

REVIEW

From Respiratory Instability to Operational Monitoring: A Deterministic Phase-Memory Framework for Chest-Mounted Wearables

Marcel Krüger^{1,*} ¹Independent Researcher, Germany

Abstract: Chest-mounted inertial sensing offers a low-cost route to continuous respiratory-pattern monitoring, but operational use is constrained by motion artifacts, baseline drift, nonspecific physiological interpretation, and the risk of treating an unvalidated deviation score as a diagnostic or readiness decision. This structured narrative review and methodological framework builds on a published proof-of-concept deterministic phase-memory operator for chest-mounted smartphone inertial measurement unit (IMU) signals. The present manuscript adds a translational validation architecture focused on high-exertion terrestrial training and occupational or first-responder heat-and-PPE tasks. It specifies signal-quality and artifact-rejection gates, a guarded dual-timescale baseline that limits absorption of progressive deterioration, pre-specified verification and validation endpoints, reference instrumentation, and governance requirements for privacy, autonomy, human oversight, and prohibition of unsupported automated command or employment decisions. Phase-memory divergence is retained as an auditable deviation observable rather than a disease classifier. A valid alert requires acceptable signal quality, no artifact-gate condition, threshold exceedance, and temporal persistence; otherwise, the required output is indeterminate/low quality. Real-time human-in-the-loop alerts and post-session trend analysis remain separate outputs. Proposed quantitative criteria are engineering targets for future preregistered studies, not achieved performance claims. The framework therefore defines a falsifiable pathway toward fit-for-purpose evaluation without claiming clinical validation, military qualification, spaceflight readiness, or deployment-level safety.

Keywords: chest-mounted IMU, respiratory monitoring, phase-memory divergence, deterministic signal processing, wearable validation

1. Introduction

Wearable respiratory and vital-parameter monitoring has become increasingly relevant across clinical, ambulatory, athletic, occupational, military, and spaceflight contexts [1–3]. Contemporary devices can acquire motion, respiratory, cardiovascular, thermal, and activity-related signals, but the translation from raw measurements to reliable decision support remains difficult when the processing chain is opaque, weakly validated, or strongly dependent on trained models [4, 5].

In this manuscript, operational vital monitoring means continuous or repeated acquisition and interpretation of physiological and kinematic signals to support human awareness during physically demanding tasks. It does not mean autonomous diagnosis, fitness-for-duty determination, command authorization, employment evaluation, or replacement of clinical assessment. Two outputs are distinguished throughout: (1) a real-time, human-in-the-loop deviation alert and (2) post-session trend analysis. These outputs require different latency, persistence, storage, and governance rules.

The present paper asks a deliberately limited question: how can a published deterministic respiratory phase-memory operator

be embedded in a reproducible, artifact-aware, ethically governed validation framework without implying that the broader system has already been validated? The proof-of-concept study by Krüger and Feeney Jr. [6] provides the signal-level methodological basis for the present work.

The revised scope is intentionally narrower than a general soldier-athlete-astronaut-firefighter platform. The core validation roadmap addresses only two priority contexts: (i) high-exertion terrestrial training, where reference respiratory instrumentation and controlled workload are feasible, and (ii) occupational or first-responder heat-and-personal protective equipment (PPE) tasks, where motion, thermal strain, and constrained equipment create realistic operational challenges. Military deployment and spaceflight are treated only as future transfer contexts requiring independent evidence and governance.

Novel contribution of the present paper. The earlier publication established the respiratory operator and a restricted proof of concept. The present paper does not claim that the operator is new. Its novel contribution is the integration of seven elements that were not established by the prior experiment: a structured evidence map, explicit signal-quality and rejection states, a motion-artifact taxonomy with deterministic gates, guarded adaptive baseline logic, measurable go/no-go validation endpoints, a reference-instrumented prospective protocol, and a

*Corresponding author: Marcel Krüger, Independent Researcher, Germany. Email: marcelkrueger092@gmail.com

governance model separating physiological monitoring from coercive or unsupported decision-making. These are architecture and validation contributions, not new clinical performance results.

2. Review Design and Evidence Synthesis

2.1. Review identity, search sources, and reproducibility

This article is a structured narrative review and methodological framework. It is not presented as a systematic review, diagnostic-accuracy meta-analysis, or complete inventory of all wearable respiratory technologies. The review objective was to identify evidence needed to define a falsifiable validation pathway for a deterministic chest-mounted IMU deviation operator. Searches covered database inception through July 14, 2026, and were executed in *PubMed/MEDLINE*, *IEEE Xplore*, Crossref-indexed publisher records, and official regulatory or standards sources. Backward citation screening was applied to directly relevant technical-validation studies and methodological frameworks.

The exact concept blocks used for the final update are reported in Table 1. Terms within a block were combined with OR and blocks were combined with AND. Spelling variants (artifact/artifactual), singular/plural forms, and title-field searches were used where supported. Search execution, eligibility assessment, and source mapping were performed manually. The review was not preregistered and did not use duplicate independent screening; therefore, no Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) flow diagram or pooled effect estimate is claimed. This limitation is stated explicitly rather than presenting the review as more comprehensive than the process supports.

2.2. Eligibility, extraction, and synthesis rules

Eligible sources were English-language peer-reviewed empirical studies, reviews, methodological frameworks, or authoritative guidance documents addressing at least one of the following: thoracic or smartphone respiratory sensing, static-versus-dynamic validation, signal quality or motion management, biometric-monitoring verification/validation, occupational heat or PPE

monitoring, usability, privacy, or workplace governance. Sources were excluded from the core synthesis when they concerned only non-wearable imaging, invasive monitoring, or algorithmic performance without sufficient information to identify sensor placement, reference measurement, validation condition, or intended use.

For each included empirical source, the following fields were extracted when available: design, sample or data source, sensor configuration, reference method, activity/context, endpoint, principal contribution, and limitation relevant to the present framework. Heterogeneity in sensor placement, reference instrumentation, activity, endpoint definition, and population precluded quantitative pooling. Methodological weight was judged using transparency of configuration, independence of the reference standard, participant-level separation, dynamic-condition testing, pre-specification, uncertainty reporting, and separation of exploratory from confirmatory analysis.

Table 2 is the requested study-level evidence map. Each row represents an individual empirical publication rather than an abstract evidence domain.

To complement the evidence map with a real-world hardware example, Figure 1 shows the actual three-IMU wearable configuration reported by Angelucci and Aliverti [7], including thoracic, abdominal, and lower-back sensor placement.

