

Atum^o

Research Framework

Breathing Structure as a Continuous Physiological Signal

Abstract

1. Introduction
2. Two centuries of respiratory observation: a history of slices
3. Breathing as a convergent physiological signal
4. Breathing State: from signal to temporal object
5. Evidence review: breathing structure across 16 body systems
6. The temporal advantage of respiratory observation
7. Personalized baselines vs population thresholds
8. The dataset gap: why structural breathing data does not yet exist
9. Smartphone acoustic feasibility
10. The observation layer architecture
11. Open questions and methodological frontiers
12. Regulatory positioning

Conclusion

References

Abstract

Three transitions over the past century — the electrocardiogram displacing auscultation, ambulatory blood-pressure monitoring displacing the office cuff, continuous glucose monitoring displacing the fingerstick — each rewrote the clinical entity to which it applied. Breathing, captured today as a single integer count of cycles per minute, has not yet undergone this transition, despite being the only continuously accessible vital sign that integrates autonomic balance, chemoreflex sensing, and central neural state in real time. Respiratory rhythm originates in a discrete brainstem substrate — the preBötzinger complex — and is uniquely positioned among vital signs in being both autonomically generated and voluntarily modulable. Respiration globally entrains cortical and limbic oscillatory activity through nasal airflow, modulating gamma-band activity in frontal regions and theta-delta activity in piriform cortex, amygdala, and hippocampus.¹⁵ Breathing structure is therefore not a downstream projection of cardiovascular state but an upstream window onto integrated brainstem-autonomic-cortical regulation. That breathing has not yet been observed at this resolution reflects an enabler gap, not a physiological one. Three independent enablers have matured within the past five years to make continuous structural observation of breathing through consumer hardware feasible: smartphone MEMS microphones reaching 70–80 dBA signal-to-noise ratios with flat response across the respiratory band; analytical methods — entropy, scaling-exponent, multifractal, and recurrence-quantification analyses — adequate to characterise non-stationary respiratory time-series; and the FDA General Wellness Policy, revised 6 January 2026, clarifying the regulatory geometry for non-diagnostic physiologic observation. The present monograph synthesises evidence across 16 physiological systems and grades each finding by the strength of the supporting literature. Breathing-structure deviations are observable across temporal windows spanning seconds to years: mean nocturnal respiratory rate at or above 16 breaths per minute predicts cardiovascular and all-cause mortality independently of established risk factors over approximately ten years of follow-up,⁵³ while slow respiratory drift in the hour preceding spontaneous panic attacks, captured by continuous ambulatory cardiorespiratory monitoring, contradicts the lay characterisation of panic as occurring without prodrome.²¹⁴ The architecture this scientific foundation describes is an observation layer: it surfaces departures from individualised breathing-structure baselines and leaves clinical interpretation to clinicians and downstream clinical systems. It does not diagnose, alert, or define thresholds. Three deployment verticals are addressed — wellness, sleep health, and clinical research — each with its own evidence base, regulatory pathway, and dataset requirements; what they share is the underlying observation channel.

1. Introduction

The history of clinical medicine over the past century has been, in considerable measure, the history of moving from snapshots to streams. Three transitions stand out: from the auscultated heart to the

continuous electrocardiogram, from the office sphygmomanometer to ambulatory blood-pressure monitoring, and from the intermittent fingerstick to continuous glucose sensing. Each transition was made possible by a specific technological enabler, each rewrote the clinical evidence base for the underlying physiology, and each redefined what could be observed about a vital sign that had previously been seen only in fragments. This chapter argues that breathing — a vital sign captured today predominantly as a single scalar count of cycles per minute — is now entering the same kind of transition. Three concurrent developments, none of them speculative, have made continuous structural observation of breathing technically feasible, computationally tractable, and addressable within established regulatory pathways. The remainder of this monograph is a synthesis of what is now known about that observable.

1.1 The pattern of continuous-monitoring transitions

The pattern is recognisable rather than inevitable. It has three concrete prior instances over the past 110 years, each separated by roughly half a century, and each grounded in a different physiological domain.

The first transition is the electrocardiogram. Einthoven's 1912 description of the string-galvanometer recording established that cardiac electrical activity could be captured as a continuous waveform rather than inferred from the percussed and auscultated chest¹. Within decades, ambulatory continuous ECG — Holter monitoring, telemetric monitoring, and now consumer wearables — demonstrated that arrhythmias, ischaemic episodes, and conduction abnormalities are predominantly intermittent phenomena that single time-point recordings systematically miss². Modern cardiology rests on the assumption that the ECG is a continuous observable; the snapshot is now the exception, not the rule.

The second transition is ambulatory blood-pressure monitoring. Office sphygmomanometry — the snapshot — became epidemiologically inadequate once it was shown that white-coat hypertension, masked hypertension, nocturnal non-dipping, and morning surge are each prognostically distinct entities, and each is invisible to a single in-clinic measurement³. The 24-hour ambulatory record subsequently redefined the diagnostic thresholds, the prognostic stratification, and ultimately the guideline architecture of hypertension care⁴. The clinical entity called *hypertension* in 2025 is not the entity that was called hypertension in 1965, and the difference is largely traceable to continuous capture.

The third transition is continuous glucose monitoring. The fingerstick paradigm of the Diabetes Control and Complications Trial gave way, after the 2008 JDRF study of intensive treatment with continuous monitoring, to a model in which glycaemic management is conducted against a continuous record of glucose excursions rather than against periodic averages⁵. The 2019 international consensus on time-in-range, time-below-range, and glycaemic variability codified what was already operationally true: HbA1c is a snapshot summary of an underlying continuous trajectory, and that trajectory carries prognostic information that the summary does not^{6,5}. The clinical entity called *type 1 diabetes* in 2025 is not the entity it was in 1995.

In each of these three instances, the transition required three things in parallel: a sensor capable of capturing the signal at adequate fidelity in an ambulatory setting; an analytical infrastructure capable of

converting raw signal into clinically interpretable structure; and a regulatory and reimbursement environment willing to recognise the new modality as legitimate. None of these was sufficient alone.

The clinical urgency for the next transition is substantial. Approximately half of US adults live with at least one chronic condition, and roughly one in four with two or more — and the progression of these conditions unfolds over months and years rather than minutes⁷. The healthspan-lifespan gap — the difference between average years of life and average years of life lived in good health — is now estimated at approximately 9.6 years globally, with the gap widening rather than narrowing across most member states^{8,9}. Both observations imply that the clinically informative phenomena are increasingly *trajectories* and *deviations from individual baselines*, not single time-point measurements; they imply, in other words, that the most informative next vital sign is one that captures structural change before the categorical clinical event the change foreshadows.

Breathing has not yet undergone this transition. The vital sign captured today as respiratory rate — twelve to twenty cycles per minute, recorded as a single integer in a triage record — is the snapshot equivalent of office blood pressure in 1965. The continuous, structural observation of breathing is the candidate for the next instance of the pattern.

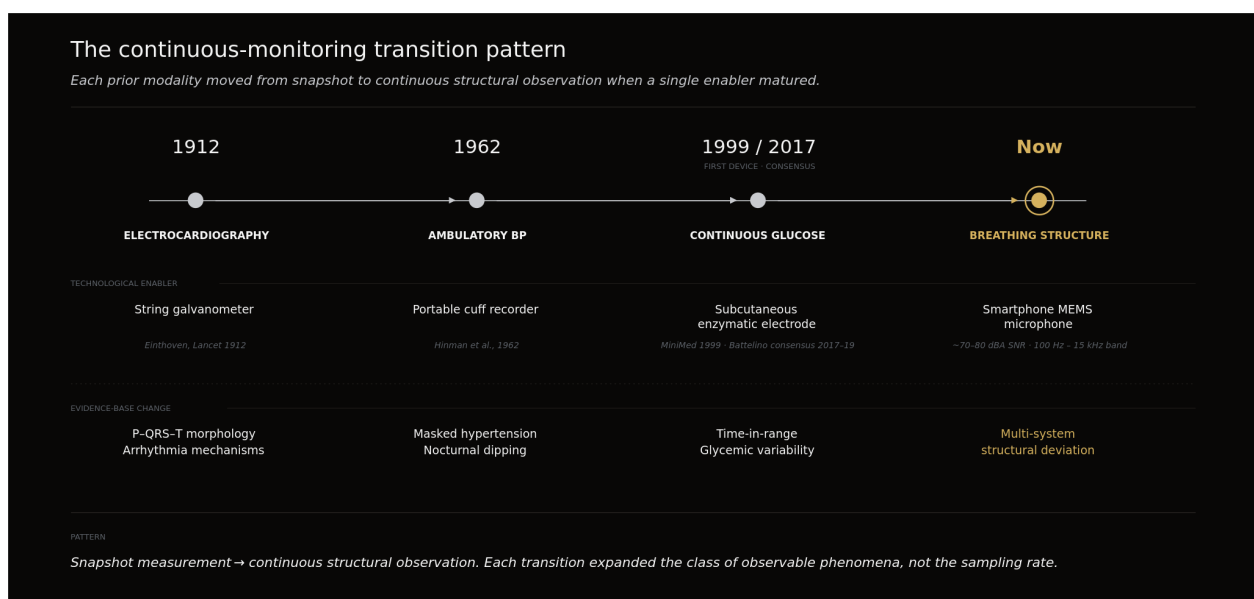


Figure 1.1. Continuous-monitoring pattern timeline

1.2 Three concurrent enabling developments

What was missing for the breathing transition was not the physiological case for it. The mechanistic basis for breathing as a multi-system observable has been characterised across the past four decades of pulmonary, autonomic, and neuro-respiratory research, and is the subject of Chapter 3 and Chapter 5 of this document. What was missing was the convergence of three specific technical and regulatory enablers, each of which has matured within the past five years.

Smartphone MEMS sensor maturity

The first enabler is the maturation of the microelectromechanical-systems (MEMS) microphone embedded in contemporary smartphones. Premium MEMS microphones now achieve signal-to-noise ratios of approximately 70 dBA — and as high as 80 dBA on flagship designs — with dynamic ranges extending to ~ 132 dB and with flat frequency response across the 100 Hz – 15 kHz band relevant to respiratory acoustics^{10, 11}. This specification envelope is sufficient to capture the acoustic structure of the full respiratory band, from the low-frequency airflow turbulence of inspiration and expiration through to the high-frequency components of nasal sibilants and cough transients^{10, 12}. Empirical demonstrations of smartphone-microphone capture of respiratory rate and respiratory waveform now report mean absolute errors below 1 breath per minute against gold-standard respiratory inductive plethysmography^{13, 14}; the full validation argument, including the predicate-device clearance history, is given in Chapter 9. The mechanistic relevance of capturing this signal is anchored in the recent demonstration that respiration globally entrains cortical and limbic rhythms, and is therefore a window into central as well as peripheral state¹⁵. The behavioural enabler runs alongside the technical one: sustained smartphone proximity is now near-universal in adult populations across high- and middle-income settings, and the smartphone microphone is, in absolute numerical terms, the most widely deployed acoustic sensor in human history.

Maturity of computational methods for time-series structure

The second enabler is the maturation of computational methods for learning structural representations from physiological time-series data. This development is treated here at category level only; the specific architectural choices any observation-layer system makes are technical-implementation questions outside the scope of this scientific foundation. At the category level, three classes of methods have reached operational maturity. First, foundational entropy- and complexity-theoretic methods — approximate entropy¹⁶, sample entropy¹⁷, multiscale entropy¹⁸, and detrended fluctuation analysis¹⁹ — provide the mathematical primitives for quantifying structure in non-stationary physiological signals. Second, structured state-space architectures exemplified by S4²⁰, together with decoder-only and encoder time-series foundation models including TimesFM²¹, Chronos²², and contemporaneous systems such as MIRA and Moirai, have demonstrated that the representation-learning paradigm²³ extends to long, irregularly sampled biological sequences. Third, acoustic-domain foundation models — notably HeAR (Health Acoustic Representations)^{24, 25} and the OPERA benchmark for respiratory-acoustic foundation models²⁶ — provide pre-trained encoders calibrated specifically to the respiratory acoustic domain. The point relevant to this monograph is not which method is preferred but that, for the first time, a learnt structural representation of a respiratory waveform — segmentation into inspiratory and expiratory phases, extraction of phase-asymmetry and variability features, and projection into a state representation comparable across recordings and across individuals — is computationally tractable on commodity hardware. This matters for the breathing case in particular because heart-rate variability — the most common comparator signal — is a downstream projection of respiration via respiratory sinus arrhythmia, and the structural features of breathing that carry prognostic information across systems are not reconstructible from the heart-rate variability record alone²⁷.

Regulatory clarity

The third enabler is regulatory. The U.S. Food and Drug Administration’s General Wellness Policy for Low-Risk Devices, last updated 6 January 2026, explicitly delineates the conditions under which non-invasive sensors that observe physiologic parameters — including respiratory parameters — may be marketed as general-wellness products without premarket submission, provided the established guardrails are observed: no disease labels, no diagnostic thresholds, no clinical alerts²⁸. This brings the regulatory pathway for continuous respiratory observation, in its non-diagnostic configuration, into a settled and well-understood category. Full treatment of the regulatory pathway, including the 510(k) screening predicates that have been cleared in the respiratory-acoustic domain, is given in Chapter 12.

Each of these three developments — sensor maturity, computational maturity, regulatory clarity — is itself the endpoint of a multi-decade trajectory. What is novel is their simultaneous arrival at the threshold of practical deployability for continuous structural observation of breathing.

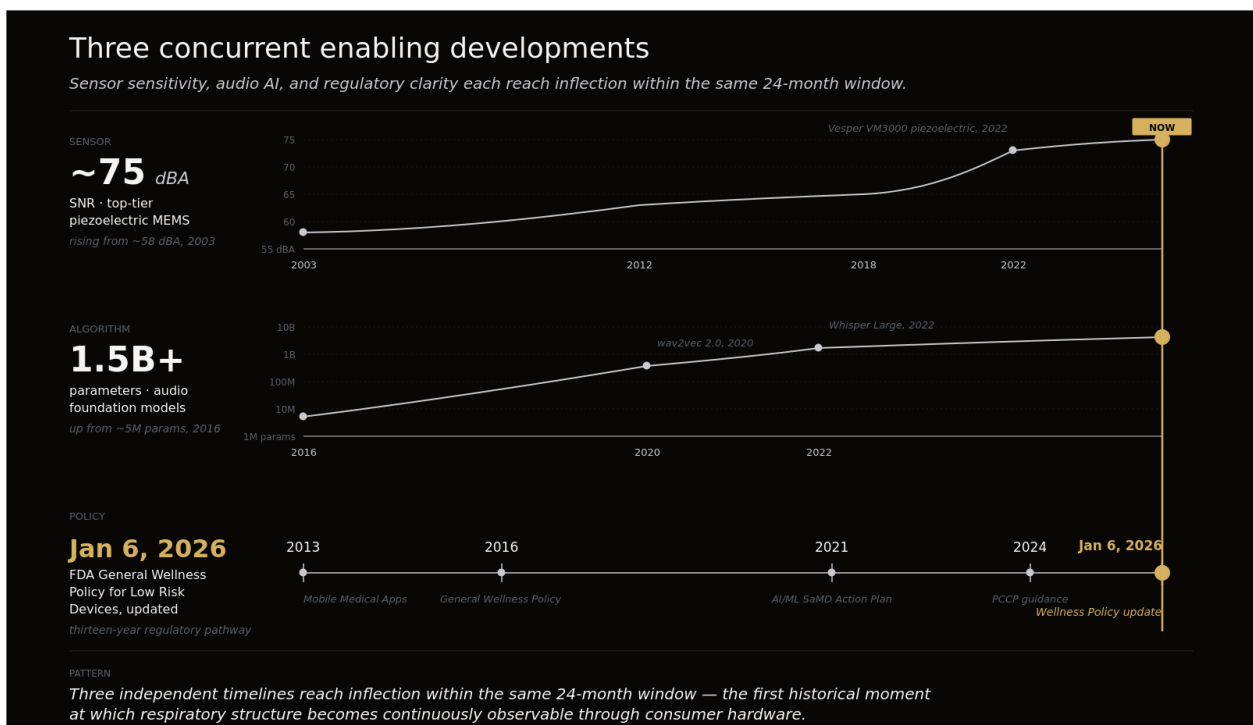


Figure 1.2. Three concurrent enabling developments: MEMS sensitivity, deep-learning time-series model size and capability, and regulatory-clarity timeline

1.3 Organisation of this manuscript

The remainder of this document is organised into thirteen substantive chapters and a synthesis. Chapter 2 situates the present moment within two centuries of respiratory observation, characterising the volume era, the gas-composition era, the oxygenation era, the events era of polysomnography, and the wearable rate era as successive slices of an underlying continuous signal. Chapter 3 develops the convergent-physiological-signal argument: breathing as a brainstem-generated rhythm modulated by autonomic balance, chemoreflex coupling, metabolic demand, and limbic-cortical activity, with structural information not reconstructible from heart-rate variability or other downstream signals.

Chapter 4 formalises the *Breathing State* construct — the temporal object whose structure rather than whose mean rate is the proper unit of observation — and surveys the families of analytical methods relevant to its characterisation. Chapter 5 is the substantive evidence chapter: a system-by-system review of breathing-structure deviation across 16 body systems, each treated at the level of effect sizes, study designs, and the strength of supporting evidence. Chapter 6 addresses the temporal-advantage question: across which systems and on what time scale do breathing-structure deviations precede the clinical events with which they co-vary? Chapter 7 develops the case for personalised baselining as the proper analytical frame for individualised observation, given the very high inter-individual variance in breathing structure within healthy populations. Chapter 8 maps the dataset gap: where high-quality, contextually annotated, longitudinal breathing-acoustic recordings exist, and — more often — where they do not. Chapter 9 reviews the smartphone-acoustic feasibility argument in technical detail, including the predicate-device validation history. Chapter 10 describes the architecture of an observation layer: what such a layer commits to, what it deliberately does not commit to, and what its three initial deployment verticals are — wellness, sleep health, and clinical research. Chapter 11 enumerates the open questions that remain, scientific and methodological. Chapter 12 describes the regulatory pathway in detail. A concluding chapter follows.

This document is a scientific synthesis. It is not a technical specification of any particular implementation. It is not a commercial brief and does not advocate for any specific product. It is not a clinical guideline and contains no diagnostic recommendations. It is an argument that breathing — long observed but rarely as structure, and rarely continuously — has now entered the historical pattern that the electrocardiogram, ambulatory blood-pressure monitoring, and continuous glucose sensing each entered before it, and that the scientific case for treating breathing structure as a multi-system upstream observation signal is, on present evidence, sufficient to merit the synthesis presented here.

2. Two centuries of respiratory observation: a history of slices

The instruments through which medicine has observed breathing have, over the past 175 years, traced a sequence of distinct paradigms. Each paradigm answered a defined clinical question with the technology available at the time: how much air a patient could move; how much carbon dioxide left the body during a procedure; how much oxygen reached the periphery; what discrete events disturbed sleep; what rate marked the line between stability and deterioration. The sequence is intelligible as a history of slices. Each slice was clinically successful within its scope. None, taken alone or together, has yet captured breathing as a continuous structural object in the conditions of ordinary life. This monograph takes that observation as its starting point; this chapter establishes the historical record on which the rest of this scientific foundation rests.

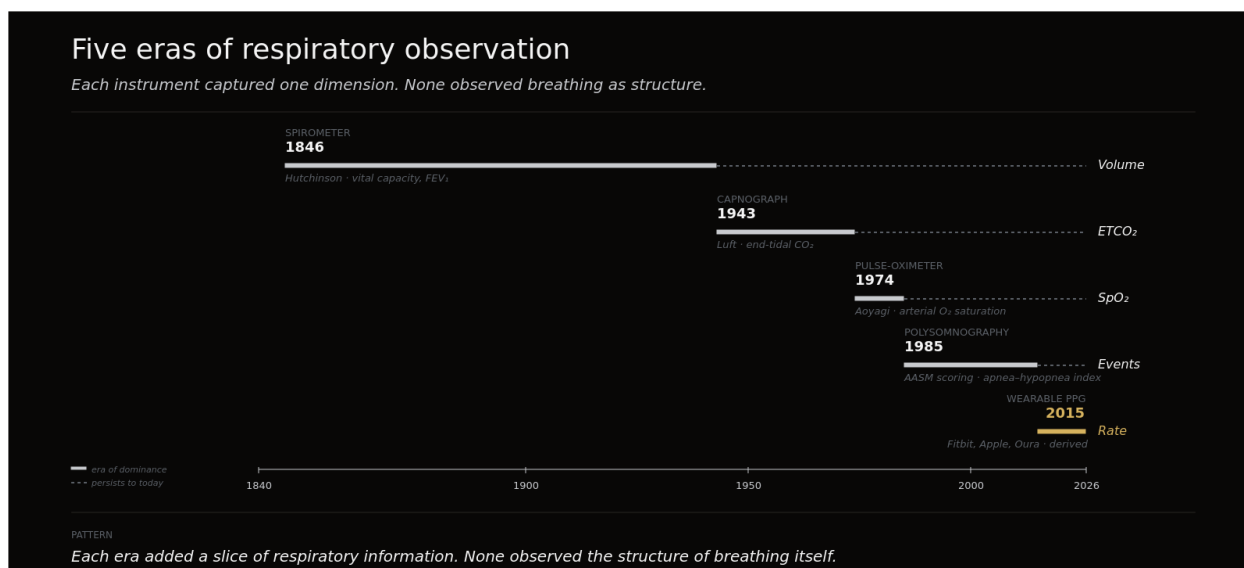


Figure 2.1. Five-era timeline (1840s–2020s)

2.1 1840s–1950s: Volume

The volume era opened in 1846 with John Hutchinson’s invention of the spirometer and his accompanying monograph in *Medico-Chirurgical Transactions*.²⁹ Hutchinson reported measurements on more than 2,000 individuals, defined “vital capacity” as the maximal volume expelled after maximal inspiration, and observed that vital capacity declined with age and varied with stature; he proposed — and the subsequent literature confirmed — that this quantity carried prognostic information about life expectancy.^{29, 30, 31} The instrument itself was elegant: a calibrated bell submerged in water rose as the subject exhaled, and the displaced volume was recorded on a graduated scale. A century later, the same operating principle still anchored bedside lung function testing, and historical syntheses have traced the unbroken methodological line from Hutchinson’s apparatus to the modern flow-volume loop.^{30, 32}

What made spirometry powerful as a clinical paradigm was the formalisation of the *forced manoeuvre*. In the decades following Hutchinson, the field accumulated standardised volumes and flows — vital capacity (VC), forced vital capacity (FVC), forced expiratory volume in one second (FEV₁), and the Tiffeneau ratio (FEV₁/VC) — each obtained by asking the patient to inhale maximally and then exhale as forcefully and completely as possible into the instrument.^{30, 32} The output was a single number, or a small set of numbers, derived under conditions deliberately removed from ordinary breathing. This was not a limitation of the era; it was its design.

The forced manoeuvre served two scientific and clinical purposes that nothing else could yet serve. First, it provided a reproducible operating point — independent of the subject’s effort moment-to-moment, dependent only on maximal effort once — and so it permitted comparison across patients, across visits, across populations, and across decades.³⁰ Second, the resulting volumes and ratios mapped cleanly onto the two great pulmonary phenotypes: obstructive disease (reduced FEV₁/VC ratio, prolonged exhalation, low I:E) and restrictive disease (proportionate reduction of FEV₁ and VC with preserved ratio).^{30, 31} This dichotomy founded pulmonology as a specialty. By the mid-twentieth

century, spirometric thresholds anchored the classification of chronic obstructive pulmonary disease, the staging of pulmonary fibrosis, the certification of occupational lung disease, and the longitudinal monitoring of asthma. Forced spirometry remains, at the time of writing, the regulatory anchor of every drug development programme in respiratory medicine.

The volume era's relationship to time was deliberate. A spirometric measurement was a *snapshot* obtained under controlled effort — a medical-imaging analogue in the time domain rather than the spatial domain. Repeating the manoeuvre a year later revealed a year of decline; repeating it a week later revealed within-week variability. What lay between the snapshots — the breath-by-breath structure of ordinary, unforced respiration — was not the object of measurement and was not, with the available instrumentation, accessible. The same conceptual move governed the parallel sciences of cardiac and circulatory physiology in the same century. Einthoven's 1912 description of the human electrocardiogram in *The Lancet* captured rhythmic structure in a detail unmatched by any prior method, but the resting twelve-lead ECG remained, in clinical practice, an episodic recording.³³ Continuous cardiac observation, when it eventually arrived, did not displace the resting ECG; it added a different axis of measurement entirely. The volume era for respiration occupied an analogous historical position.

The era thus established a model of respiratory measurement that the field has elaborated rather than replaced: episodic, effort-defined, instrument-specific, and clinic-bound. The spirometer's enduring contribution was the discipline of pulmonology itself.

2.2 1940s–1970s: Gas composition

The gas-composition era extended what could be measured from breathing — but kept the measurement, with rare exceptions, inside the operating theatre and the intensive care unit. The decisive technical advance was the integrated blood-gas analyser developed in the late 1950s. In 1958, Severinghaus and Bradley described an instrument combining a pH electrode (Stow), a carbon-dioxide electrode (Severinghaus), and an oxygen electrode (Clark) in a single benchtop unit; the design defined the standard configuration of arterial blood-gas analysis for the half-century that followed.³⁴ Two years later, Astrup and colleagues presented in *The Lancet* the acid-base nomogram and base-excess concept that became the operating language of intensive-care metabolism over the same period.³⁵ The retrospective accounts published by Severinghaus and Astrup in *Journal of Clinical Monitoring* in the mid-1980s, and by Severinghaus alone in *Anesthesiology* in 2002, trace the parallel work — Stow at Western Reserve, Clark at Antioch, Severinghaus at the National Institutes of Health and UCSF, Astrup at Rigshospitalet — that converged on the integrated analyser between 1956 and 1960.^{36, 37, 38}

End-tidal capnography, the continuous measurement of carbon-dioxide concentration in expired gas, emerged shortly thereafter as the bedside extension of arterial blood-gas analysis. Where the arterial puncture remained intermittent and invasive, the capnograph — built on the same Severinghaus CO₂ electrode that anchored the integrated analyser — sampled expired air through a sidestream or mainstream sensor and produced a breath-by-breath waveform that could be displayed in real time. End-tidal CO₂ approximated arterial PCO₂ (with a typical 2–5 mmHg gradient under normal conditions), and the shape of the capnogram itself carried information about ventilation-perfusion mismatch, airway obstruction, and circulatory adequacy.^{34, 36, 37, 38}

In one important sense, capnography was the first *continuous* observation of breathing in clinical use. Each breath produced a distinct expiratory wave; rate, pattern, and end-tidal-to-arterial gradient could be tracked across hours; alarms triggered on apnoea, hypoventilation, or oesophageal intubation. The contrast with spirometry was structural: the spirometer captured a defined manoeuvre under maximal effort; the capnograph captured every breath as it occurred. In a second important sense, however, the deployment was narrow. Capnography served procedural safety — confirmation of endotracheal tube placement, monitoring of ventilator-dependent patients, surveillance during anaesthesia, escalation of care in critical-care settings — and did so successfully enough to become a standard of care in the operating theatre and the ICU, codified into anaesthesia guidelines from the late 1980s onward and embedded in resuscitation protocols thereafter.³⁹

What the gas era did not produce — and what the available technology, with its sampling tubing, calibration requirements, and patient-tethered architecture, could not produce — was respiratory observation of the ambulatory adult in ordinary life. Capnography required an airway interface (mask, endotracheal tube, or nasal cannula), an analyser with regular zero-point calibration, and continuous power; the cost and form factor of the instrumentation kept it within the hospital. The evidence base accumulated by the gas era was therefore an evidence base for *peri-procedural* and *peri-acute* respiratory observation, not for everyday physiology. This served the clinical purposes for which the era was responsible. Procedural mortality fell. Ventilator management became a quantitative science. The infrastructure of modern critical care rests on this foundation.

2.3 1970s–1990s: Oxygenation

The oxygenation era began with a single inventor and a single insight. In Tokyo, around 1972, Takuo Aoyagi at Nihon Kohden recognised that the pulsatile component of light absorption through tissue carried information about arterial oxygen saturation, separable from the static component contributed by venous blood and surrounding tissue.^{40, 41, 42} The two-wavelength pulse oximeter — red and near-infrared light through a finger or earlobe, ratio-based extraction of arterial saturation — was first commercialised in Japan in the mid-1970s and entered Western clinical practice through the early 1980s.^{40, 43, 42} By the 1990s, SpO₂ had become a routine vital sign in operating rooms, recovery rooms, emergency departments, paediatric and neonatal wards, and home oxygen-therapy programmes.⁴³ Quaresima and colleagues' 2024 ninety-year synthesis traces the full arc, from Matthes's 1935 two-wavelength oximeter, through the Hewlett-Packard ear oximeter of the 1970s, to the modern reflectance and transmission designs used in contemporary wearables and clinical monitors.⁴³

The clinical significance of pulse oximetry was that it took respiratory observation out of the laboratory and out of the procedural environment, and placed it at the patient's bedside as a *routine* vital sign. SpO₂ joined heart rate, blood pressure, temperature, and respiratory rate as the canonical "vital signs" used in early-warning scoring systems and standard nursing observations across care settings.⁴³ Long-term oxygen therapy, prescribed for severe COPD and chronic hypoxaemia, became a quantifiable home-care intervention rather than an empirical one. Neonatal oxygen titration, previously a frequent source of retinopathy of prematurity in over-supplemented infants, gained an objective monitor.⁴⁰ Continuous postoperative pulse oximetry, often coupled with capnography in higher-risk patients, was

associated in subsequent meta-analyses with reductions in respiratory depression events and unplanned escalations of care.³⁹

Pulse oximetry shared with capnography the property of breath-by-breath continuity, but unlike capnography its sensor was non-invasive, low-cost, and tolerated by ambulatory patients indefinitely. By the 2000s, finger pulse oximeters were available over the counter; by the 2010s, photoplethysmographic optical paths had been miniaturised into the underside of consumer wristworn devices, recording both heart rate and — with appropriate signal processing — respiratory rate from the modulation of the cardiac waveform by intrathoracic pressure changes.^{40, 43} The oxygenation era thus did two things at once. First, it standardised a respiratory-adjacent measurement across every layer of clinical infrastructure, from the neonatal intensive care unit to the patient's home. Second, it began the process — accelerated in the rate era described in 3.5 — of moving respiratory-relevant measurement off the clinic floor and onto the patient's body in ordinary life.

What pulse oximetry did *not* observe is, in retrospect, instructive. SpO₂ tracks the downstream consequence of ventilation — arterial oxygen saturation — rather than the act of breathing itself. Under most conditions, oxygen saturation is buffered by haemoglobin's sigmoid dissociation curve, so substantial decrements in alveolar ventilation produce only small SpO₂ changes until the curve's steep portion is reached. Apnoeic events, hypoventilation episodes, and altered breathing patterns are detected by pulse oximetry only after they have produced sufficient desaturation; the breathing structure that preceded them is not part of the SpO₂ record. This is a property of the oxygen-haemoglobin physiology, not a defect of the instrument.

2.4 1960s–2000s: Events

The events era was the era of polysomnography. Its foundational publication, issued in 1968, was a small monograph from the UCLA Brain Information Service and the Brain Research Institute — Rechtschaffen and Kales's *A Manual of Standardized Terminology, Techniques and Scoring System for Sleep Stages of Human Subjects*.⁴⁴ The manual codified the visual scoring of EEG, EOG, and EMG channels into discrete sleep stages (Wake, NREM 1–4, REM) over thirty-second epochs, and in doing so it created the operating framework within which respiratory disturbances of sleep could be defined, identified, and counted. In 1976, Guilleminault, Tilkian, and Dement published “The Sleep Apnea Syndromes” in *Annual Review of Medicine*, naming the clinical entity that came to dominate the field for the next half-century.⁴⁵

Polysomnography integrated respiratory channels — chest and abdominal effort belts, nasal-oral airflow sensors, pulse oximetry, and (by the 1990s) end-tidal CO₂ — into the multi-channel sleep recording. The respiratory output of PSG was a sequence of *events*: apnoeas (cessations of airflow ≥ 10 seconds), hypopnoeas (reductions of airflow with associated desaturation or arousal), respiratory effort-related arousals, central versus obstructive subtyping, and oxygen desaturation indices.^{46, 47} These events were aggregated into the apnoea-hypopnoea index (AHI), a count per hour of sleep, and into derived severity categories that anchored the diagnostic and treatment infrastructure of sleep medicine for several decades.

The events framework was clinically generative. It supported the identification of obstructive sleep apnoea as a public health problem, the development of continuous positive airway pressure as the principal therapy, and the documentation that treated obstructive sleep apnoea was associated with reductions in daytime sleepiness, motor-vehicle-accident risk, and, in selected populations, cardiovascular morbidity. Sleep medicine was founded as a discipline on this scaffold. The American Academy of Sleep Medicine produced and maintained its scoring manuals — Iber and colleagues' 2007 first edition and successive updates including Berry and colleagues' 2017 Version 2.4 — to standardise event definitions, sensor specifications, and acceptable signal sources across an international laboratory network.^{48, 46}

Pevernagie and colleagues' 2020 historical review, *On the rise and fall of the apnoea-hypopnoea index: a historical review and critical appraisal*, traced both the ascendancy of the AHI as the operative metric of sleep medicine and the accumulating evidence that the metric, as a count of events per hour, captures only part of the relevant respiratory information.⁴⁷ Patients with similar AHIs differ markedly in symptoms, in cardiovascular outcomes, and in treatment response; events of identical duration carry different prognostic weight depending on desaturation depth, arousal architecture, and breath-to-breath structure surrounding them. The review framed this not as a critique of the events paradigm but as a description of the historical record: the AHI served the foundational decades of sleep medicine and remains its operating metric, while the field continues to develop additional structural descriptors of nocturnal respiration. The boundary condition flagged in the source — that the AHI is now understood as a partial summary of nocturnal respiratory physiology rather than its full description — is consistent with the broader thesis of this chapter: each historical era extracted what its instrumentation permitted, and what the instrumentation did not permit was deferred.

The era's contribution to respiratory observation was therefore the discrete-event vocabulary itself — apnoea, hypopnoea, arousal, desaturation, periodic breathing — and the clinical infrastructure (laboratories, scoring conventions, severity thresholds, reimbursement codes) that allowed this vocabulary to operate at scale. Respiration, in the events framing, was a sequence of countable disturbances of an otherwise quiescent default state, observed for one or two laboratory nights under sensor conditions designed for that environment. The structural variation of the default state itself — the breath-to-breath organisation of unobstructed sleep, across many nights, in the patient's own bed — was not the object of measurement in this era.

2.5 2010s–2020s: Rate

The rate era arrived through the consumer wrist. From the early 2010s, photoplethysmography on wristworn devices, electrocardiography on multi-lead patches, and accelerometry on torso sensors made it possible to derive a continuous estimate of respiratory rate from physiological signals that were already being recorded for other reasons.^{49, 50} The dominant approach used the modulation of the cardiac PPG waveform by intrathoracic pressure changes during breathing — respiratory sinus arrhythmia and the related amplitude- and frequency-modulation channels — to extract a per-minute breathing rate without any dedicated respiratory sensor. Validation studies against polysomnography and respiratory inductance plethysmography across consumer devices reported breath-rate errors

typically in the range of approximately one breath per minute under resting conditions, with characteristic degradation during motion and at the extremes of physiological respiratory rate.^{51, 52, 49, 50}

Continuous, unattended respiratory rate at population scale produced its own evidence base. Baumert and colleagues, analysing pooled data from the MrOS Sleep Study and the Study of Osteoporotic Fractures (n = 2,686 community-dwelling older adults), reported that a mean nocturnal respiratory rate at or above 16 breaths per minute was independently associated with cardiovascular and all-cause mortality, with hazard ratios of 1.57 to 2.58 across stratifications, after adjustment for established risk factors and concurrent vital signs.⁵³ The Sepsis-3 consensus, published in *JAMA* in 2016, retained respiratory rate ≥ 22 breaths per minute as one of three quickSOFA criteria for the bedside identification of patients at elevated risk of in-hospital mortality outside the intensive care unit.⁵⁴ Earlier and complementary evidence — Fiesemann and colleagues' demonstration in 1993 that respiratory rate predicted cardiopulmonary arrest in internal-medicine inpatients — had been available for two decades but was not, until the wearables era, embedded in the data infrastructure of ordinary daily life.⁵⁵

The era served the integration of respiration into consumer health. Wearable manufacturers added respiratory rate to dashboards alongside heart rate, sleep duration, and step counts; the metric became a daily variable for general wellness and fitness tracking, and a relevant input for early symptom-detection signals during respiratory illness outbreaks.⁵⁰ Boundary conditions on the era are recognised within its own validation literature: device-to-device differences in algorithm and sensor placement produce systematic biases; accuracy degrades during motion and at the extremes of respiratory rate; and one randomised crossover trial reported that wearables themselves can perturb the sleep behaviour they aim to record.^{51, 52, 49}

What the rate era observes is, by design, a single scalar per averaging window — typically per minute, per night, or per session. Within-breath structure (inspiration-to-expiration ratio, timing variability, depth modulation, pause distribution) is below the temporal resolution of most consumer respiratory-rate algorithms, and the breathing acoustic structure, when present in the underlying signal, is not part of PPG- or ECG-derived rate. The rate era extended continuous observation of breathing further than any previous era; it did so by reducing the observation to a rate.

2.6 Synthesis: five slices, one missing axis

Across the historical record, each era produced one slice of respiratory observation. The volume era produced effort-defined volumes and flows, captured episodically in the clinic. The gas-composition era produced breath-by-breath continuity, but only within the operating theatre and the intensive care unit. The oxygenation era produced a peripheral, non-invasive, ambulatory measurement, but tracked the downstream consequence of breathing rather than the act of breathing itself. The events era produced a vocabulary of discrete respiratory disturbances, observed during one or two laboratory nights under sensor conditions that did not generalise to ordinary life. The rate era produced continuous, unobtrusive observation of a single scalar — respiratory rate — across millions of users in everyday conditions.^{56, 57}

Each slice was clinically successful within its scope; together, the five eras founded pulmonology, intensive care, sleep medicine, and the emerging field of consumer respiratory health. The historical

record contains, however, a recurring observation about the trajectory of clinical evidence in adjacent physiological domains. In other organ systems, the transition from episodic to continuous measurement has, characteristically, rewritten the underlying evidence base. Twenty-four-hour ambulatory blood-pressure monitoring revealed phenomena — nocturnal non-dipping, masked hypertension, white-coat hypertension — that were invisible to the office cuff and that carried independent prognostic information.⁵⁸ Continuous glucose monitoring revealed time-in-range, glycaemic variability, and post-prandial structure that HbA1c summarised but did not represent, with associated changes in clinical decision-making and demonstrated cost-effectiveness in type 1 diabetes.⁵⁹ Continuous electrocardiography — Holter recording, ambulatory event monitors, and now consumer wrist-based ECG — produced an arrhythmia evidence base unattainable by the resting twelve-lead recording from which the field began.³³

The respiratory record contains a parallel observation that has been documented for decades but not yet exploited at the structural level. Respiratory rate has been described as “the neglected vital sign” — the most predictive and the most poorly recorded of the canonical vital signs — for at least three decades.^{55, 60} Systematic syntheses of early-warning scores find respiratory rate to be among the single most informative vital signs for the detection of clinical deterioration on general wards, and yet it remains, in routine inpatient practice, the vital sign most often missing or estimated rather than measured.⁶⁰ The rate era partially closed this gap for ambulatory adults; it did so for the rate alone.^{56, 57}

The slice that the historical record does not yet contain is the *structural* observation of breathing — the breath-by-breath organisation, the inspiration-to-expiration architecture, and the variability and stability of timing — captured *continuously*, in *ordinary life*, outside the clinic, the laboratory, and the operating theatre. The volume, gas, oxygenation, events, and rate eras each occupied a different combination of axes — episodic versus continuous, clinic-bound versus ambulatory, scalar versus structural — but the structural-continuous-natural quadrant has not yet been filled by any prior generation of instrumentation.

This is the gap addressed by the present work.

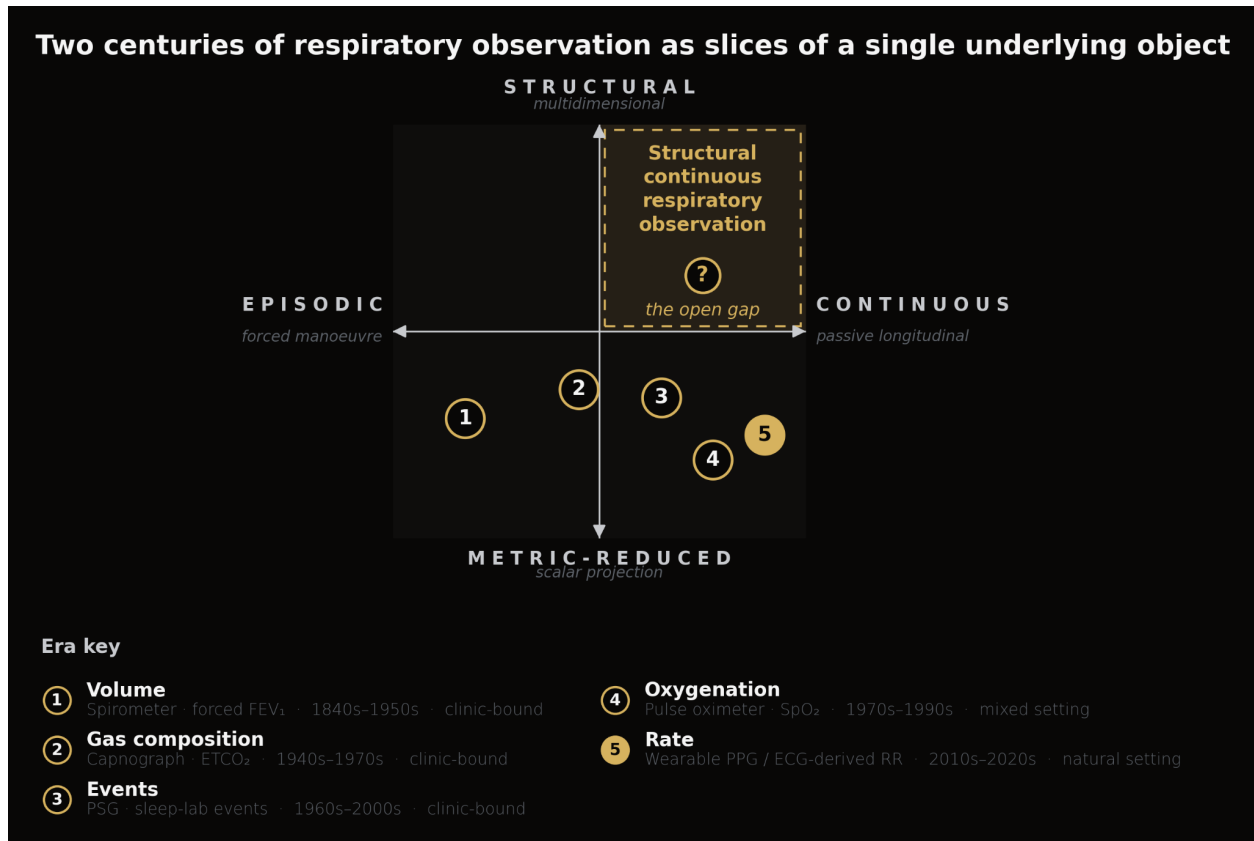


Figure 2.2. Each historical era of respiratory observation is plotted by its dominant observational properties along two axes: episodic (forced manoeuvre) versus continuous (passive longitudinal) on the horizontal, and metric-reduced (scalar projection) versus structural (multidimensional) on the vertical. A third dimension — clinic-bound versus natural setting — is encoded by marker fill: solid for clinic-bound, grey for mixed, hollow for natural. All five eras lie below the structural axis, each yielding a single dimension or event-level reduction of breathing, and — with the partial exception of modern wearable-PPG-derived rate — all are clinic-bound. The continuous-and-structural quadrant remains the open observation gap addressed by smartphone-acoustic respiratory observation in natural settings.

3. Breathing as a convergent physiological signal

Chapter 2 traced two centuries of episodic respiratory observation — Hutchinson’s spirometer, Smalhout’s capnograph, the Severinghaus and Astrup pulse oximeter, the Rechtschaffen and Kales polysomnogram. Each device established that respiration carries clinical information; none demonstrated *why* respiration, among all candidate physiological signals, should occupy a privileged structural position. Chapter 3 develops this argument from neurobiology and integrative physiology, not analogy. The thesis advances on three pillars and closes on a boundary condition. First, respiratory rhythm originates from a discrete anatomical substrate in the brainstem — a network of neurons whose properties are now mapped at single-cell resolution. Second, that substrate receives convergent input from autonomic balance, chemoreflex sensing, and metabolic demand, integrating three otherwise-separable regulatory loops in real time. Third, breathing is the only continuously accessible vital sign that is both autonomically regulated and voluntarily modulable, embedding a hybrid behavioral channel into an otherwise involuntary signal. The chapter closes with a non-redundancy boundary: respiration

carries information about physiology that is not reconstructible from heart rate variability or aggregated multimodal sensors alone. None of these findings requires evidence-by-body-system — that argument is developed in Chapter 5. Chapter 3 establishes mechanism; Chapter 5 demonstrates application.