The empirical map is complemented by reviews and fit-for-purpose frameworks. Goldsack et al. [8] separate verification, analytical validation, and clinical validation, while Bakker et al. [9] extend the framework with usability and scalability. Respiratory reviews identify motion, coupling, and activity dependence as persistent constraints [2, 3, 10, 11]. Occupational reviews support context-specific heat/PPE evaluation but do not validate the present operator [12, 13].

For continuity with the broader synthesis, Table 3 retains the evidence domains used to organize claims. Unlike Table 2, it is a framework map rather than a study map.

3. Evidence Boundary and Prior Proof of Concept

The evidence boundary must be explicit because architecture-level plausibility is not operational validation. Table 4 is introduced here to separate the published respiratory evidence from the new framework claims.

Table 1
Reproducible search log for the structured narrative review update

Source	Concept block A	Concept block B/C	Purpose
<i>PubMed/MEDLINE</i>	Wearable OR smartphone OR inertial OR IMU	Respiratory rate OR breathing OR respiratory monitoring; validation OR motion OR dynamic	Sensor and respiratory-validation evidence
<i>IEEE Xplore</i>	Chest OR thoracic; IMU OR inertial	Respiration OR breathing; estimation OR monitoring OR validation	Engineering implementations and dynamic-condition methods
Crossref/publisher records	Wearable respiratory monitoring; smartphone respiratory kinematics	Motion artifact; digital biomarker validation; signal quality	Full-text and citation-chain completion
Occupational/heat search	Wearable OR physiological monitoring	Heat stress OR PPE; worker OR occupational OR firefighter	Heat/PPE feasibility and operational context
Governance search	Wearable OR physiological data	Workplace OR occupational OR firefighter; privacy OR ethics OR surveillance	Data governance, autonomy, and surveillance evidence
Official guidance	Digital health technology; biometric monitoring	Verification; analytical validation; clinical validation; risk management	Fit-for-purpose development and regulatory terminology

Table 2
Study-level evidence map for the empirical sources most directly informing the framework

Study	Design/sample	Sensor/reference	Conditions/endpoints	Relevance boundary
Krüger and Feeney [6]	Proof of concept using five public Beth Israel Deaconess Medical Center (BIDMC) respiratory records with controlled semi-synthetic perturbations	Deterministic phase-memory pipeline; RMS and FFT comparators	Frequency drift, amplitude-suppression pause, stable controls, limited motion-gate stress	Establishes reproducible algorithmic feasibility only; no prospective human cohort, clinical label, or operational ground truth
Vignali et al. [14]	Feasibility study in 77 healthy volunteers	Smartphone accelerometer and gyroscope in five prescribed body positions	Respiratory kinematics, component/position differences, RR estimation, usability	Supports position-sensitive smartphone acquisition; lacks direct clinical-reference validation and unsupervised field testing
Valentine et al. [15]	Technical validation: 15 participants against an FDA-cleared reference and 165 participants against a study reference	Smartphone movement sensors; simultaneous reference measurement	Respiratory-rate agreement and usability in predominantly normal-rate, self-operated measurements	Supports smartphone motion sensing for RR; does not establish high-motion or deterioration detection
Angelucci and Aliverti [7]	Experimental study in 20 healthy adults	Three IMUs on thorax, abdomen, and lower back; previously validated static algorithm and activity recognition	Seven static postures and walking, running, stairs, and cycling; RR and activity recognition	Demonstrates dynamic feasibility and exposes motion-dependent failure; dynamic RR in this paper was not fully validated against an independent gold standard
Wibowo et al. [16]	Randomized crossover laboratory study in 18 participants	In-ear temperature/heart rate sensor and skin-temperature sensor	22/27 °C, with/without PPE, standardized healthcare-worker tasks	Supports controlled heat/PPE instrumentation and reveals placement-related data-quality issues; it is not respiratory-operator validation
Topazian et al. [17]	Qualitative study: four focus groups and 35 interviews, 65 fire-service participants/leaders	Interview and focus-group data; no physiological sensor performance test	Ownership, access, secondary use, trust, privacy, and organizational power	Grounds governance requirements; does not provide analytical performance evidence

The prior proof-of-concept study by Krüger and Feeney Jr. [6] provides supportive evidence for algorithmic feasibility only. Its published results are reproduced in Table 5 so that readers can see exactly what has and has not already been demonstrated. No values in this table are new results from the present manuscript.

4. Signal Model and Terminology

4.1. Chest-mounted IMU channel

Chest-mounted sensing has a direct mechanical rationale: thoracic expansion and contraction produce low-frequency motion components that can be captured by accelerometers and gyroscopes when the device is mechanically coupled to the anterior chest wall. Prior work includes contact-based respiratory techniques [10], smartphone respiratory kinematics [14], smartphone movement-sensor validation [15], and IMU respiratory estimation in static and dynamic conditions [7]. Recent reviews

continue to identify motion and unstable sensor coupling as central limitations for wearable respiratory assessment [11].

A respiration-sensitive scalar channel can be represented as

$$x(t) = \mathbf{a}(t) \cdot \hat{\mathbf{u}}(t), \quad (1)$$

where $\mathbf{a}(t) \in \mathbb{R}^3$ is accelerometer data and $\hat{\mathbf{u}}(t)$ is a gravity-aligned or principal-motion axis. This is a kinematic measurement channel, not a mechanistic lung model.

4.2. Phase-memory divergence

For a band-limited respiratory channel,

$$z(t) = x(t) + iH[x(t)] = A(t)e^{i\theta(t)}, \quad (2)$$

where H is the Hilbert transform. The instantaneous phase velocity is

$$\omega(t) = \frac{d\theta(t)}{dt}, \quad (3)$$

Figure 1

Actual IMU-based wearable system for respiratory-rate estimation: the three sensor units (top), thoracic and abdominal units worn on the anterior trunk (bottom left), and the reference unit worn on the lower back (bottom right)



and the short-term phase memory over horizon T_m is

$$\bar{\omega}(t) = \frac{1}{T_m} \int_{t-T_m}^t \omega(\tau) d\tau. \quad (4)$$

The phase-memory divergence is

$$\Delta\Phi(t) = |\omega(t) - \bar{\omega}(t)|. \quad (5)$$

Terminology is fixed as follows. Phase-memory divergence is the numerical quantity $\Delta\Phi(t)$. Respiratory instability is an operationally defined departure from a pre-specified respiratory reference pattern; it is not a diagnosis and must be tied to a protocol-specific reference label. A deviation detector is any rule that flags such a departure. The proposed operator is one deviation detector. An alert is a higher-level output that additionally requires signal quality, artifact rejection, threshold persistence, and policy constraints.

5. Deterministic Operational Architecture

Figure 2 is announced here before presentation. It extends the earlier high-level diagram by making signal-quality control, artifact rejection, an indeterminate state, baseline safeguards, and separate real-time versus retrospective outputs explicit.

5.1. Proposed implementation specification and parameter provenance

The values in Table 6 are not universal physiological constants. They are divided into three classes: (i) inherited reference settings retained to reproduce the published proof of concept, (ii) physiology-bounded settings expressed in interpretable units, and (iii) engineering safeguards that require prospective calibration. The primary configuration must be frozen before confirmatory data are opened. Sensitivity analyses are declared in advance and may characterize robustness, but they may not be used to replace a failed primary analysis.

Operationalization sequence. Parameter handling follows three stages. Stage A verifies software, sampling, filter response, and deterministic replay using synthetic and public test vectors. Stage B uses a participant-disjoint pilot cohort to select only the explicitly non-inherited engineering safeguards. Stage C freezes code, configuration, reference-onset rules, and analysis before an independent confirmatory cohort is opened. No participant-specific threshold optimization is permitted in Stage C.

Public reproducibility resource. The public GitHub repository contains the public reference implementation and reproducibility material for the earlier proof of concept, including a one-command reproduction script, dependency files, tests, validation scripts, and documented reference parameters. The repository is evidence of computational reproducibility for that implementation; it is not clinical or operational validation of the expanded architecture proposed here. A prospective study must additionally archive the exact release/commit, configuration manifest, reference labels, exclusion log, and analysis outputs used for the confirmatory report.

5.2. Signal-quality index and mathematically ordered three-state logic

A composite signal-quality index may be defined from independently interpretable components:

$$Q_{sig}(t) = w_1 S_{resp}(t) + w_2 C_{axis}(t) + w_3 A_{valid}(t) + w_4 D_{complete}(t), \quad (6)$$

where S_{resp} is spectral concentration in the declared respiration band, C_{axis} is cross-axis consistency, A_{valid} rejects saturation/floor conditions, and $D_{complete}$ is data completeness. Weights and cut-offs are locked before confirmatory testing. The quality score does not describe physiology; it determines whether physiological interpretation is admissible.

Define the admissibility gate

$$G(t) = \mathbf{1}[Q_{sig}(t) \geq q_{min}] \mathbf{1}[M_{art}(t) < m_{max}] \mathbf{1}[C_{cpl}(t) \geq c_{min}] \mathbf{1}[D_{complete}(t) \geq d_{min}] \quad (7)$$

and the deviation condition

$$E(t) = \mathbf{1}[\Delta\Phi(t) > \alpha \hat{\sigma}_\omega] \mathbf{1}[P(t) \geq p_{min}], \quad (8)$$

where $P(t)$ is the locked persistence rule. The reported state $Y(t)$ is then

$$Y(t) = \begin{cases} \text{I} & \text{if } G(t) = 0 \text{ (indeterminate/low quality),} \\ \text{D} & \text{if } G(t) = 1 \text{ and } E(t) = 1 \text{ (valid deviation),} \\ \text{N} & \text{if } G(t) = 1 \text{ and } E(t) = 0 \text{ (valid normal).} \end{cases} \quad (9)$$

Table 3
Evidence domains used to structure the narrative synthesis

Evidence domain	Question	Minimum acceptable evidence
Verification	Does the hardware/software measure and replay what it is specified to measure?	Clock integrity, calibration, deterministic replay, unit tests, documented filter response, version-locked parameters
Analytical validation	Does the derived respiratory quantity agree with a reference under intended conditions?	Synchronized reference instrument, agreement/error metrics, missing-data accounting, static and dynamic conditions
Operational association	Does the deviation measure relate to a defined task-relevant state?	Prospectively labeled events or states, preregistered thresholds, sensitivity/specificity or time-to-detection with uncertainty
Usability	Can intended users wear, place, understand, and respond to the system?	Placement repeatability, comfort, task interference, comprehension, alarm burden, subgroup assessment
Motion/environment	Does performance remain acceptable under activity and environmental stress?	Walking, running, load/PPE, posture, cough/speech, impact, coupling changes, heat/humidity where appropriate
Governance	Can the system be used without coercive or discriminatory data practices?	Purpose limitation, consent/authority analysis, local processing, role-based access, contestability, retention limits, audit trail

Table 4
Scope and evidence boundary of the revised manuscript

Category	Position used in this manuscript
Published evidence	Controlled proof-of-concept phase-memory operator evaluated on five BIDMC respiratory recordings with semi-synthetic perturbations and lightweight baselines [6]
New contribution	Structured review, architecture, quality gates, baseline safeguards, validation endpoints, instrumentation protocol, and governance requirements
Not new here	The core operator definition, prior BIDMC perturbation protocol, and prior implementation results
Not claimed	Clinical diagnosis, medical-device performance, military qualification, spaceflight readiness, autonomous readiness judgment, or field safety
Primary target contexts	High-exertion terrestrial training; occupational and first-responder heat-and-PPE tasks
Future transfer only	Military operational deployment and spaceflight, each requiring separate evidence, human-factors work, and governance approval

Table 5
Published proof-of-concept evidence anchor

Controlled condition	Phase-memory operator	Comparator	Interpretation
Frequency drift	Detection latency 0.000 ± 0.000 s	FFT peak shift 0.060 ± 0.000 s; RMS not detected	First-sample response under an injected, semi-synthetic drift; not clinical sensitivity
Amplitude-suppression pause	Detection latency 0.572 ± 0.027 s	RMS 0.000 ± 0.000 s; FFT not detected	Operator was slower than the amplitude-specific RMS baseline, demonstrating complementary rather than universal superiority
Stable control	0 false alarms	0 for RMS and FFT	Restricted stable segments only
Moderate motion stress	0.0% sustained false positive rate across 12 gated segments	Without gating, 2/12 segments showed isolated non-persistent crossings	Initial gated walking/posture stress only; cough, speech, running, impact, and loose coupling remained untested

Figure 2

Revised deterministic architecture drawn at native aspect ratio. Signal quality and artifact admissibility are evaluated before physiological deviation. Failed gates produce an indeterminate state; they cannot be reclassified as normal or deviation. Real-time alerts and retrospective trends remain separate outputs