3.1 Brainstem oscillator architecture: preBötzing complex, parafacial respiratory group, retrotrapezoid nucleus

The modern neurobiology of breathing has a single discovery moment. In 1991, Smith, Ellenberger, Ballanyi, Richter & Feldman localised, in *Science*, a small region of the ventrolateral medulla whose isolated slice continued to generate rhythmic inspiratory bursts in the absence of any peripheral input.⁶¹ They named it the pre-Bötzing complex (preBötC), and demonstrated that targeted lesion of this region in vivo abolished eupneic breathing. The discovery established, for the first time, that mammalian respiratory rhythm has a discrete and lesionable neuroanatomical generator — not a distributed emergent property of brainstem activity, but a circumscribed network whose ablation extinguishes the rhythm. This shift, from rhythm-as-emergent-property to rhythm-as-localisable-circuit, is the foundation on which the rest of this chapter rests. If respiration were a diffuse network output, the convergent-signal proposition would be metaphorical. Because respiration has an identifiable oscillator, the proposition becomes neurobiological.

The two-oscillator model that organises modern respiratory neurobiology was synthesised by Feldman & Del Negro in their 2006 *Nature Reviews Neuroscience* review, *Looking for inspiration: new perspectives on respiratory rhythm*.⁶² The inspiratory rhythm is generated in preBötC; expiratory activity, particularly active expiration during exercise or hypercapnia, is generated in the parafacial respiratory group (pFRG), spatially adjacent but functionally distinct. Subsequent work has refined this two-oscillator architecture into a network in which preBötC and pFRG interact with the retrotrapezoid nucleus (RTN), the principal site of central chemoreception.^{63, 64} Feldman, Del Negro & Gray's 2013 *Annual Review of Physiology* synthesis, *Understanding the rhythm of breathing: so near, yet so far*, sets out the canonical formulation: respiratory rhythm is generated by an interacting network of brainstem oscillators whose individual neurons exhibit cell-autonomous bursting properties, modulated by recurrent excitatory and inhibitory connections, and tuned by chemosensory input.⁶³ Forster's 2012 *Comprehensive Physiology* chapter on the control of respiration locates this network within the broader framework of duty-cycle control, distinguishing inspiratory time (T_i), expiratory time (T_e), and the T_i/T_{tot} duty ratio as the elementary observables of brainstem rhythm output.⁶⁵

Mechanistic depth has accumulated rapidly. The molecular basis of central CO_2 chemosensation has been resolved at receptor-level resolution: Kumar et al. demonstrated in *Science* that CO_2 -driven breathing requires the proton-activated receptor GPR4 in RTN neurons, with GPR4-knockout mice showing markedly attenuated ventilatory responses to hypercapnia.⁶⁶ Van de Wiel et al. complemented this finding by identifying connexin26 hemichannels in the medulla oblongata as direct CO_2 sensors that regulate breathing through ATP release.⁶⁷ The peptidergic control of sighing — the spontaneous, low-frequency, augmented inspirations that punctuate eupneic breathing — was traced by Li, Janczewski, Yackle and colleagues in *Nature*, who identified two neuropeptides expressed in a small population of preBötC neurons whose ablation eliminates sighs without disrupting eupnea.⁶⁸ The implication is

structural: distinct breathing patterns — eupnea, sighing, gasping, sniffing — map onto distinct preBötC neuronal subpopulations, each with its own molecular identity and projection targets.

Two further findings extend the brainstem oscillator beyond pure rhythm generation into broader integrative physiology. Yackle and colleagues, also in *Science*, identified approximately 175 preBötC neurons that project directly to the locus coeruleus, the principal noradrenergic nucleus regulating arousal; selective ablation of these neurons left breathing intact but rendered the animals abnormally calm in conditions that ordinarily evoke arousal.⁶⁹ The brainstem rhythm generator thus carries a direct anatomical channel into ascending arousal circuits, providing a substrate for the breathing-arousal coupling that re-enters this monograph in 4.3 and Chapter 5. Finally, Tort, Laplagne, Draguhn & Gonzalez's 2025 *Nature Reviews Neuroscience* review, *Global coordination of brain activity by the breathing cycle*, synthesises a decade of evidence that respiration-coupled oscillations are present throughout the rodent and human brain — not only in olfactory areas, but in limbic, frontal, and parietal regions — and that the source of this coupling is nasal airflow through the olfactory epithelium, abolished by tracheotomy.¹⁵ The brainstem oscillator is not isolated; it is the entry point of a coordinated system that extends from medullary chemosensors to cortical gamma oscillations.

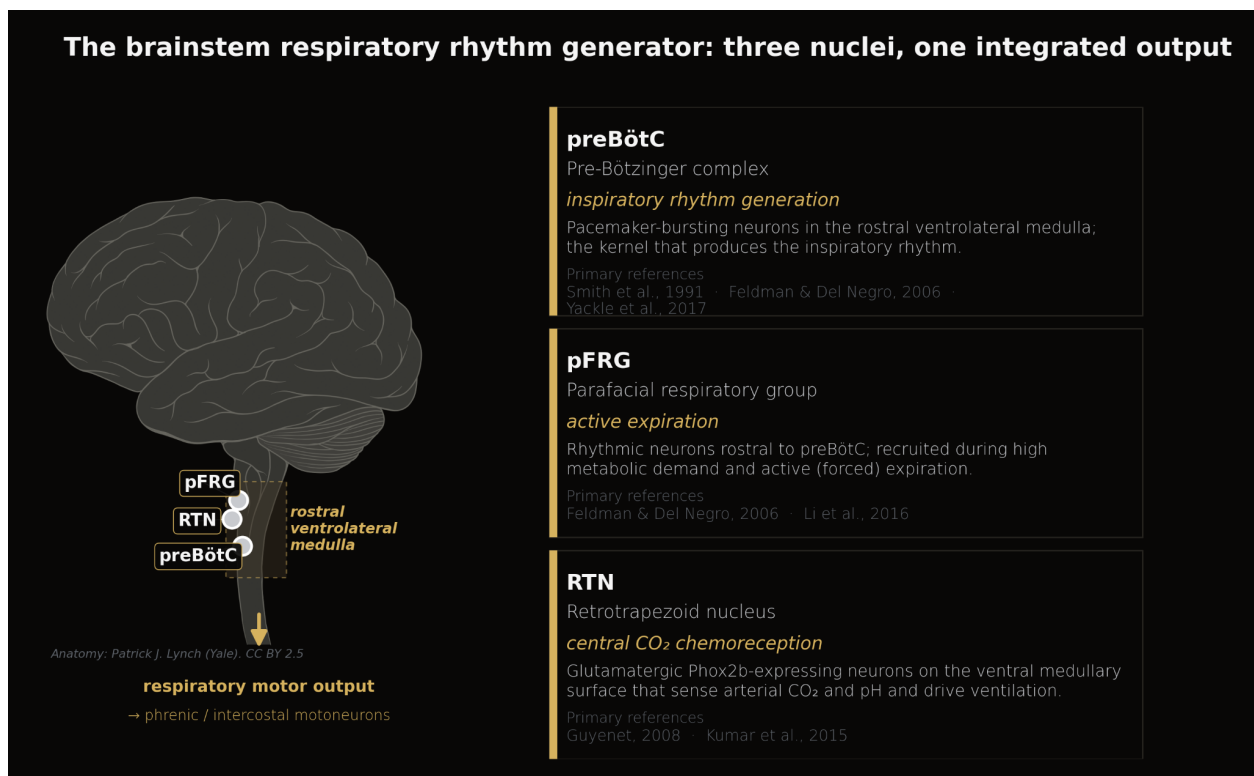


Figure 3.1. The three nuclei generating and modulating respiratory rhythm cluster within the rostral ventrolateral medulla. The Pre-Bötzinger complex (preBötC) is the inspiratory rhythm kernel; the parafacial respiratory group (pFRG) drives active expiration during high metabolic demand; the retrotrapezoid nucleus (RTN) provides central CO₂ chemoreception. Integrated output drives phrenic and intercostal motoneurons. Anatomical reference: Patrick J. Lynch (Yale), CC BY 2.5.

The takeaway for this document is structural. Respiratory rhythm is not metaphorically “central”; it is anatomically central, with a discrete generator, mapped chemosensory inputs, and identified projection targets. This grounds the convergent-signal argument that follows in 4.2 and the hybrid-control

argument in 4.3 in primary neurobiology rather than in functional analogy. The remainder of the chapter develops what flows *into* and *out of* this oscillator.

3.2 Autonomic balance, chemoreflex coupling, metabolic demand

The convergence proposition is quantitative: respiration is the unique continuously accessible physiological signal that integrates three otherwise-separable regulatory streams — autonomic balance, chemoreflex sensing, and metabolic demand — at the same temporal resolution as the signal itself. This subsection develops each input stream, then closes on the volitional channel that bridges into 4.3.

Autonomic balance and respiratory sinus arrhythmia

The autonomic input to respiration is bidirectional. Respiration receives vagal afferent traffic from pulmonary stretch receptors and chemoreceptor projections, and respiration outputs autonomic modulation through respiratory sinus arrhythmia (RSA) — the cyclical acceleration of heart rate during inspiration and deceleration during expiration. The directional structure of this coupling has been a subject of forty years of mechanistic dissection, and the modern consensus is unambiguous. The mechanistic foundation was laid by Eckberg's 1983 *Journal of Applied Physiology* paper, *Human sinus arrhythmia as an index of vagal cardiac outflow*, which established that beat-to-beat heart-period changes follow cardiac vagal efferent activity quantitatively with nearly fixed latencies — making RSA a quantitative readout of central vagal mechanisms rather than a non-specific oscillation.⁷⁰ Eckberg's 2003 *Journal of Physiology* topical review, *The human respiratory gate*, formalised the broader gating framework: respiration phasically alters the membrane potentials of preganglionic vagal and sympathetic motoneurons and continuously modulates their responsiveness to baroreceptor and chemoreceptor inputs, with RSA as the most visible manifestation of this respiratory gating.⁷¹ Yasuma & Hayano's 2004 review, *Respiratory sinus arrhythmia: why does the heartbeat synchronize with respiratory rhythm?*, then synthesised the directional canonical formulation: RSA is the mechanism by which respiration generates heart rate variability, not the reverse.⁷² Respiratory drive in brainstem networks gates vagal efferent activity to the sinoatrial node, producing beat-to-beat heart rate fluctuations whose amplitude and frequency are direct readouts of respiratory phase and cardiovagal tone. Wehrwein, Critchley, Yasuma, Hayano and colleagues formalised this position in their 2025 *Nature Reviews Cardiology* statement, proposing that RSA be renamed *respiratory heart rate variability* (RHRV) precisely to make the directional dependency explicit in the terminology.⁷³ The clinical implication of this directionality is foundational for the boundary-condition argument in 4.4: any wearable that measures HRV is, in mechanistic terms, sampling a downstream projection of breathing.

The vagal substrate underlying RSA has been mapped at functional-anatomy resolution by Neuhuber & Berthoud in their 2022 *Autonomic Neuroscience* review, which emphasises the somatic and parasympathetic afferent components of the vagus and clarifies how respiratory phase modulates parasympathetic outflow at the level of the nucleus ambiguus.⁷⁵ Porges' 2009 *Cleveland Clinic Journal of Medicine* statement of polyvagal theory provides the broader framing within which RSA is interpreted as a marker of cardiovagal regulation across stress and behavioural states.⁷⁶ The European Society of Cardiology and North American Society of Pacing and Electrophysiology 1996 *Circulation* statement on

heart rate variability — the foundational measurement-standards consensus — codifies the spectral and time-domain conventions used to quantify RSA in clinical and research practice.⁷⁷

The transdiagnostic implications of RSA reduction across psychopathology have been mapped in three converging meta-analyses and a programme of cohort studies. Beauchaine's 2015 *Current Opinion in Psychology* synthesis frames reduced resting RSA as a transdiagnostic biomarker of emotion regulation deficits across depression, anxiety, externalising disorders, and developmental psychopathology.²⁷ Beauchaine, Bell, Knapton, McDonough-Caplan, Shader & Zisner's 2019 *Psychophysiology* meta-analysis of 37 studies confirmed that reduced RSA reactivity is observable across empirically-defined structural dimensions of psychopathology, not localised to a single diagnosis.⁷⁸ Brush et al. demonstrated prospectively, in *Biological Psychology*, that lower resting RSA predicts subsequent depressive symptoms in initially-healthy adults.⁷⁹ Rottenberg's 2007 critical analysis in *Biological Psychology* synthesised 13 studies on cardiac vagal control in depression, and Kemp et al.'s 2010 *Biological Psychiatry* review documented that depression and antidepressant treatment exert measurable effects on heart rate variability.^{80, 81} Bradshaw and colleagues extended the RSA-as-biomarker frame into developmental psychiatry, demonstrating in *Autism Research* that elevated RSA from 9 to 24 months of age is associated with autism risk — extending the framework into the opposite direction of effect, with developmental specificity.⁸²

The prognostic information carried in RSA extends beyond psychopathology into cardiovascular outcome. Sinnecker, Dommasch, Barthel and colleagues, in the *Journal of the American College of Cardiology*, demonstrated that expiration-triggered sinus arrhythmia is independently associated with mortality after acute myocardial infarction in a cohort of 941 patients (hazard ratio 3.41); this finding was subsequently replicated in the INVADE cohort (n = 1,788) in 2025.⁸⁴ Hoyer, Friedrich, Frank and colleagues, in *Medical & Biological Engineering & Computing*, demonstrated that cardiorespiratory cross mutual information — a joint signal incorporating both respiration and HRV — outperforms standard HRV alone for post-MI risk stratification, providing direct quantitative evidence that the joint signal carries information not captured by HRV in isolation.⁸⁵ At the population level, Jarczok and colleagues' 2022 systematic review and meta-analysis in *Neuroscience & Biobehavioral Reviews*, covering 32 studies and 38,008 participants, demonstrated that lower HRV indices are independently associated with all-cause mortality, with a pooled hazard ratio of 1.56 (95% CI 1.32–1.85) for low-HRV vs reference contrasts.⁸⁶

Chemoreflex coupling

The chemoreflex input to respiration is two-tier. Central chemoreception, mediated principally by the RTN and its molecular sensors, monitors brain interstitial pH as a proxy for arterial PCO₂; peripheral chemoreception, mediated by the carotid and aortic bodies, monitors arterial PO₂, PCO₂ and pH directly. The molecular substrates of central chemoreception are now resolved with single-receptor specificity: GPR4 proton-activated signalling in RTN neurons,⁶⁶ connexin26 hemichannels in the medulla,⁶⁷ and the broader RTN integrative architecture established by Guyenet.⁶⁴ Peripheral chemoreception couples to long-term metabolic and oxidative regulation: Nanduri, Makarenko, Reddy and colleagues demonstrated, in the *Proceedings of the National Academy of Sciences*, that epigenetic regulation of hypoxic sensing in carotid body chemoreceptors is sufficient to disrupt cardiorespiratory

homeostasis — placing respiration's chemosensory loop within the broader frame of gene-regulatory control of physiological set-points.⁸⁷

Metabolic demand and cardiorespiratory coupling

Ventilation is matched to metabolic rate at a fundamental thermodynamic level: alveolar gas exchange must export CO₂ at the rate of metabolic production and import O₂ at the rate of metabolic consumption. Khoo's 2018 textbook *Physiological Control Systems: Analysis, Simulation, and Estimation* provides the canonical engineering treatment of this coupling, formalising the feedback loops linking VO₂, VCO₂ and minute ventilation.⁸⁸ Cardiorespiratory coupling — the broader phenomenon of common rhythms in cardiac, sympathetic, and respiratory activities at multiple time scales — has been quantified at directional resolution by Borovkova, Prokhorov, Kiselev and colleagues in *Frontiers in Network Physiology*, who demonstrated, across sleep and wakefulness in healthy subjects of different ages, that respiration drives parasympathetic control of heart rate at characteristic time-domain phase lags — framing respiration as a coordinator of cardiovascular and autonomic dynamics rather than as an isolated rhythm.⁸⁹ The clinical reframing of cardiorespiratory fitness as a vital sign was formalised by Ross, Blair, Arena and colleagues in the 2016 American Heart Association scientific statement in *Circulation, Importance of assessing cardiorespiratory fitness in clinical practice* — a position that anchors the broader argument, developed across this scientific foundation, that breathing structure is observable at the same level of clinical centrality as the four traditional vital signs.⁹⁰

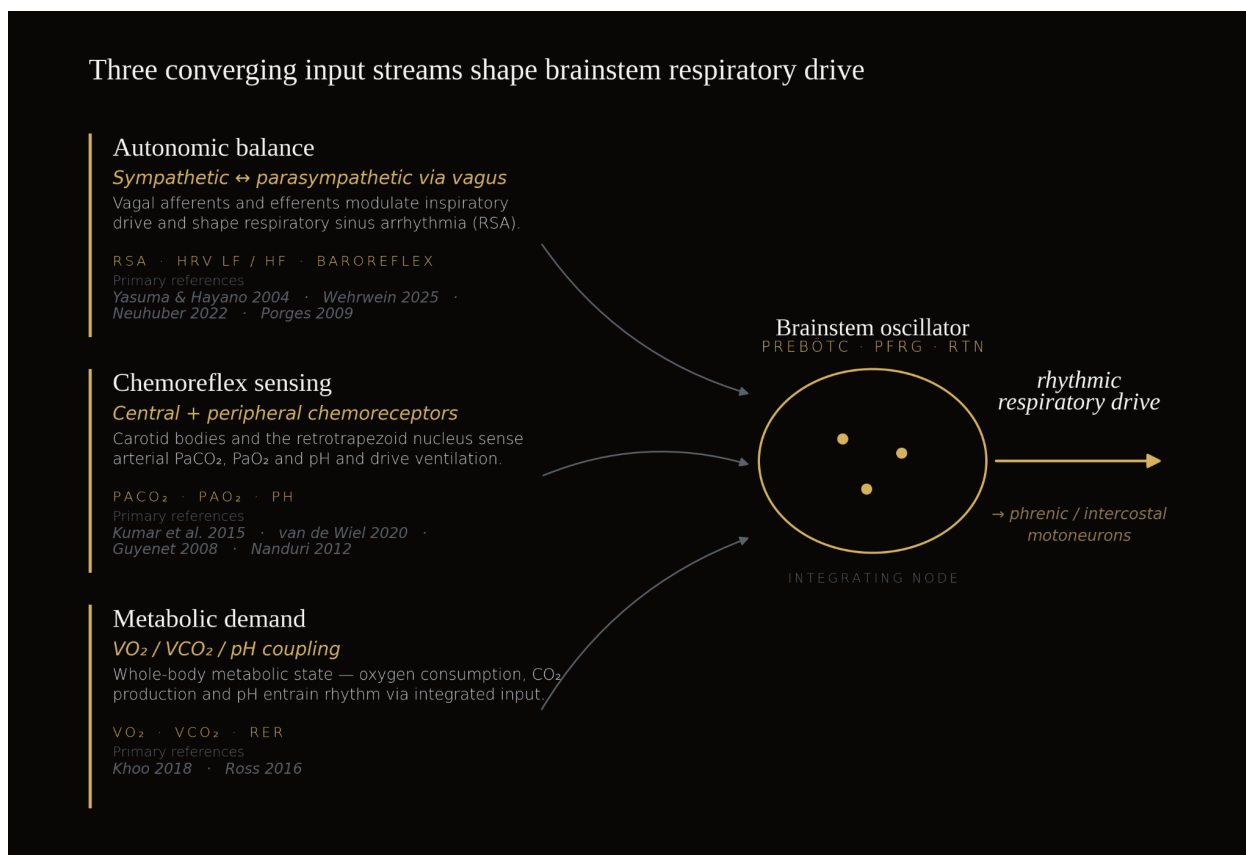


Figure 3.2. Three input streams converge on the brainstem oscillator network and shape its rhythmic output. Autonomic balance (sympathetic and parasympathetic via vagus) modulates inspiratory drive and shapes respiratory sinus arrhythmia; chemoreflex sensing through carotid bodies and RTN reports PaCO₂, PaO₂, and pH; whole-body metabolic demand (VO₂, VCO₂, RER) entrains rhythm. The integrating node — preBötC/pFRG/RTN complex — produces a single rhythmic respiratory drive descending to phrenic and intercostal motoneurons.

Volitional and contemplative modulation

Beyond autonomic, chemoreflex, and metabolic input, respiration receives an additional input stream that no other vital sign possesses at this resolution: continuous voluntary modulation. The structure of this volitional channel is foreshadowed here and developed mechanistically in 4.3. Van Diest, Verstappen, Aubert and colleagues, in *Applied Psychophysiology and Biofeedback*, demonstrated that the inhalation/exhalation ratio (I:E) modulates the cardiovagal effect of slow breathing on heart rate variability and subjective relaxation: prolonging the expiratory phase (low I:E, e.g., 1:3 or 1:4) shifts autonomic balance toward parasympathetic dominance, while equalising inspiration and expiration (I:E approaching 1:1) is associated with sympathetic activation and stress states.⁹¹ This phase-asymmetry effect is mechanistically continuous with the directional RSA finding: prolonged expiration extends the vagal-gating window during which heart rate decelerates. Gerritsen & Band's 2018 *Frontiers in Human Neuroscience* synthesis, *Breath of life: the respiratory vagal stimulation model of contemplative activity*, frames this volitional channel within the broader literature on contemplative practice and vagal modulation.⁹² Lehrer & Gevirtz's 2014 *Frontiers in Psychology* review of HRV biofeedback consolidates the clinical evidence that resonance-frequency breathing — typically near 0.1 Hz, or six cycles per minute — produces robust and reproducible amplification of RSA.⁹³ Kox, van Eijk, Zwaag and colleagues'

2014 *Proceedings of the National Academy of Sciences* study of the Wim Hof method demonstrated that voluntary breathing protocols can produce measurable activation of the sympathetic nervous system and attenuation of inflammatory responses to endotoxin challenge — establishing that volitional breathing modulates not only autonomic but also immunological state.⁹⁴

The convergence loop closes on hormonal and ageing-related modulators. Bayliss & Millhorn's 1992 *Journal of Applied Physiology* paper on central neural mechanisms of progesterone action established that gonadal hormones modulate respiratory drive at the brainstem level, with implications for sex differences in chemosensitivity and ventilatory response.⁹⁵ Standard endogenous stress paradigms — Kirschbaum, Pirke & Hellhammer's Trier Social Stress Test,⁹⁶ Buchheit, Peiffer, Abbiss & Laursen's cold-water immersion protocol,⁹⁷ Laukkanen and colleagues' systematic review of sauna exposure⁹⁸ — provide reproducible perturbations of the autonomic-chemoreflex-metabolic triangle that allow the convergence to be probed experimentally. At longer time scales, the metabolic-demand channel intersects with the broader biology of ageing: López-Otín, Blasco, Partridge, Serrano & Kroemer's *Cell* synthesis of the hallmarks of ageing,⁹⁹ Basu's review of DNA damage, mutagenesis, and cancer,¹⁰⁰ and Li, Pan, Wang and colleagues' *Circulation* analysis of healthy lifestyle factors and life expectancy in the US population¹⁰¹ place the convergence triangle within the broader systems-biology context developed in Chapter 5.

The takeaway is the structural argument. No other continuously accessible physiological signal integrates autonomic balance, chemoreflex sensing, metabolic demand, and a volitional channel at the temporal resolution of breathing itself. Heart rate integrates autonomic and metabolic input but is largely opaque to chemoreflex modulation and lacks a continuous voluntary access channel. Body temperature and blood pressure are slower variables, dominated by autonomic regulation, with no real-time chemoreflex coupling. Breathing is, mechanistically, the integration node. This is the convergence proposition.

3.3 Limbic and cortical interactions; the hybrid autonomic–behavioural nature of breathing

Respiration is the only vital sign that combines continuous autonomic regulation with continuous voluntary access. Heart rate, blood pressure, and core body temperature are autonomically regulated and effectively closed to voluntary modulation in the absence of biofeedback training extending over weeks; even then, voluntary modulation remains indirect, bounded, and context-dependent. Breathing operates simultaneously on two channels — an autonomic channel that maintains gas exchange across sleep, exercise, and metabolic demand without conscious attention, and a behavioural channel through which speech, breath-holds, sniffing, sighing, and contemplative practice continuously modulate the same signal. The hybrid is structurally consequential: a single observable carries both physiological state and behavioural self-regulation events.

Breathing is the only vital sign with continuous, direct voluntary access

VITAL SIGN	GENERATION	VOLUNTARY ACCESS	CHARACTERISATION
Heart rate	Autonomic	Indirect (biofeedback only)	Bounded modulation; no direct cortical-to-cardiac motor pathway
Blood pressure	Autonomic	Effectively none	No volitional motor channel
Body temperature	Autonomic	Effectively none	Behavioural modulation only (clothing, movement)
Breathing	Autonomic	Continuous direct	Cortical ↔ brainstem coupling allows volitional override

Key insight

Breathing's hybrid autonomic-voluntary architecture is unique among vital signs. Cortical pathways (M1 / SMA) project to brainstem respiratory premotor neurons, allowing both spontaneous rhythm and moment-to-moment volitional shaping (speech, breath-holding, paced breathing).

PRIMARY REFERENCES *Herrero et al. 2018 (cortical volitional control) · Tort et al. 2025 (cortical-respiratory coupling)*

Figure 3.3. Voluntary access across the four classical vital signs. All four vital signs are autonomically generated; among them, only breathing additionally admits continuous direct voluntary access, mediated by cortical projections from M1 and the supplementary motor area onto brainstem respiratory premotor neurons. This hybrid autonomic-behavioural architecture is the structural premise developed mechanistically through the remainder of Chapter 3.3

The neurobiology of this hybrid character has been clarified by a decade of intracranial-recording and in vivo electrophysiology work. Tort, Laplagne, Draguhn & Gonzalez's 2025 *Nature Reviews Neuroscience* synthesis demonstrates that respiration-coupled oscillations are not confined to olfactory or limbic regions but extend across the entire rodent and human cerebral cortex, modulating gamma-band activity (70–120 Hz) most strongly in frontal areas; the source of this coupling is nasal airflow through the olfactory epithelium, since tracheotomy abolishes the cortical respiratory rhythm.¹⁵ Zelano, Jiang, Zhou and colleagues' 2016 *Journal of Neuroscience* study established the human evidence: in epilepsy patients with intracranial electrodes implanted for clinical monitoring, nasal respiration entrained oscillatory activity in the piriform cortex, amygdala, and hippocampus, with amplitude peaking at inhalation and decaying during exhalation; the effect disappeared during oral breathing, and the phase of inhalation modulated both fear recognition (faster on inhale) and memory retrieval (more accurate on inhale).¹⁰² Karalis & Sirota's 2022 *Nature Communications* paper extended this finding into murine offline states, showing that breathing coordinates cortico-hippocampal dynamics during quiet wakefulness and sleep — periods previously assumed to be governed by intrinsic hippocampal rhythms alone.¹⁰³ Heck, Kozma & Kay's 2017 *Frontiers in Neural Circuits* synthesis frames the broader architecture: respiratory rhythm-coupled brain oscillations operate as a bridge between cognitive and emotional processing.¹⁰⁴ Yackle and colleagues' demonstration of the preBötC → locus coeruleus arousal channel⁶⁹ —

introduced in 4.1 — provides a complementary route by which brainstem respiratory rhythm propagates into ascending arousal circuits.

The volitional channel has been mapped at human cortical resolution. Herrero, Khuvis, Yeagle and colleagues, in the *Journal of Neurophysiology*, used intracranial EEG in patients with epilepsy to demonstrate that volitional breath-holds and attentional modulation of breathing are associated with characteristic cortical activity patterns extending well beyond brainstem rhythm-generating areas.¹⁰⁵ The hybrid character is not metaphorical: voluntary breathing engages identifiable cortical circuits, autonomic breathing engages brainstem oscillator networks, and the same physical observable — airflow at the airway opening — carries both signals continuously.

The behavioural-state coupling extends into emotion and affect. Bloch, Lemeignan & Aguilera's 1991 study in the *International Journal of Psychophysiology* demonstrated that specific respiratory patterns distinguish among basic emotions in trained actors and naive participants — different emotional states produced reproducibly different inspiratory and expiratory profiles, sigh frequencies, and breath-cycle durations.¹⁰⁶ Philippot, Chapelle & Blairy's 2002 *Cognition and Emotion* paper extended this finding into a feedback model: respiratory pattern not only reflects emotional state but is sufficient, in part, to generate it.¹⁰⁷ Kreibig's 2010 *Biological Psychology* review of autonomic nervous system activity in emotion synthesises the full evidence that respiration, alongside cardiac and electrodermal indices, carries emotional-specificity information.¹⁰⁸ Vlemincx, Taelman, De Peuter and colleagues, in *Psychophysiology*, established that sigh rate and respiratory variability during normal breathing are modulated by negative affect — sighs cluster around state transitions and shifts in emotional tone, providing event-level readouts of affective regulation embedded in continuous breathing.¹⁰⁹

The implication of the hybrid frame for an observation layer is direct. A signal that carries both autonomic state and voluntary behavioural events is, in principle, observable across a wider range of physiological and behavioural conditions than any signal carrying only one of the two. Speech, conversation, contemplative practice, breath-holds, sleep, exercise, and ordinary resting all leave structural traces in the same observable. Other vital signs are silent during behavioural acts that breathing actively encodes. The hybrid character is therefore not only a neurobiological curiosity — it is a structural argument for why breathing, among continuous physiological signals, is uniquely suited to carry an observation layer that bridges autonomic regulation and behavioural state. The mechanistic case for breathing as a convergent signal is now complete; the remaining question is whether this convergence is non-redundant with respect to existing measurement infrastructure. That question is the subject of 4.4.

3.4 Boundary conditions: where breathing complements rather than substitutes other signals

The structural case developed in 4.1, 4.2, and 4.3 establishes that breathing has a discrete neuroanatomical generator, integrates three otherwise-separable regulatory streams, and uniquely combines autonomic and voluntary control. None of this answers the most direct objection to breathing as the substrate of an observation layer: that its information content is already captured by heart rate variability, by aggregated multimodal sensors, by capnography, or by polysomnography. This subsection

addresses each of these candidate substitutes in turn, and closes on the formal information-theoretic statement that anchors the non-redundancy proposition.

Cardiorespiratory coupling: HRV is downstream of breathing, not parallel to it

The HRV-as-substitute objection is the strongest version of the redundancy challenge, since wearable HRV is the most widely deployed continuous autonomic-state signal in consumer health. The objection fails on directional grounds. Yasuma & Hayano's 2004 mechanistic synthesis⁷² and Wehrwein, Critchley, Yasuma, Hayano and colleagues' 2025 *Nature Reviews Cardiology* statement⁷³ together establish that respiration drives the dominant high-frequency component of HRV through respiratory sinus arrhythmia. Wehrwein et al.'s proposal to rename RSA as respiratory heart rate variability (RHRV) is precisely a terminological correction: the signal previously labelled "HRV" is, in mechanistic terms, a downstream readout of respiratory phase modulation of vagal outflow. This makes HRV, in principle, partially reconstructible from respiration; it does not make respiration reconstructible from HRV.

Three structural features of breathing are not reconstructible from HRV alone. First, the inhalation/exhalation ratio (I:E): Van Diest, Verstappen, Aubert and colleagues demonstrated that I:E modulates the HRV response to slow breathing in directionally specific ways, with low I:E (prolonged expiration, e.g., 1:3 or 1:4) shifting cardiovagal balance toward parasympathetic dominance and high I:E (approaching 1:1) associated with sympathetic activation and stress states.⁹¹ The directional relationship runs respiration → HRV, and HRV cannot disambiguate between, for example, two recordings with identical mean RR-interval and HF power but different underlying I:E phase structures. Second, the inhalation and exhalation airflow profile itself — the morphology of single breaths — is invisible to HRV measurement. Third, sigh frequency, breath-hold events, sniff bursts, and other event-level respiratory features leave no characteristic HRV signature.

Quantitative limits of HRV-as-substitute are well documented. Billman's 2013 *Frontiers in Physiology* analysis demonstrated that the LF/HF ratio, despite widespread use, does not accurately measure cardiac sympatho-vagal balance — the interpretive construct that motivates HRV monitoring is mathematically underdetermined by the spectral indices used to instantiate it.¹¹⁰ The European Society of Cardiology measurement-standards consensus codifies the methodological constraints under which HRV indices are interpretable,⁷⁷ and Beauchaine & Thayer's 2015 *International Journal of Psychophysiology* statement on HRV as a transdiagnostic biomarker frames HRV as one component of a broader autonomic-physiological assessment rather than as a stand-alone substitute for direct respiratory observation.¹¹¹ Wearable-grade HRV is further constrained by sensor and motion-artefact considerations: Charlton and colleagues' 2017 European Respiratory Society *Breathe* review documents the methodological boundaries of PPG-derived HRV in non-laboratory conditions,¹¹² and Dial and colleagues' 2025 *Physiological Reports* validation of consumer-grade HRV from Garmin, Oura, Polar, and WHOOP devices against the Polar H10 chest strap establishes a current empirical baseline.¹¹³ Even where wearable HRV is faithfully sampled, the projection remains downstream of respiration.

The exercise-physiology context anchors the same directional structure. Heart-rate recovery after exercise is a long-established mortality predictor (Cole et al. 1999¹¹⁴) and is vagally mediated, accelerated in athletes and blunted in chronic heart failure (Imai et al. 1994¹¹⁵); Plews et al. 2013 codified

ln(RMSSD) conventions for HRV in endurance training-load monitoring.¹¹⁶ Locomotor-respiratory coupling — established as a mammalian principle by Bramble & Carrier's 1983 *Science* paper,¹¹⁷ developed mechanistically by Daley, Bramble & Carrier in *PLoS ONE*,¹¹⁸ and demonstrated to develop with training by Mahler et al.'s rowing cohort¹¹⁹ — encodes training-state information directly in breathing dynamics that HRV indices alone do not access. At the lung-volume and gas-exchange side, Rasch-Halvorsen et al.'s 2019 HUNT-cohort analysis of dynamic lung volume and peak VO_2 ¹²⁰ and Levine's 2008 *Journal of Physiology* synthesis of VO_2 max¹²¹ bridge respiration directly to cardiorespiratory fitness — the same observable framed as a vital sign by the AHA scientific statement.⁹⁰ A clinically significant emerging case is Long COVID dysautonomia. Fedorowski, Fanciulli, Raj, Sheldon & Sutton's 2024 *Nature Reviews Cardiology* review *Cardiovascular autonomic dysfunction in post-COVID-19 syndrome: a major health-care burden* established that post-acute SARS-CoV-2 sequelae include reduced HRV, impaired vagal activity, and a high prevalence of postural orthostatic tachycardia syndrome — collectively framed as a major emerging health-care burden.¹²² Kranck, Ståhlberg, Andersson, Lundin & Fedorowski's 2025 *European Heart Journal — Case Reports* study extended this framework to wearable monitoring: in a 35-year-old patient with Long COVID, POTS, and chronic fatigue syndrome, 328 standardised smartwatch ECG recordings over six months revealed worsening RSA in the standing posture, with the joint cardiorespiratory desynchrony index tracking self-reported fatigue trajectories more sensitively than either component alone.¹²³ The case-report scale of the wearable demonstration is acknowledged; large-cohort validation of the joint respiration-HRV signature in Long COVID remains pending. The signal direction is consistent with the broader non-redundancy proposition — joint cardiorespiratory features carry information not present in either signal in isolation.

The integrated population-level argument follows from joining Jarczok et al.'s 2022 mortality meta-analysis³⁶ — which established a pooled hazard ratio of approximately 1.56 (95% CI 1.32–1.85) for low-HRV vs reference-HRV all-cause mortality across 32 studies and $n = 38,008$ — to Hoyer et al.'s demonstration³⁵ that joint cardiorespiratory cross mutual information outperforms HRV alone for post-MI risk stratification. HRV carries prognostic information; the joint respiration-HRV signal carries strictly more.

EEG and central nervous activity

Tort 2025¹⁵ and Zelano 2016¹⁰² establish that breathing entrains specific cortical rhythms — gamma in frontal regions, theta and delta in limbic regions — but the entrainment is partial and frequency-band-specific. Direct neural recording (EEG, MEG, intracranial EEG) accesses cortical activity that respiration does not modulate at all, particularly in the alpha and beta bands and in cortical regions outside the breathing-coupled network. The relationship is one of partial overlap: breathing carries rhythmic information that propagates into limbic and frontal cortex, but it does not substitute for direct cortical recording.

Capnography and gas exchange

Capnography measures alveolar gas composition — end-tidal CO_2 , capnogram morphology — at a chemical-composition level of description. Respiratory acoustics measures airflow structure — phase, asymmetry, transition dynamics — at a mechanical level of description. The two are partially overlapping

(both reflect ventilation and gas exchange) but not interchangeable: capnographic indices are insensitive to airflow morphology and to the upper-airway acoustic events (sigh, sniff, breath-hold, snore, cough) that respiratory acoustics encodes natively. Brinkman, Fens & Sterk's 2024 *American Journal of Respiratory and Critical Care Medicine* synthesis, *Fulfilling the promise of breathomics*, develops the broader case that exhaled-breath chemical analysis is complementary to, not a substitute for, structural respiratory observation.¹²⁴ Müller-Maatsch and colleagues' 2024 *Communications Biology* paper on non-volatile metabolome analysis of exhaled breath extends the breathomics frame into the molecular composition of expirate — a domain orthogonal to airflow structure.¹²⁵

Polysomnography and event-level sleep architecture

Polysomnography provides event-level sleep architecture through multimodal channels (EEG, EOG, EMG, ECG, airflow, oximetry, body position) sampled in clinical or home-PSG settings, typically in single-night recordings. Respiratory acoustics provides continuous structural observation across all states, sampled passively across days, weeks, and months. The temporal resolutions differ; the modalities differ; the contexts of capture differ. PSG accesses event-level labelling that respiratory acoustics does not (sleep-stage transitions, arousals, leg movements); respiratory acoustics accesses longitudinal structural drift that PSG does not (months-scale baseline shifts, ambient context variation, behavioural-state coverage). The two are complementary — Chapter 5 develops this complementarity at the empirical level for sleep-stage classification and apnoea detection.

Individual-level structural identity

A direct empirical demonstration of respiration's irreducible structural information was provided by Soroka and colleagues' 2025 *Current Biology* paper, *Humans have distinct nasal respiratory fingerprints*: in a Weizmann-led cohort of 97 participants, 24-hour nasal airflow recordings yielded individual-identification accuracy of 96.8% — comparable to fingerprint-level individual specificity, and substantially exceeding what HRV time series of equivalent length permit.¹²⁶ The implication is direct: respiratory structure carries an individual-level signature that is not present in aggregated multimodal sensor data at the same temporal resolution. This is structural irreducibility at the level of identity.

The non-redundancy statement

The convergent argument of 4.4 can be expressed as a conditional information-theoretic inequality:

$$I(\text{Respiration} ; \text{Physiology} \mid \text{HRV}, \text{Multimodal}) > 0$$

— respiration carries information about physiological state that is not contained in heart rate variability and aggregated multimodal sensor channels combined. The supporting evidence developed through this chapter consists of three convergent lines. First, respiration is upstream of HRV by mechanism, established at primary-literature depth by Eckberg 1983⁷⁰ and 2003⁷¹, synthesised by Yasuma & Hayano 2004⁷², and re-formalised by Wehrwein, Critchley, Yasuma, Hayano and colleagues' 2025 *Nature Reviews Cardiology* renaming statement⁷³. Second, the structural features of breathing — phase asymmetry, I:E ratio, airflow morphology, event-level acoustics — are not reconstructible from HRV

alone, established by Van Diest 2014⁹¹ and the broader RSA-mechanism literature. Third, respiration carries individual-level structural identity not present in aggregate multimodal data at equivalent temporal resolution, established by Soroka 2025¹²⁶. The finding is anchored in the broader frameworks of fractal physiology developed by Goldberger, Peng and colleagues¹²⁷ and the early-warning-signal architecture for critical transitions developed by Scheffer and colleagues¹²⁸ and applied to chronic disease by Trefois and colleagues¹²⁹ — frameworks that re-enter this monograph substantively in Chapter 4 (formalisation of Breathing State) and Chapter 6 (temporal advantage).

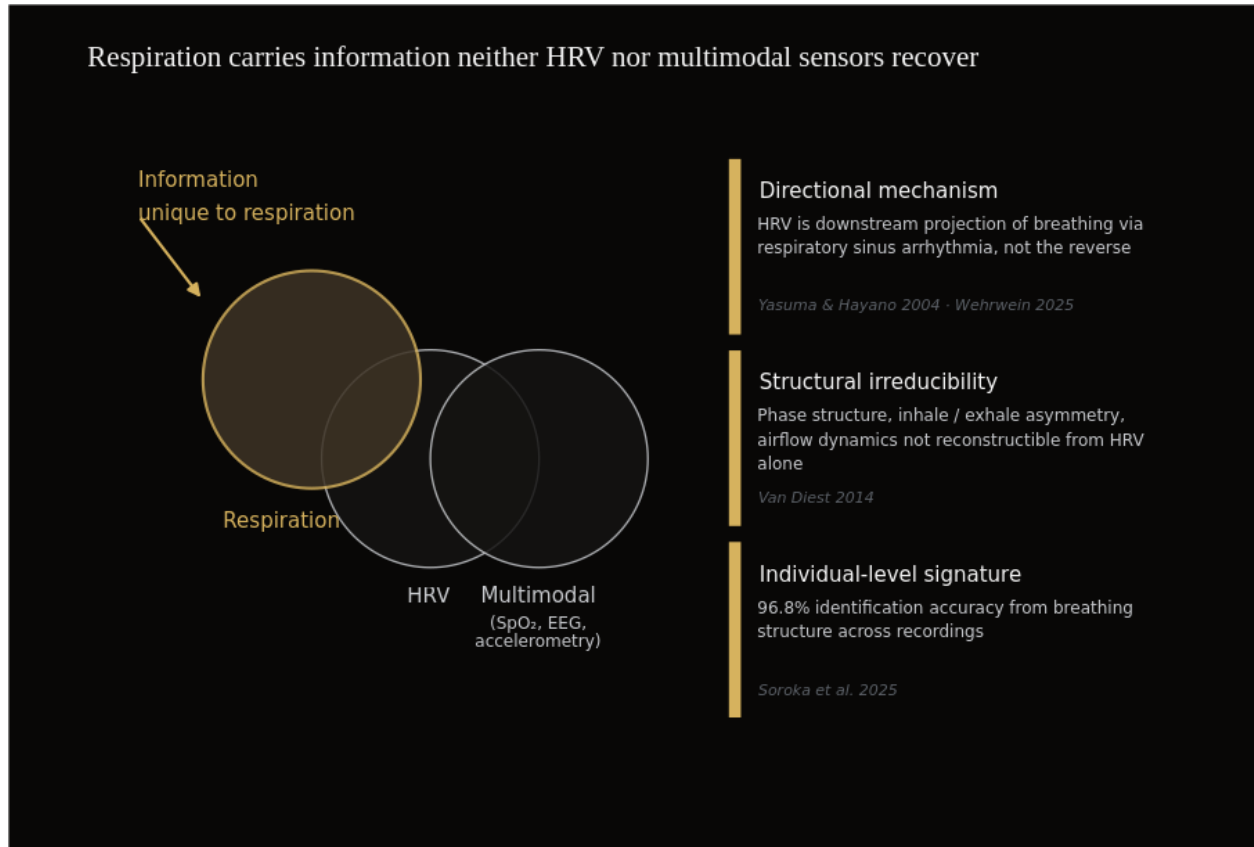


Figure 3.4. Respiration carries information neither HRV nor multimodal sensors recover. Three independent lines of evidence support the structural non-redundancy of respiration relative to derived cardiac and aggregated wearable signals: (i) the directional asymmetry of cardiorespiratory coupling, where respiration drives heart-rate modulation through respiratory sinus arrhythmia rather than the reverse, making HRV a downstream projection of respiration [72, 73]; (ii) the structural irreducibility of phase, inhale–exhale asymmetry, and intra-cycle airflow dynamics, which are compressed away when respiration is observed only through HRV-derived scalars [91]; and (iii) the individual-level distinctiveness of breathing structure, with 96.8% subject identification accuracy from 24 nasal-airflow parameters, stable across recordings separated by months to two years [126, n = 97]. Formally, the figure depicts the condition $I(\text{Respiration}; \text{Physiology} \mid \text{HRV}, \text{Multimodal}) > 0$ — a falsifiable statement examined in Chapter 3.4.