Table 6

Proposed fixed implementation, provenance, and operational lock rules

Component	Primary setting	Provenance/rationale	Lock and sensitivity rule
Acquisition	Native 50–100 Hz; resample to 50 Hz after timebase checks	50 Hz matches the published/reference pipeline and is well above the declared respiratory band	Report jitter/dropout; repeat at native rate as a secondary verification only
Respiration band	Rest: 0.1–0.5 Hz; exertion: 0.1–1.0 Hz	Corresponds to 6–30 and 6–60 breaths/min, respectively; the wider band is protocol-specific, not automatically selected	Choose the protocol band before analysis; no per-segment band optimization
Filter	Offline zero-phase fourth-order Butterworth; streaming causal matched implementation with measured delay	Zero-phase processing supports reproducible retrospective analysis; causal processing is required for real time	Compare magnitude response and latency; report both implementations separately
Memory horizon	$M = 150$ samples (3 s at 50 Hz)	Inherited from the public proof-of-concept implementation for direct reproducibility	Primary $M = 150$; declared sensitivity at 2, 3, and 5 s without retuning the threshold
Calibration interval	120 s target; minimum 60 s usable low-motion data	Engineering safeguard to estimate robust center/scale more reliably than the earlier short proof-of-concept baseline	Recalibrate after attachment change; report failure to obtain sufficient valid baseline
Threshold multiplier	$\alpha = 2.0$ applied to robust baseline scale	Matches the public reference implementation; it is a transparent sensitivity setting, not a disease threshold	Primary $\alpha = 2.0$; secondary 2.5 and 3.0 sensitivity analyses only
Persistence	1.0 s continuously above threshold, with a predeclared 70%/1.5 s robustness analysis	Engineering trade-off intended to reject isolated phase spikes	Primary rule locked before confirmatory testing; latency includes persistence time
Signal availability	At least 80% valid samples in the decision window	Initial completeness gate that prevents interpretation dominated by missing or rejected samples	Report results at 70/80/90% as sensitivity; 80% remains primary

(Continued)

Table 6
(Continued)

Component	Primary setting	Provenance/rationale	Lock and sensitivity rule
Missing samples	No interpolation across gaps > 0.5 s	Prevents phase construction across unsupported intervals	Such windows become inde- terminate and remain in availability accounting
Baseline adaptation	Fast and slow states; update only during valid, non- alert, low-drift periods	Prevents progressive deterioration from being learned as normal	Update gains and freeze thresholds tuned only in a participant-disjoint development cohort, then frozen
Audit record	Version, dependency lock, parameter manifest, sensor orientation, gate state, and decision rationale	Required for deterministic replay and failure investigation	Archive a release/commit identifier with each analysis

The order is substantive: I has priority and prevents a failed-quality segment from being counted as either normal or deviation. A human-facing alert is

$$A(t) = \mathbf{1}[Y(t) = D], \quad (10)$$

subject to the intended-use and escalation policy. This formulation corrects the earlier binary indicator representation, which did not encode the indeterminate state as a mutually exclusive mathematical outcome.

6. Motion Artifacts: Limits and Deterministic Management

No IMU-only algorithm can guarantee separation of respiratory instability from all body motion because respiratory and non-respiratory accelerations can overlap in frequency, direction, and amplitude. The physically correct response to non-identifiability is rejection or uncertainty, not confident relabeling. The proposed artifact layer therefore combines signatures rather than relying on a single gyroscope threshold.

Table 7 is announced here before presentation. It states which artifact signatures are potentially distinguishable, which require contextual sensing, and which remain fundamentally ambiguous.

The motion index should be computed from gyro RMS acceleration jerk, out-of-band energy, and orientation change. Thresholds should be normalized to a low-motion calibration distribution rather than expressed as universal absolute values. Each gate must be evaluated separately in ablation tests because a combined score can conceal failure modes.

7. Guarded Adaptive Baseline

A baseline that updates continuously can absorb progressive fatigue, heat strain, hypoxia, dehydration, or respiratory compromise. To prevent this failure, the revised framework uses two time scales and explicit freeze conditions.

Let $b_f(t)$ be a fast contextual baseline and $b_s(t)$ a slow protected baseline. For a valid, low-motion sample y_t ,

$$b_f(t) = (1 - \lambda_f)b_f(t-1) + \lambda_f y_t \quad (11)$$

$$b_s(t) = (1 - \lambda_s)b_s(t-1) + \lambda_s y_t, \quad 0 < \lambda_s \ll \lambda_f. \quad (12)$$

The update is disabled when signal quality fails, motion/coupling gates activate, an alert is active, the cumulative slope exceeds a pre-specified bound, or the fast-slow divergence

$$D_b(t) = |b_f(t) - b_s(t)| \quad (13)$$

exceeds a robust scale threshold. A cumulative-drift monitor may use a two-sided cumulative sum (CUSUM) or an equivalent locked detector. The slow baseline must also be anchored to the session-start reference and may change only within a bounded daily/session envelope.

Figure 3 is announced before display. It shows why progressive drift is retained as a deviation instead of being incorporated automatically.

This design does not prove that deterioration will always be detected. It makes the update rule falsifiable and prevents the most obvious “moving target” failure. The baseline algorithm must be tested using slow ramps in workload, temperature, hypoxia proxy, and respiratory rate, including negative controls with benign posture and task transitions.

8. Balanced Comparison with Existing Approaches

Determinism is not inherently superior to machine learning. It provides a different trade-off. Table 8 is introduced before presentation to compare deterministic, model-based, learned, and digital biomarker approaches without caricaturing any class.

A fair comparison must use identical raw inputs, synchronized labels, fixed evaluation windows, and pre-specified thresholds or model versions. The deterministic operator should be compared not only with RMS and FFT baselines but also with respiratory-rate variability, peak-interval methods, state-space estimators, and at least one properly calibrated learning-based benchmark. Performance should be reported with uncertainty and failure cases, not only mean accuracy.

9. Priority Application Contexts

9.1. Priority context 1: high-exertion terrestrial training


This context includes controlled athletic training and military-style load-carriage training but excludes combat-readiness claims. The advantage is that workload can be prescribed and reference measures can include respiratory inductance

Table 7
Operational artifact taxonomy and required system response

Condition	Likely IMU signature	Proposed deterministic handling	Residual limitation
Walking	Periodic locomotor peaks, axis-dependent harmonics, elevated gyro energy	Estimate cadence; notch/contextualize cadence harmonics; require respiratory spectral concentration and cross-axis consistency	Cadence can overlap respiratory frequency; reject if separability is poor
Running	High broadband energy, impacts, strong periodic motion	High-motion gate; no respiratory alert unless a validated dynamic model preserves a separable respiratory component	Chest IMU alone may be insufficient at high impact
Load carriage/PPE	Slow coupling changes plus motion harmonics	Attachment/coupling checks, recalibration after equipment change, task-specific validation	Mechanical transfer function depends on garment and load
Cough	Short high-kurtosis broadband impulse with abrupt amplitude/phase reset	Impulse detector; mask a pre-specified recovery interval; optionally label cough only with an independently validated channel	A cough may itself be clinically relevant; masking must be transparent
Talking	Irregular thoracic and laryngeal motion; respiration-speech coupling	Context flag or optional microphone processed locally; otherwise, indeterminate when phase pattern is inconsistent	IMU-only discrimination is unreliable
Posture transition	Step-like gravity-axis rotation and transient baseline change	Orientation-change detector; freeze baseline; require post-transition restabilization	Slow transitions can resemble drift
Impact	Very high jerk and acceleration saturation	Hard gate and recovery interval; log the event	Impact may displace the sensor and require reattachment
Loose coupling	Amplitude attenuation, axis decorrelation, intermittent contact	Coupling index from amplitude floor, axis coherence, and orientation variance; reject	Cannot be repaired algorithmically without restoring attachment

Figure 3

Guarded dual-timescale baseline logic. Progressive change is not absorbed when quality, artifact, alert, slope, or fast-slow divergence criteria fail



plethysmography, spirometry or metabolic-cart ventilation, electrocardiography (ECG)/heart rate, and task annotations. Static and dynamic respiratory monitoring differ substantially, and motion-rich evaluation is essential [7, 11].

The intended output is a deviation and recovery trend, not a medical diagnosis. Required labels include exercise stage, workload, speech, cough, posture, perceived exertion, and attachment changes. A successful study must demonstrate whether the operator adds information beyond respiratory rate and workload alone.

9.2. Priority context 2: occupational and first-responder heat-and-PPE tasks

Heat-exposed workers and first responders combine thermal strain, exertion, PPE, equipment motion, and governance concerns. Reviews emphasize that wearable occupational technologies often remain limited by short controlled evaluations and human-centered implementation gaps [12, 13]. Controlled PPE studies show that wearable temperature and cardiovascular channels can

Table 8
Balanced comparison of monitoring approaches

Approach	Strengths	Limitations	Required evidence
Simple threshold/envelope	Transparent, low compute, easy deployment	Poor specificity; threshold drift; amplitude dependence	Reference comparison, artifact challenge, threshold sensitivity
Phase-memory operator	Explicit temporal deviation, deterministic replay, low parameter count	Sensitive to preprocessing, phase quality, memory horizon, baseline and nonspecific physiology	Locked implementation, signal-quality gates, ablation, dynamic-condition validation
Model-based physiology	Can encode mechanisms and constraints	Model mismatch, parameter identifiability, calibration burden	Parameter recovery, external validation, uncertainty propagation
Machine learning	Can learn complex multimodal patterns and interactions	Dataset shift, opacity, hidden confounding, calibration and subgroup risk	External validation, calibration, subgroup analysis, explainability, locked model and data provenance
Explainable AI	May improve feature attribution and auditability	Explanations may be unstable or non-causal; underlying model may remain data-dependent	Explanation fidelity/stability plus conventional validation [4]
Digital biomarker pipeline	Explicit fit-for-purpose development from sensor to end point	Requires substantial verification, validation, usability and clinical evidence	V3/V3+ evidence, reference standards, intended-use definition [5, 8, 9]

capture thermal responses, but each device and context requires validation [16].

The respiratory operator should be one interpretable channel within a multimodal system, not a standalone heat-strain detector. The system must explicitly distinguish safety support from productivity surveillance. Fire-service research has identified data ownership, consent, secondary use, and trust as central ethical issues for occupational wearable data [17].

9.3. Future transfer contexts: military deployment and spaceflight

Military health/readiness frameworks and advanced-soldier reviews support integrated monitoring research, but they do not validate the present operator [18, 19]. Spaceflight adds microgravity, altered equipment interfaces, radiation, mission constraints, and medical-operations requirements. Therefore, these contexts are retained only as future transfer hypotheses after the two terrestrial validation pathways are completed.

10. Tailored Scientific Instrumentation and Prospective Supportive Experiment

Reviewer concerns about empirical support cannot be answered by architecture prose alone. Figure 4 is announced here before presentation and specifies a reference-instrumented rig that can generate empirical evidence while preserving the chest-mounted concept.

10.1. Protocol sequence

A first prospective study should use a repeated-measures design with a priori power analysis based on the primary endpoint

and should report all exclusions. A minimum protocol sequence is:

- 1) 5 min seated quiet breathing and 3 min paced breathing;
- 2) Standing and repeated posture transitions;
- 3) Treadmill walking at two speeds;
- 4) Running or high-intensity stepping appropriate to the cohort;
- 5) Load carriage or PPE task at two workload levels;
- 6) Scripted speech, voluntary coughs, arm-dominant tasks, and controlled impacts;
- 7) Deliberately loosened attachment under safe supervision;
- 8) Recovery period; and
- 9) For heat/PPE studies, a separately approved environmental protocol with temperature, humidity, stop criteria, and medical oversight.

The reference respiratory channel should be selected for the task: respiratory inductance plethysmography for ambulatory thoracic motion, airflow or capnography where feasible, and metabolic-cart ventilation for exercise validation. Time synchronization error must be quantified. The system should be evaluated both with and without each quality/artifact gate to show what the gates actually contribute.

10.2. Prospective target conditions and ground-truth hierarchy

The operator must not define the event that it is evaluated against. Target conditions are therefore operational protocol labels derived from independent references. Table 9 defines the initial target and non-target classes. Exact magnitude and duration cutoffs must be fixed in the protocol after pilot feasibility work and before confirmatory data are opened.

Figure 4

Redrawn reference-instrumented validation rig at native aspect ratio. The test IMU never defines its own ground truth. Target labels derive from synchronized reference channels and predeclared protocol events, followed by a frozen analysis package




Table 9

Prospective target conditions and independent ground truth

Class	Protocol definition	Ground-truth source/onset rule	Evaluation role
Acute rate transition	Preplanned paced-breathing step or workload-associated respiratory-rate change sustained for the predeclared duration	Synchronized belt/flow/metabolic-cart waveform; onset from the locked reference algorithm, not the IMU score	Primary event-detection and latency target
Brief respiratory interruption	Supervised voluntary breath hold or protocol-safe amplitude suppression; no induced pathological event	Reference airflow or belt-amplitude criterion plus task marker	Tests interruption response; safety and stop criteria required
Progressive drift and recovery	Monotonic protocol-driven change across a workload block followed by recovery	Reference RR/ventilation trajectory and workload timeline	Tests whether guarded adaptation preserves a growing deviation
Stable respiratory block	Reference signal remains within the predeclared stable band, and no target event is annotated	Reference waveform and protocol state	Denominator for false alarms per hour and specificity

(Continued)