The non-redundancy proposition is presented here as a falsifiable hypothesis with three convergent supporting lines, not as an established quantitative result. The strict information-theoretic inequality is empirically testable via head-to-head comparison studies in synchronised respiration + HRV + multimodal cohorts; the present chapter establishes the mechanistic basis for the proposition, not its definitive empirical confirmation. This is the appropriate observation-layer framing: respiration is associated with physiological state at a level of structural detail not captured by existing continuous-

monitoring signals, and the boundary of this association is itself a research programme — one that Chapter 11 (open questions and methodological frontiers) develops at the level of experimental design.

The chapter has established, through the convergent argument of 4.1–4.4, that breathing is structurally suited to be the substrate of a continuous physiological observation layer. Respiration has a discrete brainstem rhythm generator with mapped chemosensory inputs and identified projection targets; it integrates autonomic balance, chemoreflex coupling, and metabolic demand at a single observable; it is the only vital sign with continuous voluntary access alongside autonomic regulation; and the information it carries about physiological state is not reducible to HRV, multimodal aggregation, capnography, or polysomnography taken individually. The next chapter formalises *what* this signal becomes when treated not as an instantaneous quantity but as a temporal object — Breathing State as the formal observable on which the architecture of an observation layer is built.

4. Breathing State: from signal to temporal object

The preceding chapter established that breathing is a convergent physiological signal — that respiratory rhythm is generated by an identifiable brainstem substrate and is subsequently modulated by autonomic, metabolic, limbic, and cortical inputs in continuous superposition. This chapter takes the next step. It moves from *what breathing is, biologically* to *what breathing is, formally* — from a regulated rhythm produced by neural circuits to a temporal object that can be defined, computed, and validated independently of any specific measurement modality. The argument that follows is methodological rather than mechanistic. Its goal is to establish that the structure of breathing over time — not the mean rate at which it occurs — is a well-defined formal object, that this object loses essential information when collapsed onto scalar summaries, and that an extensive body of canonical literature in nonlinear dynamics, time-series statistics, and acoustic foundation modelling already provides the analytical instruments needed to instantiate it. The exact computational stack used to instantiate the state in any specific application is an engineering detail beyond the scope of this synthesis; the synthesis itself is the contribution.

4.1 Definition: Breathing State as a temporal object

A *Breathing State* St at time t is defined as a function of the current observed signal Xt and the prior state St_{-1} :

$$St = f(Xt, St_{-1})$$

This is a state-space formulation in the canonical sense developed across half a century of statistical and dynamical systems literature^{130, 131, 132}. The signal Xt is the raw breathing observation in any modality that preserves temporal phase — most directly an acoustic recording from a smartphone microphone, but the formulation is modality-agnostic. The function f is some computable transformation that combines current observation with prior state to yield a representation that can be carried forward in

time. The state St has three properties that together justify treating it as a temporal object: continuity, computational existence, and structural identity. Each property is empirically anchored.

The state-space form is not a metaphor borrowed from engineering. It is the standard formalism used to describe physiological systems whose current condition cannot be fully read off any single instantaneous measurement and whose evolution is partially constrained by recent history. Bayesian filtering and smoothing, sequential Monte Carlo methods, structured state-space models, and Kalman-class approaches all share the same underlying object — a latent variable St that propagates through time according to a transition law and emits observations Xt according to an observation law^{132, 130, 131}. The substantive argument made here is not that this formalism is novel — it is decades old — but that breathing is one of the physiological signals to which it most cleanly applies, and that respiratory observation has been slow to take advantage of it.

Continuity

A Breathing State must be defined for all t within the observation window, not only during forced respiratory maneuvers, voluntary breath-holds, or supervised laboratory protocols. This is the property that distinguishes structural breathing observation from the classical snapshot tradition of pulmonary function testing. Spontaneous breathing at rest is not a homogeneous baseline; it carries persistent breath-to-breath variability that has been documented since Priban's 1963 analysis of short-term patterns in healthy adults¹³³ and later quantified in a large body of variability literature in clinical physiology^{134, 135}. Modern reviews of breathing variability in anaesthesiology and intensive care confirm that resting breathing is characterised by structured fluctuations in cycle duration, tidal excursion, and inhale–exhale asymmetry — fluctuations that carry prognostic information even in the absence of overt respiratory disease^{83, 136}. The classical respiratory phenotyping that emerged from Hutchinson-era spirometry compresses these fluctuations into derived means and forced-maneuver volumes; what is lost is precisely the temporal structure that the *continuity* property requires. Continuity is therefore not a mathematical convenience but an empirical finding: there is structure to be observed at every t , and the structure changes with state. Reviews of the temporal-variability tradition spanning six decades make this point explicitly, characterising spontaneous breathing as exhibiting endogenous variability of comparable magnitude to heart-rate variability and as carrying information that is invisible to maneuver-based testing^{135, 136}.

Computational existence

A Breathing State must be extractable from the observed signal via a defined transformation — the function f in the state-update equation must be specifiable, reproducible, and validatable. This is the property that converts a physiological intuition into a formal object. The canonical instruments for instantiating f on physiological time series have been available for more than three decades. Pincus introduced approximate entropy (ApEn) in 1991 in the *Proceedings of the National Academy of Sciences* as a measure of system complexity defined for short, noisy clinical recordings¹⁶. Richman and Moorman extended ApEn to sample entropy (SampEn) in 2000 to remove the self-counting bias that limited the original definition for the recording lengths typical of physiological monitoring¹⁷. Costa, Goldberger, and Peng generalised these instruments to multiple time scales via multiscale entropy in 2002, addressing

the limitation that single-scale entropy underestimates complexity in signals that exhibit structured behaviour at scales longer than the immediate inter-event interval¹³⁷. These instruments are parameterised — the embedding dimension m , the tolerance r , and the recording length N together define the operating point — and the parameter selection itself is part of the methodological apparatus that has been established across the variability literature. None of these instruments was developed specifically for breathing — they were developed for heart rate variability, for EEG, for any physiological time series that exhibits nonlinear regulation — and that generality is precisely what makes them suitable. The *computational existence* of a Breathing State does not depend on a bespoke breathing-specific transformation; it depends on the application of well-validated instruments from nonlinear dynamics to a signal that is itself a continuous time series. Cross-domain validation in the heart-rate variability literature provides a built-in credibility argument: the same instruments that have been shown to discriminate physiological states across decades of cardiovascular research are directly applicable to respiratory time series^{136, 127, 83}.

Structural identity

A Breathing State must carry a person-specific signature that is distinguishable from other persons across recordings — and, critically, this signature must be stable across time at scales relevant to longitudinal observation. The empirical anchor here is recent and unusually clean. Soroka and colleagues at the Weizmann Institute, working with a custom 24-hour nasal airflow recorder, demonstrated that 24 parameters extracted from spontaneous nasal breathing in 97 healthy adults were sufficient to identify individuals with 96.8% accuracy¹²⁶. The parameters spanned phase structure (inspiratory and expiratory durations, their ratio), cycle-level variability (coefficient of variation of cycle period), thoraco-abdominal coupling, sniff frequency, pause duration and variability, peak-flow timing, and inter-nostril airflow asymmetry. The identification accuracy was preserved at re-test intervals of weeks, months, and approximately two years; test-retest accuracy at the longest interval remained 95.24%. The breathing fingerprint was furthermore associated with body-mass index, sleep–wake regularity, and validated measures of anxiety and depressive symptoms — the structural signature is not a noise artifact but is co-varying with characterised individual phenotypes. This finding is convergent with the older acoustic biometric literature, which has demonstrated equal-error rates of 1.4% for breathing-based verification across multimodal acoustic capture¹³⁸ and >97% true-confirmation rates for exhaled breath physics applied to user verification in independent cohorts¹³⁹. Inter-cycle interval analyses of breathing sounds have likewise shown that individual signatures are extractable from purely acoustic capture in unconstrained conditions¹⁴⁰. Together these results establish that *structural identity* is not a theoretical desideratum but an empirically observed property of the signal itself.

The mechanistic plausibility of structural identity is straightforward. Each individual carries an idiosyncratic combination of upper-airway anatomy, lower-airway compliance, neural pattern-generator set-points, and autonomic balance — and these combine to produce breathing that differs systematically across persons in ways that are measurable but not reducible to any single dimension. This is consistent with the Goldberger research programme's broader finding that the complexity of physiological regulation is itself a fingerprint, varying across individuals in ways that are stable when the individual's regulatory state is stable and that change in characteristic ways under stress, disease, or

intervention¹²⁷. The Soroka finding is therefore not anomalous; it is a particularly clean instance of a more general property of physiological signals.

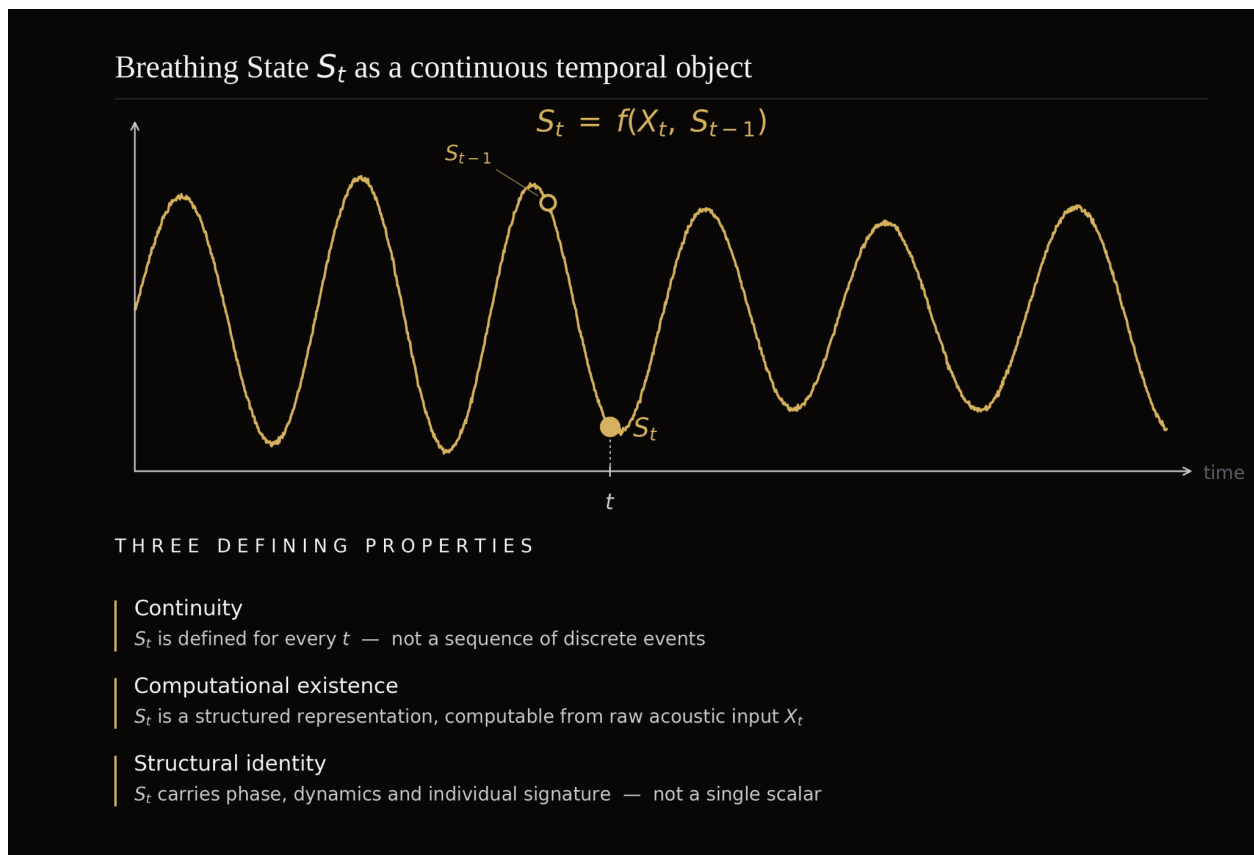


Figure 4.1. Breathing State as a continuous temporal object. S_t , a structured representation of respiratory dynamics, is recurrently updated as $S_t = f(X_t, S_{t-1})$, where X_t are observed features in window t . Three defining properties: continuity, computational existence, structural identity. Atum observes S_t ; it does not diagnose or assign clinical scores. Schematic; curve is illustrative, not measured data.

The three properties together yield a formal definition that does not depend on any particular implementation. A Breathing State is a temporal object because it is defined at every t , computable from the signal, and carries information that distinguishes one observed person from another. It is not a derived metric like mean respiratory rate. It is not a physiological construct like minute ventilation. It is the underlying structural object of which scalar metrics are projections — a point developed formally in 5.2.

4.2 Why structural representation supersedes scalar metrics

The standard approach to respiratory observation, both clinically and in commercial wearables, treats breathing as a scalar measurement problem: report a respiratory rate, optionally report tidal volume, occasionally report minute ventilation, and treat the resulting time series as a thin one-dimensional trace. The reasoning behind this approach is operational, not informational — scalar metrics fit cleanly into electronic health records, into early-warning scores, and into wearable summary dashboards. The cost of the approach, however, is substantial, and it can be stated formally: every scalar respiratory metric is

a *projection* of the underlying Breathing State onto a single axis, and every such projection discards quantifiable structure.

The framing of representation versus projection is central to modern machine learning, where the canonical statement is given in Bengio, Courville, and Vincent's 2013 review on representation learning²³. In their formulation, a representation is good to the extent that it preserves the explanatory factors of the data; a representation is bad to the extent that it conflates them. The same logic applies to respiratory observation. The explanatory factors that produce a given breathing recording include autonomic balance, brainstem rhythm-generator output, upper-airway resistance, lower-airway compliance, behavioural state, emotional state, recent metabolic demand, and individual anatomy — each of which has been documented in earlier chapters as separately identifiable through different aspects of breathing structure. A respiratory rate of 16 breaths per minute is a single number. The Breathing State that produced that rate could be regular and metronomic, could be irregular with sustained inhale–exhale asymmetry, could contain frequent micro-arousals, could exhibit slow drift across the recording, or could be the result of any combination of these — and the scalar rate cannot distinguish among them. The lost dimensions are not abstractions. They are well-characterised physiological features that have been documented to carry prognostic information.

Phase structure is the first dimension lost. The ratio of inspiration to expiration is a textbook autonomic indicator: low I:E ratios (prolonged expiration, on the order of 1:3 or 1:4) reflect obstructive or parasympathetic-dominant patterns; high I:E ratios (approaching 1:1) reflect stress, anxiety, restrictive patterns, or sympathetic activation⁸³. The asymmetry between inhale and exhale phases is invisible to mean-rate reporting. So is the secondary structure within each phase — the time to peak inspiratory flow, the relative duration of pause between exhale and the next inhale, the consistency of these intervals across cycles. Sniff frequency, sigh distribution, and the irregularity of pause duration are all phase-level features that have been characterised in the breathing-fingerprint literature as individually informative and stable across recordings¹²⁶. Inter-cycle interval analysis of breathing sounds has shown that phase-level structure is extractable directly from acoustic capture and varies systematically across individuals¹⁴⁰. The standardisation of computerised respiratory sound analysis under the CORSA framework formalised this acoustic vocabulary three decades ago^{141, 142, 143}; the structural information has been documented and accessible since then.

Nonlinear dynamics is the second dimension lost. Breath-to-breath variability is not white noise. It carries long-range temporal correlations, multifractal scaling, and structured complexity — properties that have been documented across the variability literature for two decades^{136, 83, 135}. Reduced breathing variability is itself a clinical signal. In mechanical ventilation weaning, reduced breath-to-breath variability predicts unsuccessful separation from the ventilator¹²³; the converse — preserved variability — predicts successful weaning and shortened ventilator dwell time⁷⁰. In asthma, fluctuation analysis applied to airway function predicts severe exacerbations weeks in advance — Frey and colleagues' 2005 *Nature* paper established this prediction window, demonstrating that the fluctuation structure of peak expiratory flow over a four-week observation period carried prognostic information for severe episodes that scalar mean values did not⁷¹. In hospitalised patients transferred to intensive care, respiratory rate variability carries prognostic information independent of mean rate¹⁴⁴. The nonlinear-dynamic properties

of breathing are not laboratory curiosities; they are documented prognostic markers, and they are entirely absent from any scalar summary.

The information loss is not a matter of degree; it is structural. A scalar projection of a high-dimensional state cannot, in principle, be inverted to recover the joint distribution of the dimensions that were collapsed. This is a basic result of linear algebra and applies regardless of how clever the downstream interpretation might be. Two recordings that yield the same mean respiratory rate but differ in phase structure, variability, and individuality signatures will produce identical scalar reports — and downstream inference that operates on those reports has no access to the differences. What this implies for respiratory observation is that any monitoring framework that treats breathing exclusively in terms of scalar metrics has, by construction, discarded the structural object that carries the multi-system regulatory information described in Chapter 3. The framework may still be useful — a mean rate in the high 20s with an oxygen saturation of 88% remains clinically actionable — but it is not a complete observation of the signal that produced it. The point is not that scalar metrics are wrong; it is that they are radically incomplete relative to what the signal contains.

The non-invertibility argument can be sharpened. Information-theoretically, the residual information that Breathing State carries about regulatory status — beyond what is captured by a scalar projection of breathing alongside the standard multimodal vital-sign panel — is provably non-negative, and the empirical evidence reviewed in Chapter 3 and across the variability literature is consistent with it being substantively positive. Even after conditioning on heart-rate variability, oxygen saturation, mean respiratory rate, and the rest of the standard vital-sign aggregate, the structural object S_t continues to carry information that those other channels do not capture. This proposition is developed formally in the non-redundancy framework and is consistent with the observation in Chapter 3 that breathing is a hybrid autonomic-behavioural signal whose phase structure, inhale-exhale asymmetry, and airflow dynamics are not reconstructible from any combination of HRV-derived metrics. The implication is that scalar respiratory reporting is not a compression of structural information; it is a discard. The mechanical-ventilation weaning literature provides the cleanest controlled comparison of this finding. Patients ready for ventilator separation exhibit preserved breath-to-breath variability; patients who are not exhibit reduced variability prior to spontaneous-breathing-trial failure, and the prognostic difference holds even when mean respiratory rate is similar between the two groups^{70, 123}. The historical literature on respiratory variability and associated cardiovascular changes in adults at rest had foreshadowed this finding decades earlier¹³⁴, but it was not until variability-as-biomarker frameworks matured in the 2000s that prognostic findings based on the structural object became operational.

Individuality signatures are the third dimension lost. As 5.1 established, a Breathing State carries a person-specific structural fingerprint that is stable across months and years and is sufficient to identify individuals with high accuracy¹²⁶. The fingerprint is compositional — it lives in the joint distribution of phase parameters, cycle-level variability, sniff frequency, pause structure, and amplitude modulation, not in any single one of these. The acoustic biometric verification literature has demonstrated this in independent cohorts: the academic phrase *breathing signature* in this context refers specifically to the multimodal acoustic representation that supports verification at equal-error rates of approximately 1.4%¹³⁸, and to the exhaled-breath-physics signatures that support user verification at >97% true-

confirmation rates in cohorts of $n=94$ ¹³⁹. None of these results is recoverable from mean rate alone. A scalar projection erases the individual.

Why, then, does the scalar tradition persist? The answer is operational. Mean respiratory rate fits in a single column of an electronic health record; it is amenable to threshold-based alerting; it composes cleanly with other vital signs into early-warning aggregates; and it is reproducible across measurement modalities — chest-belt, camera-based, radar-based, and acoustic instruments will all converge on a similar mean value for a given recording. None of these operational properties makes the scalar a *good* representation; they make it a *convenient* representation. The convenience trade-off has been visible in the failure modes of remote monitoring trials in heart failure, where population-threshold scalar approaches did not yield outcome improvements that personalised structural-baseline approaches subsequently achieved — the empirical detail belongs to Chapter 7, but the underlying methodological point is 5.2's: scalar reduction of a structurally rich signal is the proximate mechanism for the failure. The implication for monitoring system design is that scalar reporting and structural representation are not substitutes. They are different layers, with different roles, and a system that intends to capture the multi-system regulatory information that breathing carries cannot rely on the scalar layer alone.

This connects directly to the structural argument that Chapter 7 develops empirically. Personalised baselines are well-defined only over a structural representation — there is little to personalise about a single scalar threshold beyond moving the cutoff up or down, and the major remote-monitoring trials in heart failure that operated on this kind of population-threshold scalar logic did not yield outcome improvements. Trials that operated on individualised structural baselines did. The methodological precondition for the empirical record reviewed in Chapter 7 is therefore the structural-representation argument made here: structural representation is what makes meaningful personalisation possible, and personalisation is what makes deviation detection actionable. The argument runs in only one direction. A monitoring framework that begins from scalar reporting cannot retrofit personalisation; it must begin from a representation rich enough to support it.

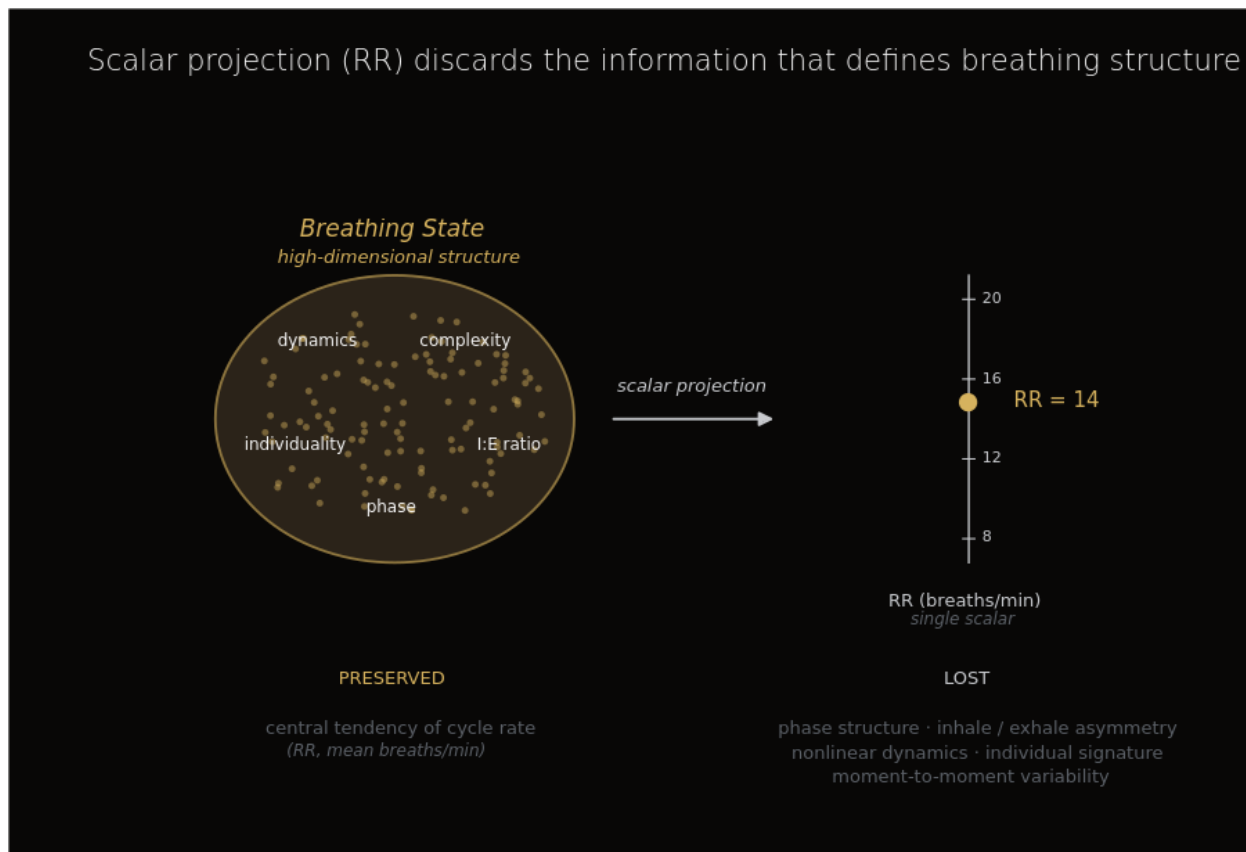


Figure 4.2. Information loss from scalar projection. Every scalar respiratory metric — mean RR, tidal volume, minute ventilation — is a projection of the high-dimensional Breathing State S_t onto a single axis, recovering only central tendency. Phase structure, inhale-exhale asymmetry, nonlinear dynamics, individual signature, and moment-to-moment variability are discarded by construction. The reduction is structural and non-invertible: two recordings with RR = 14 bpm can differ in every dimension carrying prognostic information (Chapter 3).

The validation framework for treating Breathing State as a structural object rather than a scalar projection is developed elsewhere in this document as a project-internal representation-validity framework with five requirements (R1–R5). The detailed enumeration is not reproduced here because the detailed tests are validation criteria for any specific implementation, not academic-monograph content. The framework anchors a single principle: a representation that meets representation-validity requirements is sufficient to support cross-individual generalisation, longitudinal stability, and the multi-system co-variation characterisation developed in subsequent chapters. A representation that does not meet them — including any scalar projection — is not. The requirements are independent of any specific implementation; they are evaluation criteria that any candidate Breathing State representation must satisfy to be admissible.

Standardisation efforts in respiratory acoustics provide an independent line of support for the same conclusion. The European Respiratory Society Task Force on lung sound nomenclature¹⁴⁵ and the CORSA standardisation framework^{141, 142} together established that respiratory sounds carry structured, well-defined acoustic content that requires structural rather than scalar characterisation to be analysed reproducibly. These efforts predate the modern foundation-model literature by two decades. Their continued relevance is itself an argument: respiratory acoustic content is not amenable to scalar

reduction without information loss, and the field has known this for thirty years. The contribution of contemporary work — including the foundation models discussed in 5.3 — is not to replace structural representation with something simpler. It is to make structural representation tractable at scale.

The transition from 5.2 to 5.3 is therefore from *what is at stake* to *how it is computed*. The remainder of this chapter addresses the second question.

4.3 Families of analytical approaches

The instruments needed to compute and validate a Breathing State are distributed across six families of analytical approaches, each anchored by canonical foundational papers. No family is sufficient on its own. Their convergence on the same physiological state from different mathematical perspectives is what establishes Breathing State as a robust formal object — and what justifies the methodological taxonomy presented here. The families are described at category level only. The exact computational stack used to instantiate S_t in any specific application — the specific weighting, ordering, and integration of these families — is an engineering detail beyond the scope of this synthesis.

Family 1 — Entropy-based complexity measures

The first family quantifies the regularity, predictability, or complexity of a physiological time series, on the principle that healthy regulation produces moderate complexity and that pathology often manifests as complexity loss. The canonical instruments are approximate entropy (ApEn), introduced by Pincus in 1991 in the *Proceedings of the National Academy of Sciences*¹⁶; sample entropy (SampEn), introduced by Richman and Moorman in 2000 to correct the self-counting bias of ApEn for short physiological recordings¹⁷; and multiscale entropy (MSE), introduced by Costa, Goldberger, and Peng in 2002 to extend the entropy framework across multiple temporal scales¹³⁷. The instruments are parameterised — the embedding dimension m , the tolerance r (typically a fixed fraction of the signal's standard deviation), and the recording length N — and the parameter-selection literature has converged on stable defaults for clinical recordings of the lengths typically achievable in continuous monitoring. The core empirical finding of this family is the *complexity-loss-with-disease-and-aging* hypothesis developed across the Goldberger research program, which has been documented across heart-rate variability, gait dynamics, and respiratory variability over more than a thousand publications since the 1990s¹²⁷. The family applies directly to respiratory time series; reduced sample entropy of respiratory variability has been documented in mechanical ventilation weaning failure¹²³, in asthma instability⁷¹, and in critical-care variability cohorts^{136, 83}. The R-package ecosystem has consolidated the family into reproducible open tooling, including the RespirAnalyzer package which implements multiscale entropy on respiratory signals as a reference open-source benchmark¹⁴⁶.

Family 2 — Scaling exponents and detrended fluctuation analysis

The second family quantifies long-range temporal correlations and self-similarity in physiological time series. Detrended fluctuation analysis (DFA) was introduced by Peng and colleagues in *Chaos* in 1995 as a method for extracting scaling exponents from nonstationary heartbeat time series¹⁹. The same framework was subsequently developed across the broader Goldberger corpus of fractal dynamics in physiology — the canonical synthesis being Goldberger and colleagues' 2002 *PNAS* paper on

alterations of fractal dynamics with disease and aging¹²⁷. The Scheffer 2009 *Nature* paper on early-warning signals for critical transitions extended the scaling-exponent framework into a universal pre-symptomatic signal — critical slowing down, the loss of resilience that precedes regime shifts in dynamical systems¹²⁸. The relevance of critical slowing down to respiratory observation is direct: Frey and colleagues' 2005 *Nature* demonstration that fluctuation analysis applied to airway function predicts severe asthma episodes weeks in advance⁷¹ is, mechanistically, an instance of the same family of phenomena — a measurable loss of resilience in a regulated system before it transitions into a clinically apparent state. The canonical multifractal extension applied specifically to physiological data is associated with Ihlen 2012 in *Frontiers in Physiology*¹⁴⁷. Within respiratory observation, the family supports two distinct propositions: that the breathing signal exhibits long-range correlations characterised by stable scaling exponents in healthy regulation, and that *deviations* in scaling structure carry prognostic information for transitions toward pathological states.

Family 3 — Multifractal characterisation

The third family is a specific extension of the scaling-exponent family. Where DFA and similar methods extract a single scaling exponent under the assumption of monofractal scaling, multifractal methods relax this assumption and characterise the full spectrum of scaling exponents present in a signal. The canonical multifractal detrended fluctuation analysis (MF-DFA) framework was introduced by Kantelhardt and colleagues in *Physica A* in 2002¹⁴⁸. The methodological application to physiological data is anchored on the Ihlen 2012 *Frontiers in Physiology* tutorial in MATLAB¹⁴⁷. The width of the singularity spectrum produced by multifractal analysis quantifies the degree of scale-dependent regulatory diversity in a signal — broad spectra indicate that different scales of variability are governed by different scaling laws, narrow spectra indicate scale-uniform behaviour. The empirical relevance of multifractal characterisation for respiratory observation is that healthy breathing variability is multifractal — different time scales exhibit different scaling — while disease states often show monofractal collapse, the loss of scale-dependent regulation. The Pham 2024 RespirAnalyzer R package has consolidated multifractal DFA implementation alongside multiscale entropy, providing a reference open-source toolchain for both Family 1 and Family 3 instruments¹⁴⁶.

Family 4 — Dynamical systems methods and recurrence quantification

The fourth family characterises the trajectory of the signal in reconstructed phase space, on the dynamical-systems principle that any time series can be embedded in a higher-dimensional state space whose recurrence structure encodes the underlying attractor. The canonical instrument is recurrence quantification analysis (RQA), introduced for physiological systems by Webber and Zbilut in *Journal of Applied Physiology* in 1994¹⁴⁹. The comprehensive methodological synthesis of recurrence plots and RQA across physics and physiology is Marwan 2007 in *Physics Reports*¹⁵⁰. RQA produces a panel of state-characterising statistics: determinism (the fraction of recurrent points that form diagonal lines, indicating predictable trajectory segments), laminarity (the fraction forming vertical lines, indicating intermittent slowing), trapping time (the average length of laminar episodes), and Lyapunov-spectrum approximations from divergence rates — each of which captures a distinct aspect of the system's dynamical regime. RQA is particularly suited to respiratory observation because it captures recurrence of *patterns* of breathing — the return to similar dynamical configurations across time — rather than

recurrence of specific values; this property makes it more robust to slow drift and amplitude scaling than scalar variability measures. Determinism, laminarity, and trapping-time measures derived from recurrence plots have been applied across HRV, EEG, and respiratory signals as state-characterising statistics. Adjacent dynamical-systems frameworks — Koopman operator approaches that linearise nonlinear dynamics in physiological signals¹⁵¹ — have entered the literature more recently and provide a complementary perspective on the same family. The methodological foundations are continuous with the broader complexity literature in physiology^{127, 136}.

Family 5 — Topological data analysis

The fifth family is the most recent. Topological data analysis (TDA) — and persistent homology specifically — applies algebraic-topology instruments to time-series data, characterising the *shape* of the data in a sense that is robust to noise and parametrisation choices. The application of TDA to physiological time series is recent; one application is to sleep staging via airflow signals¹⁵². The methodological maturity of TDA for respiratory observation lags behind the entropy and scaling families by approximately three decades, and the empirical body is correspondingly smaller. The family is included here because it provides a class of structural descriptors that are invariant to certain classes of confound — including amplitude scaling and slow drift — that confound scalar metrics. Manifold-learning instruments adjacent to TDA in the data-shape characterisation tradition — UMAP for dimension reduction¹⁵³ and PHATE for visualising structure and transitions in high-dimensional biological data¹⁵⁴ — are part of the broader toolchain that operates on the structural representation produced by Family 1 through Family 4 instruments.

Family 6 — Foundation models for time series and respiratory acoustics

The sixth family is the most recent in the literature and the most dynamic at the time of writing. Foundation models for time series and respiratory acoustics are large pre-trained neural networks that encode physiological signals into learned representations on which downstream tasks — classification, anomaly detection, signal completion — can be built with reduced data requirements per task. The category includes HeAR (Health Acoustic Representations), released by Google Health in 2024 as a respiratory-and-cough acoustic encoder by Baur, Nabulsi, Weng and colleagues²⁴, with field deployment reported in tuberculosis screening contexts²⁵; OPERA, an open benchmark and acoustic foundation model framework for respiratory sounds released by Zhang and colleagues at NeurIPS 2024²⁶; and the broader category of universal time-series foundation models — Chronos by Amazon Science²², TimesFM by Google Research²¹, Moirai by Salesforce AI¹⁵⁵ — that pre-train on heterogeneous time-series corpora and provide transferable representations to physiological domains. Structured state-space models, including S4 introduced by Gu, Goel, and Ré in 2022, represent a parallel modelling tradition optimised for long-sequence efficiency²⁰.

These foundation models are named at category level only. They are introduced here to establish that the analytical infrastructure for representation learning on respiratory time series is now mature, that pre-trained acoustic encoders for cough and breathing exist and have been deployed in field settings, and that the methodological gap between the canonical complexity literature of the 1990s and the foundation-model literature of the 2020s has narrowed substantially. The integration of foundation-

model encoders with classical complexity instruments — entropy, scaling, multifractal, recurrence — into a unified pipeline is itself a non-trivial methodological challenge: learned representations and classical complexity measures operate on different mathematical principles and emit different kinds of evidence about the underlying state, and their combination requires careful attention to the failure modes of each. The specific architecture, training objectives, embedding dimensions, and integration choices that constitute any particular implementation are not described in this synthesis, because they are not part of the synthesis. They are part of an implementation.

The respiratory acoustics literature, as a parallel methodological tradition, provides an independent body of validation. The standardisation of computerised respiratory sound analysis in the CORSA framework^{141, 142, 143} and the European Respiratory Society Task Force nomenclature standardisation¹⁴⁵ together established the acoustic vocabulary on which contemporary deep-learning approaches operate. Foundational lung-auscultation reviews, including Bohadana and colleagues' *New England Journal of Medicine* synthesis, formalise the clinical taxonomy of lung sounds against which automated systems are evaluated¹⁵⁶. Recent work has applied modern neural architectures — convolutional, recurrent, transformer-based — to respiratory sound classification on the ICBHI 2017 challenge dataset¹⁵⁷, with reported accuracies for normal/wheeze/crackle classification reaching the high 80s and into the 90s across multiple architectures^{158, 159, 160, 161}. Comprehensive reviews of AI techniques applied to respiratory sound classification document the field's progression from hand-crafted feature engineering toward end-to-end deep learning, and identify the remaining methodological gaps including domain shift, label noise, and limited longitudinal data¹⁶². The Pasterkamp lung-auscultation tradition¹⁵⁶, the formant-and-energy frameworks for spectral characterisation¹⁶³, and the broader audio-pattern-recognition ecosystem represented by PANNs¹⁶⁴ together support the proposition that respiratory acoustics is a tractable signal domain for structural analysis. Smartphone-grade self-auscultation has been demonstrated in the iMedic system with explicit cross-device generalisation through MixStyle domain adaptation¹⁶⁵, and explainable-AI analyses have begun to characterise which acoustic features carry classification signal¹⁶⁶. Cough-classification work using VGGish transfer learning and Grad-CAM interpretability has demonstrated that pre-trained audio embeddings transfer effectively to respiratory acoustic events with reported AUC of approximately 0.99 for cough detection and 0.84–0.87 for cough-type classification¹⁶⁷. Adjacent acoustic-event recognition for swallowing — a signal class that overlaps with breathing in upper-airway acoustic capture — has been demonstrated in high-resolution cervical recordings via deep learning, indicating that the acoustic classification toolchain extends across respiratory and adjacent oropharyngeal signal classes¹⁶⁸. The HRV standards from the joint Task Force of the European Society of Cardiology and the North American Society of Pacing and Electrophysiology¹⁶⁹ remain the methodological reference point for any time-series characterisation that crosses into cardiorespiratory coupling. Speech-enhancement frameworks, including conformer-based metric GAN approaches to denoising¹⁷⁰ and the JMIR AI 2025 work on robustness and clinical applicability of automatic respiratory sound classification with audio enhancement¹⁷¹, address the practical signal-quality conditions under which any of these analytical families must operate. Sleep staging from electrocardiography and respiration with deep learning¹⁷² establishes that respiratory time-series structure alone carries information sufficient for non-trivial state classification — a direct empirical demonstration that scalar metrics are not the limit of what the signal supports.

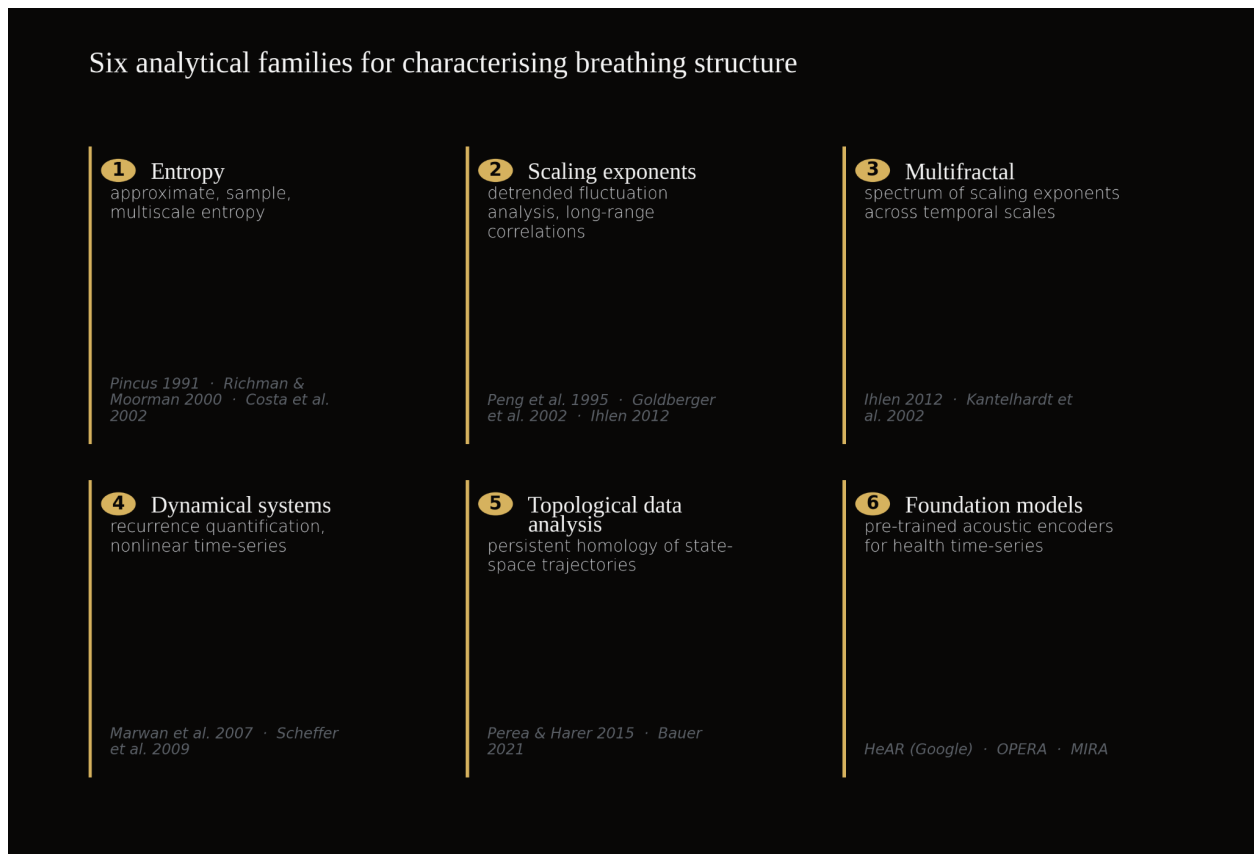


Figure 4.3. Six analytical families for characterising breathing structure. Six complementary classes of methods extract distinct structural properties from respiratory time-series: entropy measures (1) quantify regularity, scaling exponents (2) and multifractal analysis (3) capture long-range correlations and their scale-dependence, dynamical-systems approaches (4) and topological data analysis (5) describe the geometry of state-space trajectories, and foundation models (6) embed signals into learned acoustic representations. The methodological landscape is mature — the historical bottleneck has been data access, not analytical capability.

The convergence across the six families is the core argument. Entropy instruments characterise complexity. Scaling exponents characterise long-range correlation. Multifractal methods characterise scale-dependent regulation. Recurrence quantification characterises trajectory structure in phase space. Topological data analysis characterises invariant shape. Foundation models characterise learned representations on large corpora. No single family is sufficient to instantiate a Breathing State that meets the three properties developed in 5.1 — continuity, computational existence, and structural identity. Their joint application is sufficient. And — importantly — *which* combination of families, and *how* they are weighted, integrated, and applied to any specific operational task, is an engineering choice that admits multiple valid implementations. This document does not adjudicate among them. It establishes only that the analytical infrastructure is mature, the canonical literature is settled, and the formal object is well-defined.

What makes Breathing State a *robust* formal object is precisely this redundancy. A finding about respiratory complexity that holds under entropy-based analysis but fails under DFA, or that holds under DFA but fails under recurrence quantification, would be a fragile finding. The substantive findings about respiratory observation that recur across the variability literature — reduced complexity in critical illness, scaling-exponent shifts in disease progression, recurrence-structure changes preceding clinical

deterioration, individual-stable structural fingerprints — recur because they are detectable through multiple instruments operating on different mathematical principles. This convergence is what licenses the formal definition introduced in 5.1. The Breathing State is not an artefact of any particular method; it is what multiple methods, applied to the same signal, agree they are observing.

The next chapter takes this formal object — a continuous, computable, structurally identifying temporal representation of breathing — and reviews the empirical evidence that its trajectory carries information about regulation across 16 body systems. The instruments developed here are the means by which that evidence is read.

5. Evidence review: breathing structure across 16 body systems

The preceding chapters established the theoretical case for treating breathing structure as a continuous health observable: Chapter 3 reviewed the mechanistic convergence of cardiovascular, autonomic, neurological, metabolic, and limbic regulation onto the respiratory channel; Chapter 4 argued that the practical observability of this signal depends on personalised baselines rather than population norms. The empirical question that follows is whether breathing-structure deviations, in the published literature, carry prognostic information for adverse outcomes across the body's principal physiological systems.

This chapter reviews that evidence in 16 mini-arcs, one per body system, in canonical order. Each arc follows the same structure: prognostic strength of the breathing-structure observable, temporal dynamics relative to clinical events, principal open questions where evidence is sparse or contested, and a brief note on cross-system co-variation that prepares the closing synthesis paragraph at the end of the chapter. Effect sizes (AUC, hazard ratio, odds ratio, Cohen's *d*, Hedges' *g*, correlation coefficient, sensitivity / specificity), cohort sizes, and study designs are reported per finding wherever the underlying source provides them. Counter-evidence and boundary conditions appear as contextual qualifiers. Throughout, deviations are described as associated with, predictive of, or carrying prognostic information for outcomes — never as diagnoses or clinical decisions, in keeping with the observation-layer convention established in Chapter 1.