Table 9
(Continued)

Class	Protocol definition	Ground-truth source/onset rule	Evaluation role
Artifact-only block	Speech, cough, posture change, arm task, impact, or loosened coupling without a reference-defined target transition	Task marker plus valid-reference respiratory channel	Negative control; expected output is normal or indeterminate, not physiological deviation
Reference failure	Reference dropout, desynchronization, or uninterpretable waveform	Predeclared reference-quality gate	Excluded from physiological accuracy; retained in availability and failure reporting
Heat/PPE context	Controlled environmental/PPE exposure with independent temperature, heart rate (HR), workload, and stop criteria	Contextual measurements only; not a respiratory diagnosis label	Stratification/context analysis, not a standalone target disease state

Ground-truth construction follows a hierarchy: (1) quantitative reference waveform and locked onset algorithm, (2) synchronized protocol marker, and (3) blinded review by two assessors when the reference algorithm is ambiguous, with a third assessor resolving disagreement. The test-IMU output is hidden during adjudication. Inter-rater agreement and all unresolved/excluded events are reported. No label such as “fatigue,” “heat strain,” “hypoxia,” or “respiratory compromise” is permitted unless an independent, condition-specific clinical or physiological definition is prospectively supplied.

10.3. Primary hypotheses

The study should preregister the following directional hypotheses: (H1) the phase-memory operator detects injected or independently labeled respiratory-pattern transitions faster than a spectral-only baseline under low-motion conditions; (H2) the full quality-gated system reduces sustained false alarms during artifact tasks relative to the ungated operator; (H3) guarded baseline adaptation retains sensitivity to progressive ramps better than unrestricted adaptation; and (H4) the operator adds incremental information beyond respiratory rate alone. Failure to support these hypotheses would constrain or falsify the proposed implementation class.

11. Measurable Validation Endpoints and Statistical Validation Plan

The targets below are proposed engineering gates, not achieved performance claims. The prospective protocol must identify one primary endpoint for power calculation and treat the remaining endpoints as ordered secondary or exploratory outcomes.

11.1. Analysis populations and parameter lock

The development and confirmatory samples must be participant-disjoint. All parameter choices, reference-onset rules, comparator versions, exclusion rules, and output definitions are frozen after the development stage. Repeated windows from one participant are not treated as independent participants. The confirmatory analysis is performed once under the locked primary specification; sensitivity analyses are reported separately and cannot replace it.

11.2. Endpoints and initial go/no-go targets

Table 10 summarizes the proposed validation layers, associated metrics, and pre-specified go/no-go targets.

11.3. Uncertainty, repeated measures, and comparator tests

RR agreement is reported using mean absolute error/root mean square error (MAE/RMSE) and Bland–Altman analysis with a repeated-measures or mixed-effects treatment of participant clustering; correlation alone is insufficient. Event-level sensitivity is the proportion of independently labeled target events detected within the predeclared horizon. False alarms are counted per valid target-negative hour. Latency distributions retain missed events as right-censored at the analysis horizon or report them explicitly as misses; reporting only detected events is prohibited.

Confidence intervals are participant-clustered, obtained through a suitable hierarchical bootstrap or mixed-effects model. Comparisons with root mean square (RMS), fast Fourier transform (FFT), respiratory rate (RR) variability, peak-interval, state-space, and at least one locked learning-based benchmark use paired participant-level contrasts on identical raw inputs and labels. The primary hypothesis and metric are specified before data collection. Secondary multiplicity is controlled with a declared method such as Holm correction, while explicitly labeled exploratory analyses are reported without confirmatory language.

11.4. Missing data, indeterminate intervals, and denominators

No gap longer than the declared limit is silently interpolated. Physiological performance is reported on valid-reference, admissible intervals; system availability is reported on all scheduled time. Two denominator views are mandatory: (i) performance conditional on valid admissible time and (ii) end-to-end coverage including indeterminate time. Reference failures, test-sensor failures, quality rejections, participant withdrawals, and protocol deviations are tabulated separately. This prevents high apparent accuracy obtained by rejecting difficult segments without accounting for the loss of coverage.

Table 10
Proposed validation layers, metrics, and pre-specified go/no-go targets

Layer	Metric	Initial target	Failure interpretation
Verification	Deterministic replay	Identical outputs for identical input, release, and configuration within the declared numeric tolerance	Software/version or platform equivalence failure
Verification	End-to-end processing delay	< 100 ms excluding the declared persistence interval on target hardware	Real-time configuration unsupported
Availability	Valid/indeterminate time	$\geq 80\%$ valid time per protocol block; indeterminate fraction reported by activity	Placement, sensor, or gate configuration not fit for task
RR agreement	MAE, bias, and limits of agreement	MAE ≤ 2 breaths/min at low motion and ≤ 3 during walking, with activity-specific limits	Derived respiratory channel lacks acceptable agreement
Event detection	Event-level sensitivity and target-negative specificity	Both ≥ 0.85 with 95% confidence intervals under the frozen event definition	Threshold/operator insufficiently discriminative in that context
Alarm burden	Sustained false alarms per valid hour	≤ 1 per valid hour in target-negative blocks; rejected/indeterminate spikes reported separately	Operational alarm burden excessive
Latency	Time from reference onset to sustained deviation	Median ≤ 5 s, including persistence; censored misses retained	Detection too slow or frequently absent
Motion robustness	Within-participant degradation	Sensitivity decrease < 10 percentage points from low motion to walking; running reported separately	Dynamic use unsupported
Reliability	Test-retest ICC and within-subject error	ICC ≥ 0.75 for predeclared stable session summaries	Poor repeatability
Incremental value	Locked comparison with RR/workload and simpler baselines	Improvement with confidence interval and calibration/error analysis	Phase-memory score adds no demonstrated value
Subgroups	Performance and indeterminate fraction	Report by sex, body size/attachment group, fitness/task group, and relevant interface factors	Generalizability or attachment equity remains unknown

11.5. Sample-size principle

The current framework does not assert a universal sample size. The prospective protocol must calculate sample size from one primary endpoint using the participant as the independent sampling unit, the expected within-participant event count, clustering/intraclass correlation, anticipated indeterminate fraction, and attrition. A study powered on windows while ignoring participant clustering would be methodologically invalid.

12. Regulatory and Quality-Development Pathway

The framework is not currently a medical device, but any future diagnostic or clinical-monitoring claim would require a fit-for-purpose evidence pathway. The V3 framework separates verification, analytical validation, and clinical validation, while V3+ adds usability and scalability [8, 9]. FDA guidance on digital health technologies for remote data acquisition emphasizes fit-for-purpose selection, verification/validation, usability, data handling, and risk management [20]. ISO 14155 defines good clinical practice for medical-device investigations, and ISO 14971 provides a risk-management framework [21, 22].

These sources do not imply that the present system is regulated as a medical device today. They define the evidence discipline required if the intended use changes from research or informational monitoring to medical decision support.

13. Governance, Privacy, Surveillance, and Responsible Deployment

Physiological monitoring in workplaces, fire services, military structures, and elite sport creates risks that cannot be solved by informed consent language alone. Occupational wearable data may influence perceived readiness, performance judgment, access to duties, insurance, discipline, or employment. Research in fire-service settings emphasizes concerns about ownership, secondary use, trust, and organizational power [17]. Broader workplace-surveillance research associates perceived surveillance with poorer well-being outcomes [23].

Table 11 is announced here before presentation. It converts ethical concerns into enforceable design requirements.

A research prototype should therefore keep all decision authority outside the algorithm. Even a technically accurate

Table 11
Minimum governance requirements for operational physiological monitoring

Principle	Required control	Prohibited practice
Purpose limitation	Written intended use, defined alert recipients, separate research and operational policies	Reusing physiological data for discipline, productivity scoring, or unrelated profiling
Data minimization	Prefer on-device processing; transmit alerts/features only when raw data are unnecessary	Continuous raw-data upload by default
Human oversight	Alerts support trained human judgment; documented escalation and override	Autonomous diagnosis, command decision, or employment action
Contestability	Wearer can review, correct context, and challenge adverse interpretation	Treating sensor output as unquestionable ground truth
Role separation	Clinical/safety reviewers separated from line management where feasible	Unrestricted supervisor access to raw physiological histories
Consent and coercion	Assess legal/organizational authority and power imbalance; provide alternatives where possible	Calling participation “voluntary” when refusal has adverse consequences
Retention	Predefined shortest retention; automatic deletion and access logging	Indefinite retention or undisclosed secondary datasets
Security	Encryption, signed software, role-based access, breach response, provenance	Unauthenticated dashboards or shared credentials
Fairness	Subgroup performance and attachment usability evaluated; accommodations provided	Penalizing users for device failure, body morphology, disability, pregnancy, or health status

signal can be used unethically. Governance is part of system validity, not a post hoc policy appendix.

14. Limitations and Falsification Criteria

The principal limitations are no new human-subject data in this manuscript, dependence on the prior five-record proof of concept, nonspecificity of phase deviation, possible phase-estimation failure at low amplitude, sensitivity to filter and baseline choices, incomplete motion separability, and proposed numerical targets that remain to be empirically justified.

The framework should be rejected or narrowed for a target context if any of the following occur in a properly powered, preregistered evaluation:

- 1) Respiratory agreement fails under the intended activity despite acceptable attachment;
- 2) Quality gates reject so much data that monitoring is not practically available;
- 3) Motion artifacts produce an unacceptable sustained false-alarm burden;
- 4) Guarded adaptation still absorbs progressive deterioration ramps;
- 5) Phase-memory divergence adds no incremental value over respiratory rate, workload, or simpler baselines;
- 6) Subgroup performance is materially unequal and cannot be mitigated; or
- 7) Users cannot understand or safely respond to the alert state.

A null or negative result would be scientifically informative. The operator is not protected from falsification by redefining thresholds after test-set inspection.

15. Conclusion

This revised manuscript positions the work as a structured narrative review and methodological framework rather than a new empirical validation study. The published chest-mounted smartphone IMU proof of concept by Krüger and Feeney Jr. [6] established a deterministic respiratory phase-memory operator under restricted semi-synthetic conditions. The present contribution is the explicit translational architecture needed to test whether that operator can become fit for a defined operational purpose.

The revised framework narrows primary validation to high-exertion terrestrial training and occupational and first-responder heat-and-PPE tasks; treats military deployment and spaceflight only as future transfer contexts; specifies signal-quality, artifact, coupling, persistence, and indeterminate states; introduces guarded baseline adaptation; defines measurable go/no-go endpoints; proposes tailored reference instrumentation; compares deterministic and learned approaches fairly; and makes privacy, autonomy, contestability, and human oversight part of the technical requirements.

Nothing in this paper establishes clinical sensitivity, diagnostic utility, readiness assessment, or deployment safety. The next publication-level advance must be empirical: a preregistered, reference-instrumented, motion-rich prospective study with locked parameters, uncertainty estimates, failure analysis, and transparent negative results. Until then, the system remains a candidate deviation-monitoring framework, not a qualified operational device.

AI Statement

No generative AI system is part of the proposed monitoring method, signal-processing pipeline, data generation, or

decision logic. Language and document-formatting assistance was used during manuscript revision. The author reviewed and approved all scientific statements, equations, references, and proposed validation criteria and retain full responsibility for the manuscript.

Acknowledgement

The author gratefully acknowledges Don Michael Feeney Jr. for substantial contributions during the development of this work, including software and implementation context, methodology, technical review, validation-roadmap input, and writing – review and editing.

Ethical Statement

This study does not contain any studies with human or animal subjects performed by the author.

Conflicts of Interest

The author declares that he has no conflicts of interest to this work.

Data Availability Statement

No new empirical dataset was generated or analyzed for this framework article. The public repository at <https://github.com/dfeen87/Smartphone-Based-Chest-Monitoring> contains the reference implementation and reproducibility materials for the previously published respiratory proof-of-concept article, including documented parameters, tests, validation scripts, dependency files, and a one-command reproduction workflow. The repository does not constitute clinical or operational validation of the broader architecture, baseline safeguards, proposed go/no-go targets, or future application contexts described here. Any prospective confirmatory analysis must cite and archive the exact release/commit and configuration manifest used.

Author Contribution Statement

Marcel Krüger: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Resources, Data curation, Writing – original draft, Writing – review & editing, Visualization, Supervision, Project administration.