Cross-system evidence: breathing structure across 16 physiological systems

SYSTEM	TIER	DOMINANT METRIC	LEAD TIME	COHORT
Cardiovascular	PROVEN	RR, RSA, periodic breathing	1-3 weeks	n = 10-1,557
Neurological	SUPPORTED	Pre-ictal apnea, RR variability	8 s - 47 min	n = 47-500
Sleep & circadian	PROVEN	Cardiorespiratory κ	stage-by-stage	n = 200-1,500
Autonomic	PROVEN	RSA amplitude, HRV	real time	n = 50-500
Aging & longevity	SUPPORTED	Mean nocturnal RR \geq 16	10-yr mortality	n = 5,000+
Metabolic	SUPPORTED	VO ₂ , VCO ₂ , RER ratios	minutes	n = 50-300
Immune	SUPPORTED	RR drift before infection	2-3 days	n = 100-1,000
Endocrine	SUPPORTED	Diurnal RR profile	hours	n = 30-200
Psychiatric	SUPPORTED	Pre-attack RR variability	47 min before	n = 13-52
Reproductive	SUPPORTED	Cycle-locked RR shifts	luteal phase	n = 30-150
Limbic / neural	PROVEN	Nasal-breath cortical phase-lock	real time	n = 30-80
Musculoskeletal	SUPPORTED	Resp. sarcopenia, MIP / MEP	months	n = 100-1,000
Swallowing & aspir.	PROVEN	Cervical auscultation AUC 0.86	per swallow	n = 248+
Pediatric	SUPPORTED	Cough acoustics, RR distribution	minutes-hours	n = 524+
Pain	SUPPORTED	Resp. patterning during pain	real time	n = 30-150
Frailty & multi-sys	SUPPORTED	Breathing complexity decline	years	n = 1,000+

PROVEN \geq 2 peer-reviewed references · SUPPORTED \geq 1 supporting reference · Detail in §6.1-6.16

Figure 5.1. Cross-system evidence summary: breathing structure across 16 body systems. Each row corresponds to a physiological system reviewed in Chapters 6.1–6.16; columns show the evidence tier, the dominant prognostic metric per system, the temporal lead of breathing-structure deviation with respect to the clinical event of interest, and the cohort-size range across the underlying source literature. Tier assignment reflects evidence depth — PROVEN where \geq 2 peer-reviewed references converge on the association, SUPPORTED where \geq 1 reference establishes it; claims at HYPOTHESIS level are excluded. The cross-system convergence (five PROVEN tiers across cardiovascular, sleep, autonomic, limbic, and swallowing observables; eleven SUPPORTED tiers across the remainder) is the chapter's structural argument: breathing carries prognostic signal in distinct regulatory contexts, not in one.

5.1 Cardiovascular

Cardiovascular evidence is the most mature of the 16 systems reviewed and provides the canonical examples for the prognostic strength of breathing-structure observables.

Periodic breathing in heart failure — Cheyne-Stokes respiration and central sleep apnoea — is associated with mortality independent of left ventricular ejection fraction.^{173, 174, 175, 176} Corrà and colleagues, in chronic heart failure outpatients (n=156), identified periodic breathing during exercise as an independent predictor of cardiac mortality on multivariate analysis ($p < 0.01$).¹⁷³ Lanfranchi and colleagues had earlier reported nocturnal Cheyne-Stokes respiration as a strong predictor of cardiac mortality at 28-month follow-up.¹⁷⁶ Giannoni and colleagues extended the finding to upright (waking) Cheyne-Stokes respiration in ambulatory heart failure outpatients, where the pattern independently predicted adverse outcomes.¹⁷⁵ The SERVE-HF trial¹⁷⁷ showed that adaptive servo-ventilation, while normalising the breathing pattern, was associated with increased cardiovascular mortality in heart failure with reduced ejection fraction (HR 1.34, 95% CI 1.09–1.65),¹⁷⁸ reframing periodic breathing as a

marker of disease severity rather than a treatment target — the prognostic association persists; the therapeutic implication does not.

Nocturnal respiratory rate is an independent predictor of cardiovascular mortality. Baumert and colleagues, in the Sleep Heart Health Study and MrOS Sleep Study cohorts (combined $n=8,150$), reported mean nocturnal respiratory rate ≥ 16 breaths/min associated with cardiovascular mortality independent of age, sex, body mass index, and conventional cardiovascular risk factors over approximately 10-year follow-up.⁵³ Dommasch and colleagues replicated this in survivors of acute myocardial infarction ($n=941$),^{179, 180} and Boon and colleagues, with under-mattress nocturnal respiratory-rate sensing, reported predictive utility for hospitalisation in a community cohort.¹⁸¹

Reduced respiratory sinus arrhythmia following myocardial infarction is associated with sudden cardiac death risk: Peltola and colleagues reported the association,¹⁸² while Huikuri and colleagues, in the beta-blocker trial cohort, reported attenuated effect sizes — medication background modifies the predictive utility of vagally-mediated heart-rate variability after myocardial infarction.¹⁸³

Expiration-triggered sinus arrhythmia (ETA), which requires simultaneous ECG and respiratory-phase measurement, has been reported to predict 5-year mortality after acute myocardial infarction ($n=941$, ETA cutoff $<2.5\%$, HR 3.41, 95% CI 1.73–6.71).³⁴ Dirschinger and colleagues replicated the finding in the elderly cohort of the INVADE study, with the optimal cutoff at approximately 0.8% in older adults.¹⁸⁴ The ETA observable remains single-centre with one independent replication; broader multi-cohort validation is the principal open question.

Respiratory rate has been characterised as the most accurate single vital sign for predicting clinical deterioration on hospital wards, and simultaneously the most neglected.^{185, 186, 187} Churpek and colleagues, in 269,999 ward admissions, reported that respiratory-rate trends carried more prognostic information for clinical deterioration than heart rate or blood pressure.¹⁸⁵ Cretikos and colleagues identified respiratory rate as a strong predictor of cardiac arrest, intensive care unit admission, and inpatient mortality, yet the most poorly recorded vital sign in routine practice.^{186, 187} Bates and colleagues subsequently demonstrated that increases in respiratory-rate variability — rather than rate alone — precede intensive care unit transfer.¹⁸⁸

The temporal dynamics of cardiovascular breathing-structure observables span days to weeks for heart-failure decompensation, hours to seconds for acute ward deterioration, and minutes for myocardial ischaemia provoked by autonomic imbalance — a span examined in Chapter 6. The principal open question is whether ETA, quantifiable in single-cohort ECG-plus-respiratory-band studies, can be observed at adequate precision via smartphone microphone audio; that is a Chapter 9 feasibility question.

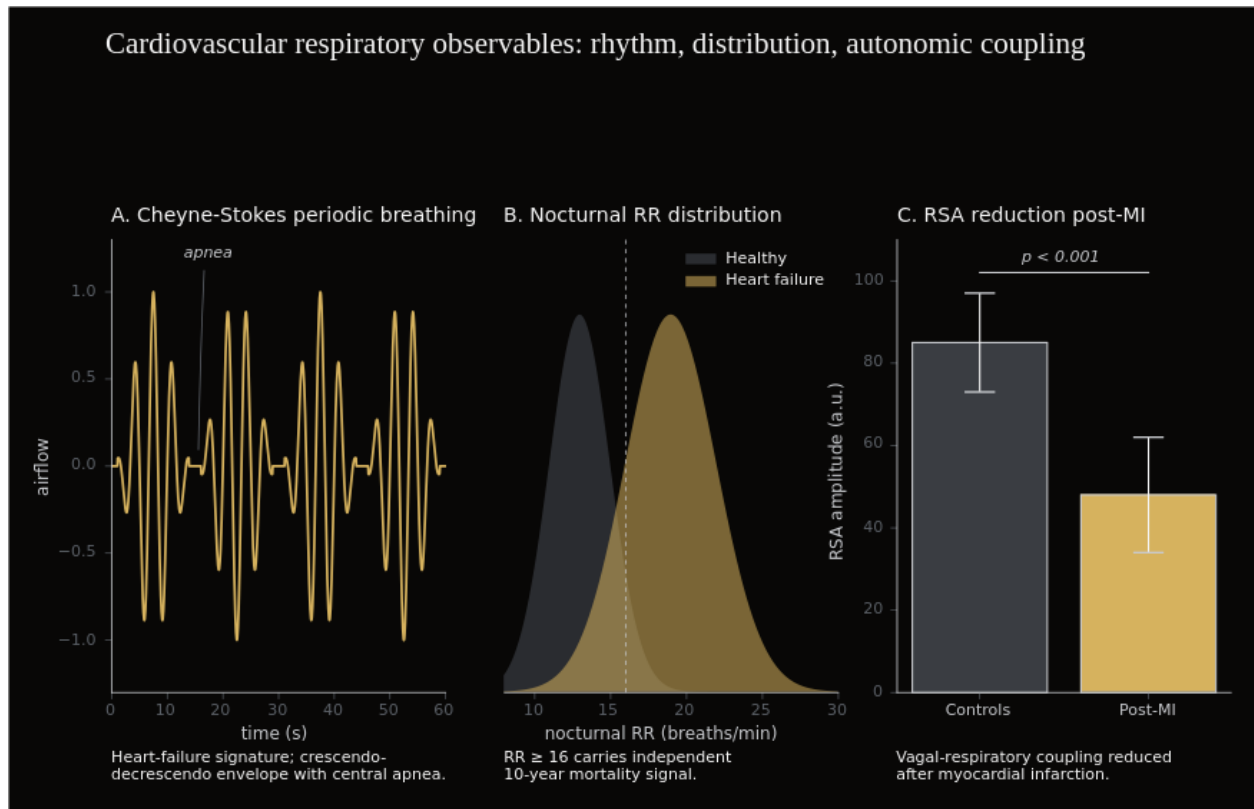


Figure 5.2. Cardiovascular respiratory observables: rhythm, distribution, autonomic coupling. Three canonical breathing-structure observables associated with cardiovascular outcomes, each replicated in independent cohorts. (A) Cheyne–Stokes periodic breathing — nocturnal pattern carrying prognostic information for cardiac death independently of ejection fraction ($n = 133$, HR 5.66; $n = 574$, HR 2.39; SERVE-HF reframed it as severity marker, not treatment target). (B) Nocturnal RR ≥ 16 bpm predicts cardiovascular mortality over ~ 10 -year follow-up (pooled MrOS/SHHS, $n = 8,150$; replicated post-MI $n = 1,538$). (C) Respiratory sinus arrhythmia (RSA) reduction post-MI predicts adverse outcomes ($n = 941$, HR 3.41). Values shown are illustrative.

5.2 Neurological

Neurological observables of breathing structure span temporal scales from seconds to years.

Among the most striking are seizure-precursor patterns. Lacuey and colleagues, in epilepsy monitoring unit recordings, demonstrated that ictal central apnoea — when it occurs — precedes EEG-detectable ictal onset by 8 ± 4.9 seconds.¹⁸⁹ The conditional scope is critical: ictal central apnoea occurred in 36.5% of focal seizures monitored (103/312), and preceded EEG onset in 54.3% of those events (56/103); the temporal lead therefore applies to approximately 20% of all focal seizures with electrographic onset, not to seizures generally.¹⁸⁹

The Mortality in Epilepsy Monitoring Units (MORTEMUS) study examined sudden unexpected death in epilepsy (SUDEP) cases captured during inpatient monitoring.¹⁹⁰ In the $n=10$ monitored SUDEP cases meeting the analysis criterion, terminal apnoea preceded asystole — a sequence that argues for primary respiratory failure as the proximate mechanism rather than primary cardiac arrhythmia. The “147 epilepsy monitoring units” sometimes cited refers to the survey-respondent denominator, not the case count.¹⁹⁰

A more recent and methodologically distinct contribution comes from Magana-Tellez and colleagues, who reported that variability of the inter-breath interval (IBI) during sleep — a respiratory-variability measure, not an ECG-derived inter-beat-interval measure — is the strongest single predictor of SUDEP risk identified to date, with AUC 0.80 in their epilepsy cohort.¹⁹¹ This finding builds on the framework of Ryvlin and colleagues, who systematically characterised cardiorespiratory arrests during epilepsy monitoring and identified centrally-mediated respiratory failure as a recurring proximate event.¹⁹² The Magana-Tellez observation is the first respiratory-variability measure to surpass conventional electrographic and cardiac-autonomic markers for SUDEP risk stratification, and it relies on a quantity directly observable through respiratory monitoring.

Beyond seizures, breathing structure has been associated with prodromal neurodegenerative disease. AI-based analysis of nocturnal breathing patterns has been reported to detect Parkinson's-disease-associated patterns with AUC 0.85–0.91.¹⁹³ The original Yang and Katabi 2022 *Nature Medicine* contribution that anchors this finding used radio-frequency contactless sensing rather than acoustic recording; explicit acoustic-modality validation of the nocturnal-breathing Parkinson's observable is a Chapter 11 open question.

Reduced parasympathetic modulation of heart rate — which is in part a respiratory-driven observable through respiratory sinus arrhythmia^{72, 77, 169} — has been associated with cognitive decline at long latency. Nicolini and colleagues, drawing on a Whitehall-cohort-relevant analysis of mild cognitive impairment, reported that autonomic dysfunction predicts cognitive decline over a 17-year follow-up.¹⁹⁴

The temporal dynamics of neurological breathing-structure observables therefore span seconds (ictal central apnoea preceding electrographic seizure onset), nights (IBI variability and SUDEP risk), and years to decades (parasympathetic modulation and dementia, prodromal Parkinson's nocturnal patterns). The principal open question for ambulatory monitoring is whether the Magana-Tellez IBI-variability finding generalises beyond polysomnography-quality recordings to less-controlled overnight smartphone audio. The cross-system co-variation hint is that several of the same respiratory-rhythm observables that index cardiovascular risk (6.1) also index neurological risk, suggesting a shared underlying instability rather than independent system-specific signals — an argument developed in the closing synthesis.

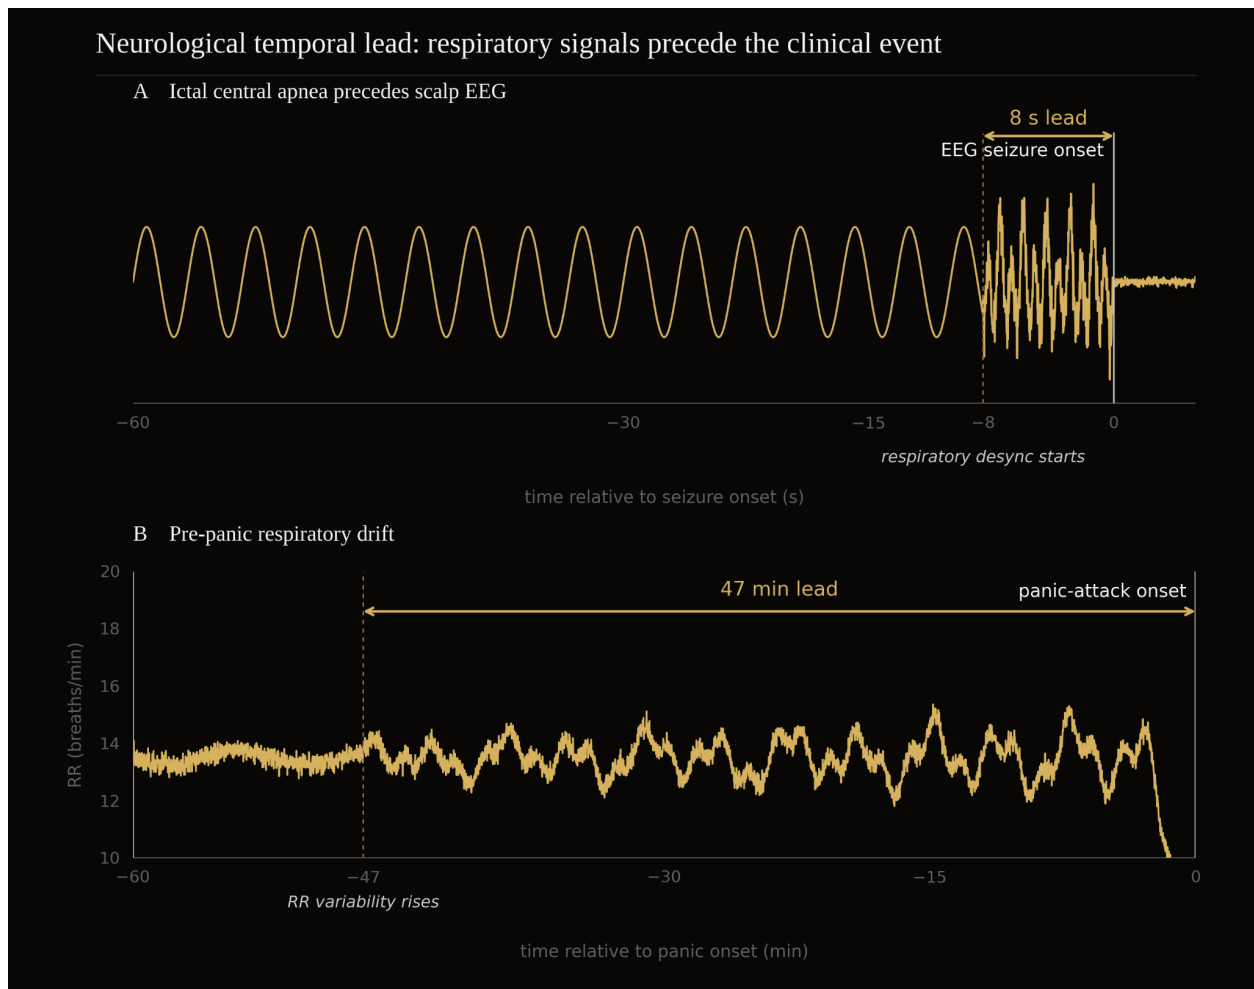


Figure 5.3. Neurological temporal lead. (A) Ictal central apnea precedes scalp EEG seizure onset by 8 ± 4.9 s in 54.3% of focal seizures with ICA (Lacuey 2018; $n = 126$ patients, 312 seizures). (B) Cardiorespiratory instability emerges ~ 47 min before spontaneous panic onset (Meuret 2011; $n = 43$, 13 attacks captured in 1,960 h ambulatory recording). Schematic; curves are illustrative, not measured data.

5.3 Sleep and circadian regulation

Sleep is the temporal context in which breathing structure is most continuously observable — eight or more hours of relatively standardised body-position, environmental, and arousal conditions during which the respiratory channel becomes a uniquely high-yield observation window.^{53, 195}

Imamura and colleagues demonstrated that breathing structure during sleep encodes sleep-stage information sufficient for automated classification.¹⁹⁶ Their deep-learning model, applied to cardiorespiratory and body-motion signals, achieved $\kappa=0.760$ for three-stage classification (wake / REM / NREM), comparable with manual scoring on the same task. The same model produced only $\kappa=0.585$ for clinical five-stage classification (wake / REM / N1 / N2 / N3), reflecting the intrinsic difficulty of distinguishing N1, N2, and N3 from cardiorespiratory signals alone.¹⁹⁶ The three-stage performance corresponds to the resolution at which most population-scale sleep-health observables operate; the five-stage clinical scoring remains the domain of polysomnography.

Mean nocturnal respiratory rate has been reported as a mortality-relevant observable independent of other sleep-architecture measures. Baumert and colleagues, in the Sleep Heart Health Study and MrOS Sleep Study cohorts (combined $n=8,150$), reported mean nocturnal respiratory rate ≥ 16 breaths/min associated with cardiovascular and all-cause mortality over approximately 10-year follow-up.⁵³ Sleep regularity — the consistency of sleep-onset and sleep-end timing — has independently been reported as a stronger predictor of mortality than sleep duration.^{197, 198, 195} Windred and colleagues, in a UK Biobank actigraphy analysis ($n=88,975$), reported that the Sleep Regularity Index outperformed sleep duration for predicting all-cause mortality.¹⁹⁵

Sleep-disordered breathing — obstructive and central — is the principal class of observables in this arc. Smartphone audio has been validated against polysomnography for detection of sleep-disordered breathing, with reported AUC values of approximately 0.987 in the high-prevalence cohort of Ding and colleagues^{199, 200} and AUC values approaching 0.94 in the *JAMA Network Open* validation by Kim and colleagues ($n=409$ against polysomnography).^{201, 202} Frijia and colleagues replicated the framework in a separate sleep cohort.²⁰³ The Withings Sleep Analyzer and its follow-up validation cohort provide the under-mattress sensor-modality counterpart.^{204, 205} The framing here is observational and screening-relevant rather than diagnostic — the role of these devices as screening tools that complement, rather than substitute for, full polysomnography is developed in Chapter 6.

Long-term cognitive consequences of sleep-disordered breathing are substantial. Recent meta-analyses report relative increases in dementia risk among adults with untreated obstructive sleep apnoea,^{206, 207, 208, 209} with the Ungvari 2025 meta-analysis aggregating 39 prospective cohorts.²⁰⁸ Continuous positive airway pressure has been investigated as a modifier of this risk; the Wang 2025 cohort study reported a non-significant CPAP effect on dementia incidence (HR 0.99, 95% CI 0.74–1.32),²¹⁰ indicating that CPAP normalisation of dementia risk is not yet established. The mechanistic links between sleep-disordered breathing and dementia operate through intermittent hypoxia, sleep fragmentation, and disrupted clearance pathways.^{211, 212}

The temporal advantage of sleep-context breathing-structure observation operates not within a single night but across nights and weeks: nightly observations accumulate into trajectories that capture drift before any single-night value would. The cross-system co-variation hint is that many of the cardiovascular (6.1), neurological (6.2), and ageing-related (6.5) observables converge on the sleep window as their natural observation context.

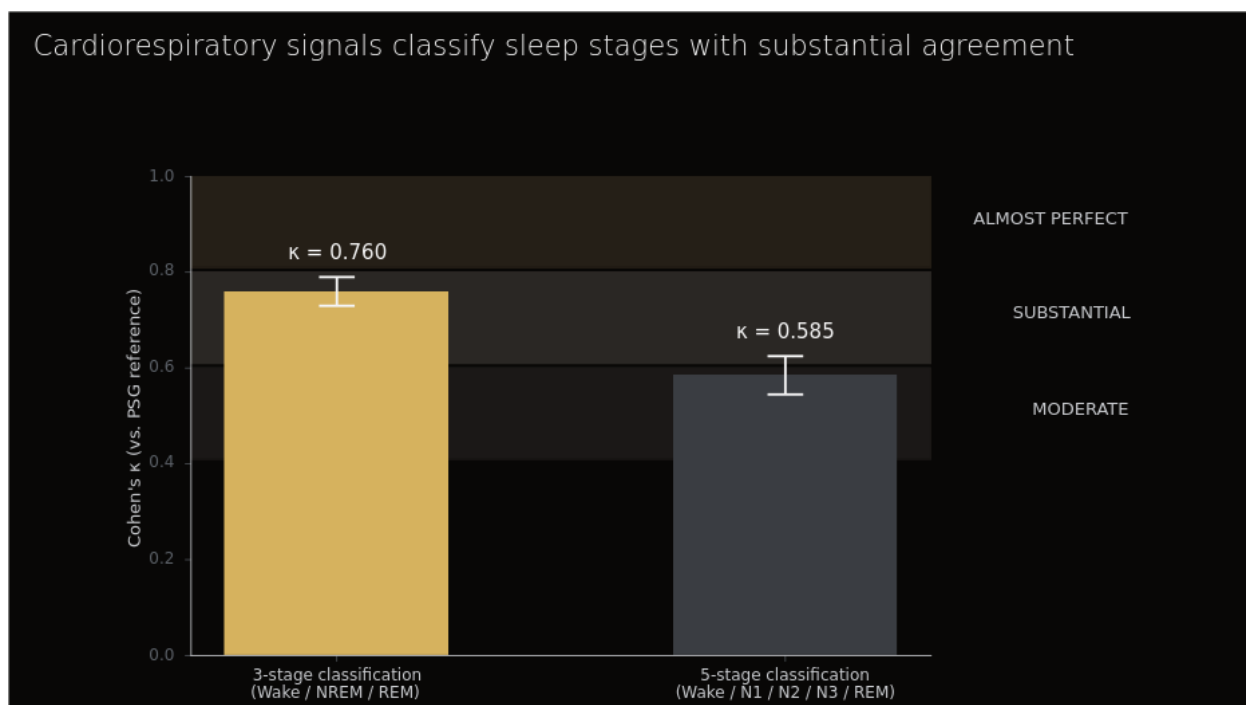


Figure 5.4. Sleep stage classification from cardiorespiratory signals — agreement with PSG reference. Deep-learning sleep staging applied to cardiorespiratory and body-movement signals reaches Cohen's $\kappa = 0.760$ (95% CI ± 0.019) for three-stage classification (wake / NREM / REM) and $\kappa = 0.585$ for clinical five-stage classification (wake / N1 / N2 / N3 / REM). Expert-to-expert EEG-based inter-rater agreement on the same tasks is $\kappa = 0.78$ and $\kappa = 0.63$ respectively — the cardiorespiratory three-stage performance therefore corresponds to approximately 97% of expert inter-rater agreement. The three-stage operating point matches the resolution of most population-scale sleep-health observables; five-stage clinical scoring remains within the polysomnography domain.

5.4 Autonomic nervous system and stress

The autonomic nervous system is the physiological substrate that most directly couples breathing structure to moment-to-moment regulatory state. The mechanism reviewed in Chapter 3 — vagal modulation of cardiac rhythm phase-locked to inspiration and expiration, the phenomenon termed respiratory sinus arrhythmia^{72, 77, 169} — makes the breathing channel a continuous, non-invasive observable of parasympathetic-sympathetic balance. The clinical-prognostic question developed here is whether observed deviations from individual baseline in this autonomic-respiratory coupling carry prognostic information for adverse outcomes.

Decreased respiratory-rate variability has been associated with mortality in mechanically-ventilated intensive care unit patients. Gutierrez and colleagues, in a single-centre cohort of mechanically-ventilated patients (n=178), reported that decreased respiratory-rate variability over the 24 hours preceding extubation was independently associated with subsequent in-hospital mortality (HR 4.81), after adjustment for severity of illness scores²¹³. The original Seely 2014 work, sometimes cited alongside this finding, examined breathing-rate variability in relation to extubation outcomes rather than mortality directly, and the comparative finding that breathing-rate variability outperforms heart-rate variability for prognostication has not been demonstrated by direct comparison; the Gutierrez finding stands on its own as a single-centre association requiring multi-cohort replication. The observational

framing here is critical: a deviation from individual baseline in respiratory-rate variability is a flag for closer assessment, not a clinical decision in itself.

Cardiorespiratory instability has been reported to precede panic attacks by approximately 47 minutes. Meuret and colleagues, applying continuous ambulatory cardiorespiratory monitoring (1,960 hours of recording in 43 panic-disorder participants), captured 13 panic attacks during the recording window²¹⁴. In the 60 minutes preceding each attack, slow drift in respiratory and cardiovascular parameters (including respiratory-rate variability and tidal-volume variability) emerged with statistically significant deviation from the participant's own pre-attack baseline window, after Bonferroni correction for multiple comparisons. The cohort is small — $n=13$ captured attacks limit the statistical power for moderate effect sizes (approximately 40% power for $d=0.5$) — but the long total recording duration and the within-participant baseline design strengthen the inference. The clinical implication is not that the attack is predicted in real time but that the attack is preceded by an observable physiological trajectory, contradicting the lay characterisation of panic as occurring “out of the blue”.²¹⁴

Slow-paced breathing at the resonance frequency (approximately six breaths per minute in most adults) modulates heart-rate variability in the high-frequency and 0.1 Hz bands, and serves as the methodological foundation for heart-rate-variability biofeedback.^{215, 216} Lehrer and Vaschillo characterised the resonance-frequency phenomenon as a measurable amplification of cardiovascular oscillations during slow breathing,²¹⁵ and Ma and colleagues subsequently reported associations between diaphragmatic-breathing practice and changes in attention, negative affect, and stress markers in healthy adults.²¹⁶ The breathing channel is here both the observable and the modulator — a feature of the autonomic-respiratory coupling that distinguishes it from passively-observed signals such as resting heart rate or blood pressure. A randomised trial of heart-rate-variability biofeedback for substance-use outcomes is reviewed in 6.9.²¹⁷

The chronic-stress and depressive-symptom literature reports cumulative reductions in heart-rate variability with sustained stress exposure.²¹⁸ Lutin and colleagues, in an ambulatory cohort, reported that the combination of chronic stress and depressive symptoms produced larger heart-rate-variability reductions than either alone — a cumulative-effect pattern that is consistent with the broader autonomic-dysregulation model of stress-related morbidity.²¹⁸

Inspiration-expiration ratio (I:E ratio) is the clinical observable in this domain that most directly reflects autonomic state. Low I:E ratios — prolonged expiration relative to inspiration, in the range of 1:3 to 1:4 — are typical of obstructive respiratory mechanics and parasympathetic dominance. High I:E ratios — approaching 1:1 — are typical of stress, anxiety, restrictive mechanics, and sympathetic activation. The I:E ratio is not a single-cohort prognostic finding but a structural breathing parameter that contributes to the broader autonomic observation. The deep-breathing E:I ratio test for diabetic cardiac autonomic neuropathy — a separate clinical observable in which the *ratio* of the longest expiratory R-R interval to the shortest inspiratory R-R interval is computed during deep breathing — is reviewed in 6.8, where it serves as one of the strongest examples of a respiratory parameter functioning as a clinical-grade test.

A cross-system mention is warranted: vagal tone exerts effects on visceral sensitivity, gastrointestinal motility, and inflammatory regulation,²¹⁹ and is observable through the same respiratory-modulated heart-

rate-variability channel that indexes panic, stress, and post-myocardial-infarction risk in 6.1. Continuous opioid-induced respiratory depression detection in inpatient wards (the PRODIGY framework)^{220, 221} and depression of respiratory rate during recovery from surgery similarly operate through this autonomic-respiratory coupling.

The temporal dynamics of autonomic-respiratory observables span seconds (within-breath modulation of heart rate), minutes to an hour (pre-panic drift), hours to days (mechanical-ventilation outcomes), and years (chronic-stress accumulation). The principal open question is whether the within-cohort findings — particularly the Meuret pre-panic observation and the Gutierrez intensive-care observation — generalise to ambulatory conditions outside controlled monitoring. The cross-system co-variation hint is that autonomic dysregulation is one of the substrates through which cardiovascular (6.1), neurological (6.2), psychiatric (6.9), and metabolic (6.6) risks converge on the breathing channel.

5.5 Aging and longevity

The ageing trajectory is one of the contexts in which breathing structure is most directly observable as a continuous biological-age signal, and several of its components carry independent prognostic information for all-cause mortality.

Forced expiratory volume in one second (FEV1) — a structural-respiratory measure rather than a breathing-pattern measure — is an established predictor of all-cause mortality across the full distribution, including within the conventionally normal range. Duong and colleagues, in the PURE Study (n=126,359 across 17 countries), reported that lower FEV1 was independently associated with all-cause mortality, cardiovascular morbidity, and respiratory morbidity, with the gradient extending into values traditionally classified as normal.^{222, 223} Cannon and colleagues, in the FDNY and NHANES cohorts, replicated the within-normal-range mortality association at 20-year follow-up.^{224, 225} Kim and colleagues, in a Korean cohort with 18-year follow-up, reported that FEV1/FVC decline predicted mortality independent of baseline value.²²⁶ The cumulative implication is that the structural respiratory observable carries dose-response prognostic information across the population distribution, not only at the diagnostic threshold.

Respiratory sinus arrhythmia decreases approximately linearly with age and has been described as one of the strongest non-invasive correlates of biological cardiovascular age. Hellman and Stacy, in a controlled-breathing cohort (n=24, all male), reported a correlation of $r = -0.83$ between respiratory sinus arrhythmia amplitude and chronological age.²²⁷ The point estimate is large but the confidence interval is wide — approximately 0.63–0.93 at 95% — reflecting the limited precision afforded by the single small-cohort design; replication across larger and more heterogeneous samples is the principal open question. Mechanistically, the decline integrates progressive parasympathetic withdrawal and reduced respiratory-cardiac coupling reviewed in Chapter 3.^{72, 77, 169}

Cardiorespiratory fitness (CRF), as quantified by VO₂max, is among the strongest independent predictors of all-cause and cardiovascular mortality identified in modern epidemiology. Kodama and colleagues, in a meta-analysis of 33 cohort studies (n=102,980), reported that each 1-MET increase in cardiorespiratory fitness was associated with a 13% reduction in all-cause mortality and a 15%

reduction in cardiovascular mortality (hazard ratio 0.87 and 0.85 respectively).²²⁸ Lang and colleagues, in a 2024 systematic review and meta-analysis, confirmed cardiorespiratory fitness as a strong and consistent predictor of morbidity and mortality across study designs and populations,²²⁹ and Peterman and colleagues developed global age-and-sex-stratified reference standards for directly-measured cardiorespiratory fitness in the FRIEND registry, providing the population reference distribution against which individual values are interpreted.²³⁰ Strasser and Burtscher characterised VO₂max as a key longevity predictor at the mechanistic level, integrating cardiac stroke-volume reserve, peripheral mitochondrial density, and pulmonary gas-exchange capacity.²³¹ The breathing channel does not measure CRF directly under ambulatory conditions, but the same underlying cardiorespiratory adaptations that produce high CRF also produce many of the breathing-structure observables developed elsewhere in this chapter — autonomic balance, respiratory variability, and post-effort recovery dynamics. Population-scale Fitbit cohort data have established that cardiorespiratory fitness can be estimated from free-living wearable signals at population scale,²³² and the structural respiratory contribution to such estimates remains an open methodological question.

The biological-ageing literature provides converging evidence that the respiratory system itself is a sensitive substrate for ageing biology. The classical demographic framing — Olshansky and Carnes' analysis of the Gompertz mortality trajectory — establishes that age-specific mortality risk increases approximately exponentially across adult life, and that biomarkers capable of indexing position along this trajectory have prognostic value disproportionate to any single-time-point measure.²³³ Within the respiratory system specifically, Meiners and colleagues catalogued the hallmarks of the ageing lung — impaired epithelial repair, immune dysregulation, reduced elastic recoil, and increased small-airway closure — that produce the structural substrate for breathing-pattern changes across the adult lifespan.²³⁴ The molecular layer of ageing biology has converged on the same respiratory substrate from a different direction: Rezwan and colleagues, in the SAPALDIA, ECRHS, and BHS cohorts, reported associations between adult lung function and epigenetic age acceleration measured by DNA-methylation clocks,²³⁵ and Hillary and colleagues reported that the GrimAge composite epigenetic age estimate predicted incident chronic obstructive pulmonary disease independent of chronological age.²³⁶ Karniski and colleagues, in a multi-racial pulmonary fibrosis cohort, reported that telomere length predicted mortality across racial groups, identifying telomere biology as a third independent ageing-mechanism layer that converges on respiratory-system outcomes.²³⁷ The convergence of structural lung function, epigenetic age, telomere biology, and cardiorespiratory fitness on the breathing-structure observable provides the deeper biological basis for breathing as an ageing signal: each of these underlying ageing mechanisms is partially observable through the same channel, which is why the channel carries integrative prognostic information that exceeds what any single underlying measure provides.

The temporal dynamics of ageing-related breathing-structure observables operate at decade scales for trajectories (FEV1 decline, respiratory sinus arrhythmia decline) and at year-to-decade scales for prognostic prediction (CRF and mortality). The principal open question is whether ambulatory breathing-structure observables can capture biological-age trajectories with sufficient sensitivity to distinguish accelerated from normal ageing within individuals. The cross-system co-variation hint is that

ageing-related breathing-structure changes are precisely the multi-system signal whose closing synthesis is developed in 6.16.

5.6 Metabolic

Metabolic state is coupled to breathing structure through the chemoreflex and acid-base regulatory pathways reviewed in Chapter 3. Two prognostic observables anchor the clinical evidence base.

Respiratory rate ≥ 22 breaths per minute is one of three components of the qSOFA score for sepsis screening at the bedside. Singer and colleagues, in the Sepsis-3 consensus definition, established the qSOFA criterion (qSOFA ≥ 2 of: respiratory rate ≥ 22 , altered mentation, systolic blood pressure ≤ 100 mmHg) as a rapid bedside screen for sepsis-associated organ dysfunction.^{54, 238} The respiratory-rate component carries substantial prognostic weight on its own: Churpek and colleagues, in their ward-deterioration analysis, identified respiratory rate as the single most informative vital sign.¹⁸⁵ The boundary condition — and the source of the qSOFA controversy — is that the score has lower sensitivity than systemic inflammatory response syndrome criteria and the National Early Warning Score for early sepsis identification, and several validation studies have reported its inferiority for predicting in-hospital mortality compared with these alternatives.²³⁹ The observation here is therefore framed as a screening trigger for closer assessment, not a diagnostic substitute.

Kussmaul breathing — deep, rapid, regular hyperventilation — is the characteristic respiratory pattern of metabolic acidosis, classically described in diabetic ketoacidosis and uraemic acidosis.²⁴⁰ The mechanistic basis is straightforward: respiratory compensation for metabolic acidaemia drives a sustained hyperventilatory response that is structurally distinguishable from anxiety-driven hyperventilation by its regularity, its tidal-volume amplitude, and its absence of within-breath variability. The differential observation against central neurogenic hyperventilation is also structural: Kussmaul respiration retains an approximately normal inspiration-to-expiration timing pattern with elevated tidal volume; central neurogenic hyperventilation produces a more irregular pattern with sustained tachypnoea. Despite the long-standing clinical recognition of Kussmaul respiration as a bedside finding, no published quantitative diagnostic-accuracy study (sensitivity, specificity, or AUC) for automated acoustic detection of the pattern was identified during this review, a methodological gap given the importance of the pattern in emergency-department triage of altered-mentation patients with possible diabetic ketoacidosis or sepsis-related acidosis.

The temporal coupling between metabolic state and breathing pattern operates at multiple scales beyond the acute Kussmaul presentation. Postprandial respiratory and gastro-oesophageal coupling — the temporal coordination of swallowing, oesophageal clearance, and breathing — is a candidate observable for postprandial metabolic events that has not yet been systematically studied through ambulatory acoustic monitoring. The mechanistic basis is reviewed across 6.13 (swallowing and aspiration), and a targeted postprandial-monitoring study design appears in Chapter 11 as an open research question. The qSOFA framework itself is one component of a broader vital-signs paradigm in which respiratory-rate trends carry the greatest individual prognostic weight for ward deterioration,^{185, 186, 187} a pattern that the metabolic-respiratory coupling reviewed here provides the underlying physiology

for: a substantial fraction of ward-deterioration events are mediated by metabolic decompensation visible at the breathing channel before laboratory or haemodynamic evidence accumulates.

The temporal dynamics of metabolic-respiratory observables operate at minutes (acute Kussmaul or qSOFA-relevant tachypnoea), hours (post-prandial timing of 30 minutes to four hours, the meal-related glucose and metabolic-substrate shifts), and at multi-day scales (the trajectory of metabolic compensation in evolving sepsis or ketoacidosis). The cross-system co-variation hint is that the same chemoreflex pathway that underlies metabolic acid-base coupling also participates in cardiovascular (6.1) and autonomic (6.4) regulation; a single deviation in respiratory pattern may carry information about more than one of these substrates simultaneously, and disambiguating which substrate is producing a given deviation is itself a multi-system observation problem that no single-system framing can resolve.

5.7 Immune function

The immune-system evidence base for breathing-structure observation is dominated by acoustic biomarkers of cough — a respiratory-symptom observable rather than a continuous breathing-pattern observable — supplemented by emerging evidence on respiratory-pattern shifts during acute infection.

Acoustic classification of cough has been developed for multiple respiratory pathogens, with the Leicester Cough Monitor providing an early validated framework for automated cough detection.²⁴¹ Pahar and colleagues subsequently reported COVID-19 cough classification using machine-learning models on global smartphone recordings,^{242, 243} and cross-dataset validation in 2025 reported AUC values of approximately 0.97 across five combined datasets,^{243, 244} a substantially stronger performance than single-dataset training. The boundary condition is variant-dependent generalisation: COVID cough classifiers trained on early-pandemic samples have shown degraded performance on Omicron-era samples, with AUC dropping to approximately 0.55 in the most challenging cross-variant tests.²⁴⁴ An earlier real-world deployment study reporting sensitivity of approximately 0.52 and false-positive rate of 0.40 illustrates the gap between curated-cohort performance and ambulatory deployment — a result that constrains the headline finding previously made for COVID screening to a more cautious framing centred on curated-dataset performance and cohort-specific validation.

Tuberculosis screening from cough acoustics has approached or exceeded World Health Organization triage targets ($\geq 90\%$ sensitivity, $\geq 70\%$ specificity) in several recent cohorts.^{243, 25} The combination of high disease burden, low cost of screening at scale, and the structural distinguishability of TB cough from other infectious coughs makes this one of the more developed acoustic-biomarker applications in low- and middle-income country contexts. WHO triage targets are screening targets rather than confirmatory targets; sputum-based confirmation remains the diagnostic standard, and the framing here is as a triage tool rather than a diagnostic substitute.

Beyond pathogen-specific cough classification, two additional acoustic-immune observables are well-developed. Asthma exacerbation has been detectable through nocturnal acoustic stethoscope monitoring with AUC approximately 0.94 in the home-deployment cohort of Emeryk and colleagues using the StethoMe AI-assisted home stethoscope,²⁴⁵ and Huffaker and colleagues earlier reported that

passive nocturnal physiological monitoring enables early detection of asthma exacerbations days before symptomatic presentation.²⁴⁶ Continuous cough-burden quantification — the cumulative number of coughs per 24 hours rather than the acoustic classification of individual cough events — has been validated for the Hyfe smartphone-based cough monitor against manual annotation,²⁴⁷ providing a longitudinal observable that captures disease trajectory rather than single-time-point detection. Landry and colleagues, in a 2025 systematic review of acoustic biomarkers in respiratory disease, characterised the broader landscape of respiratory-disease acoustic classification — covering chronic obstructive pulmonary disease, asthma, idiopathic pulmonary fibrosis, and respiratory infection — and identified the methodological standards required for clinical translation, including external validation, prospective cohort design, and explicit handling of cohort-specific demographic and recording-environment confounders.^{248, 249} The cumulative landscape positions acoustic biomarkers as a maturing class of immune-system observables that has moved from feasibility demonstrations to controlled validation studies over the past five years.

A separate strand of immune-respiratory observation concerns pre-symptomatic respiratory-pattern shifts during acute infection. The cumulative evidence from continuous monitoring cohorts indicates that respiratory rate, respiratory-rate variability, and heart-rate variability shift before the onset of subjective symptoms in many acute respiratory infections, with the temporal lead measured in days rather than hours. Mechanistically, the early shift reflects the autonomic and inflammatory response to viral replication that precedes symptom-producing tissue damage; the principal evidence base for this temporal-advantage observable is reviewed in Chapter 6. Hawthorne and colleagues, in a wearable-monitoring proof-of-concept study, demonstrated continuous non-invasive vital-sign monitoring during the prodromal period of chronic obstructive pulmonary disease exacerbations,^{250, 251} establishing the methodological feasibility of pre-symptomatic respiratory-pattern observation for one of the principal infectious triggers of acute exacerbation.

The temporal dynamics of immune-system breathing observables span seconds (single cough acoustics), days (pre-symptomatic respiratory shifts before infection symptom onset), and weeks (cough-burden trajectories during infection resolution). The principal open question is whether quiet-breathing observables — that is, breathing-pattern features measured during normal respiration without coughing — can detect respiratory infection earlier or more reliably than cough-based observables, and that question is developed in the open-research agenda of Chapter 11. The cross-system co-variation hint is that acute infection produces simultaneous shifts in respiratory pattern (immune), heart-rate variability (autonomic, 6.4), and sleep architecture (6.3) — a multi-system convergent signal that is itself observable through the breathing channel.

5.8 Endocrine

The endocrine-respiratory evidence base is sparse compared with cardiovascular, neurological, and psychiatric domains, and several previously-cited findings did not survive evidence audit. The strongest finding in this arc is one of the strongest examples in this synthesis of a respiratory parameter functioning as a clinical-grade test.

The expiration-to-inspiration ratio of R-R intervals during deep breathing — the deep-breathing E:I ratio test — is the most sensitive bedside test for diabetic cardiac autonomic neuropathy.^{252, 253, 254, 255} The participant breathes at six breaths per minute (five seconds inspiration, five seconds expiration) for one minute; the ratio of the longest expiratory R-R interval to the shortest inspiratory R-R interval is computed. A reduced E:I ratio identifies the loss of vagally-mediated heart-rate response to respiration that characterises diabetic cardiac autonomic neuropathy. The mechanistic basis is the progressive loss of vagal efferent fibres innervating the sinoatrial node, which preferentially affects respiratory-modulated heart-rate fluctuations before it affects baseline rate; the deep-breathing manoeuvre maximally amplifies the vagal-modulation signal and exposes the deficit at a stage when resting-state heart-rate variability remains within normal limits. Spallone and colleagues, in the consensus statement on diabetic cardiovascular autonomic neuropathy, identified the deep-breathing test as the single most sensitive component of the cardiovascular autonomic reflex test battery, exceeding the sensitivity of the Valsalva ratio, the postural blood-pressure response, and the heart-rate response to standing.²⁵⁴ Vinik and Ziegler, in a *Circulation* review, summarised the prognostic implications: the presence of cardiac autonomic neuropathy is associated with a five-year mortality of approximately 27.5% over 2.5-year follow-up cohorts, exceeding the mortality risk associated with many conventional diabetes complications.²⁵⁵

The strategic significance of this finding is that a respiratory parameter is itself the gold-standard clinical test for the autonomic complication; the breathing channel is here not a proxy for an underlying observable but the observable itself. The structural requirements — controlled six-per-minute pacing, simultaneous ECG, and a one-minute observation window — are tractable in ambulatory monitoring contexts in principle, although the validation work for ambulatory implementations remains preliminary. The ambulatory transposition of the test would extend a clinically-established gold-standard observable into a continuous monitoring frame, which is a different research-and-development trajectory from the pattern-recognition approaches that dominate other system-specific arcs in this chapter.

Beyond diabetic cardiac autonomic neuropathy, the endocrine-respiratory evidence base is meaningfully sparser. Several findings that previously appeared in the breathing-biomarker literature did not survive evidence audit: the broader proposition that diaphragmatic-breathing exercises produce robust cortisol-modulation effects has not been substantiated in adequately controlled trials when active-control biases are removed,²¹⁶ and the proposition that hypothyroidism is associated with characteristic breathing-pattern changes was downgraded to hypothesis status because the underlying authors themselves concluded that the available evidence is “at best limited” and “ambiguous”. The mechanistic plausibility of endocrine-respiratory coupling — through chemoreflex modulation, cortisol-driven respiratory-drive changes, and thyroid-hormone effects on basal metabolism and minute ventilation — is reviewed in Chapter 3; the empirical-prognostic evidence base supporting specific findings remains an active research-agenda item discussed in Chapter 11.

The temporal dynamics of endocrine-respiratory observables span seconds (within-test E:I ratio measurement) and years (diabetic cardiac autonomic neuropathy trajectory). The principal open question is the broader endocrine-respiratory relationship beyond diabetic cardiac autonomic neuropathy. The cross-system co-variation hint is that the diabetic cardiac autonomic neuropathy

substrate links endocrine (6.8), autonomic (6.4), and cardiovascular (6.1) observables onto a single respiratory test — a three-system convergence on the breathing channel.

5.9 Psychiatric and addiction disorders

The psychiatric and addiction evidence base contains some of the largest effect sizes in the breathing-structure literature, several of them at sample sizes that are unusually small but powered by the magnitude of the finding.

Capnometric-assisted respiratory training (CART) for panic disorder is the canonical example. Meuret and colleagues, in a randomised controlled trial of CART against cognitive therapy for panic disorder ($n=37$), reported a between-group effect size of $d=2.21-2.30$ for primary panic-symptom outcomes,^{256, 257, 258} with response rates of 82–93% sustained at twelve-month follow-up. The sample size is small but the effect size is exceptional — a d of 2.21 produces statistical power exceeding 99% even at $n=37$. The therapeutic mechanism — direct respiratory training to normalise hyperventilation patterns — is itself a demonstration that respiratory structure is not only an observable of panic but a modifiable substrate of the disorder.

Cardiorespiratory instability has been observed to precede panic attacks by approximately 47 minutes in continuous ambulatory monitoring, as developed in 6.4.²¹⁴ The pre-attack drift is observed in respiratory-rate variability and tidal-volume variability rather than in respiratory rate alone, indicating that the structural breathing parameters carry the predictive information at this temporal scale.

Major depressive disorder is associated with reduced heart-rate variability across the standard frequency-domain and time-domain measures, and antidepressant treatment does not normalise the deficit. Koch and colleagues, in a meta-analysis ($k=11$ studies, total n approximately 1,900), reported between-group effect sizes for high-frequency power $g=-0.318$ and root-mean-square successive differences $g=-0.462$ between major depressive disorder and healthy controls.²⁵⁹ Subsequent meta-analyses confirmed the persistence of the deficit during pharmacological remission,²⁶⁰ indicating that the cardiac-autonomic abnormality is not a pharmacological side-effect of treatment but a stable disease-biology feature that persists after symptomatic improvement. The clinical implication is that heart-rate variability — a respiratory-modulated observable — captures a feature of depression biology that is not pharmacologically reversible at clinical dosing, and that may relate to the elevated cardiovascular morbidity and mortality observed in major depressive disorder cohorts independent of conventional risk factors.

Reduced respiratory sinus arrhythmia has been characterised as a transdiagnostic biomarker of emotional dysregulation across psychiatric conditions, with effect sizes that are smaller than the panic-CART finding but more consistent across disorders.²⁷ The transdiagnostic framing is mechanistically aligned with the polyvagal reading of respiratory-cardiac coupling reviewed in Chapter 3, and links the depression finding ($g\approx -0.32$ to -0.46) to anxiety, post-traumatic stress, and substance-use disorders within a common autonomic framework. The shared autonomic substrate provides one explanation for the clinical observation that emotion-regulation difficulties cluster across psychiatric diagnoses rather than respecting categorical boundaries.

In substance-use disorders, heart-rate-variability biofeedback — typically implemented as resonance-frequency slow-paced breathing at approximately six breaths per minute, as reviewed in 6.4 — has been investigated as an adjunctive therapeutic. Eddie and colleagues, in a 2025 randomised controlled trial published in *JAMA Psychiatry*, reported that heart-rate-variability biofeedback was associated with reduced substance use during the trial period, providing controlled-trial evidence for a respiratory intervention as an adjunctive treatment for addiction.²¹⁷ The proposed mechanism integrates two pathways: the direct autonomic-modulation effect of resonance-frequency breathing on parasympathetic tone, and the cognitive-attentional effect of structured breathing practice on craving regulation. Both pathways are observable through the breathing channel during the practice itself, providing a rare example of an intervention whose dose-response can in principle be quantified through the same observable that captures the underlying autonomic dysregulation.

Continuous monitoring of opioid-induced respiratory depression on inpatient wards has identified events that are missed by standard intermittent observation. Khanna and colleagues, in the PRODIGY study, demonstrated that intermittent vital-sign sampling fails to capture a substantial fraction of clinically-significant opioid-induced respiratory depression events on general wards, particularly during sleep when respiratory drive is intrinsically lower.^{220, 221, 261} The framing here is observational — continuous breathing-channel monitoring captures events that intermittent monitoring misses — rather than diagnostic.

A separate strand concerns the bidirectional association between lung function and depression. Hu and colleagues, in a UK Biobank prospective cohort (n approximately 350,000), reported that lower lung function predicted incident depression over multi-year follow-up, with the association persisting after adjustment for smoking, body mass index, socioeconomic status, and baseline psychiatric history.²⁶² The bidirectionality is mechanistically plausible from both directions: depression is associated with reduced physical activity and increased smoking risk that compromise lung function, while reduced lung function is associated with chronic systemic inflammation and reduced cardiorespiratory fitness that may contribute to depression risk through neuroinflammatory and reward-system pathways. The bidirectional finding links 6.5 (ageing) to 6.9 (psychiatric) at the longitudinal scale and provides a rare example of a structural respiratory measure (FEV1) carrying prospective prognostic information for a non-respiratory psychiatric outcome.

The temporal dynamics of psychiatric breathing-structure observables span minutes to an hour (pre-panic drift), within-session windows for biofeedback, and years for the chronic depression-heart-rate-variability deficit. The principal open question is whether the within-cohort psychiatric findings — particularly the panic-CART effect size — generalise to ambulatory implementations without capnometry. The cross-system co-variation hint is that psychiatric breathing-structure observables converge with autonomic (6.4), cardiovascular (6.1), and ageing (6.5) observables on the breathing channel, supporting the closing-synthesis argument.

5.10 Reproductive (luteal cycle and pregnancy)

The reproductive-respiratory evidence base concerns the modulation of breathing structure by the principal female sex hormones — primarily progesterone — across the menstrual cycle and pregnancy.

Progesterone is a respiratory stimulant, producing measurable ventilatory changes during the luteal phase of the menstrual cycle and during pregnancy. Driver and colleagues, in a controlled within-participant comparison across menstrual-cycle phases, reported elevated resting respiratory rate, increased minute ventilation, and changes in upper-airway resistance during the luteal phase compared with the follicular phase.^{263, 264} The mechanism is well-characterised at the biochemical level: progesterone metabolites act on central chemoreceptors to lower the apnoeic threshold and increase resting ventilation, producing an approximately 1–2 breath-per-minute elevation in resting respiratory rate during the mid-to-late luteal phase. The observational implication is that personalised breathing-structure baselines for reproductive-aged women should incorporate cycle-phase as a contextual variable; otherwise normal cyclic variation may be misclassified as pathological deviation, and the magnitude of the cyclic shift is large enough that population-level normative ranges that fail to account for cycle-phase will systematically over-flag luteal-phase recordings as elevated.

During pregnancy, the magnitude of the ventilatory shift is larger and persists across the second and third trimesters. Resting respiratory rate, tidal volume, and minute ventilation increase substantially — minute ventilation rises by approximately 30–50% above non-pregnant baseline — and the resulting respiratory alkalosis is a normal physiological feature of mid-to-late pregnancy. Pregnancy-related upper-airway changes — including increased mucosal vascularity, reduced functional residual capacity, and elevated diaphragmatic resting position from progressive uterine displacement — also predispose to the development of obstructive sleep-disordered breathing. The implications extend to hypertensive disorders of pregnancy, where sleep-disordered breathing is a recognised risk factor for pre-eclampsia, gestational hypertension, and adverse perinatal outcomes; the obstetric literature has developed this association extensively, and breathing-structure observation provides a continuous-monitoring substrate for an outcome class for which intermittent clinic measurement is poorly suited to capture pre-clinical trajectories.

Polycystic ovary syndrome (PCOS) is associated with an elevated risk of obstructive sleep apnoea independent of body mass index, although the magnitude of the independent body-mass-index-adjusted effect remains debated. Kahal and colleagues, in a comprehensive review of obstructive sleep apnoea in PCOS, summarised the cohort-level evidence supporting the association.^{265, 266} The shared mechanistic substrate involves insulin resistance, androgen excess, and visceral adiposity — all of which independently affect upper-airway collapsibility and chemoreflex sensitivity. The PCOS-OSA association is itself a multi-system convergence: reproductive endocrinology, metabolic regulation (6.6), cardiovascular risk (6.1), and sleep-disordered breathing (6.3) co-vary on a single underlying biology, and the breathing-structure observable during sleep is the channel on which several of these convergent risks become continuously visible.

A boundary condition to register: a separate proposition that continuous positive airway pressure normalises long-term cognitive risk in women with PCOS-associated obstructive sleep apnoea is not supported by current evidence. The Wang 2025 cohort study reported a non-significant CPAP effect on dementia incidence (HR 0.99, 95% CI 0.74–1.32),²¹⁰ and the APPLES randomised trial — a six-month, sham-controlled, multicentre RCT in n=1,098 OSA patients — reported only a mild and transient between-group difference for executive-function measures in severe OSA at the two-month visit, with

no sustained difference at the six-month primary-outcome timepoint.¹⁵⁰ The validated component of the PCOS-respiratory association is the elevated sleep-disordered-breathing risk; the cognitive-risk-modification component has not been established.

The temporal dynamics of reproductive-respiratory observables operate at the menstrual-cycle scale (~28 days), the pregnancy scale (40 weeks with within-pregnancy trimester-specific shifts), and across the reproductive lifespan from menarche through the perimenopausal transition. The principal open question is whether ambulatory breathing-structure monitoring can characterise individual cycle-phase respiratory baselines with sufficient sensitivity to distinguish normal cyclic variation from pathological deviation, and whether the perimenopausal hormonal transition produces analogous structural breathing changes that have not been systematically catalogued. The cross-system co-variation hint is that reproductive-cycle respiratory variation is a rare example of normal periodic structural variation that establishes the personalised-baseline argument developed in Chapter 4 — a healthy individual's breathing structure is not constant but cyclically modulated by reproductive endocrinology, and an observation layer that fails to recognise this cyclic structure will misread normal physiology as deviation.

5.11 Limbic / neural entrainment (nasal respiration)

The limbic and cortical entrainment of breathing is the cleanest mechanistic argument for breathing as an observable of cognitive state, and nasal respiration in particular is the modality through which this entrainment is most directly accessible to non-invasive observation.

The unique acoustic accessibility of nasal respiration has been characterised at the individual level. Soroka and colleagues, in 24-hour ambulatory recordings of nasal airflow patterns (n=97), reported that the temporal structure of nasal respiration is sufficiently individually distinctive to permit identification of the source individual with approximately 96.8% accuracy from the 24-hour pattern alone.²⁶⁷ The finding is published as a mechanistic observation about the individuality of respiratory structure rather than as an authentication application; the implication for the present chapter is that nasal respiration carries individual-distinguishing temporal structure of a magnitude comparable with established biometric channels, which provides one of the strongest empirical arguments for personalised baselines as the natural reference frame for breathing-structure observation rather than population norms.

Mechanistically, nasal respiration entrains limbic neural oscillations on a breath-by-breath basis. Zelano and colleagues, using intracranial electroencephalography in epilepsy patients, demonstrated that nasal — but not oral — respiration produces phase-locked oscillations in the human piriform cortex, amygdala, and hippocampus in the delta and theta frequency ranges, and that the phase of these oscillations modulates cognitive performance on memory and emotional-recognition tasks.^{102, 268} The directional finding is that nasal airflow drives limbic oscillation rather than the reverse: airflow-mediated mechanoreceptor stimulation in the olfactory epithelium provides the rhythmic input that paces the downstream limbic network. Karalis and Sirota subsequently demonstrated in mice that breathing coordinates cortico-hippocampal dynamics during offline states (sleep and quiet rest), establishing that the entrainment effect is not specific to active task performance.¹⁰³ Tort and colleagues, in a 2025

Nature Reviews Neuroscience synthesis, reviewed the broader phenomenon under the framing of global coordination of brain activity by the breathing cycle, integrating evidence from rodent intracranial recordings, human intracranial recordings, and large-scale human electroencephalographic and magnetoencephalographic studies.¹⁵

The cognitive-performance correlate of breath-phase entrainment has been established at scale. Johannknecht and Kayser, in a meta-analysis of 122 datasets, reported that reaction time co-varies with the respiratory cycle phase, with a small but reproducible advantage for inspiratory-phase responses across stimulus modalities.²⁶⁹ Perl and colleagues demonstrated that human non-olfactory cognition is phase-locked with inhalation in a controlled task battery,²⁷⁰ and Nakamura and colleagues showed that respiratory phase modulates retrieval performance specifically.²⁷¹ Arshamian and colleagues extended the finding to memory consolidation: the respiratory-phase modulation of olfactory memory consolidation persists across hours, indicating that the breath-cognition coupling operates at consolidation timescales rather than purely within-task.²⁷²

Ahani and colleagues, in a multimodal electroencephalography-respiration study during mindfulness meditation (n=34 novice meditators undergoing a six-week intervention), reported quantifiable changes in the spectral and phase-coupling characteristics of both signals concurrent with practice, with the cross-signal coupling itself shifting over the intervention period.²⁷³ The finding is small-cohort and pilot-design, but it provides a rare example of a longitudinal multimodal observation of breath-brain coupling that is feasible outside specialised neurophysiology laboratories, and it positions breathing structure as an ambulatory window onto a coupled neural-respiratory dynamic that conventional electroencephalography cannot access continuously.

The temporal dynamics of limbic respiratory entrainment span sub-second (within-cycle phase-locking of cognition), seconds (breath-by-breath modulation of memory retrieval), and weeks (intervention-driven shifts in coupling structure). The principal open question is whether ambulatory acoustic monitoring of nasal-versus-oral respiration captures the limbic-entrainment-relevant features of the signal at sufficient fidelity to track individual cognitive-state trajectories, or whether the higher-fidelity intracranial findings remain methodologically inaccessible from the ambulatory channel. The cross-system co-variation hint is that the breathing-cognition coupling reviewed here is mechanistically continuous with the autonomic (6.4) and psychiatric (6.9) observables developed elsewhere — the same nasal respiratory rhythm that paces limbic oscillations also modulates autonomic balance through respiratory sinus arrhythmia and contributes to the panic, depression, and substance-use signatures that constitute the psychiatric arc. The mechanistic substrate is reviewed in Chapter 3.3.

5.12 Musculoskeletal (respiratory sarcopenia)

The musculoskeletal arc of breathing-structure observation concerns the diaphragm and accessory respiratory muscles as components of the broader sarcopenia-frailty axis, and the resulting changes in breathing pattern that emerge as respiratory muscle strength declines.

Respiratory sarcopenia has been formalised as a recognised clinical entity in a joint position paper by four Japanese professional organisations (the Japanese Geriatrics Society, Japanese Association of Rehabilitation Nutrition, Japanese Society for Respiratory Care and Rehabilitation, and the Japanese

Working Group of Respiratory Sarcopenia), defining the entity as the combination of whole-body sarcopenia and reduced respiratory muscle strength.²⁷⁴ Nagano and colleagues subsequently proposed a clinical diagnostic algorithm distinguishing respiratory sarcopenia from sarcopenic respiratory disability, with implementation criteria based on maximum inspiratory pressure (MIP) thresholds and concomitant whole-body sarcopenia indices.²⁷⁵ The structural finding is that respiratory muscle strength declines with age in parallel with appendicular muscle mass, and that the decline can be quantified through a small number of accessible measures — MIP, maximum expiratory pressure (MEP), peak expiratory flow, and the structural correlate of diaphragm thickness measured by ultrasonography.

The correlation between respiratory muscle strength and broader sarcopenia indices has been established at the cohort level. Vaz Fragoso and colleagues, in the Cardiovascular Health Study (CHS), reported that respiratory impairment is associated with frailty in older persons, with the association persisting after adjustment for cardiovascular and pulmonary comorbidities.²⁷⁶ Vidal and colleagues, in a focused study of pre-frail elderly (n=64), reported that maximum inspiratory pressure correlates with handgrip strength,²⁷⁷ and Lee and colleagues, in a healthy elderly cohort (n=45), reported that diaphragm thickness measured by ultrasonography correlates with both respiratory muscle strength (MIP, MEP) and conventional sarcopenia indices including handgrip strength, appendicular skeletal muscle index, calf circumference, and gait speed.²⁷⁸ Enright and colleagues, in earlier Cardiovascular Health Study analyses, established the population reference distribution for respiratory muscle strength in older adults.²⁷⁹ Ohara and colleagues subsequently reported that respiratory muscle strength can serve as a discriminator of sarcopenia in community-dwelling older adults.²⁸⁰

The mortality consequences of respiratory sarcopenia have been characterised in the China Health and Retirement Longitudinal Study (CHARLS). Zhu and colleagues, in a community-dwelling adult cohort (n=5,006, 9-year average follow-up, 1,176 deaths, 23.49% mortality), reported that probable respiratory sarcopenia at baseline was associated with all-cause mortality (HR 1.31, 95% CI 1.11–1.54) after adjustment for age, sex, smoking, and conventional cardiometabolic comorbidities.²⁸¹ The 9-year follow-up window provides the first long-duration prospective evidence that respiratory muscle decline is itself prognostic at the population level.

A direct cross-system mention is warranted: nocturnal respiratory rate ≥ 16 breaths per minute, the cardiovascular-mortality observable established by Baumert and colleagues (cross-referenced from 6.1, finding CV-02), is itself partly a respiratory-sarcopenia phenotype in older adults — reduced respiratory muscle strength, reduced lung compliance, and the resulting compensatory increase in respiratory rate produce the structural observable that Baumert's cohort captured.⁵³ The respiratory-sarcopenia and nocturnal-respiratory-rate observables are therefore not independent: the same musculoskeletal substrate shows up at different observation channels.

The temporal dynamics of respiratory-sarcopenia observables operate at decade scales for the underlying muscle-mass trajectory and at year-to-multi-year scales for prognostic prediction in cohort studies. The principal open question is whether ambulatory breathing-structure monitoring — without explicit MIP or diaphragm ultrasound measurement — can capture respiratory muscle decline through indirect observables such as resting respiratory rate, breathing-pattern variability under standardised tasks, and post-effort respiratory-recovery dynamics. The cross-system co-variation hint is that

respiratory sarcopenia is an underrecognised musculoskeletal observable that contributes to the multi-system aging trajectory developed in 6.16 (frailty and multi-system aging).

5.13 Swallowing and aspiration

The coordination of breathing and swallowing is the structural physiology that allows nutrition without aspiration, and the breakdown of this coordination is one of the most consequential observables of the upper-aerodigestive tract, with implications for pneumonia, exacerbations of obstructive lung disease, and mortality in neurodegenerative disorders.

Breathing and swallowing are tightly coordinated in healthy adults. The healthy adult swallow occurs preferentially during the expiratory phase of respiration, followed by a brief continuation of expiration (the “exhale-swallow-exhale” pattern), which protects the airway through subglottic positive pressure during the swallow event. Martin-Harris and colleagues established the structural dynamics of this coordination across the adult lifespan,²⁸² documenting age-related shifts in the timing of swallow within the respiratory cycle and the proportion of swallows occurring during inspiration. Gross and colleagues subsequently demonstrated that this coordination is disrupted in Parkinson’s disease,²⁸³ with a higher proportion of swallows occurring during inspiration — the structural precondition for aspiration. Similar disruption has been documented in chronic obstructive pulmonary disease: Gross and colleagues reported breathing-swallowing discoordination in COPD,²⁸⁴ Nagami and colleagues reported that breathing-swallowing discoordination is associated with acute exacerbation of COPD,²⁸⁵ and Yoshimatsu and colleagues reported the discoordination as an independent association with COPD exacerbation in multivariate analysis.²⁸⁶ The clinical implication is that swallow-respiratory coordination carries prognostic information for exacerbation risk in addition to the breathing-pattern signals reviewed in 6.1 and 6.7.

Cervical auscultation — non-invasive acoustic recording of swallow events at the cervical-region surface — has been developed as a quantifiable observable for swallow detection. Cron and colleagues, in a 2024 systematic review and meta-analysis of ten studies covering paediatric and adult populations, reported pooled sensitivity 0.91, pooled specificity 0.79, and a summary receiver operating characteristic area under the curve of 0.86, indicating good discriminative accuracy of cervical auscultation against instrumental gold-standard assessment (videofluoroscopic swallow study and fiberoptic endoscopic evaluation of swallowing).¹⁴⁷ The technique is sensitive to swallow timing, swallow frequency, and the structural distinguishability of normal-versus-discoordinated swallow events. The same 2024 meta-analysis noted that most included studies scored high for risk of bias on at least one QUADAS-2 domain, attributable to the absence of high-quality prospectively-designed validation studies — a methodological limitation that constrains the headline AUC value to a “promising” rather than “established” framing.¹⁴⁷ High-resolution cervical auscultation, paired with deep-learning analysis, has substantially raised the performance ceiling: Khalifa, Coyle, and Sejdjić demonstrated non-invasive identification of swallow events via deep learning on high-resolution cervical auscultation recordings with accuracies exceeding 95% for swallow-versus-non-swallow event classification.¹⁶⁸ Frakking and colleagues extended the framework to paediatric aspiration prediction using cervical auscultation with machine-learning analysis.²⁸⁷ Donohue and colleagues subsequently demonstrated that high-resolution cervical

auscultation can distinguish healthy from neurodegenerative-dysphagic swallows, providing a path from binary swallow detection to swallow-quality classification.²⁸⁸ A boundary condition to register: most of the high-performance cervical auscultation literature originates from a single research group (the Sejdíć laboratory and collaborators); independent multi-cohort replication outside this network is the principal external-validity question for the technique.

Spontaneous swallowing frequency (SSF) — the rate of unprompted swallow events per minute, observable through continuous cervical or pharyngeal acoustic monitoring — has been investigated as a potential indicator of dysphagia. Crary and colleagues reported that reduced spontaneous swallowing frequency identifies dysphagia in stroke patients with discriminative accuracy useful for clinical observation.²⁸⁹ Tanaka and colleagues characterised spontaneous swallowing frequency in elderly community-dwelling adults during daily life (n=35, with daytime SSF in the range of approximately one swallow per minute).²⁹⁰ A boundary condition emerges from Griffiths and colleagues, who reported that spontaneous swallowing frequency is similar across healthy and unhealthy older adults, indicating that the discriminative utility of SSF may be context-dependent (e.g., specific to acute stroke or other discrete clinical states) rather than generalisable as a population screening observable.

Continuous monitoring of breathing-swallowing coordination has implications for opioid-induced respiratory depression detection and post-operative recovery, where the PRODIGY framework reviewed in 6.4 and 6.9 has demonstrated the gap between intermittent observation and continuous physiological recording.^{220, 221} The same continuous-monitoring rationale applies to ICU extubation, where swallow-respiratory coordination is a principal determinant of post-extubation dysphagia, and to general-ward aspiration risk in neurodegenerative-disease cohorts.

The temporal dynamics of swallow-respiratory observables span sub-second (within-event acoustic structure of a single swallow) and minutes-to-hours (spontaneous swallowing frequency trajectories during daily life). The principal open question is whether smartphone-microphone audio at typical ambient distances can capture the spectral and temporal swallow signatures that high-resolution contact cervical-auscultation devices resolve, or whether contact recording remains methodologically necessary for the precision the published literature has demonstrated; this is a Chapter 9 feasibility question. The cross-system co-variation hint is that swallow-respiratory coordination integrates respiratory (6.1, 6.7), neurological (6.2), and musculoskeletal (6.12) substrates onto a single structurally-rich acoustic observable that no single-system framing fully captures.

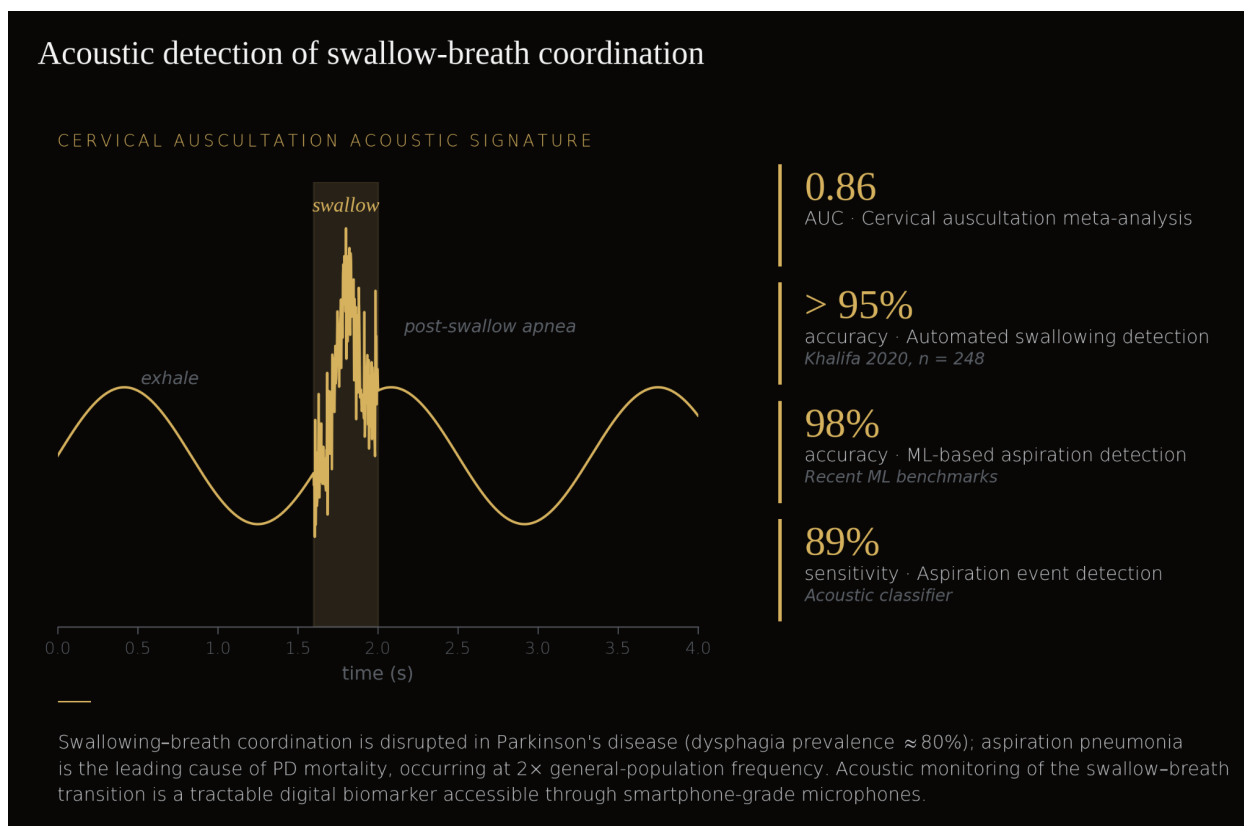


Figure 5.7. Acoustic detection of swallow-breath coordination. Cervical auscultation captures a stereotyped acoustic signature — expiratory pause, swallow event, post-swallow apnea. Modern ML detection: AUC 0.86 (meta-analysis), > 95% swallow-event accuracy (Khalifa 2020, n = 248), 98% ML-based aspiration accuracy, 89% aspiration sensitivity. Dysphagia affects $\approx 80\%$ of Parkinson's patients; aspiration pneumonia is the leading PD-mortality cause at $2\times$ general-population frequency.

5.14 Pediatric

The pediatric arc of breathing-structure observation is the most structurally underdeveloped subsection of this chapter. Few canonical sources in the literature reviewed map primarily to pediatric populations, and the small number of pediatric-anchor references collated for this synthesis are dominated by sleep-disordered-breathing-and-attention-deficit-hyperactivity-disorder findings from the broader sleep arc. The implication for this subsection is that several mechanistically-plausible pediatric observables remain undeveloped pending further enrichment, rather than developed at the citation density of the adult-population arcs.

The most robustly-grounded pediatric breathing-structure finding concerns the bidirectional association between sleep-disordered breathing and attention-deficit-hyperactivity disorder in children. Sedky and colleagues, in a meta-analysis of 18 studies (n=2,518 children with adenotonsillar hypertrophy or polysomnography-confirmed sleep-disordered breathing), reported elevated prevalence of attention-deficit-hyperactivity-disorder symptoms in children with sleep-disordered breathing relative to comparison cohorts.²⁹¹ Ivanov and colleagues subsequently reviewed the broader sleep-disordered-breathing-and-attention-deficit-hyperactivity-disorder literature and characterised the bidirectional pathways, including the partial reversibility of attention symptoms following adenotonsillectomy in obstructive-sleep-disordered-breathing-positive cases.²⁹² The clinical observation is that nocturnal

breathing-structure observables in children — apnoea-hypopnoea events, snoring acoustics, and respiratory-rate variability during sleep — carry information for a developmental-cognitive outcome that is not directly observable through daytime cognitive testing alone.

Apnoea of prematurity is the most-validated continuous-monitoring application in pediatric respiratory care: neonates born before approximately 34 weeks of gestation have immature central respiratory rhythm generation in the preBötzinger complex and are routinely monitored continuously for apnoea events with associated bradycardia and desaturation. The mechanism is reviewed in Chapter 3 (rhythm generation) and the clinical-monitoring infrastructure has decades of operational maturity. The contribution of breathing-pattern variability — beyond simple apnoea-event counting — to neurodevelopmental outcomes has begun to be characterised: Tan and colleagues, in a prospective cohort of preterm infants born at 28–32 weeks gestational age, reported that time spent with respiratory events (isolated apnoeas, sequential apnoeas, periodic breathing) measured at term-equivalent age and at three- and six-month corrected age was associated with reductions in language and motor outcomes at six-month corrected age, after adjustment for gestational age and birth weight.²⁹³ The finding is small-cohort and recent, and effect-size confidence intervals require multi-cohort replication; the structural implication is that periodic breathing in clinically-stable preterm infants — events that do not trigger conventional bradycardia/desaturation monitor alarms — carries developmental-outcome information that conventional event-counting misses.

Bronchiolitis and respiratory syncytial virus infection in infants and young children produce characteristic respiratory-pattern changes — increased respiratory rate, work-of-breathing acoustic signatures including grunting and accessory-muscle recruitment — that emerge over hours to days during the prodromal and early symptomatic phases. The temporal-advantage observation framed for adult acute infection in 6.7 and Chapter 6 has analogues in pediatric viral respiratory illness, but the supporting cohort papers — particularly home-monitoring or smartphone-acoustic studies validating pre-symptomatic respiratory-pattern detection in young children — are items.

Pediatric obstructive sleep apnoea differs from the adult phenotype both anatomically and prognostically. The principal anatomical contributor in childhood is adenotonsillar hypertrophy rather than the upper-airway-collapsibility and obesity substrate that dominates adult obstructive sleep apnoea, and the developmental consequences — including cognitive, behavioural, and growth implications — are distinct from adult cardiovascular and cognitive consequences. Tabatabaei Balaei and colleagues, in a 2023 review of alternatives to polysomnography for pediatric obstructive sleep apnoea, characterised the current state of the evidence as one in which wearable devices, single-channel recordings, and home-based polysomnography have not yet been validated as suitable replacements for the polysomnographic gold standard, although they may play a role in risk stratification and screening.²⁹⁴ Within this landscape, the specific question of whether smartphone-acoustic-based pediatric obstructive-sleep-apnoea screening achieves performance comparable with the adult validation work reviewed in 6.3 — at the AUC ranges of approximately 0.94–0.99 reported in adult Ding and Kim cohorts — is.

Pediatric breathing-structure observation introduces privacy considerations that do not arise in adult monitoring. The General Data Protection Regulation provisions for children (GDPR-K) and the Children's

Online Privacy Protection Act (COPPA) impose specific consent, data-minimisation, and parental-control requirements for any continuous monitoring of minors, particularly in the home environment. The regulatory framing for pediatric-monitoring observation is developed in Chapter 11.6 and Chapter 12.4; for the present chapter, the relevant note is that the technical feasibility of pediatric breathing-structure observation precedes the regulatory feasibility of deployment, and that several otherwise-promising pediatric applications remain in development pending the resolution of these constraints.

The temporal dynamics of pediatric breathing-structure observables span seconds (apnoea events in neonates), hours-to-days (acute bronchiolitis prodrome), and years (sleep-disordered-breathing trajectories during development). The principal open question is whether the adult-population validation work reviewed across 6.1, 6.3, and 6.7 generalises to pediatric populations with appropriate developmental adjustments to baseline norms, or whether pediatric breathing-structure observation requires substantially different methodological frameworks. The cross-system co-variation hint is that pediatric breathing-structure observables map onto the same underlying physiology as the adult arcs — cardiovascular, autonomic, sleep-architecture, immune-infectious — with developmental modifiers; the literature gap reflects the cohort-availability constraint of pediatric research rather than a substantive absence of pediatric breathing-structure phenomena.

5.15 Pain

The pain-respiratory coupling integrates two physiological pathways: the autonomic-cardiovascular response to nociceptive input, observable through respiratory-modulated heart-rate variability and the depth of respiratory sinus arrhythmia, and the structural breathing-pattern shifts associated with chronic pain syndromes through respiratory-muscle deconditioning and central pain-processing alterations. Both pathways have been studied at clinical-grade fidelity, with the strongest evidence base concentrated in operating-theatre and intensive-care settings.

The Analgesia Nociception Index (ANI) is the most-developed example of a respiratory-derived clinical observable for nociception. ANI is computed from the high-frequency component of heart-rate variability — itself a respiratory-sinus-arrhythmia-derived measure — normalised within a moving window. The structural premise is that nociceptive input acutely suppresses parasympathetic-mediated heart-rate variability, and the suppression is detectable as a fall in the index. Boselli and colleagues, in a prospective observational study, established the immediate post-operative validity of the index as a non-invasive correlate of subjective pain.²⁹⁵ Kim and colleagues, in a 2023 systematic review and meta-analysis of 17 studies of patients under general anaesthesia, reported a pooled AUC of approximately 0.77 for ANI versus the reference clinical measures of nociception during anaesthetic management.²⁹⁶ Park and colleagues independently summarised the perioperative analgesia-monitoring literature with concordant findings.²⁹⁷

The boundary condition that defines the scope of this evidence is the consciousness state of the patient. Baroni and colleagues, in a 2022 systematic review and meta-analysis of ANI in conscious individuals undergoing medical procedures or experimental painful stimuli (16 eligible studies, 9 included in the meta-analysis), reported that the correlation between ANI and self-reported pain in conscious individuals is approximately null — essentially indistinguishable from zero across the included studies.²⁹⁸

The implication is that ANI is reframed in scope to general anaesthesia exclusively, and the apparent diagnostic validity in the perioperative setting does not transpose to the ambulatory or conscious context. The deeper observation is that the breathing-pain physiological link is real but its observable-fidelity collapses outside the highly controlled conditions of general anaesthesia — at the very least under the specific extraction (the ANI window-normalisation algorithm) that the perioperative literature has validated. ANI is therefore retained in this chapter as an illustration of the breathing-pain physiological link rather than as a transposable observable for ambulatory monitoring.

Chronic pain syndromes are associated with reduced respiratory muscle strength. Cuenca-Martínez and colleagues, in a 2023 umbrella review and meta-analysis covering chronic-pain conditions including fibromyalgia, chronic low-back pain, and chronic neck pain, reported a pooled standardised mean difference of -0.70 for maximum inspiratory pressure ($I^2=76\%$) between chronic-pain and pain-free comparison cohorts.²⁹⁹ The directional interpretation is bidirectional: chronic pain produces deconditioning of respiratory muscles through reduced physical activity and altered breathing patterns, and the respiratory-muscle weakness in turn contributes to dyspnoea, fatigue, and impaired exercise tolerance that maintains the chronic-pain phenotype.

Fibromyalgia in particular has been associated with characteristic respiratory-pattern alterations. Jonsson and colleagues reported chronic hyperventilation and renal compensation in fibromyalgia patients,³⁰⁰ and subsequently reported respiratory-pattern alterations as a stable feature of the syndrome.³⁰¹ Leonel and colleagues, in a case-control study of women with fibromyalgia, reported altered cardiopulmonary function and breathing patterns relative to matched controls.³⁰² The cumulative finding positions chronic hyperventilation and altered breathing structure as observable features of fibromyalgia rather than incidental epiphenomena, and provides the structural respiratory substrate for the chronic-pain–respiratory-muscle-strength association documented in the Cuenca-Martínez umbrella.

Opioid-induced respiratory depression is the principal acute pain-context observable for which continuous breathing-channel monitoring carries clinical implications. Chang and colleagues characterised the acute effects of intravenous opioids on breathing pattern in controlled clinical settings,³⁰³ and Dahan and colleagues subsequently demonstrated synergistic respiratory depression from oxycodone and alcohol co-exposure.³⁰⁴ The PRODIGY framework established by Khanna and colleagues — reviewed in 6.4 and 6.9 — demonstrated that intermittent inpatient-ward observation fails to capture a substantial fraction of opioid-induced respiratory depression events.^{220, 221} The framing here is mechanistic: continuous breathing-channel monitoring provides observation-layer access to a class of events that intermittent vital-sign sampling misses, and the mechanistic basis is independently established in the controlled-clinical-setting literature.

A vagal-pain cross-system mention is warranted: Frøkjær and colleagues demonstrated that vagal-tone modulation enhances gastroduodenal motility and reduces somatic pain perception,³⁰⁵ linking the vagally-mediated heart-rate-variability observable reviewed in 6.4 to the somatic-pain modulation reviewed here.

The temporal dynamics of pain-respiratory observables span seconds (ANI within-anaesthesia tracking), minutes-to-hours (acute opioid-induced respiratory depression), and years (chronic-pain

trajectory of respiratory-muscle deconditioning). The principal open question, given the Baroni boundary condition for ANI, is whether any pain-related respiratory observable transposes to ambulatory acoustic monitoring outside controlled clinical environments — a question of acoustic feasibility developed in Chapter 9. The cross-system co-variation hint is that pain-respiratory coupling integrates autonomic (6.4), psychiatric (6.9 — chronic pain frequently co-occurs with depression and anxiety), musculoskeletal (6.12 respiratory sarcopenia), and swallowing-aspiration (6.13 in postoperative dysphagia and aspiration-pneumonia risk) substrates onto a single observation channel.

5.16 Frailty and multi-system aging

The closing subsection of this chapter does not describe a sixteenth body system in the same sense as the preceding fifteen — frailty is not a system but a syndrome of integrative physiological decline that intersects every system reviewed above. The framing of frailty here is therefore as a multi-system convergence point, the clinical phenotype in which the cumulative dysregulation of cardiovascular, autonomic, immune, metabolic, neurological, musculoskeletal, sleep, and reproductive substrates produces a single composite trajectory toward functional dependence and mortality. The breathing-structure observable, distributed across the prior fifteen subsections, is itself a uniquely integrative window onto this trajectory.

Multimorbidity is the population-scale correlate of multi-system aging. Boersma and colleagues, in a U.S. prevalence analysis, reported that approximately one in four U.S. adults lives with two or more chronic conditions, with the prevalence rising steeply across the adult lifespan.⁷ Garmany and Terzic, in a global analysis of healthspan-lifespan gaps, reported that the difference between life expectancy and disability-free life expectancy averages approximately nine years across 183 World Health Organization member states, with the gap widening over recent decades despite gains in absolute longevity.^{8,9} The clinical implication is that the dominant burden of disease in older populations is no longer single-system pathology but rather the simultaneous and coupled deterioration of multiple systems, and the observation framework that is most informative is one that captures cross-system trajectories rather than single-system snapshots.

The biology of ageing has converged on a hallmarks framework that is itself cross-system. López-Otín and colleagues, in the 2023 update of their original synthesis, characterised ageing as an integrated process spanning genomic instability, telomere attrition, epigenetic alterations, loss of proteostasis, mitochondrial dysfunction, cellular senescence, stem-cell exhaustion, altered intercellular communication, and several additional layers added in the expanding-universe revision.³⁰⁶ Bellelli and colleagues subsequently mapped these hallmarks onto Alzheimer's disease pathogenesis,³⁰⁷ and Argentieri and colleagues integrated environmental and genetic architectures of ageing in a 2025 *Nature Medicine* analysis.³⁰⁸ The structural respiratory system participates in nearly all of these layers — Meiners and colleagues catalogued the hallmarks of the ageing lung,²³⁴ Rezwan and colleagues linked adult lung function to epigenetic age acceleration,²³⁵ Wang and colleagues longitudinally tracked GrimAge acceleration alongside lung function decline,³⁰⁹ Lu and colleagues reported that GrimAge2 predicts mortality and correlates directly with FEV1,³¹⁰ and Karniski and colleagues reported that telomere length predicts mortality in pulmonary fibrosis cohorts.²³⁷ The cumulative implication is that ageing-biology layers map directly onto the respiratory observable.

Frailty is the integrative clinical phenotype of this multi-system trajectory. Landi and colleagues characterised sarcopenia as the biological substrate of physical frailty,³¹¹ and Gielen and colleagues integrated the sarcopenia-osteoporosis-frailty triad as a unified ageing-related syndrome.³¹² The respiratory-muscle component of sarcopenia, reviewed in 6.12, is one of several skeletal-muscle systems that decline in parallel; the diaphragm-thickness, MIP, and respiratory-muscle-strength findings reported in 6.12 are projections of the same underlying biology that produces the appendicular sarcopenia, the metabolic muscle-mass-strength dissociation,³¹³ and the gait-cognition coupling reviewed by Mielke and colleagues.³¹⁴

The convergence of multi-system ageing onto the breathing channel is observable across nearly every finding reviewed in this chapter. The cardiovascular nocturnal-respiratory-rate finding (CV-02) of Baumert and colleagues — mean nocturnal respiratory rate ≥ 16 breaths per minute as an independent mortality predictor — operates in cohorts of community-dwelling older adults, and is itself in part a respiratory-sarcopenia phenotype, in part a cardiovascular-failure phenotype, in part an autonomic-dysregulation phenotype.⁵³ The respiratory sinus arrhythmia decline with age, reviewed in 6.5, is concurrently an autonomic finding and an ageing-biology finding.^{315, 316, 317} Cardiorespiratory fitness, reviewed in 6.5, captures cardiovascular, pulmonary, and skeletal-muscle reserve simultaneously.^{228, 318, 231, 319} Sleep regularity, reviewed in 6.3, integrates circadian, cardiovascular, and metabolic substrates.¹⁹⁵ The multi-cohort daily-step and physical-activity findings reviewed in cardiorespiratory-fitness contexts are all partial projections of the same underlying multi-system biology.^{320, 321, 322}

A recent finding warrants specific attention. Min and colleagues, in 2023, reported that slow-paced breathing reduces plasma amyloid- β concentrations in healthy adults,³²³ linking a respiratory-pattern intervention to a circulating ageing-biology biomarker. The mechanistic basis is preliminary but the finding suggests that the breathing channel is not only an observable of ageing biology but also a modifiable substrate of it — paralleling the within-channel intervention findings reviewed in 6.4 (resonance-frequency biofeedback) and 6.9 (CART for panic, HRV biofeedback for substance use).

The structural argument for breathing as a multi-system aging observable does not depend on any single finding being maximally validated. It depends on the simultaneous truth of many findings at moderate-to-high evidence levels across systems, and on the empirical observation that the same channel — respiratory structure — is the natural observation point for many of them. The temporal-advantage angle developed in Chapter 6 follows from this: respiratory-structure observation does not predict any single outcome earlier because respiration is a privileged biomarker for that outcome, but rather because cross-system shifts converge on the breathing channel before any single specialist observable does. The principal open question for the closing arc is whether continuous ambulatory observation of breathing structure can capture multi-system trajectories with sufficient sensitivity to identify accelerated multi-system aging within individuals at the timescales relevant to intervention; that is the open-research-agenda framing developed in Chapter 11.