References

- [1] Dunn, J., Coravos, A., Fanarjian, M., Ginsburg, G. S., & Steinhubl, S. R. (2024). Remote digital health technologies for improving the care of people with respiratory disorders. *The Lancet Digital Health*, 6(4), e291–e298. [https://doi.org/10.1016/S2589-7500\(23\)00248-0](https://doi.org/10.1016/S2589-7500(23)00248-0)
- [2] Vitazkova, D., Foltan, E., Kosnacova, H., Micjan, M., Donoval, M., Kuzma, A., ..., & Vavrinsky, E. (2024). Advances in respiratory monitoring: A comprehensive review of wearable and remote technologies. *Biosensors*, 14(2), 90. <https://doi.org/10.3390/bios14020090>
- [3] Hussain, T., Ullah, S., Fernández-García, R., & Gil, I. (2023). Wearable sensors for respiration monitoring: A review. *Sensors*, 23(17), 7518. <https://doi.org/10.3390/s23177518>
- [4] Abdelaal, Y., Aupetit, M., Baggag, A., & Al-Thani, D. (2024). Exploring the applications of explainability in wearable data analytics: Systematic literature review. *Journal of Medical Internet Research*, 26, e53863. <https://doi.org/10.2196/53863>
- [5] Daniore, P., Nittas, V., Haag, C., Bernard, J., Gonzenbach, R., & von Wyl, V. (2024). From wearable sensor data to digital biomarker development: Ten lessons learned and a framework proposal. *NPJ Digital Medicine*, 7(1), 161. <https://doi.org/10.1038/s41746-024-01151-3>
- [6] Krüger, M., & Feeney Jr, D. M. (2026). A proof-of-concept deterministic phase-memory operator for respiratory instability detection using chest-mounted smartphone IMU signals. *Smart Wearable Technology*. <https://doi.org/10.47852/bonviewSWT62029484>
- [7] Angelucci, A., & Aliverti, A. (2023). An IMU-based wearable system for respiratory rate estimation in static and dynamic conditions: Angelucci and Aliverti. *Cardiovascular Engineering and Technology*, 14(3), 351–363. <https://doi.org/10.1007/s13239-023-00657-3>
- [8] Goldsack, J. C., Coravos, A., Bakker, J. P., Bent, B., Dowling, A. V., Fitzer-Attas, C., ..., & Dunn, J. (2020). Verification, analytical validation, and clinical validation (V3): The foundation of determining fit-for-purpose for Biometric Monitoring Technologies (BioMeTs). *NPJ Digital Medicine*, 3(1), 55. <https://doi.org/10.1038/s41746-020-0260-4>
- [9] Bakker, J. P., Barge, R., Centra, J., Cobb, B., Cota, C., Guo, C. C., ..., & Goldsack, J. C. (2025). V3+ extends the V3 framework to ensure user-centricity and scalability of sensor-based digital health technologies. *NPJ Digital Medicine*, 8(1), 51. <https://doi.org/10.1038/s41746-024-01322-2>
- [10] Massaroni, C., Nicolò, A., Lo Presti, D., Sacchetti, M., Silvestri, S., & Schena, E. (2019). Contact-based methods for measuring respiratory rate. *Sensors*, 19(4), 908. <https://doi.org/10.3390/s19040908>
- [11] Pecik, M., Vavrinsky, E., Vitazkova, D., Kosnacova, H., Nevrela, J., & Foltan, E. (2026). Respiratory monitoring in motion: An overview of wearable methods and algorithmic approaches for reliable assessment. *Biosensors*, 16(6), 306. <https://doi.org/10.3390/bios16060306>
- [12] Cannady, R., Warner, C., Yoder, A., Miller, J., Crosby, K., Elswick, D., & Kintziger, K. W. (2024). The implications of real-time and wearable technology use for occupational heat stress: A scoping review. *Safety Science*, 177, 106600. <https://doi.org/10.1016/j.ssci.2024.106600>
- [13] Mendes, D., Terradillos, E., Navas, H. V., Costa, O., Matias, J., & Soares, V. (2026). Wearable technologies in occupational safety and health: A systematic review and a human-centered implementation model. *Applied Sciences*, 16(10), 4715. <https://doi.org/10.3390/app16104715>
- [14] Vignali, E., Gasparotti, E., Miglior, L., Gervasi, V., Simone, L., Haxhiademi, D., ..., & Celi, S. (2024). A new smartphone-based method for remote health monitoring: Assessment of respiratory kinematics. *Electronics*, 13(6), 1132. <https://doi.org/10.3390/electronics13061132>
- [15] Valentine, S., Cunningham, A. C., Klasmer, B., Dabbah, M., Balabanovic, M., Aral, M., ..., & Plans, D. (2022). Smartphone movement sensors for the remote monitoring of respiratory rates: Technical validation. *Digital Health*, 8, 20552076221089090. <https://doi.org/10.1177/20552076221089090>
- [16] Wibowo, R., Do, V., Quartucci, C., Koller, D., Daanen, H. A., Nowak, D., ..., & Rakete, S. (2024). Effects of heat and personal protective equipment on thermal strain in healthcare workers: Part B—application of wearable sensors

- to observe heat strain among healthcare workers under controlled conditions. *International Archives of Occupational and Environmental Health*, 97(1), 35–43. <https://doi.org/10.1007/s00420-023-02022-2>
- [17] Topazian, R. J., Wee, A., Ali, J., Frattaroli, S., Locke, P., Taylor, J. A., & Crifasi, C. K. (2025). Ethical use of wearable device data in occupational settings: Perspectives from the fire service. *Journal of Occupational and Environmental Medicine*, 67(9), e621–e629. <https://doi.org/10.1097/JOM.0000000000003446>
- [18] de Vries, H. J., van der Wal, S. J., Delahaij, R., Venrooij, W., & Kamphuis, W. (2025). Real-time monitoring of military health and readiness: A perspective on future research. *Frontiers in Digital Health*, 7, 1542140. <http://dx.doi.org/10.3389/fdgth.2025.1542140>
- [19] Sipos, D., Vészi, K., Bogár, B., Pető, D., Füredi, G., Betlehem, J., & Pandur, A. A. (2025). Physiological state monitoring in advanced soldiers: Precision health strategies for modern military operations. *Sci*, 7(4), 137. <https://doi.org/10.3390/sci7040137>
- [20] US Food and Drug Administration. (2023). *Digital health technologies for remote data acquisition in clinical investigations*. <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/digital-health-technologies-remote-data-acquisition-clinical-investigations>
- [21] International Organization for Standardization. (2026). *ISO 14155:2026, Clinical investigation of medical devices for human subjects - Good clinical practice*. <https://www.iso.org/standard/83968.html>
- [22] International Organization for Standardization. (2019). *ISO 14971:2019, Medical devices - Application of risk management to medical devices*. <https://www.iso.org/standard/72704.html>
- [23] Glavin, P., Bierman, A., & Schieman, S. (2024). Private eyes, they see your every move: Workplace surveillance and worker well-being. *Social Currents*, 11(4), 327–345. <https://doi.org/10.1177/23294965241228874>

How to Cite: Krüger, M. (2026). From Respiratory Instability to Operational Monitoring: A Deterministic Phase-Memory Framework for Chest-Mounted Wearables. <i>Smart Wearable Technology</i>. https://doi.org/10.47852/bonviewSWT620210391