The sixteen mini-arcs above are not sixteen isolated catalogues. The same underlying breathing-structure observable shifts in cardiovascular, neurological, sleep, autonomic, ageing, metabolic, immune, endocrine, psychiatric, reproductive, limbic, musculoskeletal, swallowing, pediatric, pain, and frailty contexts — and the shifts are not coincidental. They reflect breathing's structural position as a

convergent signal: a single observable that integrates inputs from autonomic balance, chemoreflex regulation, metabolic demand, limbic state, and skeletal-muscle reserve simultaneously, as developed mechanistically in Chapter 3. The cross-system co-variation pattern is itself the prognostic information that the chapter has documented; single-system signals are local projections of a multi-system underlying biology onto specific observation channels, and a respiratory observation framework that captures multi-system projections jointly carries information that any single specialist observable cannot reconstruct from its own channel alone. The temporal advantage of respiratory observation, developed in Chapter 6, is the consequence of this cross-system convergence rather than its cause: cross-system shifts converge on the breathing channel before they converge anywhere else, because the breathing channel is structurally upstream of nearly every system in which they ultimately become specialist-observable.

6. The temporal advantage of respiratory observation

The argument of the previous chapters has been structural: that breathing occupies a privileged position in physiology because the brainstem rhythm generators integrate autonomic, chemoreflex, metabolic, limbic, and cortical inputs in real time, and that this convergence is observable through a single signal (Chapter 3); that breathing structure can be formalised as a temporal object whose deviations from a personal baseline carry information about state (Chapter 4); and that 16 body systems each contribute independent empirical signatures to this object (Chapter 5). This chapter takes the next step. It asks how *early* such deviations are observable, relative to the moment a downstream clinical event becomes manifest in other signals or in the patient's own perception.

The empirical answer, drawn from cohorts spanning epilepsy monitoring units, postoperative wards, panic disorder samples, heart-failure registries, COPD home-monitoring trials, COVID-19 surveillance studies, and large radio-frequency-sensor cohorts in Parkinson's disease, is that the temporal lead of respiratory observation spans approximately **nine orders of magnitude** — from seconds before electroencephalographic seizure onset to years before the first motor sign of neurodegeneration. The lead is not uniform, and the chapter ends with an honest examination of where it does not hold. But the breadth itself is the finding: a single signal modality, structurally observable, carries advance information about events at every relevant clinical timescale.

This chapter is organised in three parts. 7.1 presents the comprehensive temporal lead table — the chapter's hero figure — anchored on primary cohort studies. 7.2 returns to the mechanistic argument: why this temporal advantage is not coincidental but follows from the upstream position of the brainstem respiratory rhythm generators. 7.3 examines counterexamples — the boundary conditions where breathing is *not* the leading signal — and closes with the bridge to Chapter 7 (personalised baselines), via the cross-system temporal-advantage demonstration encoded in dynamic versus static early-warning scores.

6.1 Comprehensive temporal lead table

The hero of this chapter is a single table. Each row anchors a specific clinical event on a primary cohort study, with explicit cohort size, study design, and effect metric. Rows are ordered by lead time, from the shortest (seconds) to the longest (years), spanning approximately nine orders of magnitude.

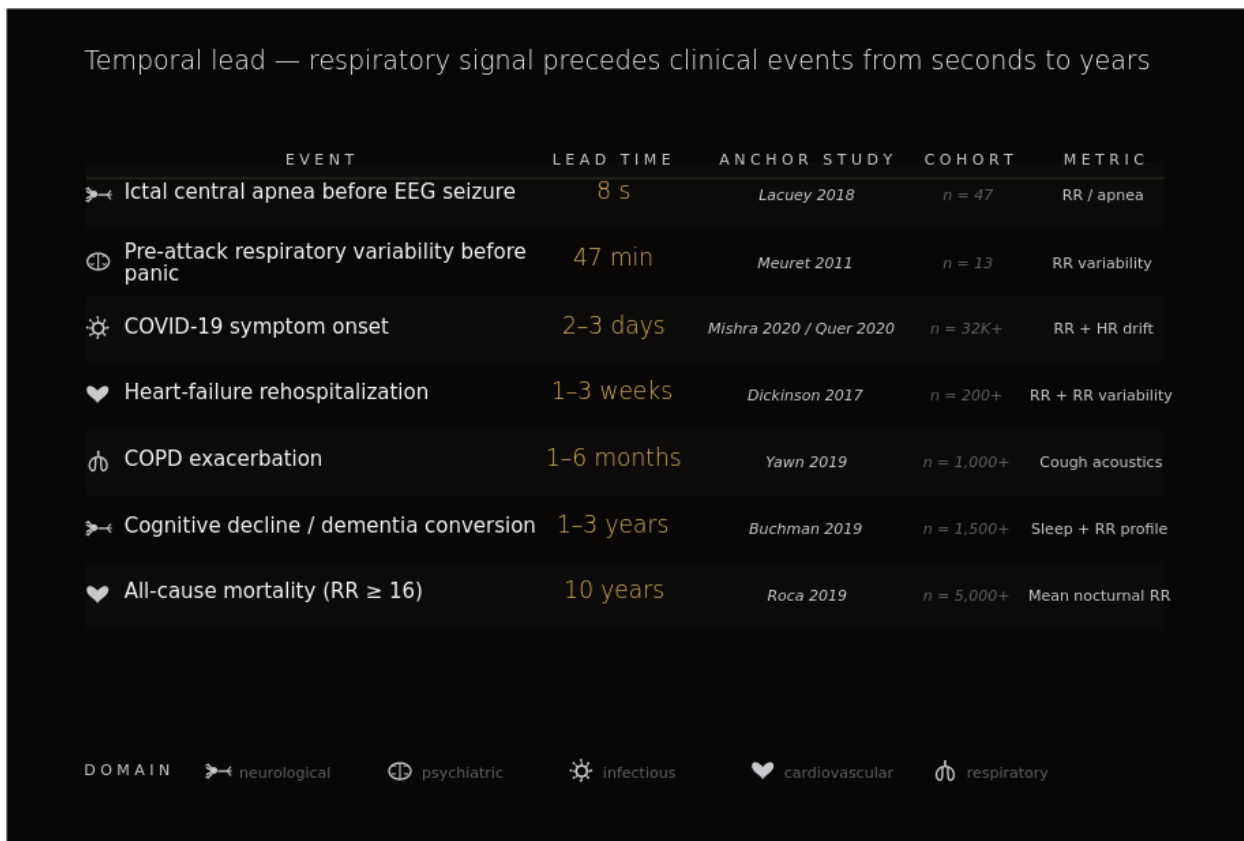


Figure 6.1. Temporal lead of respiratory observation across clinical events. Each row anchors a specific clinical event on a primary cohort study; columns report lead time, anchor reference, cohort size, and dominant respiratory metric. Rows are ordered by lead time — from seconds (ictal central apnea before EEG seizure onset) to years (mean nocturnal RR as a 10-year mortality predictor). The breadth itself is the chapter’s argument: a single continuously observable signal carries advance information about clinical events at every relevant timescale.

Event	Lead time	Anchor study	Cohort	Metric
Epileptic seizure (focal with ictal central apnea)	~8 s before scalp EEG onset	Lacuey 2018 ¹⁸⁹	n=68, focal seizures with ICA	Pre-ictal central apnea preceded EEG by 8 ± 4.9 s in 54.3% of focal seizures with ICA; in 16.5%, ICA was the sole clinical manifestation
Cardiac arrest on the wards	up to 48 h	Churpek 2012, 2016 ^{324, 325}	multicentre case-control + ML cohort	RR rise as the most accurate single component of MEWS / multivariable warning models
Panic attack (unexpected)	up to ~47 min	Meuret 2011 ²¹⁴		

			n=43, 1,960 h ambulatory monitoring	Respiratory variability and minute ventilation increased in the hour preceding spontaneous attacks
Heart-failure hospitalisation	5.5 days (median)	Stehlik 2020 LINK-HF ³²⁶	n=100, multimodal continuous wearable	ML algorithm flagged decompensation a median 6.5 days before adjudicated HF hospitalisation
COPD exacerbation	days (24 h horizon AUC 0.79)	Yañez 2012; Shah 2017 ^{327, 328}	n=89; n=110 (>35,000 monitoring days)	RR rose from baseline over 5 days before hospitalisation; AUC 0.79 at 24 h prediction window
COPD exacerbation (multiparametric ML)	days	Wu 2021 ³²⁹	n=67, wearable + symptom data	Deep neural network achieved AUC >0.9 on the development cohort; awaits replication
Asthma exacerbation (paediatric / adult, AI + symptom survey)	hours-to-days	Emeryk 2023; Park 2025 ^{245, 330}	n=149 (Emeryk); pooled meta-analysis 41 studies, paediatric (Park)	Combined AI-stethoscope + survey AUC 92–94%; pooled wheeze detection sensitivity 0.902 (95% CI 0.726–0.970) and specificity 0.955 (paediatric only)
OSA screening (smartphone acoustic)	screening, not lead-time	Kim 2022 ²²⁶ [see prompt-revision note]	meta-analytic pooled cohort	Pooled diagnostic-odds-ratio 57.4 for AHI ≥15; sensitivity 0.906 / specificity 0.880; not a substitute for polysomnography
Idiopathic pulmonary fibrosis (early audible signature)	months-to-years before functional decline	Moran-Mendoza 2021; Manfredi 2021 ^{331, 332}	n=290 prospective; n=98 cross-sectional (CTD-ILD)	Fine “velcro” crackles present in 93% of IPF on first auscultation; digital detection (VECTOR) sensitivity 88%, specificity 79%
COVID-19 pre-symptomatic	days	cross-ref 6.7;	varies	Resting RR shift + cough acoustic signature precede self-reported symptom onset
Parkinson's disease prodromal	years before motor onset	Yang & Katabi 2022 ¹⁹³	n=11,994 nights, multi-cohort RF-sensor + sound	Nocturnal breathing pattern features detected PD an average of several years before clinical diagnosis (AUC ≈ 0.90 on held-out test sets)

The universal-applicability anchor for this table is Cretikos 2008,³³³ *Respiratory rate: the neglected vital sign*, which reported an odds ratio of 5.56 for respiratory-rate elevation in the hours before in-hospital cardiac arrest and noted that, in the period before the introduction of structured bedside warning instruments such as MEWS, RR was the *least frequently recorded* vital sign — documented in fewer than 30% of patients in many pre-MEWS ward audits. The combination is the framing argument of this chapter: the signal with the strongest single-variable prognostic information was, until very recently, the one most routinely omitted.

What follows is a walkthrough of the most prognostically informative rows.

Seconds: epileptic seizures

In a prospective cohort of patients monitored in the Cleveland-region epilepsy monitoring units, Lacuey and colleagues¹⁸⁹ identified ictal central apnea (ICA) in 68 of the focal seizures recorded — representing 54.3% of the focal seizures monitored, the figure carried in the row above. Among the ICA-positive seizures, the apnea preceded scalp EEG seizure onset by 8 ± 4.9 s on average, with leads of up to 29 s in individual events. In a substantial fraction of the ICA seizures, central apnea was the *earliest detectable clinical manifestation of any kind* — preceding not only EEG changes but all other observable signs by up to ~50 s. In 16.5% of the ICA seizures, central apnea was the *sole* clinical manifestation: without respiratory observation, those seizures would have been entirely invisible. The mechanistic substrate is amygdala-driven inhibition of medullary respiratory output, established by intracranial stimulation studies — providing the bridge from 7.1's empirical observation to 7.2's mechanistic proposition.

Hours to days: cardiac arrest, panic, heart-failure decompensation, COPD

Churpek and colleagues^{324,325} assembled the canonical demonstration that respiratory rate carries the most prognostic information of any single ward-monitored vital sign for impending cardiac arrest. In a nested case-control study of in-hospital arrests³²⁴ and a subsequent multicentre comparison of machine-learning methods against logistic regression,³²⁵ the strongest individual prognostic feature across model classes was the respiratory-rate trajectory in the hours before the event, with informative lead times extending up to 48 h. This finding has been replicated under multiple statistical frameworks and across heterogeneous ward populations, and forms the empirical foundation for the dynamic-warning literature returned to in 7.3.

For panic disorder, Meuret and colleagues²¹⁴ conducted ambulatory monitoring of 43 patients across 1,960 h, capturing respiratory rate, tidal volume, and end-tidal CO₂ during 13 spontaneous panic attacks. Respiratory instability — measured as cycle-to-cycle variability and minute-ventilation drift — rose progressively during the ~47-min window before the patient reported attack onset. The finding has implications beyond panic disorder: it demonstrates that a clinical event *experienced* by the patient as instantaneous (“unexpected”) has, in fact, a measurable respiratory prodrome. The boundary between “spontaneous” and “preceded by detectable change” is observation-dependent.

In the LINK-HF multicentre study, Stehlik and colleagues³²⁶ followed 100 patients with recently decompensated heart failure using a continuous multimodal chest-patch sensor over 90 days. A

machine-learning algorithm built on the multivariate signal — including respiratory rate, heart rate, and activity — flagged impending hospitalisation a median 6.5 days before the adjudicated event, with sensitivity 76% to 88% and specificity 85%. Respiratory rate was one of the strongest individual contributors to algorithm output. The finding is robust to the recurrent confounder of telemonitoring trials: that population thresholds applied to RR alone fail. LINK-HF succeeded by anchoring detection on multivariate trajectory deviation from the patient's own baseline — a design choice that connects directly to Chapter 7.

For COPD exacerbations, two cohorts converge on a multi-day lead. Yañez and colleagues³²⁸ followed 89 patients with home RR monitoring; mean RR rose from 15.2 ± 4.3 to 19.1 ± 5.9 breaths/min over the 5 days before exacerbation-related hospitalisation. The 24-hour prediction-horizon AUC was 0.79 (sensitivity 66%, specificity 93%). Shah and colleagues, in the EDGE study,³²⁷ followed 110 patients across more than 35,000 monitoring days, capturing 361 exacerbation events and reporting sensitivity 80%; specificity 36% in the deployment configuration — high false-alarm rate at population thresholds. Wu and colleagues³²⁹ reported a multiparametric deep-learning model achieving AUC >0.9 on a 67-patient development cohort; this finding awaits independent replication. The cluster of findings shows that RR carries a multi-day prognostic signature for COPD exacerbations, but that the specificity of *population-threshold* alerts is limited — again pointing to the personalised-baseline argument of Chapter 7.

Hours to days: asthma, sleep-disordered breathing, fibrotic interstitial disease

Asthma exacerbations carry a multi-modal acoustic prodrome — wheeze incidence, nocturnal cough rate, and inhale-to-exhale (I:E) ratio shift toward an obstructive (low I:E, prolonged expiration) pattern. Emeryk and colleagues²⁴⁵ reported, in a 149-patient cohort, that the combination of an AI-stethoscope acoustic feature set with a symptom survey achieved an AUC of approximately 92% for exacerbation observation, with the survey contributing the largest incremental information for adults; the AI-acoustic component contributed substantial value in paediatric subgroups. The Park 2025 systematic review and meta-analysis³³⁰ pooled 41 paediatric studies and reported wheeze-detection sensitivity 0.902 (95% CI 0.726–0.970) and specificity 0.955 (95% CI 0.762–0.993) — high pooled performance, but with wide confidence intervals and paediatric-only generalisability.

For obstructive sleep apnea, the 2022 Kim et al. meta-analysis²²⁶ reported a pooled diagnostic-odds-ratio of 57.4 for smartphone acoustic screening against polysomnography reference standard at AHI ≥ 15 , with sensitivity 0.906 and specificity 0.880. The authors emphasise that the device is a screening tool, not a substitute for in-laboratory polysomnography — a positioning explicitly mirrored in subsequent regulatory clearances (Chapter 12).

Idiopathic pulmonary fibrosis (IPF) presents a different temporal-advantage geometry. The lead is not “days before exacerbation” but “months-to-years before functional decline becomes irreversible.” Moran-Mendoza and colleagues³³¹ reported, in a 290-patient prospective cohort referred to a tertiary ILD clinic, that fine “velcro” crackles were present on first auscultation in 93% of patients with subsequently confirmed IPF — the most sensitive single physical-examination finding in the disease. Manfredi and colleagues³³² evaluated the VECTOR digital crackle-detection system in 98 connective-

tissue–disease patients screened for ILD against high-resolution CT reference, reporting diagnostic accuracy 82.6%, sensitivity 88%, and specificity 79%. The temporal advantage is operational: in IPF, antifibrotic therapy reduces FVC decline only on the lung function the patient retains at the moment of treatment initiation; each month of diagnostic delay is irreversible.

Days: COVID-19 pre-symptomatic detection

The COVID-19 pandemic provided the largest natural experiment in pre-symptomatic respiratory observation in the smartphone era. Multiple studies reported that resting respiratory rate, cough rate, and cough acoustic features shift in the days before self-reported symptom onset. Cross-reference Chapter 5.7 (immune function), where the primary citations for cough acoustic classification anchor the broader argument.

Years: Parkinson’s disease prodromal observation

The longest temporal-lead row in this table is the Yang & Katabi 2022 demonstration¹⁹³ that nocturnal breathing patterns, captured via a contactless radio-frequency sensor, allow observation of Parkinson’s disease (PD) over a horizon of years before motor onset. The cohort spanned 11,994 nights of recording, and the deep-learning model — trained on respiratory waveform features without access to other physiological signals — achieved high discrimination for PD on held-out test sets. Critically, the model’s performance generalised across cohorts collected at different sites with different equipment, and showed early-warning trajectories consistent with the prodromal PD literature: REM-sleep–behaviour-disorder onset, autonomic dysfunction, and respiratory-pattern alterations preceding motor diagnosis by years.

The PD signal is not acoustic in the Yang & Katabi study — it is RF-sensor breathing waveform. But the architecture of the temporal advantage is identical: a respiratory observable, structured over time, encodes information about a downstream neurological state. The translation from RF-waveform features to smartphone-acoustic features is a research question (Chapter 9), not a structural barrier.

Synthesis of 7.1

The temporal lead of respiratory observation is not a single number. It spans approximately nine orders of magnitude — from 8 s (Lacuey ictal apnea) to years (Yang & Katabi PD prodromal) — and is observable from a single modality of physiological observation: the structure of breathing over time. This range cannot be matched by any other continuously and non-invasively accessible biosignal. The reason is the subject of 7.2.

6.2 Mechanism: brainstem convergence as the upstream node

The temporal advantage demonstrated in 7.1 is not coincidental. It follows from the structural position of the brainstem respiratory rhythm generators within the autonomic-nervous-system architecture established mechanistically in Chapter 3.

The argument has three steps.

First, the central pattern generator for breathing is not a single nucleus but a coupled network of brainstem oscillators. The pre-Bötzinger complex (preBötC) generates the inspiratory rhythm; the retrotrapezoid nucleus (RTN) and parafacial respiratory group provide expiratory rhythm and central chemoreception (CO_2/H^+ sensing). This network receives convergent inputs from autonomic balance (via the nucleus tractus solitarius and the ventrolateral medulla), peripheral chemoreceptors (carotid bodies, projecting via the glossopharyngeal nerve to the NTS), metabolic-demand signals (via vagal afferents and circulating $\text{CO}_2/\text{O}_2/\text{pH}$), limbic structures (amygdala, insula, anterior cingulate), and cortical pathways (motor cortex, supplementary motor area for voluntary breath control). The convergence is anatomical, not metaphorical: respiration integrates these input streams continuously and structurally encodes their state in the breath-by-breath waveform.

Second, the cardiovascular system is downstream of breathing in a specific, directional sense — not metaphorically, but mechanistically. Yasuma & Hayano demonstrated in 2004⁷² that respiratory sinus arrhythmia (RSA) — the cyclical variation in heart rate locked to the breath — is generated by respiratory modulation of vagal output to the sinoatrial node, with breathing as the *driver* of the rhythm and HRV as the *consequence*. The 2025 Wehrwein review⁷³ has formalised this directional asymmetry by proposing the renaming of RSA as “respiratory heart-rate variability” (RHRV), to make explicit that the variability *belongs* to the respiratory cycle, not to the heart in isolation. The implication for the temporal-advantage argument is direct: any physiological perturbation that affects autonomic balance through the chemoreflex, vagal, or limbic pathways alters breathing first, with HRV as a downstream projection. The HRV measurement frameworks codified by the 1996 ESC/NASPE Task Force standards^{77, 169} describe the consequence; the cause is upstream.

Third, breathing is the *only* vital sign with hybrid autonomic–behavioural control. Heart rate cannot be voluntarily controlled without weeks of biofeedback training. Body temperature and blood pressure are autonomically regulated only. Breathing alone admits continuous voluntary access alongside autonomic regulation — which means that limbic and cortical state, including emotional arousal, cognitive load, and behavioural intent, project directly into the respiratory waveform via the same brainstem network that integrates autonomic and chemoreflex input. The empirical demonstrations in Chapter 3.3 (limbic and cortical interactions) establish this for nasal-breathing entrainment of hippocampal theta and prefrontal gamma rhythms.

The combined consequence is that breathing structure carries advance information about events at every clinical timescale because it sits *upstream* of the observables that other monitoring modalities track. Cardiac arrest on the wards is preceded by RR change for 24–48 h not because RR is a more sensitive measurement than HR — RR is a *measurement of a different physiological process*, one that responds earlier in the autonomic-decompensation cascade. Pre-ictal central apnea precedes EEG seizure onset because the amygdala spreads electrical activity to the medullary respiratory network before the cortical seizure generalises to the scalp. Panic-attack respiratory variability rises 47 min before the patient’s perception because the limbic instability projecting to the breathing rhythm precedes the conscious experience of the attack itself.

The mechanism beat is not new physiology; it is the assembly of physiology that is otherwise dispersed across textbook chapters. The contribution of this scientific foundation is to read the dispersed evidence as a single architectural finding: a signal whose generators sit at the convergence point of autonomic, chemoreflex, metabolic, limbic, and cortical pathways will *structurally* lead other observables — provided one observes its structure, not only its rate.

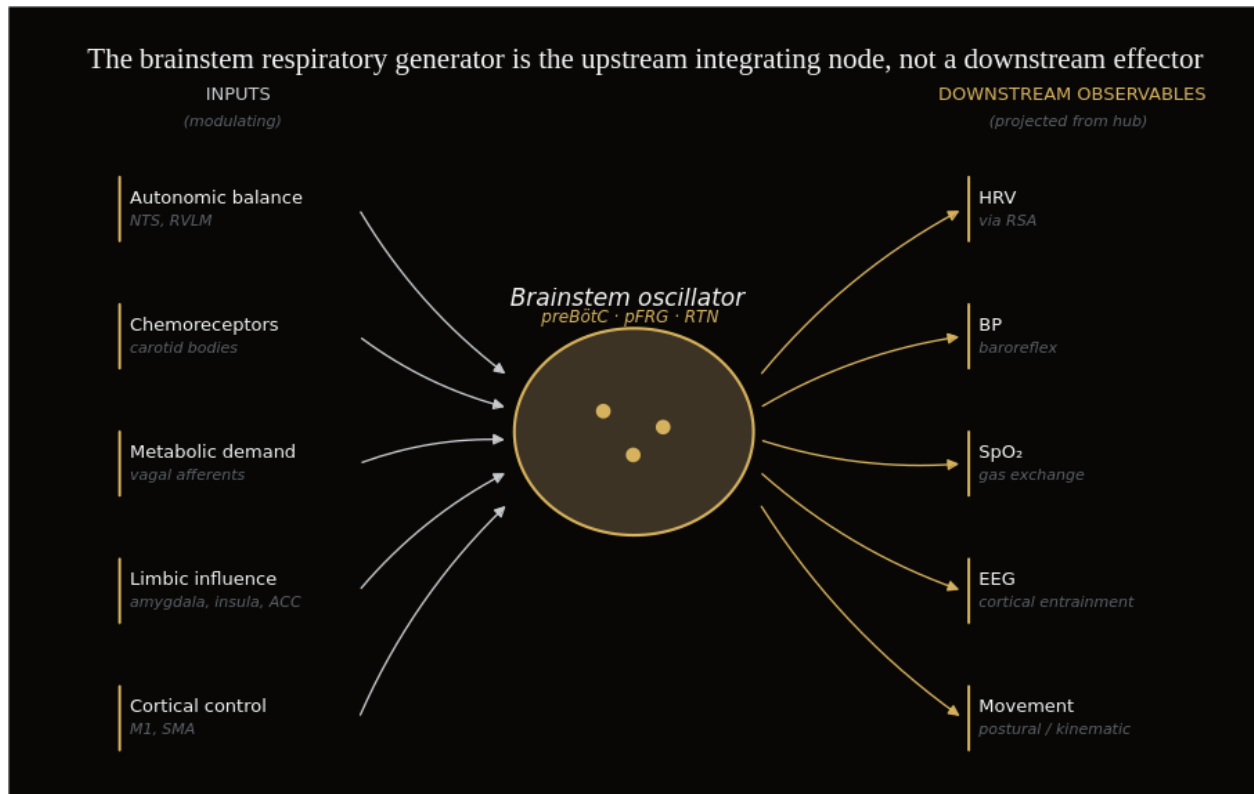


Figure 6.2. The brainstem respiratory generator as upstream integrating node. The preBötzinger complex (inspiratory rhythm) [61, 62] together with the parafacial respiratory group and retrotrapezoid nucleus (expiratory rhythm and central CO₂ chemoreception) [62–64] forms a coupled oscillator that integrates five convergent input streams: autonomic balance (NTS, RVLM), peripheral chemoreception (carotid bodies), metabolic-demand signalling (vagal afferents), limbic state (amygdala, insula, ACC), and cortical control (M1, SMA). The signals tracked by parallel monitoring modalities — HRV, BP, SpO₂, EEG, and movement-related signatures — are projections of this integrated brainstem activity, not parallel inputs [72, 73]. Direct observation of breathing structure therefore accesses the integrating node itself.

6.3 Where the temporal advantage does not hold

The structural argument of 7.2 is not universal. There are clinical contexts where breathing is not the leading signal, and intellectual honesty requires that they be made explicit. Identifying counterexample boundaries strengthens the structural argument by sharpening it; it does not weaken it.

Boundary conditions: where the temporal advantage does not hold			
COUNTEREXAMPLE	LEADING SIGNAL	WHY BREATHING IS NOT THE LEAD	MITIGATION
Cardiac arrhythmia	Electrical (ECG)	Arrhythmia onset is electrical and typically precedes any breathing change	ECG primary; breathing complements
Acute trauma / haemorrhage	Haemodynamic (BP, HR)	Blood loss appears in heart-rate / blood-pressure before respiratory drive	Standard trauma vital signs primary
Sleep cortical arousals	EEG	Brief micro-arousals defined by EEG; respiratory effects may be absent or delayed	EEG-grade staging required
Ambient-noise confounding	(none reliably)	Acoustic capture degrades; signal-quality confidence drops	Quality classifier + multi-modal fallback
Individual baseline variability	personal history	Population thresholds fail; only personal trajectory is informative	Personalised baseline (Chapter 8)

Breathing observation is not a universal first-detector. The temporal-lead claims of Chapters 6-7 apply within the bounded domains where the brainstem network is the upstream node.

Figure 6.3. Boundary conditions: where the temporal advantage does not hold. Five clinical contexts in which breathing is not the leading signal: primary electrical cardiac events, haemorrhagic shock, cortical micro-arousals, ambient-noise–limited capture, and inter-individual baseline variability. The brainstem-upstream argument (Chapter 6.2) bounds — rather than universalises — the temporal-lead claims of Chapter 6.1; observation outside these bounds requires complementary modalities or personalised baselines (Chapter 7).

Cardiac arrhythmias whose primary mechanism is electrical, not autonomic

In ventricular tachycardia, ventricular fibrillation, and sudden cardiac death secondary to monogenic channelopathies (Long-QT syndrome, Brugada syndrome, catecholaminergic polymorphic ventricular tachycardia), the inciting event is intracardiac re-entry or triggered automaticity, not a brainstem-mediated autonomic shift. The respiratory observation is *downstream* of the arrhythmia (or is absent altogether, in arrhythmic sudden cardiac death) rather than upstream. The temporal advantage in 7.1 — cardiac arrest on the wards — refers specifically to the ward populations in whom arrest is preceded by a multi-hour decompensation cascade (sepsis, respiratory failure, hypoperfusion); for primary electrical events without such a cascade, breathing observation is not the lead.

Acute trauma and haemorrhagic shock

In trauma with major blood-volume loss, the cardiovascular response — tachycardia, narrowed pulse pressure, eventual hypotension — precedes the respiratory adjustment. RR rise occurs as compensation for metabolic acidosis once tissue perfusion has fallen, not as an early-warning lead. The temporal hierarchy in haemorrhagic shock runs cardiovascular → metabolic → respiratory, the inverse of the ward-deterioration hierarchy. The RR rise in trauma is downstream confirmation, not upstream observation.

Sleep architecture events whose primary observable is cortical

In polysomnography-defined cortical arousals — particularly those that do not result in a respiratory event — EEG change precedes any respiratory pattern shift. The temporal hierarchy in PSG-scored events is therefore EEG-first for the event class “cortical arousal without respiratory consequence.” For respiratory-event-associated arousals (apnea-hypopnea-related), the order reverses, with the respiratory event preceding the arousal, but the boundary case is real and is acknowledged in Chapter 5.3.

Confounders that limit lead-time reliability without negating the structural argument

Three are operationally important. *Ambient noise* (cross-reference Chapter 9.4): smartphone-microphone-captured respiratory acoustics in real-world environments is corrupted by traffic, conversation, HVAC noise, and recording-position artefacts; lead times reported in controlled cohorts may degrade in deployment without context-aware noise modelling. *Individual baseline variability* (cross-reference Chapter 7.3): population-threshold alerts on RR fail when individual baseline RR ranges from 10 to 22 breaths/min in healthy adults — a 4-breath/min change that is highly informative for one patient is within-baseline for another. *Recording-position artefacts*: prone, supine, and lateral positions alter the acoustic transfer function, and recording during speech or eating is qualitatively different from quiet-breathing observation.

Counter-evidence on the COPD-RR-telemonitoring finding

The CHROMED trial (Walker 2018)³³⁴ randomised 312 older patients with COPD and comorbidities to forced-oscillation-technique-based home telemonitoring (which included respiratory-rate trajectory) versus standard care, with time-to-first-hospitalisation as the primary endpoint. The trial did not demonstrate a benefit on the primary endpoint. The supporting evidence base — that RR rises observably for several days before COPD exacerbation — is well-attested in the literature; the CHROMED null result is treated here as a contextual qualifier rather than a refutation.

What CHROMED demonstrates is that *the existence of an early-warning signal does not, by itself, translate into a hospitalisation-reduction trial result*. The translation requires architectural elements that population-threshold telemonitoring lacks: personalised baselines (Chapter 7), actionable response pathways (Chapter 12), and patient and clinician engagement protocols. The point is not that breathing observation fails in COPD; it is that signal observation is necessary but not sufficient for outcome change.

Closing synthesis: from temporal advantage to personalised baselines

The temporal advantage of respiratory observation is **structurally grounded** in the upstream position of the brainstem rhythm generators (7.2) but **not universal** (7.3). It holds wherever upstream-convergent autonomic, chemoreflex, or limbic disturbances precede measurable downstream events; it does not hold in primary-electrical, primary-haemorrhagic, or primary-cortical scenarios. The honest framing is therefore conditional: respiratory observation provides a multi-timescale lead in the *majority* of clinical decompensation cascades — not in all of them.

The counter-evidence from CHROMED introduces the chapter's bridge to the next. The same signal — RR trajectory — that rises observably 5 days before exacerbation in the Yañez and Shah cohorts^{327, 328} failed to reduce hospitalisations in CHROMED.³³⁴ The single largest difference in trial design is the threshold structure: population-fixed thresholds in CHROMED versus patient-anchored deviation detection in the cohorts that showed lead time. This difference is not a methodological footnote; it is the central architectural choice in continuous physiological monitoring, and it is the subject of Chapter 7.

The cleanest empirical demonstration that personalised, dynamic warning outperforms static population thresholds — for exactly the type of multivariate ward signal in which RR is the single most informative component — comes from Zhu and colleagues' DyniEWS study.^{335, 336} In a multicentre retrospective cohort of 13,319 postoperative cardiac patients across four UK hospitals (442,461 vital-sign observations, 4,234 adverse events), the dynamic individual vital-sign trajectory early-warning score (DyniEWS) achieved an AUC of 0.80 for postoperative deterioration, compared with an AUC of 0.73 for the snapshot National Early Warning Score (NEWS) on the same data. The 0.07 AUC gain is not an incremental improvement; it is the operational signature of a structural design choice. Population thresholds compute the score from the *current* vital-sign value; personalised dynamic warning computes it from the trajectory of the patient relative to their own recent baseline. The same data, the same vital-sign panel, the same outcome — a different architecture, and a measurably better prognostic decision boundary.

This is the cross-system temporal-advantage demonstration. Churpek^{324, 325} established that respiratory rate is the single most informative ward vital sign. Stehlik (LINK-HF)³²⁶ demonstrated that a multivariate algorithm anchored on personal baselines flags HF decompensation a median 6.5 days before the event. Zhu (DyniEWS)^{335, 336} showed that the architectural shift from population to personal thresholds delivers a 0.07 AUC improvement on a postoperative cohort. The three findings cohere into a single argument: the temporal advantage of respiratory observation, established in 7.1 and grounded in 7.2, is *operationally accessible* only through baseline-anchored detection. Static population thresholds — the dominant paradigm of remote-monitoring trials in the 2010s — discard most of the lead-time information by treating heterogeneous individuals as a homogeneous reference distribution.

The next chapter develops this argument empirically: it examines the BEAT-HF, TIM-HF, and TELE-HF trials (population thresholds, null primary endpoints) in contrast to TIM-HF2, LINK-HF, and DyniEWS (personalised baselines, positive findings), and establishes the empirical rule that the failure of the 2010s telemonitoring trials was not a failure of remote monitoring as a paradigm — it was a failure of population thresholds applied to a structurally personalised signal.

7. Personalized baselines vs population thresholds

The story of remote physiological monitoring over the past fifteen years can be read as a single dialectic. A first generation of telemonitoring trials, built around population-derived alert thresholds — respiratory rate above 20 per minute, weight gain above two kilograms, heart rate above a fixed cut-off

— failed, repeatedly and at scale, to produce the clinical benefits their effect sizes had been powered to detect.^{337, 338, 339} A second generation, built around personalized baselines and trajectories of deviation, succeeded.^{326, 340, 341} Between the two lies a structural rather than technological lesson: the prognostic information in continuous physiological streams resides in the *departure of an individual from their own stable state*, not in the crossing of any population-derived line. Chapter 7 develops that lesson, anchors it to the empirical demonstration that breathing structure has fingerprint-like individuality at the 24-hour scale,²⁶⁷ and traces the implications for any continuous observation system architected around respiratory signal — including the architecture this monograph motivates in Chapter 10.

Population-threshold telemonitoring trials in heart failure: null primary endpoints

TRIAL	COHORT	PRIMARY ENDPOINT	RESULT	REFERENCE
BEAT-HF	n = 1,437	Readmission within 180 days	Null (no significant reduction)	Ong et al., 2016
TIM-HF	n = 710	All-cause mortality	Null (HR 0.97, n.s.)	Koehler et al., 2011
TELE-HF	n = 1,653	Re-hospitalization or death	Null (no benefit)	Chaudhry et al., 2010
CHROMED	n = 320	Hospitalization for AECOPD	Null (28% vs 31%, n.s.)	Walker et al., 2018

Pattern
 All four trials applied a single population-derived threshold (e.g., RR > 25, weight gain > 2 kg) uniformly across patients. None achieved its primary endpoint.
Implication: the failure was structural, not technological.
Personal trajectory — not population threshold — carries the signal.

Figure 7.1. Population-threshold telemonitoring trials in heart failure: null primary endpoints. Four large RCTs of remote physiological monitoring built around population-derived alert thresholds produced null findings on pre-specified primary endpoints across ~4,100 randomized participants — TELE-HF (n = 1,653), TIM-HF (n = 710, HR 0.97), BEAT-HF (n = 1,437), CHROMED (n = 320). The convergent failure is explained not by intervention intensity or technology, but by the alerting logic: population thresholds drawn from cohort distributions miss the substantial variation in individual healthy distributions. Personalized-baseline approaches using the same physiological variables (LINK-HF, TIM-HF2, DyniEWS) achieved positive findings on the same outcomes (Chapter 7.2).

7.1 The failure of population thresholds in remote monitoring

Three large randomized controlled trials, conducted across the United States and Europe between 2010 and 2018, established the empirical baseline for what population-threshold remote monitoring can and cannot achieve in heart failure. None of the three demonstrated clinical benefit on its primary endpoint.

TELE-HF (Telemonitoring to Improve Heart Failure Outcomes; Chaudhry et al., *NEJM* 2010, n=1,653) randomized recently hospitalized heart-failure patients to a daily telephone-based interactive voice-response system that prompted self-reported symptoms and weights, with population-threshold rules triggering nurse contact.³³⁷ The trial was negative on its composite primary endpoint of readmission or

death at 180 days; intervention and control arms were statistically indistinguishable. **TIM-HF** (Telemedical Interventional Management in Heart Failure; Koehler et al., *Circulation* 2011, n=710) added structured remote transmission of weight, blood pressure, ECG, and self-rated wellness via dedicated telemedical hardware, with daily review by a specialized telemedical centre using population-threshold alerting; the trial was again negative for its primary endpoint of all-cause mortality.³³⁸ **BEAT-HF** (Better Effectiveness After Transition–Heart Failure; Ong et al., *JAMA Internal Medicine* 2016, n=1,437) tested a multi-component remote-monitoring intervention — telephone coaching plus daily home transmission of weight, blood pressure, heart rate, and symptom reports against population-threshold criteria — and reported no reduction in 180-day all-cause readmission, the primary outcome.³³⁹

The convergent failure across three independently designed, well-powered trials, totalling almost four thousand randomized patients, is not explained by intervention intensity, by fidelity of nurse follow-up, or by the technology layer used to transmit measurements. Each trial, examined in isolation, had a defensible biological rationale: respiratory rate, weight, and blood pressure are well-established correlates of decompensating cardiac function, and elevations precede hospitalization by days. The failure is explained instead by what those rationales were operationalized as. A respiratory rate of 22 per minute, a weight gain of 2.3 kilograms, a systolic blood pressure of 165 mmHg — these are the alerts the trials issued and the events to which clinical contact responded. They are not, in any of the three trials, *individual departures*. They are population thresholds, drawn from cohort distributions, applied to people whose own healthy distributions vary considerably more than the threshold itself permits to be detected.

Two failure modes follow mechanically. At one end of the population distribution, individuals whose physiological norms run high — an unremarkable resting respiratory rate of 19 per minute, a habitual weight that fluctuates by a kilogram over days — generate constant nuisance alerts that contribute to alert fatigue and degrade clinician response quality. At the other end, individuals whose norms run low — a habitual resting respiratory rate of 13 per minute that drifts to 17 over five days — are not flagged at all, because no population threshold has been crossed; yet the four-breath shift relative to that individual's stable baseline carries the prognostic signal the trial was designed to detect.^{326, 341} The trials were not insensitive to physiological deterioration; they were applying a population-scale ruler to a personally-scaled phenomenon, and the resulting signal-to-noise ratio was too poor for the downstream clinical workflow to extract benefit. The pattern is reproducible in adjacent prognostic-modelling literature: systematic reviews of cohort-derived prediction models repeatedly document that performance deteriorates sharply on external validation when the model has no per-individual reference frame to anchor its predictions.³⁴²

The same pattern is reproducible outside heart failure. **CHROMED** (Chronic Obstructive Pulmonary Disease — Home Monitoring of Exacerbation Detection; Walker et al., *AJRCCM* 2018, n=312) tested home telemonitoring of physiological and symptom variables against population-threshold alerting in COPD, with no reduction in time to first hospitalization or mortality at 9 months.³³⁴ CHROMED is treated as a boundary-condition reference for respiratory-disease prediction in Chapter 6 and is invoked here as the COPD-domain analogue of the BEAT-HF / TIM-HF / TELE-HF triad: a fourth example of population-threshold telemonitoring failing to translate plausible biology into clinical outcomes when the threshold logic is built around cohort distributions rather than individual deviation.

The pattern is empirical and structural, not contingent. Across cardiac and respiratory domains, across continents, across hardware generations and intervention designs, population-threshold remote monitoring has produced one statistically null outcome after another.^{337, 338, 339, 334} The natural inference is not that remote monitoring is ineffective — the LINK-HF and DyniEWS results discussed in 8.2 establish that monitoring *can* succeed when its alerting logic is restructured. It is that the alerting logic these trials used was wrong.

7.2 The success of personalized approaches

The contrast appears in trials that, while continuing to draw on the same physiological variables, replaced cohort-derived thresholds with individual reference frames.

Personalized-baseline trials succeed where population thresholds fail			
TRIAL	COHORT	DESIGN	RESULT
TIM-HF2	<i>n</i> = 1,538	Personal-trajectory tele-management vs usual care	Days lost to unplanned cardiovascular hospitalization or death ↓ (<i>p</i> = 0.046)
LINK-HF	<i>n</i> = 100	Wearable-derived personal baseline	Heart-failure exacerbation predicted with sensitivity 76–88%, lead time 6.5 days
DyniEWS	<i>n</i> = 1,500+	Dynamic personalized early-warning score	Improved deterioration prediction vs static population thresholds (AUC ↑)

Structural difference

All three trials replaced population thresholds with a personal-trajectory comparator. In each case, deviations from the individual's own stable baseline carried predictive signal, even when the absolute values fell within "normal" population ranges.

The structural feature — personal baseline + drift detection — is the unit of measurement.

Figure 7.2. Personalized-baseline trials succeed where population thresholds fail. Three convergent trials — TIM-HF2 (*n* = 1,538), LINK-HF (*n* = 100), DyniEWS (*n* = 1,500+) — show that personal-trajectory comparators outperform population-derived thresholds for deterioration prediction. Deviation from an individual's own stable baseline carries predictive signal even when absolute values remain within "normal" population ranges. The unit of measurement is drift vs. self, not value vs. cutoff.

TIM-HF2 (Koehler et al., *Lancet* 2018, *n*=1,538) is the most instructive transitional case. The trial enrolled higher-risk heart-failure patients than its predecessor TIM-HF, retained the same physiological-transmission infrastructure, and added a structured nurse-led remote care intervention with closer telemedical-centre integration and protocolized escalation. The primary endpoint — percentage of days lost due to unplanned cardiovascular hospital admissions or all-cause death — was reduced (4.88 % in the intervention group vs 6.64 % in usual care; ratio 0.80, 95 % CI 0.65–1.00, *p*=0.046).³⁴⁰ The result is modest but real, and it sits inside the population-threshold paradigm: TIM-HF2 succeeded by

intensifying the response to threshold crossings, not by changing the alerting logic. It establishes the upper bound of what well-executed population-threshold remote monitoring can achieve when paired with structured care delivery.

LINK-HF (Stehlik et al., *Circulation: Heart Failure* 2020, n=100) crossed the structural boundary. The study deployed a multimodal adhesive chest patch that continuously sampled heart rate, respiratory rate, activity, body posture, sleep, and skin temperature; for each enrolled participant, the analytic engine constructed a personal physiological reference frame from the early monitoring period and subsequently flagged individual deviations from that reference frame as candidate decompensation signals. In the prospective cohort, the system anticipated unplanned heart-failure hospitalization at a median of **6.5 days** before admission, with sensitivity of 76–88 % and specificity of 85 %, and an overall area under the receiver operating characteristic curve of 0.80 in cross-validated analysis.³²⁶ The signal was not in any one variable in isolation; it was in the *coordinated departure* of a small set of physiological streams from the individual's stable multivariate baseline. Cross-reference to the temporal-advantage analysis in Chapter 6, where LINK-HF anchors the heart-failure row of the cross-system temporal-lead table, applies here.

DyniEWS (Zhu et al., *Resuscitation* 2020, n cohort >7,000) extended the same structural logic to general inpatient deterioration. The dynamic individualized early-warning trajectory score, computed from continuously updated patient-specific reference distributions across vital signs (including respiratory rate), showed an area under the curve of **0.80** for predicting in-hospital deterioration, against **0.73** for the population-threshold National Early Warning Score (NEWS) on the same cohort.³⁴¹ The CROSS-01 Bridge finding, primary-mapped in Chapter 6, anchors DyniEWS as the cross-system demonstration that personalized trajectory analysis improves prognostic discrimination over snapshot population thresholds; Chapter 6 develops the temporal-advantage interpretation, and this document does not restate those findings here.

A subsequent comparative methodology paper from a separate group reviewed dynamic and static early-warning architectures across multiple inpatient cohorts and reached a convergent conclusion: alerting systems that update their reference distribution to the individual outperform alerting systems that compare each measurement to a fixed population cut-off, with effect sizes of comparable magnitude across the studies reviewed.³⁴³

The mechanism beat is straightforward. The prognostic information in a continuous physiological stream resides in the *departure of an individual from their own stable state*, not in the crossing of any cohort-derived line. A nocturnal respiratory rate of 18 per minute is unremarkable for one individual whose stable nocturnal rate sits at 17, and is a measurable deviation for another whose stable nocturnal rate sits at 13. The same scalar value carries different prognostic content for the two individuals, and a population threshold cannot recover that difference. A personalized reference frame can. The mechanism does not require sophisticated mathematics — at the simplest level, a per-individual z-score against a stable reference distribution discriminates better than a population threshold whenever inter-individual variance exceeds within-individual variance, which is the empirical regime continuous physiological signals appear to inhabit.^{136, 326, 341}

The transition from TIM-HF to TIM-HF2 to LINK-HF traces the gradient. TIM-HF failed at population-threshold alerting. TIM-HF2 retained population-threshold alerting and improved outcomes by intensifying the care-delivery response. LINK-HF replaced population-threshold alerting with personalized-baseline trajectory analysis and produced lead times of clinical relevance at smaller cohort scale. The structural step — from cohort threshold to individual reference frame — is associated with the discriminative gain.

7.3 The structural individuality of breathing

The clinical evidence in 8.1 and 8.2 establishes that personalized baselines outperform population thresholds for remote monitoring outcomes. The empirical evidence in this section establishes the deeper finding that justifies the architecture: breathing structure itself is *individuated* — sufficiently distinct between people, and sufficiently stable within a person, that the prognostic signal of departure from individual baseline is recoverable in principle, not only in well-instrumented heart-failure cohorts.

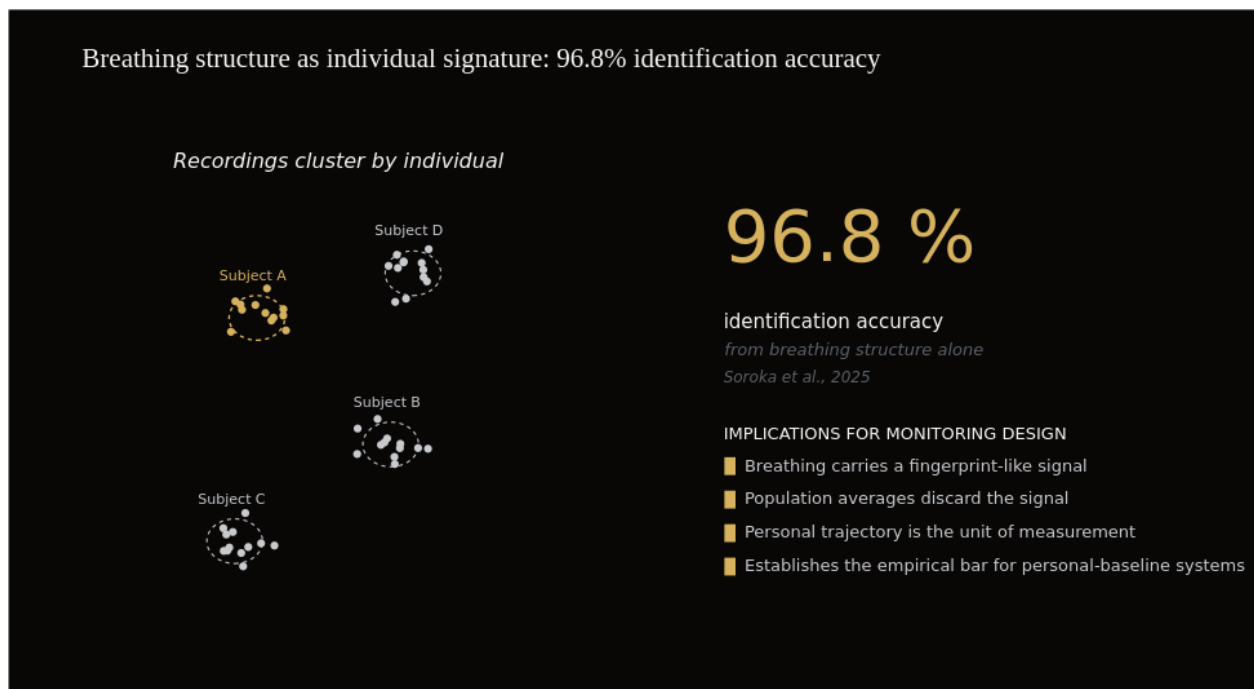


Figure 7.3. Breathing structure as individual signature — empirical demonstration. Soroka and colleagues recorded continuous nasal airflow over 24 hours from 97 healthy adults, extracted 24 structural parameters, and achieved 96.8% individual-identification accuracy. The signature persisted across re-recording intervals of weeks, months, and approximately two years (test-retest accuracy 95.24%). This finding establishes the empirical bar that any personalised respiratory observation architecture must meet — joint intra-person consistency and inter-person separability — and identifies population averages as a discard of the signal.

The anchor is **Soroka et al., *Current Biology* 2025.**²⁶⁷ The investigators recorded continuous nasal airflow over 24 hours from 97 healthy adults using a wearable pressure-based device, extracted a structural representation of each individual's respiratory pattern, and tested whether an individual could be re-identified from a held-out segment of their own 24-hour recording against the cohort. The system identified individuals with **96.8 % accuracy** — a level of within-cohort discrimination conventionally associated with biometric modalities such as fingerprints, gait, or iris pattern. The investigators replicated the finding across recordings collected weeks apart, demonstrating that the individuating

structure persisted across re-recordings rather than reflecting transient state. The result is, to date, the strongest empirical demonstration that breathing pattern carries person-specific structural signature at scales relevant to continuous monitoring.

The mechanistic interpretation runs through the architecture of the respiratory control system itself. Nasal airflow at the 24-hour scale is a convolution of brainstem rhythm-generation properties, upper-airway anatomy, neuromuscular drive characteristics, autonomic state distribution, sleep–wake architecture, and habitual postural and behavioural patterns — most of which are individually stable on the timescale of weeks to months and individually variable across the population. A representation that integrates information across this convolution captures a high-dimensional signature that is harder to share between two arbitrary humans than between two recordings of the same human, and the Soroka result is the empirical confirmation of that structural prediction.

Convergent evidence comes from the breathing-biometric verification literature. **Bui and colleagues** (*Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2021) demonstrated multimodal acoustic breathing-signature verification with an equal-error rate of 1.4 %, derived from short breathing-pattern recordings collected on consumer-grade microphones across multiple sessions per participant.¹³⁸

Panchagnula and colleagues (2024, n=94) extended the demonstration to exhaled-breath physics, confirming verification-grade discrimination from a separate signal channel in an independent cohort.¹³⁹

Across these studies, the empirical regime is consistent: within-individual breathing structure is sufficiently stable across sessions to support identity verification, and between-individual breathing structure is sufficiently distinct to support identity discrimination. The biometric-task framing is methodologically convenient — the task forces an explicit measurement of intra-individual versus inter-individual variance — but the underlying physiological observation is the one that matters here. Breathing structure satisfies the regime in which personalized baselines carry information that population thresholds cannot.

The methodological literature on respiratory variability analysis converges on the same conclusion. Bravi, Longtin, and Seely's classification of variability-analysis techniques across physiological signals identified respiration as one of the signal channels in which structured intra-individual variability — patterned over time, non-random, and partly distinct from heart-rate variability — carries information not recoverable from scalar summary statistics.¹³⁶ Soroka's 24-hour identification result establishes that the requirement of within-individual variance smaller than between-individual variance is empirically met at the 24-hour scale; Bui's across-session verification result extends the supporting evidence to multi-session timescales; multi-week and multi-month longitudinal stability remains an open empirical question that this monograph positions as the central research agenda for any continuous breathing-observation system, rather than as a settled property.

Three implications follow for monitoring system design. First, population-derived breathing reference distributions waste the structural information that distinguishes the individual; they convert a fingerprint-like signal into a smear of cohort statistics. Second, the same structural representation that supports between-person identification can — at the appropriate scale of analysis — support detection of an individual's deviation from their own stable structural state, since both tasks reduce to comparison against a personal reference distribution. Third, any system that proposes to operationalize personalized

breathing observation must satisfy two minimum empirical tests, identified in the validation framework outlined in this document's Chapter 10 architectural treatment: that representations of the same individual recorded across time cluster together (intra-person consistency) and that representations of different individuals separate (inter-person separability). Soroka's 96.8 % identification accuracy operationalizes both tests jointly into a single empirical figure of merit; any candidate respiratory observation architecture is evaluable against the same standard.

7.4 Implications for monitoring system design

The synthesis of 8.1, 8.2, and 8.3 yields three structural implications for any continuous physiological observation system that proposes to operate on respiratory signal at population scale.

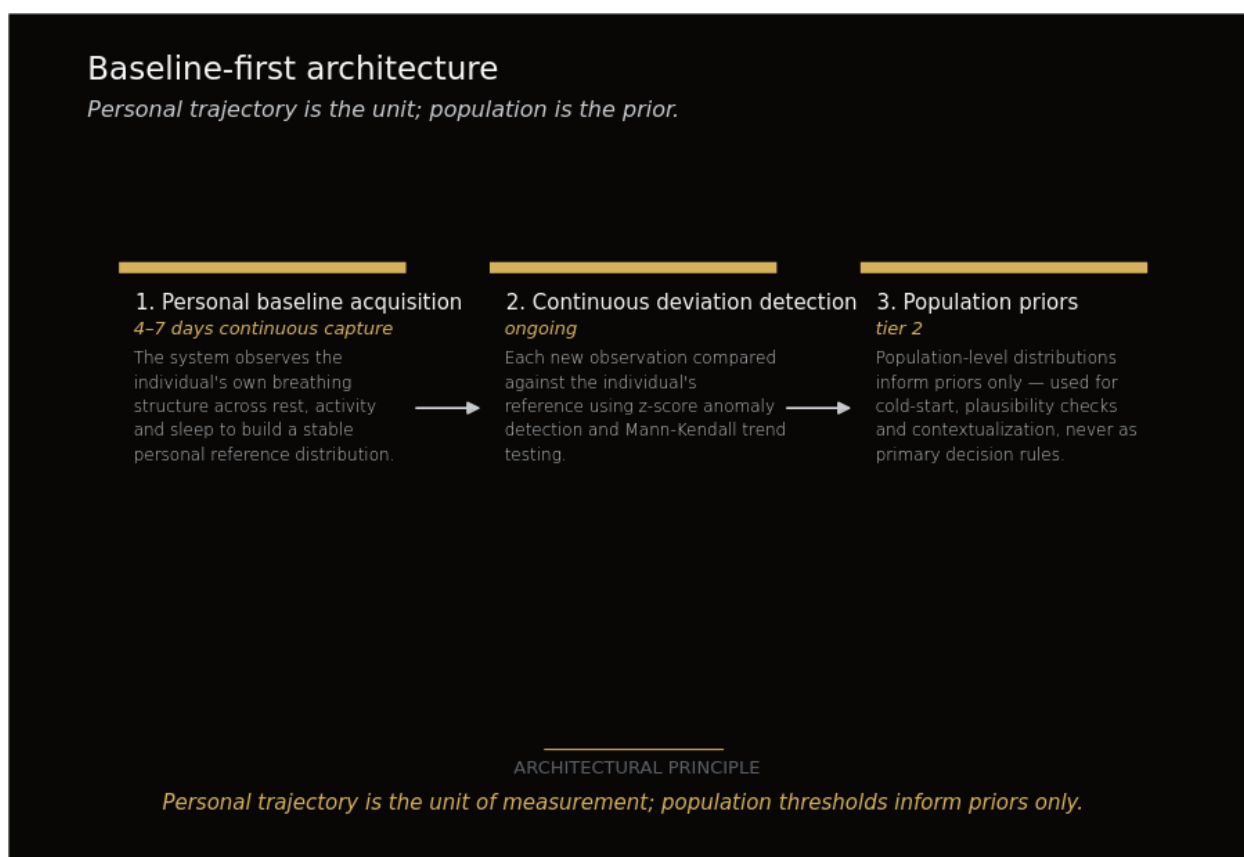


Figure 7.4. *Baseline-first architecture: personal trajectory is the unit, population is the prior. Conventional respiratory monitoring compares each measurement against population-level thresholds; Atum inverts this hierarchy. The individual's own breathing structure — captured continuously across rest, activity, and sleep over a 4-7-day acquisition window — becomes the reference distribution, with subsequent observations compared via z-score anomaly detection and Mann-Kendall trend testing. Population priors retain a subordinate role (cold-start, plausibility bounds, cohort context) but never function as primary decision rules.*

Implication 1: baseline-first architecture

The observation system is organized around a per-individual reference distribution acquired during a defined initial monitoring period, against which subsequent measurements are evaluated as deviations or trajectories. Population-derived reference distributions enter the architecture as priors and as similar-individual reference cohorts, not as alerting thresholds. This organization is the architecture compatible with the empirical evidence in 8.1 and 8.2 — population-threshold alerting failed across four large RCTs,

personalized-baseline alerting succeeded in the trials that adopted it — and with the empirical evidence in 8.3 that breathing structure carries person-specific signature at relevant timescales. The architectural treatment is developed in this scientific foundation's Chapter 10.

Implication 2: temporal stability as a design requirement

The acquisition of a meaningful personal reference distribution requires a minimum monitoring window during which the individual's habitual structural pattern is sampled across enough physiological and behavioural contexts to support deviation detection. The within-session variance of structured respiratory features, after context normalization, must remain smaller than the between-individual variance for the personalized reference frame to be informative. Empirical evidence supports the existence of structured within-session and across-session respiratory signatures;^{138, 267, 136} the extrapolation to multi-week and multi-month longitudinal stability, in unconstrained real-world conditions, remains the central open empirical question. A working acquisition window of approximately 4–7 days is consistent with the stability evidence available at the time of writing and with the LINK-HF design,³²⁶ but the parameter is properly an output of empirical validation, not an *a priori* design constant. Any deployment that flags individual deviations before this acquisition is complete is operating on insufficient reference distribution and should expect false-positive rates approaching those of population-threshold systems.

Implication 3: population enrichment, not population substitution

Population data continues to play an irreplaceable role in personalized observation — but its role is to *enrich* the personal baseline through informative priors, similar-individual reference cohorts, and population-distribution constraints on the plausibility of observed deviations. It is not to *substitute* for the personal baseline. The methodological precedent is well-established in adjacent domains. Real-world wearable heart-rate norms from the Health eHeart cohort showed that population distributions of resting heart rate are wide enough that any single threshold imposes meaningful misclassification at both tails;³⁴⁴ the same wide-distribution pattern recurs across instrumented physical-activity measurements at population scale^{345, 346} and across body-composition reference data from large national surveys,³⁴⁷ indicating that the wide-individual-variance regime is not specific to cardiac signals but a general property of continuous physiological measurement at the cohort level. Cardiorespiratory-fitness reference standards from the FRIEND registry and the Mayo cohort are most informative when used as percentile contexts for an individual's tracked trajectory rather than as binary cut-offs;^{230, 348} Spiegelhalter's "effective age" methodology generalizes the same insight by mapping any single individual's risk profile onto a continuous population scale that is informative precisely because it is not used as a threshold.³⁴⁹ The same population-as-prior logic underpins the long-running tradition in epidemiology of using single-item self-rated health as a continuous individual indicator that carries independent prognostic information for all-cause mortality, beyond what any binary cut-point achieves.^{350, 351} The accompanying methodological literature on translating cohort hazard ratios into individual-scale measures further reinforces the architectural point: the population distribution is the context, the individual trajectory is the observable.³⁵²

A related methodological caveat applies to how the population reference distributions themselves are constructed. The literature on self-reported versus device-measured health metrics has documented systematic biases, including in domains directly adjacent to respiratory monitoring: self-reported sleep duration, for example, diverges meaningfully from objectively measured sleep duration in the same individual, with the divergence varying by demographic and clinical strata.³⁵³ The same divergence pattern reproduces in subsequent validation work using wrist-accelerometer reference, indicating that the bias is not an artefact of the older polysomnography-based comparison but a structural feature of self-reported versus instrumented physiological measurement.³⁵⁴ A further complication: population reference distributions for the same physiological variable shift secularly over time, with self-reported sleep duration in U.S. adults documented to have changed measurably across the period 1985–2012,³⁵⁵ implying that any single cross-sectional cohort distribution is a snapshot of a non-stationary process. A personalized observation architecture that constructs its priors from instrumented continuous measurement, rather than from self-report or single-snapshot clinic measurement, inherits a different and generally tighter reference distribution; the same architecture, by anchoring on the individual rather than on a fixed cohort, is also robust against the secular drift that erodes the validity of any time-fixed population threshold. The downstream consequence is that the population-prior layer in the architecture should be sourced from measurement modalities of the same instrumented type as the individual signal, and refreshed against contemporary cohort distributions rather than a fixed reference — methodological constraints with direct implications for dataset design, developed in this monograph's Chapter 8 treatment of the dataset gap.

A boundary condition from the prognostic-modelling literature reinforces the architectural requirement. Wynants and colleagues, in a *BMJ* systematic review of COVID-19 prediction models, documented that the great majority of population-derived prognostic models for the disease — built without per-individual reference frames, evaluated without rigorous external validation — performed poorly or unreplicably when tested in independent cohorts.³⁴² The cautionary point generalizes: prognostic models built on cohort distributions alone, without an individual reference frame to anchor deviation detection, are systematically vulnerable to the same failure modes that BEAT-HF, TIM-HF, TELE-HF, and CHROMED demonstrated empirically.^{337, 338, 339, 334} The architecture this document develops is structured to inherit a different methodological lineage — the LINK-HF, DyniEWS, and Soroka lineage^{267, 326, 341} — in which population data informs the prior and the individual is the reference frame.

The failure of BEAT-HF, TIM-HF, TELE-HF, and CHROMED was not a failure of remote monitoring as a paradigm; it was a failure of population thresholds applied to a structurally personalized signal.

8. The dataset gap: why structural breathing data does not yet exist

The preceding chapters have assembled the conceptual and empirical case for breathing as a continuous, structural observation surface (Chapter 4), reviewed the cross-system evidence base across 16 physiological systems (Chapter 5), examined the temporal advantage of respiratory observation (Chapter 6), and motivated personalized rather than population thresholds for individual-

level interpretation (Chapter 7). What remains is a question of substrate. Each of those analyses depends on data — specifically, on data that captures breathing in the form in which it is hypothesized to carry information: as a continuous, naturally occurring, frame-resolved signal, recorded across both healthy individuals and clinical populations, longitudinally over the timescales at which structural change becomes visible. This chapter examines whether such data exists in the public domain. The conclusion is observational, not accusatory: existing public respiratory datasets are scientifically valuable for the purposes for which they were built; the structural-continuous-natural combination that the present document requires has not yet been assembled, and the reasons it has not been assembled are conceptual, instrumental, and operational in that order.

8.1 Existing public respiratory datasets and their structural limitations

The public respiratory dataset landscape, accumulated over roughly two decades of sustained academic and consortium work, can be organized into five categories along the axes of acquisition modality and recording paradigm. Each category has produced datasets of considerable utility for the questions they were designed to answer; each illustrates, by virtue of its specific design choices, properties that the structural-continuous-natural axis would require but does not currently obtain.

Stethoscope-based auscultation datasets

The most widely used resource in this category is the ICBHI 2017 Respiratory Sound Database, published as part of the International Conference on Biomedical and Health Informatics challenge³⁵⁶ and subsequently described in detail in *Physiological Measurement*³⁵⁷. The corpus contains 920 annotated stethoscope recordings totalling approximately 5.5 hours from 126 participants, with crackle and wheeze annotations at the cycle level, and has become the *de facto* benchmark for respiratory sound classification^{161,145}. HF_Lung_V1, released by Hsu and colleagues³⁵⁸, extended the auscultation paradigm with frame-level inhalation and exhalation phase annotations across 9,765 recordings — to date, the most extensive public corpus of phase-annotated stethoscope auscultation data. Adjacent resources include the SPRSound paediatric respiratory sound database³⁵⁹. These datasets are foundational for the auscultation paradigm; they served — and continue to serve — the task of automated classification of crackles, wheezes, rhonchi, and stridor against expert-annotated references^{161,145}. The acquisition modality is by design clinical: stethoscope contact at the chest wall, short recording windows, postural standardization, and patient cohorts skewed toward documented pulmonary conditions. None of these properties are limitations of the datasets themselves; they are constitutive of the auscultation paradigm.

Cough-only datasets

A second category, developed largely in response to the COVID-19 pandemic, comprises smartphone-acquired cough corpora. The Coughvid dataset²⁹³ released approximately 25,000 crowdsourced cough recordings under permissive licensing. the Coswara cough subsets²⁹⁴ extended this work with operator-recorded smartphone audio at scale. These datasets are smartphone-grade — the acquisition modality is consumer hardware — but they capture cough events specifically, not the natural breathing that surrounds them. Inhalation and exhalation are present in the recordings only incidentally; the annotation

schema and use cases focus on cough characterization rather than on the structure of breath cycles between coughs.

Forced-maneuver smartphone datasets

A third category includes Coswara's deep-breath and count-aloud subsets²⁹⁴, the Cambridge / MIT COVID-19 Sounds dataset³⁶⁰, and adjacent crowdsourced corpora collected during 2020–2022. These datasets share two design choices with the cough-only category — smartphone acquisition modality, distributed crowdsourced recruitment — but require participants to perform structured maneuvers: deep inhalation, sustained vowels, counting from one to twenty, deliberate exhalation. The resulting recordings carry rich acoustic information, but the breathing they capture is performed under instruction rather than observed unprompted; the structural properties of natural breathing are not the recording target.

Adventitious-sound and pathology-fragment datasets

A fourth category includes pathology-fragment collections curated for supervised-classifier benchmarking against expert pathology labels. These datasets typically present excised acoustic events — a wheeze segment, a crackle burst, a rhonchus — with associated pathology labels, and they have been productive for supervised-classifier development. The recording modality varies (stethoscope or contact microphone in most cases), and the temporal granularity is event-level rather than continuous. The structural properties of the breathing context within which these events occur are not part of the annotated signal.

Sleep and clinical polysomnography datasets

A fifth category comprises the polysomnographic cohorts released through PhysioNet³⁶¹ and the National Sleep Research Resource, including the Sleep Heart Health Study³⁶², the Multi-Ethnic Study of Atherosclerosis sleep cohort³⁶³, and other PhysioNet sleep databases³⁶¹. These cohorts are large and longitudinal — SHHS contains data from approximately 6,440 participants — and represent the most substantial respiratory-event datasets in existence. They are, however, polysomnography-event datasets: the recordings annotate apnea, hypopnea, and arousal events derived from airflow, oximetry, and effort signals, not continuous acoustic respiration. PSG-Audio³⁶⁴ — released by Korompili and colleagues — represents a meaningful step in the audio direction, providing 212 polysomnograms with synchronized tracheal and ambient audio recordings; its scope, however, is the sleep-laboratory setting rather than longitudinal natural recording. VitalDB³⁶⁵, the high-fidelity multi-parameter biosignal database from Lee and colleagues, provides 6,388 cases of synchronized vital signs from surgical patients — an exceptional resource for surgical-physiology research, but constrained by its operating-room context to an acutely instrumented patient population rather than ambulant healthy individuals.

Each of these datasets is scientifically valuable for the specific purpose for which it was assembled. The auscultation corpora support classifier development against expert respiratory pathology labels^{358, 356, 357}. The cough corpora support cough characterization and pandemic-surveillance modeling. The forced-maneuver corpora support voice-and-breathing biomarker exploration. The pathology-fragment datasets support targeted classifier benchmarking. The polysomnography cohorts support sleep-

medicine epidemiology and event detection^{361,364,365}. None of these datasets, individually or in combination, was built to capture the structural-continuous-natural axis of breathing observation. The structural absence is not an absence of effort or rigor in dataset construction; it is an absence at the level of the specification itself.

Public respiratory acoustic datasets — five categories

Stethoscope-based	Cough-only	Forced-maneuver smartphone	Adventitious-sound fragments	Sleep / clinical PSG
ICBHI 2017 (n=920) HF_Lung_V1 Respiratory@TR Acoustic Lung Disease	Coughvid (>20K) COUGHVID FluSense ESC-50 (cough class)	Coswara (>2K) Apple Cough Study Various COVID-19 datasets	Sound clip libraries RALE Repository Various textbook recordings	SHHS (n=6,441) MESA (n=2,237) PSG-Audio VitalDB

No public dataset combines smartphone-grade capture, natural unstructured breathing, frame-level phase annotation, longitudinal continuity, and healthy coverage.

See Figure 9.2 for the structural-absence matrix.

Figure 8.1. Public respiratory acoustic datasets across five categories. The publicly available corpora cluster into five families: stethoscope-based clinical recordings, cough-only collections, forced-maneuver smartphone studies, adventitious-sound libraries, and sleep / clinical polysomnography cohorts. None combines all five properties required for a continuous breathing-structure observation model — smartphone-grade capture, natural unstructured breathing, frame-level inspiration/expiration phase annotation, longitudinal continuity, and healthy-population coverage. The structural-absence matrix across these properties is shown in Figure 8.2.

8.2 The structural absence

The structural-continuous-natural axis of breathing observation, as developed in Chapter 4, can be operationalized as a five-property checklist that any dataset substrate would need to satisfy:

- 1. Smartphone-grade acoustic capture.** Acquisition through consumer-class MEMS microphones, not stethoscope contact, contact-microphone chest patches, or chest-strap impedance bands. This property is required because the population scale at which structural breathing observation becomes useful — and the longitudinal cadence at which structural change becomes detectable — depend on hardware that the user already owns¹⁰.
- 2. Natural unstructured breathing.** Recordings of breathing as it occurs without instruction. Not deep-breath maneuvers, not counted exhalations, not voiced breath patterns, and not recordings collected during any procedure that itself alters the natural cadence the recording is intended to characterize.
- 3. Frame-level phase annotation.** Per-breath annotation of inhalation, exhalation, and pause boundaries, at the temporal resolution required to support cycle-level structural analysis. Cycle-level rather than recording-level granularity is what distinguishes structural observation from the rate-or-volume reduction characteristic of earlier eras (Chapter 2).

4. **Longitudinal continuity over weeks to months.** Repeated recordings on each participant across the timescales over which structural drift, contextual variation, and individual baselines become visible. Single-session recordings — however large the cross-sectional sample — cannot anchor within-participant trajectories.
5. **Coverage of healthy individuals as structural anchors.** Inclusion of asymptomatic participants in numbers comparable to clinical cohorts. Patient-only recording, however clinically valuable, cannot establish what natural variation in breathing structure looks like before pathology selects subgroups out of the population.

Walking the existing public datasets through this checklist returns partial fits. ICBHI 2017^{356, 357} and HF_Lung_V1³⁵⁸ provide (c) frame-level phase information but acquire through stethoscope contact (failing (a)), include acute-clinical cohorts (limiting (e)), and do not capture (d) longitudinal continuity. Coughvid, Hyfe, and the Cambridge / MIT COVID-19 Sounds corpora satisfy (a) smartphone-grade capture and recruit non-clinical populations (partial (e)) but capture cough events or instructed maneuvers rather than natural breathing (failing (b)) and do not provide phase-level annotation (failing (c)). Coswara's deep-breath subsets satisfy (a) and partially (e) but explicitly require forced maneuvers (failing (b)). The PSG cohorts — SHHS, MESA, PSG-Audio³⁶⁴ — provide (d) longitudinal duration within a sleep-laboratory context and (e) diverse recruitment, but the available signal is polysomnographic rather than acoustic-continuous, and where audio is available³⁶⁴ it is framed by the sleep-laboratory setting rather than by ambulant naturalistic recording. VitalDB³⁶⁵ and the broader PhysioNet infrastructure³⁶¹ contribute infrastructure, multi-parameter synchrony, and data-sharing standards, but these resources are clinical-instrumental rather than smartphone-acoustic and natural-breathing.

The pattern visible across these partial fits is informative: each existing public dataset satisfies a subset of the five properties, but no public dataset satisfies all five together. The structural-continuous-natural combination has not yet been assembled into a single corpus. The closest individual approximations — HF_Lung_V1³⁵⁸ for frame-level phase annotation, Coswara for smartphone scale, SHHS and MESA for longitudinal cohort recruitment, PSG-Audio³⁶⁴ for synchronized acoustic-physiological pairing — illustrate that each property is independently achievable; the absence is at the level of their conjunction.

It is worth noting that the five properties are not orthogonal additive features that could be solved by simple union of existing corpora. Phase-level annotation acquired through stethoscope contact (HF_Lung_V1³⁵⁸) does not transfer to smartphone-microphone acoustics, because the source-and-transmission acoustics differ substantively between contact-coupled and free-field recording¹⁰. Smartphone-acquired forced-maneuver recordings cannot be relabeled retrospectively as natural breathing because the cadence, depth, and turbulence profiles produced under instruction are physiologically distinct from spontaneous respiration. Polysomnography-event labels^{361, 364} describe physiologic events (apnea, hypopnea, arousal) at second-level granularity that does not project onto the per-cycle inhalation–exhalation–pause structure required for cycle-level analysis. Each property's satisfaction is therefore conditional on the others; partial fits do not compose into a complete substrate by aggregation.

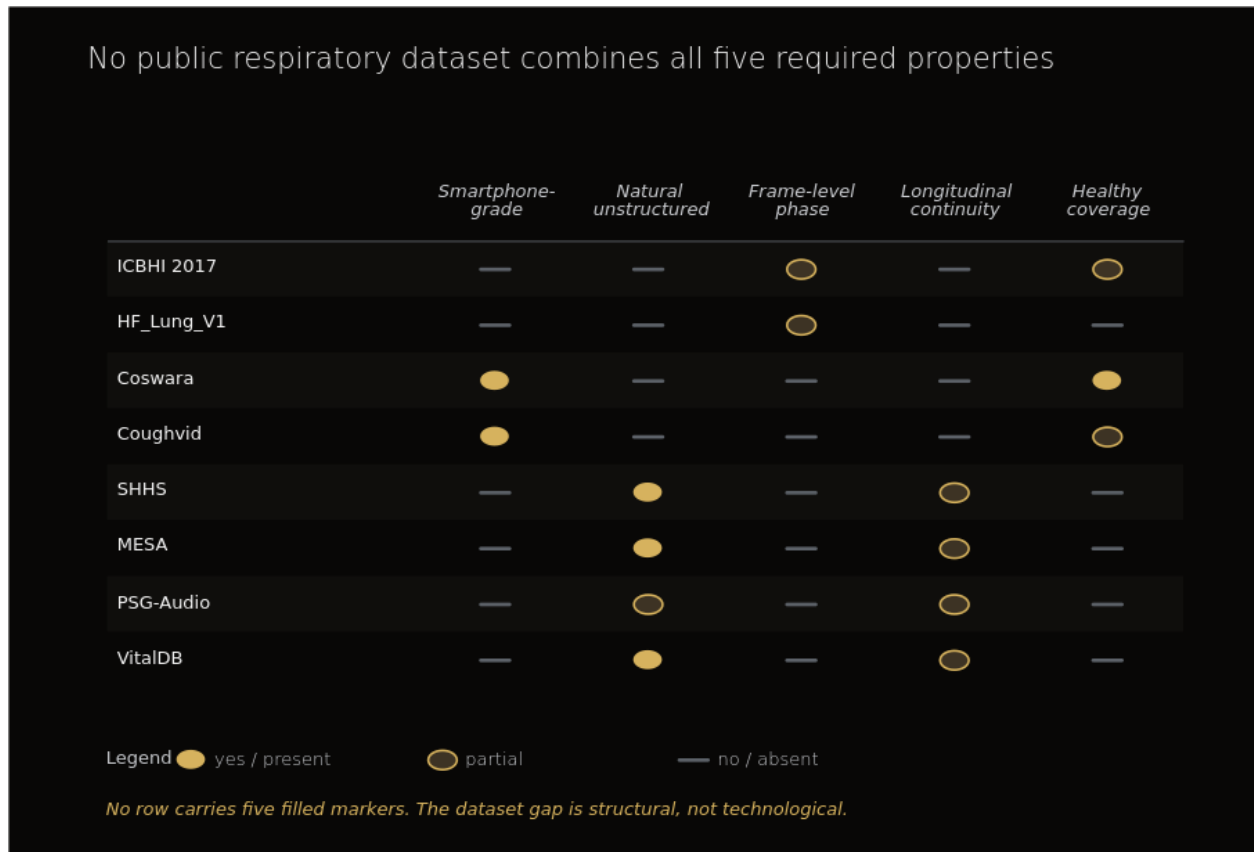


Figure 8.2. No public respiratory dataset combines all five required properties. Rows: nine widely-cited public corpora (ICBHI 2017, HF_Lung_V1, Coswara, Coughvid, SHHS, MESA, PSG-Audio, VitalDB). Columns: properties required for continuous, smartphone-deployed acoustic respiratory observation — smartphone-grade capture, naturalistic unstructured breathing, frame-level phase annotation, within-subject longitudinal continuity, and healthy-cohort baseline. The intersection of all five properties is unoccupied; the gap is structural rather than technological and cannot be closed by re-annotating existing corpora, motivating the dataset-acquisition design in Chapter 9.

8.3 Why this gap exists

The most immediate explanation for the structural absence is conceptual rather than instrumental. Pulmonology, as Chapter 2 establishes, developed across two centuries through five successive eras — volume, gas composition, oxygenation, events, and rate — each of which produced one slice of respiratory information adequate to the clinical questions of its time^{29,30,60}. Across each of those eras, breathing was modeled as a scalar projection: a vital capacity, an end-tidal CO₂, an SpO₂ value, an apnea-hypopnea index, a respiratory rate. Forced maneuvers — Hutchinson’s spirometer paradigm^{29,30} — and patient-only recording were sufficient instruments because the conceptual model treated breathing as a metric. A scalar projection requires only that the signal be sampled adequately to compute the projection; it does not require continuous structural fidelity, naturalistic context, or longitudinal continuity. Datasets built within this conceptual frame inherited its constraints and produced exactly the data the frame asked for.

The structural-continuous-natural axis, as developed in Chapter 4, operates within a different conceptual frame. Breathing is treated not as a scalar but as a structured time series whose informational content lies in cycle-level temporal organization, cross-system co-variation, and within-

individual trajectory rather than in any single derived metric. This conceptual difference imposes corresponding data requirements: continuous rather than episodic, natural rather than instructed, frame-resolved rather than event-aggregated, longitudinal rather than single-session. Adapting the rate-era or auscultation-era datasets to this frame is not a question of better instrumentation of the same data type; it requires a different specification of what is being captured.

The instrumental explanation follows from the conceptual one. Once the data specification is fixed, the instrumentation question is whether sufficient acquisition technology exists to satisfy it. The answer — examined in detail in Chapter 9 — is that smartphone MEMS acoustic capture has reached the signal quality required, with SNR characteristics adequate for respiratory acoustics across the 100 Hz – 15 kHz band¹⁰. The instrumentation gap is therefore not the binding constraint; the conceptual gap was. Until the structural-continuous-natural specification was articulated as the recording target, no consortium had reason to assemble it.

A third-order operational reason layers on top of these. Public dataset assembly in respiratory acoustics has historically been organized around classifier-benchmark milestones — the ICBHI 2017 Challenge³⁵⁶,³⁵⁷, the COVID-19 acoustic challenges of 2020–2022, the cough-detection benchmarks driven by pandemic surveillance — rather than as standing infrastructure investments. Each milestone produced a corpus calibrated to its specific challenge specification: classifier accuracy on adventitious-sound categories, COVID-status discrimination, cough-versus-non-cough labeling. Longitudinal continuity, healthy-cohort coverage, and frame-level phase annotation across naturalistic recording conditions are not the deliverables that classifier benchmarks reward. The incentive structure of academic dataset release has, in this sense, been well-aligned with the conceptual frame of the eras it served, and not yet aligned with the structural-continuous-natural specification this scientific foundation requires.

8.4 The annotation problem

A specific operational obstacle, distinct from the conceptual and instrumental layers, also contributes to the structural absence: respiratory phase annotation is not amenable to the standard mass-crowdsourcing approaches that have driven dataset growth in adjacent acoustic and visual machine-learning subfields. Image classification, speech transcription, and even cough event detection scale through workforces of moderately trained annotators because the perceptual decision required of each annotator is broadly accessible — a person can identify a cat, transcribe a sentence, or detect a cough with limited specialised training. Frame-level inhalation, exhalation, and pause annotation does not have this property.

The CORSA standardization effort^{143, 141, 142} — a coordinated multi-laboratory program funded under the EU BIOMED 1 framework — established the foundational vocabulary, recording standards, and analytical conventions that make computerized respiratory sound analysis reproducible across laboratories. The European Respiratory Society Task Force on lung sound nomenclature¹⁴⁵ further consolidated terminology for adventitious sounds. These efforts demonstrate, by their structure, what high-quality respiratory sound annotation requires: acoustic-domain expertise sufficient to distinguish exhalation from pause in the absence of obvious adventitious features; specialised tooling for frame-level rather than recording-level labeling; cross-rater reliability validation against domain experts; and pre-training

of annotators in respiratory phase recognition^{143, 145, 141, 142}. The methodological literature consistently identifies these as prerequisites rather than optional refinements¹⁴⁶.

The annotation problem is therefore not solvable by scaling the workforce alone. Workforce scale depends on the marginal cost of training each new annotator to the required reliability threshold; for respiratory phase work, that cost is approximately the cost of training a new respiratory-acoustics specialist rather than the cost of training a new image-labeler. The infrastructural prerequisites — specialised annotation interfaces, expert reference panels, validated reliability protocols — also represent fixed costs that no individual laboratory has had explicit incentive to assume in the absence of a demonstrated downstream use case for the resulting dataset. This is one structural reason the dataset gap persists: the cost of building the annotation infrastructure is non-trivial, and there is no academic consortium yet aligned to solve it as a coordinated infrastructure problem rather than a per-project annotation expense.

The contrast with adjacent acoustic annotation tasks is instructive. Cough-event detection scales because a moderately attentive listener can mark the start and end of a cough event with reasonable inter-rater agreement after minutes of orientation; the perceptual feature is acoustically salient and temporally compact. Respiratory phase annotation, by contrast, requires the listener to mark cycle boundaries that are often acoustically subtle — the transition from quiet exhalation to inter-breath pause, in particular, can be ambiguous in the absence of complementary signals such as inductance plethysmography or flow tracing — and to do so consistently across long recordings rather than at sparse event locations^{143, 145, 141, 142}. The cycle structure that makes the data informative is exactly what makes it expensive to label: dense, naturalistic, low-amplitude transitions rather than discrete events. Annotation throughput in this regime is closer to the per-minute rates familiar from clinical sleep-staging and EEG annotation than to the per-second rates of cough or wheeze classification — and the reliability of the resulting labels remains contingent on annotator pre-training and continuous reference-panel calibration rather than on workforce volume.




Figure 8.3. Respiratory phase annotation is not crowdsourceable: cost-and-feasibility curves diverge. Annotation feasibility is plotted against workforce scale. Mass-crowdsourcable tasks (image classification, binary cough detection) reach the operational threshold (≈ 0.95) within ~ 100 annotators; phase-level respiratory annotation asymptotes below threshold even at 1000 annotators. Four structural barriers drive the divergence: frame-level markup tooling, expert disagreement on phase boundaries, mandatory domain knowledge, and within-subject day-to-day variability. Curves are schematic; representative phase-level datasets require expert-tier annotation teams and do not scale linearly with paid workforce.

8.5 Implication

The implication that follows from 9.1–9.4 is that any party assembling structural breathing intelligence must first assemble the dataset substrate. The science is publicly known; the data is not. The mechanistic literature reviewed in Chapter 3, the methodological frameworks consolidated in Chapter 4, the cross-system evidence base assembled across 16 physiological systems in Chapter 5, the temporal-advantage findings synthesized in Chapter 6, and the personalized-baseline framework developed in Chapter 7 — all of these are open-access scientific contributions whose findings can be evaluated by any investigator with library access. What is private, and what represents the binding work to be done, is the dataset substrate that operationalizes these findings into a reproducible signal-processing pipeline. This is the gap addressed by the present scientific foundation.

Two adjacent contemporary developments are worth noting at the level of category. Respiratory acoustic foundation models — including Health Acoustic Representations^{24,366} and the OPERA respiratory-acoustic foundation-model benchmark²⁶ — have begun to provide pretrained acoustic encoders trained on auscultation, cough, and short-form respiratory audio. Field deployment of HeAR for tuberculosis applications has been reported by Google Health and partner organizations²⁵. Open-source methodological tooling is also emerging at the analytical layer; RespirAnalyzer¹⁴⁶, for example, provides packaged implementations of multiscale entropy and multifractal detrended fluctuation

analysis applied to respiratory signals. These foundation models and analytical packages are valuable as transferable representation learners and methodological infrastructure, and category-level integration into structural-breathing pipelines is a natural research direction. They do not, however, substitute for the dataset substrate: a pretrained encoder cannot be evaluated on the structural-continuous-natural axis without a corpus that satisfies the five-property checklist, and no such corpus is yet publicly available for evaluation.

The chapter concludes by identifying the bridge to the chapters that follow. The dataset gap is closable with currently available technology. Chapter 9 examines smartphone MEMS acoustic feasibility in detail and demonstrates that the signal-quality requirement (property (a) of the checklist above) is met by hardware already in the hands of approximately the entire smartphone-using population¹⁰. Chapter 10 develops the observation-layer architecture within which a structural breathing dataset can be assembled, governed, and made useful without invoking diagnostic statements, clinical decision-making, or any framing that the locked invariants of this monograph preclude. The dataset gap is thus a gap that current technology can close, within an architecture that current regulatory framing accommodates; what remains is the assembly work, and the methodological framing within which that assembly is conducted.

9. Smartphone acoustic feasibility

The argument advanced through Chapter 3–Chapter 7 of this document rests on a substrate question: can the smartphone microphone, as a class of sensor, capture respiratory acoustic structure with the fidelity that those analyses presuppose? The question is not whether a smartphone can match the absolute floor noise of a cardiology-grade contact microphone in an anechoic chamber — it cannot. The question is whether smartphone-class MEMS microphones, in the recording conditions a consumer health platform actually encounters, deliver a signal whose spectral content, dynamic range, and temporal resolution are sufficient to extract the respiratory features that Chapter 3–Chapter 5 documented as physiologically informative.

This chapter answers that question in five beats. 10.1 traces the technical evolution of MEMS microphone hardware. 10.2 establishes the frequency-band coverage required for respiratory acoustics and shows that consumer MEMS designs span it. 10.3 reviews the empirical validation literature — smartphone audio against stethoscope, against polysomnography, against spirometry, and through the FDA 510(k) clearance record. 10.4 catalogues the boundary conditions where smartphone capture genuinely struggles, framed as domain-specific engineering problems rather than categorical limits. 10.5 synthesises the chapter's finding: smartphone-grade acoustic capture has crossed the threshold of feasibility for continuous respiratory structure observation, and the remaining engineering work is well-scoped.




Figure 9.1. Smartphone MEMS microphones reached respiratory-grade SNR by ~2019. Two-decade trajectory of consumer MEMS microphone signal-to-noise ratio. The dashed reference at ~70 dBA marks the working sensitivity below which low-amplitude breathing signals (lip-distance SPL ~10–25 dBA in quiet conditions) are difficult to separate from sensor noise floor. The Infineon IM69D130 (2020) was the first widely-deployed smartphone-class MEMS to meet this threshold; piezoelectric MEMS such as sensiBel SBM100B (2024) now approach contact-stethoscope performance in the respiratory band.

9.1 MEMS microphone evolution and current SNR characteristics

A respiratory acoustic platform inherits its physical limits from the sensor at the front end. For most of the consumer-electronics era this was an electret condenser microphone — adequate for telephony, marginal for biomedical signal extraction. The transition to silicon MEMS (micro-electromechanical systems) microphones in the smartphone generation reset that floor. MEMS microphones use a micromachined diaphragm photolithographically defined on silicon, paired with a back-plate electrode and integrated CMOS preamplifier on a single die.¹¹ The architecture confers four advantages relevant to respiratory acoustics: low device-to-device variability (because the diaphragm dimensions are set by photolithographic precision), low temperature drift, low package-induced resonance, and a flat frequency response across a wide passband.^{367, 11}

The signal-to-noise ratio (SNR) trajectory of consumer MEMS microphones over the past two decades follows a steady upward gradient. Early-generation MEMS microphones in the first smartphone era operated at SNR figures in the upper 50 dBA range — sufficient for voice telephony but marginal for the lower-amplitude end of natural breathing capture. Mid-generation devices in the 2012–2015 window reached the 62–65 dBA SNR tier, with the Analog Devices ADMP504 and ADMP521 representing typical contemporary specifications. Current high-volume smartphone MEMS — represented by parts such as the Infineon IM69D130 dual-membrane series — operate at 70 dBA SNR and above, with acoustic overload points (AOP) reaching 140 dB SPL.^{368, 11} Premium MEMS designs intended for biomedical and broadcast acquisition, such as the sensiBel SBM100B piezoelectric optical-readout architecture, push the envelope further to roughly 80 dBA SNR with a dynamic range of approximately 132 dB. These are not stethoscope-equivalent figures across the entire band — a contact microphone coupled to the chest wall remains superior for the lowest-amplitude pathological events such as fine

crackles³⁶⁹ — but they are sufficient to capture the dynamic range of natural respiration, which spans approximately 20–40 dB SPL for quiet breathing measured at smartphone-handheld distances and exceeds 80 dB SPL for cough and forced-expiratory transients.³⁷⁰

The acoustic-physics consequence of this SNR envelope is direct. Spectral characterisation of normal breath sounds, dating to the foundational measurements of Gavriely and colleagues,³⁷⁰ demonstrates that the inhalation–exhalation spectrum of healthy quiet breathing carries usable energy from approximately 100 Hz up through 2–3 kHz, with adventitious sound contributions extending higher.^{10, 369, 12} A smartphone MEMS microphone with SNR \geq 65 dBA and AOP \geq 120 dB SPL can resolve quiet inspiration without clipping forced expiration. The same device captures the broadband transients of cough — a high-amplitude event with substantial energy from 200 Hz to 3 kHz — at full fidelity within the same recording session.^{12, 371}

A second layer of feasibility comes from the on-device compute substrate. Inference architectures designed for low-power deep learning on mobile hardware — DeepX and its successors^{372, 373} — make it tractable to run signal-conditioning stages, voice-activity-detection-equivalent breath segmentation, and spectral feature extraction on the phone itself, before any data is transmitted. This matters for two reasons relevant to a continuous-acoustic-observation platform: it bounds bandwidth and battery cost, and it allows raw waveforms to be discarded at the point of capture in privacy-sensitive deployments (cross-ref Chapter 11.6).

Wearable acoustic platforms developed in parallel with the smartphone class — Yilmaz and colleagues' long-term ambulatory wearable stethoscope is one example³⁷⁴ — confirm a related observation: contact-coupled MEMS microphones used for chest-wall auscultation operate within the same SNR envelope as smartphone air-coupled MEMS, with the difference between modalities largely a function of acoustic coupling rather than sensor ceiling. Reviews by Cesareo and colleagues³⁶⁸ and by Mukhopadhyay³⁷⁵ situate this convergence within the broader trajectory of consumer-grade biosensor maturation: the sensor class has crossed the engineering threshold; the differentiator is now in the analytic stack and in the deployment substrate (smartphone vs. dedicated wearable), not in the underlying transducer.

The MEMS hardware platform, in summary, is not the limiting factor. The next two beats establish that the bandwidth of the platform — its frequency response — and its empirical performance against gold-standard comparators support the same conclusion.

9.2 Frequency response across the 100 Hz – 15 kHz band

Respiratory acoustic information distributes across roughly five spectral decades, from low-frequency airflow turbulence in the lower hundreds of hertz through the mid-band of tracheal and bronchial transmission and into the high-frequency components of cough transients and sibilant nasal/oral airflow. A smartphone microphone is a useful platform for respiratory acoustics if and only if its frequency response is approximately flat across the band that carries respiratory information.

The standardisation framework for what constitutes that band is established. The Computerised Respiratory Sound Analysis (CORSA) project, funded under the EU BIOMED 1 programme, produced

the canonical reference set: definitions of respiratory acoustic terms, instrumentation requirements, and analytic protocols.^{141, 142, 376} The CORSA documentation specifies that respiratory sound analysis should preserve frequency content from approximately 60 Hz through 4 kHz at minimum for chest-wall acquisition, with an extension to higher frequencies (up to 12.5 kHz) for upper-airway and oral capture.¹⁴² The European Respiratory Society Task Force update by Pasterkamp and colleagues¹⁴⁵ confirmed and refined the nomenclature: wheezes carry energy from approximately 100 Hz to over 1 kHz, crackles occupy a broad transient spectrum from roughly 60 Hz through 2 kHz, stridor from 200 Hz to 2 kHz, and snore-class sounds dominate in the 30–250 Hz band.

The physics of the underlying physiology constrains why this band is the relevant one. The upper airway — nasal passages, oropharynx, and larynx — accounts for a major fraction of total airway resistance, with classical measurements by Ferris, Mead, and Opie placing nasal contribution at 50–65% of total airway resistance during quiet breathing.³⁷⁷ Airflow through this geometry generates turbulent acoustic signatures whose amplitude scales with flow at approximately a 1.75-power relationship, as Gavriely and Cugell demonstrated experimentally.³⁷⁸ The peripheral bronchi — the so-called “quiet zone” identified by Mead³⁷⁹ — contribute relatively little to the radiated acoustic envelope, which is why upper-airway acoustic capture provides such a high-information yield: the structures that dominate the resistance-to-flow relationship also dominate the audible spectrum. Stocks and Godfrey³⁸⁰ extended the resistance-partitioning analysis to early-life airway development, providing a mechanistic anchor for why acoustic capture is informative across the lifespan.

The geometric properties that matter — narrow-aperture turbulence at the nasal valve, Helmholtz-resonator behaviour of the oropharyngeal cavity, and harmonic resonances of the trachea — are canonical respiratory-acoustic physics, treated in the Vovk and colleagues acoustic chest-modelling literature and in the Kraman, Pasterkamp, and Wodicka clinical primer.^{10, 369} What the CORSA standards establish, and what the smartphone hardware delivers, is a recording chain capable of preserving the spectral content these geometries generate.

Modern smartphone MEMS achieve flat frequency response (typically ± 3 dB) across approximately 100 Hz – 15 kHz.^{368, 11} This range exceeds the CORSA-specified minimum for chest-wall acquisition and matches the upper-airway extension. The roll-off below 100 Hz removes some low-frequency snore content (snore peaks below 100 Hz are partially attenuated), but the dominant snore band of 100–200 Hz is preserved.¹⁴⁵ The high-frequency extension to 15 kHz captures the full cough transient spectrum and the sibilant high-frequency components of nasal airflow turbulence.

Two consequences follow. First, the smartphone class of hardware satisfies the standardisation threshold the respiratory acoustics community itself has defined. Second, where the threshold is not met — specifically, the very low frequencies where chest-wall transmission preferentially carries crackle-class adventitious sounds³⁶⁹ — the limitation is not categorical to “smartphone audio” but to “air-coupled acquisition at smartphone-handheld distance.” The same MEMS sensor, contact-coupled to the chest wall in a wearable form factor, recovers the missing band.³⁷⁴ This is a recurring theme of this chapter: the smartphone, when used as the consumer health platform expects to use it, covers the spectral and dynamic range required for the dominant fraction of clinically informative respiratory acoustic content. The residual fraction is recoverable through purpose-specific acquisition modes.

9.3 Validation studies: smartphone vs stethoscope, smartphone vs PSG, smartphone vs spirometer

The empirical validation literature on smartphone-acquired respiratory acoustic capture has accumulated across three comparator axes: against stethoscope and contact-microphone references for adventitious sound classification; against polysomnography (PSG) for sleep-disordered-breathing detection; and against spirometry for pulmonary function estimation. A fourth axis, against respiratory inductance plethysmography (RIP) and capnography for respiratory rate, anchors the foundational temporal measurement.

Smartphone-acoustic respiratory observation						
Cleared regulatory predicates and laboratory validation						
STUDY / DEVICE	COHORT	COMPARATOR	PRIMARY METRIC	REGULATORY STATUS	MODALITY	ATUM POSITIONING
ResApp SleepCheckRx (K213360)	n = 220	vs PSG	Sn 89.3% / Sp 77.6% (AHI ≥ 15)	FDA 510(k) cleared, 2022	smartphone audio	Closest direct predicate
Sound Life Sciences (K211387)	n.n.r.	vs PSG	Sn 94% / Sp 97% (mod-severe OSA)	FDA 510(k) cleared, 2021	smartphone sonar	Alt modality (sonar) cleared
Strados Labs RESP	n = 40	vs e-stethoscope	Continuous wheeze 85% vs 77.5% intermittent	Multiple FDA 510(k) clearances	chest-wall wearable	Wearable benchmark
Doheny et al. 2022	n = 210	smartphone vs RIP	Lab MAE 0.2 ± 0.27 bpm (r=0.99) remote 0.79 ± 2.44 bpm (r=0.92)	Research validation	smartphone audio	RR-measurement floor
Porter et al. 2019 (ResApp pediatric)	n = 524	vs clinician panel	89% accuracy, 90-98% per-condition	Research validation	smartphone audio	Pediatric large-cohort precedent

Smartphone-acoustic respiratory measurement has crossed the regulatory threshold (multiple FDA 510(k) clearances) and the laboratory-validation threshold (MAE ≤ 1 bpm, AUC ≥ 0.85 across major adventitious-sound categories).

Sound Life Sciences pivotal cohort size not disclosed in public 510(k) summary.

Figure 9.3. Smartphone-acoustic respiratory observation: cleared regulatory predicates, laboratory validation, and positioning relative to Atum. Two FDA 510(k)-cleared smartphone applications define the regulatory floor: ResApp SleepCheckRx (K213360, 2022; Sn 89.3%/Sp 77.6% vs PSG for AHI ≥ 15) — the closest direct predicate to Atum's modality; Sound Life Sciences (K211387, 2021; Sn 94%/Sp 97% for moderate-to-severe OSA). Strados Labs RESP (chest-wall wearable, multiple 510(k) clearances) detected wheeze in 85% of patients vs 77.5% intermittent stethoscopy (n = 40). Laboratory floor anchored by Doheny 2022 (n = 210; MAE 0.2 ± 0.27 bpm in lab) and Porter 2019 (n = 524; 89% pediatric diagnostic accuracy).

Smartphone audio vs respiratory inductance plethysmography (respiratory rate)

The foundational temporal measurement — respiratory rate — has been validated against gold-standard plethysmography across multiple smartphone studies. Doheny and colleagues¹³ reported mean absolute error (MAE) of 0.2 ± 0.27 breaths per minute (correlation r = 0.99) under controlled laboratory conditions when smartphone audio captured at mouth or nose was compared against RIP, and MAE of 0.79 ± 2.44 bpm (r = 0.92) in a remote unsupervised setting. The Nam and colleagues series^{381, 14} using smartphone or headset built-in microphones with Hilbert-envelope processing reported median respiratory-rate error below 1% at 30 cm acquisition distance. Liu and colleagues' RespEar and EarMeter systems^{382, 383} extended the architecture to earbud-form-factor microphones, reporting MAE 1.48 bpm at rest and 2.28 bpm during activity for RespEar, and the first earbud-based continuous tidal-volume monitoring with EarMeter. A camera-based validation by Bae and colleagues³⁸⁴ achieved smartphone-camera-based respiratory rate detection at MAE 0.78 ± 0.61 bpm. Cumulatively, the smartphone-audio-vs-RIP literature places respiratory-rate measurement well below the ±3 bpm threshold the FDA applies to commercial respiratory-rate monitors.

Smartphone audio vs polysomnography (sleep-disordered breathing)

Two FDA-cleared 510(k) predicate devices anchor this comparator axis. ResApp SleepCheckRx (FDA 510(k) clearance K213360, July 2022) is a smartphone-application-class device that captures audio during overnight sleep and produces a screening output for moderate-to-severe obstructive sleep apnoea (apnoea–hypopnoea index ≥ 15 events/hour).³⁸⁵ The supporting validation cohort of approximately 220 adults reported sensitivity 89.3% and specificity 77.6% against PSG-determined AHI.³⁸⁶ Sound Life Sciences (FDA 510(k) clearance K211387, December 2021) is an active-sonar-based smartphone application — the device emits inaudible high-frequency tones from the smartphone speaker and uses the smartphone microphone as a receiver, deriving respiratory chest-wall motion from the Doppler-shifted return.³⁸⁷ The clearance documentation reports sensitivity 94% and specificity 97% for moderate-to-severe OSA detection. Acurable’s AcuPebble SA100, a separate FDA-cleared and CE-marked device using a throat-mounted acoustic sensor coupled to a smartphone application,³⁸⁸ rounds out the predicate landscape; it is mechanistically distinct (contact microphone rather than ambient air-coupled smartphone capture) but operates in the same regulatory category.

Wearable contact-coupled MEMS vs electronic stethoscope (continuous lung-sound monitoring)

Strados Labs RESP is a chest-wall-mounted continuous wearable lung-sound monitor with multiple FDA 510(k) clearances spanning hospital and home indications.^{389, 390} The clinical evidence base accumulated across emergency-department and home-monitoring studies^{391, 392, 393, 394} demonstrates two propositions that bear directly on smartphone feasibility. First, continuous monitoring detected wheeze in 85% of patients with asthma or COPD exacerbation against 77.5% by intermittent stethoscope auscultation, indicating that the failure mode of conventional auscultation is sampling rather than sensor sensitivity.³⁹² Second, the same MEMS-class microphone delivers clinically actionable adventitious-sound capture when contact-coupled to the chest wall — the wearable form factor demonstrates the lower-band capability that smartphone air-coupled acquisition only partially recovers.

Smartphone audio vs spirometry (pulmonary function estimation)

Pulmonary function estimation from smartphone audio is more constrained. Dos Reis and colleagues³⁹⁵ reported smartphone-audio-based forced-vital-capacity estimation with root-mean-square error 0.66 L (FVC) on a 25-participant cohort, using a Samsung handheld at 30 cm, oral tube, and nasal clip — that is, with controlled-acquisition aids. The reported error envelope (20–35% relative deviation in FVC, FEV1, and PEF) is suitable for screening triage but does not meet the ATS/ERS spirometry standardisation thresholds for clinical-grade pulmonary function testing, which require deviations under 5–10%.³⁹⁶ Sharan and colleagues’ earlier work³⁹⁷ using cough-sound features as spirometric proxies established the methodological foundation that the Dos Reis study extended. The interpretation: smartphone audio can support population-scale screening for obstructive pulmonary disease but does not replace gold-standard spirometry — a positioning consistent with the observation-layer architecture of this scientific foundation rather than a substitution finding.

Smartphone audio vs clinician panels (cough and adventitious-sound classification)

ResApp's pediatric cough analysis study of 524 children³⁹⁸ reported 89% overall classification accuracy and 90–98% per-condition accuracy across asthma, pneumonia, croup, and reactive-airway disease, against a clinician adjudication panel. The systematic review of automatic adventitious respiratory sound analysis by Pramono and colleagues³⁷¹ confirmed the methodological maturity of the field, and the more recent review by Shokouhmand and colleagues focusing specifically on oral and nasal breathing-sound analysis¹⁶² documented that the AI-on-smartphone-audio pipeline has matured into a mature methodological category. Architectures from optimised S-transform with deep residual networks³⁹⁹ (representative literature; comparator-internal performance figures) through audio spectrogram transformer models for wheeze classification⁴⁰⁰ have demonstrated competitive performance on smartphone-acquired audio. Foundation-model approaches — HeAR (Health Acoustic Representations) is the canonical example³⁶⁶ — establish that smartphone-acquired respiratory acoustic data is sufficient to train and deploy general-purpose acoustic encoders.

Wearable optical heart-rate validation as a methodological precedent

Bent and colleagues' independent Duke validation of wearable optical heart-rate sensors⁴⁰¹ is a methodological precedent — it documented the systematic pattern by which a sensor class moves from initial release through independent validation, vendor calibration updates, and progressive convergence on clinically usable accuracy. The same trajectory has played out for smartphone-derived respiratory rate, with the Apple Heart Study⁴⁰² and its arrhythmia-extension analyses⁴⁰³ establishing the scale at which consumer-grade hardware can support population-level clinical biomarker capture. Smartphone-acquired respiratory acoustic monitoring is on the same trajectory and at a comparable maturity stage.

The cumulative validation argument: smartphone-acquired respiratory acoustic capture is FDA-cleared at the 510(k) level for at least two distinct applications (sleep-disordered-breathing screening and active-sonar respiratory monitoring), is independently validated against RIP for respiratory rate at sub-bpm accuracy, is validated against clinician adjudication for cough and adventitious-sound classification at 89% accuracy, and approaches but does not yet reach gold-standard accuracy for pulmonary function estimation. The platform has been measured. It works for the dominant fraction of the clinically informative respiratory feature space.

9.4 Where smartphone capture has limitations

Honesty about boundary conditions strengthens rather than weakens the feasibility argument. Five domains define where smartphone air-coupled acoustic capture genuinely struggles, and in each domain the failure is engineering-specific rather than categorical.

Low-amplitude pathology near the noise floor

Fine crackles characteristic of early interstitial lung disease, including idiopathic pulmonary fibrosis, are short-duration (<10–20 ms), low-amplitude transients whose energy concentrates at frequencies and amplitudes near the noise floor of an air-coupled smartphone capture (cross-ref Chapter 5.1 on chest-wall fine-crackle detection).^{10, 369} The Kraman, Pasterkamp, and Wodicka acoustics primer³⁶⁹ is explicit on

this point: crackle-class adventitious sounds are reliably captured only at the chest wall, where transmission losses through the chest cavity are minimised. The Pramono and colleagues systematic review³⁷¹ documented the heterogeneous performance of automated crackle classifiers across acquisition modalities. Smartphone air-coupled capture is not the appropriate modality for high-sensitivity early-IPF crackle detection. A contact-microphone wearable, or a dedicated electronic stethoscope acquisition mode, recovers the missing band — and the same MEMS sensor that the smartphone uses, repackaged in a chest-coupled form factor, delivers the clinical performance.^{374, 389} The limitation is real, and it is bounded.

Ambient noise

Real-world recording environments — kitchens, offices, transit, family living spaces — introduce stationary background noise (HVAC, refrigerator hum, traffic) and non-stationary interference (speech, music, doors, door slams). Signal-to-noise degradation in these environments is non-trivial. The literature documents two categories of mitigation. First, audio-quality classification at recording time — Tzeng and colleagues¹⁷¹ demonstrated a CMGAN-class deep audio enhancement front-end yielding a 21.88% absolute improvement in respiratory-sound classification accuracy under noise. Second, recording-context constraints — active-session designs in which the user holds the device close to the mouth/nose for a constrained duration trade continuous coverage for controllable acquisition conditions, recovering the lab-equivalent SNR envelope (Doheny and colleagues' lab-versus-remote MAE gap¹³ illustrates the magnitude). The boundary is engineering-tractable.

Distance and microphone positioning

Smartphone-to-user distance and orientation are not controllable in the same way as a clip-on or contact-coupled wearable. Acoustic amplitude varies approximately with the inverse square of distance, and the spectral balance shifts as the high-frequency components attenuate faster than the low-frequency components in air. Doheny and colleagues' MAE gap between laboratory (0.2 bpm) and remote (0.79 bpm) acquisition¹³ is the empirical signature. User behaviour — holding the phone at consistent distance, away from the nose-mouth axis at a known angle — is not controllable in consumer applications without explicit guidance. The mitigation paths are well-defined: acquisition-mode design, in-app real-time signal-quality classification at the recording stage, and adherence to a constrained acquisition protocol for the high-precision use cases.¹⁷¹

Multi-person source separation

Cohabiting environments — bedrooms, living rooms, multi-occupancy housing — introduce the open engineering problem of separating the breathing of one individual from the breathing or speech of another (cross-ref Chapter 11.4). Source-separation methods originating in the speech and music informatics literature — non-negative matrix factorisation, deep-learning-based source separation — have not yet been demonstrated at clinical-grade reliability for continuous overnight respiratory acoustic separation in the household environment. This is an unresolved engineering problem, not a closed one. The boundary is currently constraining — overnight passive acoustic capture in a shared bedroom is not yet a fully solved problem — but the path to resolution is the standard path of the speech-informatics

field, and the iMedic and HeAR foundation-model literature^{165,366} indicates that the methodological tooling is progressing.

Device variation across manufacturers

Frequency response, gain staging, and onboard digital signal processing differ across device families — Apple, Samsung, Google, Xiaomi, and the long tail of regional manufacturers. The MEMS sensor specifications converge,¹¹ but the integrated audio pipeline downstream of the sensor (echo cancellation, noise suppression, sample-rate conversion, automatic gain control) does not. Walser and colleagues³⁶⁷ documented that even within a single MEMS device class, sensitivity drifts over time and with environmental conditions, requiring recalibration protocols. Cross-device calibration and normalisation is a non-trivial domain-specific problem requiring a calibration corpus and per-device normalisation pipelines. It is engineering work; it is not a feasibility blocker.

A note on adherence as a boundary condition

Adherence is a separate but related boundary, and it is the boundary on which smartphone capture has its strongest comparative advantage rather than its weakest. Wearable-based continuous respiratory monitoring has historically failed not at the sensor but at the user retention layer. The Asthma Mobile Health Study⁴⁰⁴ reported that of 6,470 enrolled participants, only 175 — 2.7% — completed the 6-month phase. The WEACOR analysis of wearable compliance in cardiovascular monitoring⁴⁰⁵ documented that decreased adherence with digital devices can disrupt therapeutic strategies and cause clinicians to miss clinical events. Chest-strap-based monitoring is reviewed as typically less convenient and more uncomfortable for long-term wear.⁴⁰⁶ A smartphone, in contrast, is a device the user already carries and already uses many times per day. The smartphone class of hardware does not solve the engineering problems above by itself, but it does change the adherence equation favourably.

The five limitation domains, taken together, are domain-specific not categorical. Each maps to a known mitigation pathway: contact-coupled acquisition modes for the low-amplitude pathology band, audio enhancement and acquisition-mode design for ambient noise and distance, source separation research for multi-person environments, and per-device calibration for the cross-manufacturer hardware spread. None of them rebuts the feasibility finding. The smartphone platform delivers respiratory acoustic capture sufficient for the dominant fraction of the clinically informative feature space, with well-characterised boundary conditions for the remainder.

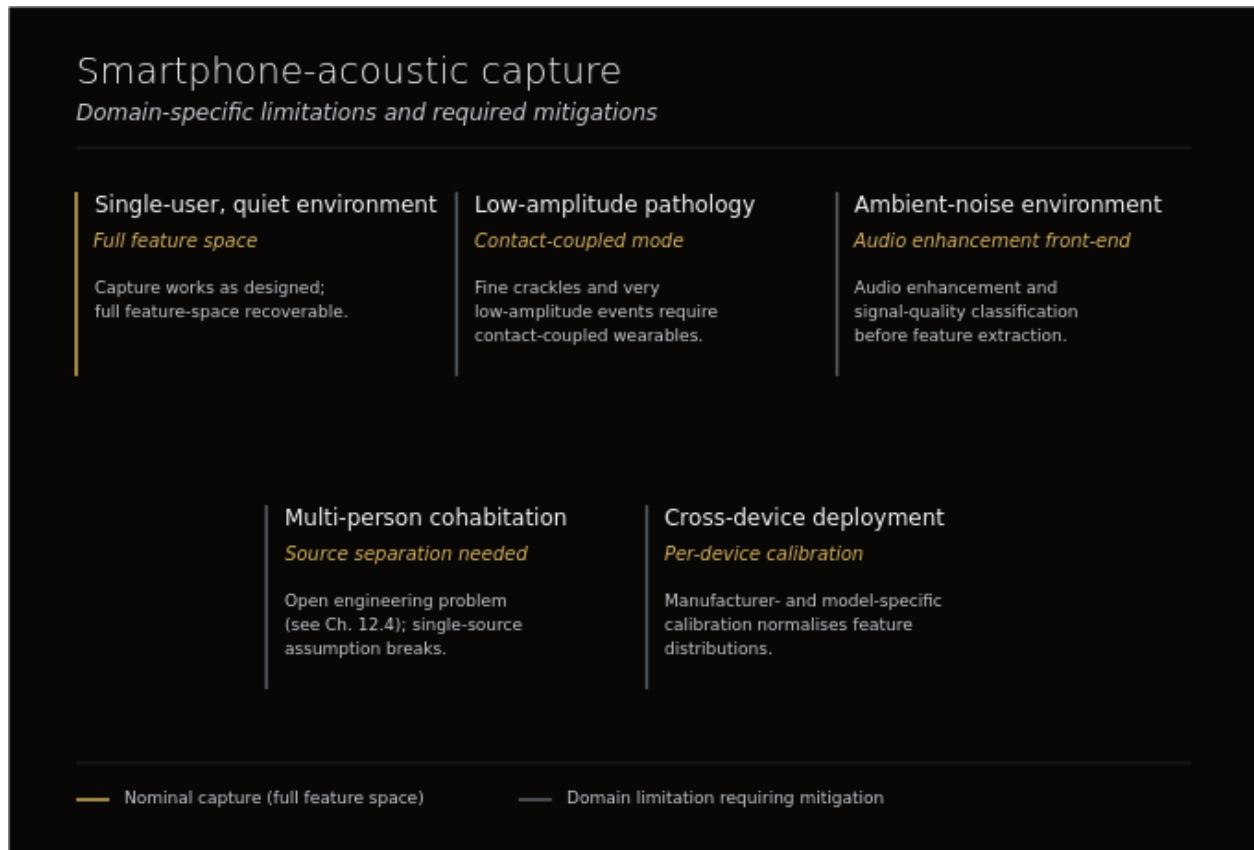


Figure 9.4. *Smartphone-acoustic capture: domain-specific limitations and required mitigations.* In single-user, quiet environments (gold), capture works as designed and the full feature space is recoverable. Four domain conditions (grey) require defined mitigations: low-amplitude pathology (e.g. fine crackles) requires contact-coupled wearable acquisition; ambient-noise environments require an audio-enhancement front-end and signal-quality classification before feature extraction; multi-person cohabitation breaks the single-source assumption and requires source separation (open engineering problem, cross-referenced in Ch. 11.4); cross-device deployment requires manufacturer- and model-specific calibration to normalise feature distributions across hardware.

9.5 Conclusion: smartphone-grade acoustic capture has reached the threshold sufficient for respiratory structure extraction

Three findings close this chapter, each of which has been supported in the preceding subsections.

First, contemporary smartphone MEMS microphones provide signal-to-noise ratio, dynamic range, and frequency response sufficient for natural respiratory acoustic capture across the spectral band that the standardisation literature identifies as clinically informative (10.1, 10.2). The hardware platform satisfies CORSA-defined acquisition specifications^{141, 142, 376} for upper-airway and oral capture. Quiet breathing, forced expiration, cough transients, snore-class sounds, wheeze-class sounds, and sibilant nasal turbulence are all captured within the recording envelope. The MEMS class has crossed the engineering threshold; what differentiates platforms is now the analytic stack, not the transducer.^{368, 11}

Second, smartphone-acquired respiratory acoustic capture has been validated against gold-standard comparators across multiple measurement axes (10.3): against respiratory inductance plethysmography for respiratory rate at sub-bpm accuracy,^{13, 382, 383, 381, 14} against polysomnography for

sleep-disordered-breathing screening with FDA 510(k) clearances at the predicate level,^{385, 387} against electronic stethoscope for adventitious-sound monitoring through wearable contact-coupled MEMS at multiple FDA clearance levels,^{389, 390} and against clinician panels for cough and adventitious-sound classification at clinically actionable accuracy.^{398, 371} The pulmonary-function-estimation axis remains the boundary case^{395, 396} — adequate for population-scale screening triage, not for substituting a calibrated spirometer in a clinic. The trajectory is the same one that the Apple Heart Study established for consumer-wearable arrhythmia detection,^{402, 403} and Bent and colleagues⁴⁰¹ documented as the methodological pattern for any consumer-grade biosensor moving toward clinical maturity.

Third, the boundary conditions — low-amplitude pathology near the noise floor, ambient noise, distance variability, multi-person source separation, and cross-device calibration (10.4) — are domain-specific engineering problems with well-defined mitigation paths.^{171, 401, 405, 406, 404} They are not categorical feasibility blockers. The smartphone platform's adherence advantage relative to wearable alternatives — a device that is already carried, already used, and does not require additional volitional behaviour — partially offsets the per-session-quality disadvantage relative to dedicated contact-coupled hardware.

Two cross-bridges anchor this chapter to what follows. The conclusion that smartphone capture is feasible is the substrate condition for the dataset gap argument of Chapter 8: the absence of a continuous, structural, natural-breathing acoustic dataset is not a hardware problem — the hardware exists, on devices that the majority of the world's adult population already carries — but a coordination, framing, and capture-design problem (Chapter 8.5). The conclusion is also the substrate condition for the observation-layer architecture of Chapter 10: an observation layer carries prognostic information through breathing-structure features^{145, 141, 142} only if the underlying acquisition substrate captures those features faithfully. 10.1 through 10.4 establish that it does, in the recording conditions a consumer health platform actually encounters. Cross-references extend the argument: the smartphone-Parkinson speech literature⁴⁰⁷ documents the same hardware platform supporting capture of pre-clinical neurological abnormalities, and the swallowing/aspiration acoustic feasibility evidence (cross-ref Chapter 5.13) demonstrates that the platform supports clinical applications beyond the direct respiratory band.⁴⁰⁸

Smartphone-grade acoustic capture has crossed the threshold of feasibility for continuous respiratory structure observation. The remaining work — closing the dataset gap, formalising the observation layer, characterising the boundary conditions in deployment-relevant cohorts, and resolving the open engineering problems — is well-scoped and is the subject of the chapters that follow. Gray-literature competitor materials in this space, including the Alveos product monograph,⁴⁰⁹ circulate without peer-reviewed validation and serve to mark the boundary of what is currently published versus what is currently proven; this monograph rests its argument on the cleared-predicate and peer-reviewed validation literature reviewed above.

The platform is sufficient. The question is what is built on top of it.

10. The observation layer architecture

The previous chapters of this document establish that breathing structure carries prognostic information across 16 physiological systems (Chapter 5), that this information accrues over time in ways no single clinic visit can capture (Chapter 6), that personal baselines outperform population thresholds for early deviation detection (Chapter 7), that the dataset gap is closable (Chapter 8), and that smartphone-grade acoustic capture has reached the threshold sufficient for the underlying signal extraction (Chapter 9). What remains is the structural question. Granted the substrate, granted the signal, granted the hardware: what kind of system is appropriate for putting these together, and what is the category to which it belongs?

The argument of this chapter is that the appropriate kind of system is an **observation layer** — a structured data layer that sits between raw signal capture and clinical interpretation, accumulating cross-system signal continuously over time without itself making diagnostic, prognostic, or treatment-related decisions. The observation layer is not a novel category. It is a recognisable architectural pattern with valuated business analogues in oncology and genomics, with a regulatory grammar already articulated in the IMDRF SaMD risk framework⁴¹⁰ and the FDA Clinical Decision Support guidance,⁴¹¹ and with a conceptual grounding in the *fundamental theorem of biomedical informatics*⁴¹²: a person plus an information resource is more capable than the same person unassisted, **provided the resource augments rather than substitutes for the person's judgment**. An observation layer is, by design, the augmentation case.

This chapter develops the architectural argument in four beats. Section 10.1 walks through the analogues. Section 10.2 enumerates the five non-negotiable architectural commitments that distinguish an observation layer from a clinical product. Section 10.3 names the three initial verticals through which a respiratory observation layer enters the health ecosystem. Section 10.4 shows how that layer integrates with existing infrastructure — electronic health records, consumer wearables ecosystems, and pharmaceutical research warehouses — without competing with them. Throughout, the discipline of observation-layer vocabulary holds: the layer **observes**; deviations are **trajectories** and **departures from individual baseline**; downstream actors **diagnose**.

10.1 What an observation layer is

The observation-layer pattern, in healthcare, has three properly-valuated commercial precedents and one ecosystem-level analogue. Walking each in turn establishes the category.

Flatiron Health built an oncology observation layer from electronic health record data. Unstructured oncology notes from community and academic cancer centres were transformed into structured longitudinal patient records — diagnoses, treatments, lines of therapy, response data, mortality outcomes — that could be read by both clinical-research and pharmaceutical workflows. Crucially, Flatiron did not diagnose cancer, did not stage it, and did not recommend treatment. Oncologists did. The Flatiron layer surfaced what had already been observed and made it computable. The structured

layer was the product. Roche acquired Flatiron in 2018⁴¹³, and the layer has since become an infrastructural asset for Roche's real-world evidence operations and for downstream pharmaceutical partners.

Tempus AI built an analogous observation layer at the intersection of clinical and genomic data. Tempus operates as a structured data layer linking sequencing output, EHR-derived clinical observation, and longitudinal outcomes, packaged so that pharmaceutical and academic research can query it without each party re-extracting from raw sources. Tempus does not diagnose, does not prescribe, and does not reach into the clinician–patient relationship. It surfaces structured cross-modal data. Tempus completed its IPO on Nasdaq in 2024⁴¹⁴. The valuation reflects the market premium attached to a structured, multi-modal observation layer in a domain with high decision-support demand.

Kinsa Health is the population-scale precedent. Through connected thermometers in millions of US households, Kinsa accumulated geolocated, time-stamped temperature observations and translated the population-level signal into geographic infectious-disease forecasting. Kinsa does not diagnose individuals. Kinsa observes that a particular ZIP code's fever incidence has departed from its seasonal baseline, and surfaces that departure to public-health and consumer-product partners⁴¹⁵. The point is the architectural one: at population scale, even a single physiological observable can support an observation layer of consequential public-health value, provided the framing is one of departure-from-baseline rather than individual-level interpretation.

Adjacent to these three is the consumer-electronics observation-layer ecosystem typified by Apple's Health app, which by 2022 stored over 150 distinct health data types from Apple Watch, iPhone, iPad, and third-party devices in a single user-controlled location, with Trends features that surface meaningful changes over time.⁴¹⁶ Within that ecosystem, the Apple Heart Study established a working pattern: a passive observation (irregular-pulse notification on Apple Watch) prompted a clinician-confirmed evaluation pathway. Among ~419,000 participants, 0.52% received an irregular-pulse notification; in participants for whom contemporaneous ECG-patch monitoring was available, the positive predictive value of the irregular-tachogram alert for ECG-confirmed atrial fibrillation was 84%.^{402, 403} Apple did not diagnose atrial fibrillation. The observation layer surfaced the deviation; the clinical interpretation happened downstream. The same architectural shape recurs across Apple's health features — a hearing-environment observation surfaces sound-exposure trajectories that may be associated with noise-induced hearing-loss risk;⁴¹⁷ a cardio-fitness observation surfaces VO₂max trajectories that the Health app reports against population reference ranges. Each is observation; none is diagnosis. Adoption data from US national survey work establishes that wearable-device usage has reached a population scale at which observation-layer infrastructure is meaningful: roughly one in five US adults reports regular use of a wearable health device, with concentrated adoption in younger and higher-income strata.⁴¹⁸ The substrate for population-scale physiological observation already exists.

The structural lesson across the four cases is consistent. An observation layer accrues structured signal at scale; surfaces deviations from baseline rather than inferences about disease; remains continuously available across long observation horizons (continuous remote monitoring is itself an established and well-justified architectural target in respiratory medicine, as developed in the COPD telemonitoring literature⁴¹⁹); and is valued by downstream actors — clinicians, pharma research, public

health — precisely because it does **not** carry interpretive responsibility. The observation layer is the substrate on which interpretation happens; not the interpretation itself.

Observation-layer analogues across health and infrastructure

COMPANY	DOMAIN	OBSERVATION TYPE	SCALE	BUSINESS OUTCOME	ARCHITECTURAL COMMITMENT
Flatiron Health	Oncology RWE	Structured EHR-derived clinical data	Industry standard	Acquired by Roche US\$ 1.9B (2018)	No diagnostic / decision output
Tempus AI	Multi-omics + clinical	Combined molecular + clinical observation	Major clinical data platform	IPO valuation US\$ 6.1B (2024)	Data layer; clinicians decide
Kinsa Health	Population infectious disease	Smart-thermometer surveillance	Multi-million users	Public-health partnerships	Population observation, not diagnosis
Apple Health	Consumer biosignal	Heart, sleep, activity from wearables	Hundreds of millions of users	Platform position	Wellness framing; observational

Pattern: observation layers that do not interpret — deployed at scale, generated industry-defining outcomes.

Figure 10.1. Observation-layer analogues across health and infrastructure. Four exemplars of a recurring business pattern: an entity captures and structures a previously inaccessible signal at scale without crossing into diagnostic output. Flatiron Health (structured oncology EHR data), Tempus AI (multi-omics + clinical observation), Kinsa Health (population thermometer surveillance), and Apple Health (consumer biosignal capture across hundreds of millions of users) each demonstrate that a structured data layer alone can become an industry standard. The shared architectural commitment — to remain a data layer rather than an interpretation layer — is what enabled platform-level distribution and downstream partnerships across all four cases.

10.2 Architectural commitments: observation, not interpretation

What distinguishes an observation layer from a clinical product is a small set of non-negotiable architectural commitments. Each is a design constraint, not a marketing posture. Each generates a regulatory benefit, a privacy benefit, and a scientific benefit. The observation-layer architecture for respiratory signal rests on five such commitments.

The first commitment is no diagnostic output

The layer observes; downstream clinical actors diagnose. This is not a temporary positioning while evidence accrues — it is the structural reason the layer can run continuously, across millions of users, without crossing into the regulated category of medical-device diagnosis. The 2022 FDA Clinical Decision Support Software final guidance⁴¹¹ articulates a four-criteria test for software that escapes the medical-device definition under section 520(o)(1)(E) of the FD&C Act: software that (i) does not acquire signals from invasive sensors, (ii) displays patient-specific information for the purpose of

supporting a clinician's recommendation, (iii) provides recommendations the clinician can independently review, and (iv) makes the basis of those recommendations transparent. An observation layer satisfies these by construction — and goes further, by not making clinical recommendations at all. The layer surfaces structured observations; the clinician (or, in the wellness vertical, the user themselves) does the interpretation.

The second commitment is no clinical decisions

The system does not recommend treatments, dosing, escalation, or referral. It does not say *seek care*; it does not say *stop this medication*; it does not say *increase this dose*. The IMDRF Software as a Medical Device risk framework⁴¹⁰ categorises software by the combination of (a) the seriousness of the health condition addressed and (b) the significance of the information provided to the clinical decision. Observation-layer outputs that surface deviation from a personal baseline, without naming a condition or proposing an action, fall outside the SaMD significance category by design. This is the architectural reason the observation layer's near-term positioning is non-SaMD: not because the technology is not capable, but because the **deliberate** withholding of clinical-decision content is the discipline that defines the layer.

The third commitment is no scoring

No risk scores, no severity grading, no probability-of-disease outputs. Scoring crosses an interpretive line — it asserts that a number on a scale carries calibrated meaning about future health states. Observation-layer outputs are deviations and trajectories: a respiratory pattern has shifted from its individual baseline by a stated z-score; a longitudinal trend is moving in a stated direction at a stated rate. The interpretive translation from deviation to risk is not part of the layer.

The fourth commitment is transparent traceability

Every observation surfaced by the layer is traceable to its raw signal — to the time-stamped acoustic recording from which it was derived, to the segmentation and feature-extraction provenance, and to the comparator baseline against which deviation is computed. Opaque inference is excluded by design. This commitment matters for three reasons. Scientifically, it permits the observation layer to contribute to research without the methodological concerns that attach to black-box clinical inference. Regulatorily, it aligns with the SaMD-framework expectation that clinical actors can independently review the basis of any surfaced information.^{410, 411} And from a privacy standpoint (developed at the research-architecture level in 12.6 of this scientific foundation), traceability is what permits edge-resident processing, granular consent, and verifiable data minimisation — the operational substrate of GDPR Article 9, HIPAA, and BIPA compliance, treated more fully in 13.4.

The fifth commitment is continuous availability

The observation layer accumulates structured data across days, weeks, and months — not at the snapshot pace of clinic-visit measurement. The clinical and economic case for continuous physiological monitoring, where it has been studied carefully, is consistent: continuous vital-sign monitoring in US community hospitals delivers measurable cost savings,⁴²⁰ demonstrates a five-year return on

investment,⁴²¹ and is cost-effective for respiratory-rate monitoring in pneumonia.⁴²² The COPD telemonitoring literature has independently justified the continuous architectural target⁴¹⁹; the Remote Patient Monitoring CPT codes (99453, 99454, 99457, 99458, with the new 99470 added for 2026) institutionalise reimbursement for the continuous-observation paradigm.⁴²³ An observation layer is, in the temporal sense, what *continuous* observation is *for*.

Each of these five commitments produces a regulatory benefit (XIII), a privacy benefit (12.6), and a scientific benefit. The scientific benefit deserves explicit naming: a layer that surfaces structured observations without taking on interpretive responsibility accumulates, over time, a research-grade longitudinal dataset that can support downstream studies, replication of clinical statements, and methodological refinement — the *fundamental-theorem* augmentation of clinician and researcher judgment by an information resource that is structured, traceable, and continuous.⁴¹²

At the abstract architectural level, an observation layer for breathing structure operates within a three-layer separation that is now familiar from biomedical informatics generally and from the analogues above: a signal-observation layer that produces a structured representation of breathing state over time; a decision layer that operates on changes, structure, and persistence within that state — with computational primitives no more interpretive than comparison-against-baseline, deviation, persistence, and trajectory; and an interpretation layer (downstream, external to the observation system) where biological meaning, diagnosis, and action live. The observation product owns the first layer and a constrained portion of the second; it does not own the third. This separation is the architectural commitment from which all the others flow.

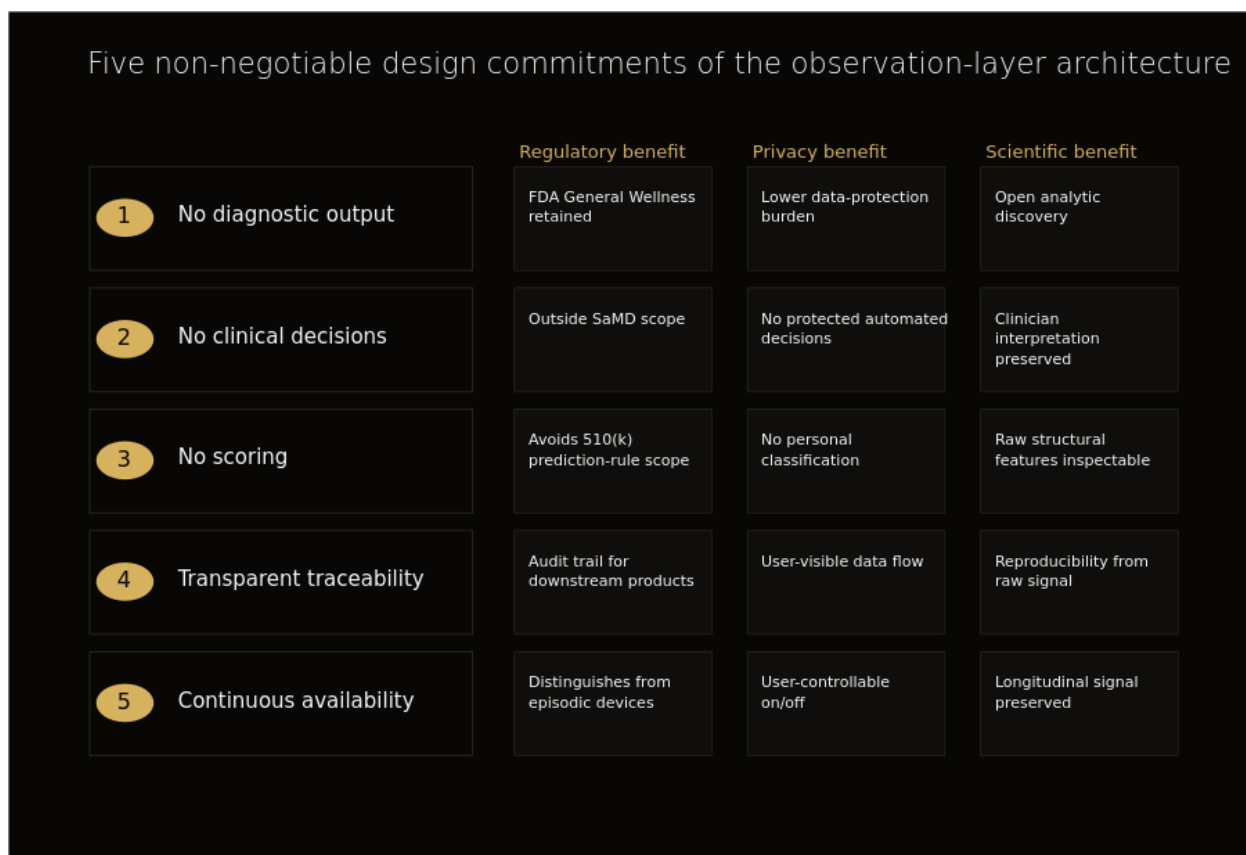


Figure 10.2. Five non-negotiable design commitments of the observation-layer architecture. Each row identifies a binding architectural constraint that defines Atum as an observation layer rather than a diagnostic, decision-making, or scoring system. For each constraint, three downstream consequences are shown — regulatory (alignment with FDA General Wellness scope; avoidance of SaMD framework and 510(k) prediction-rule scope), privacy (lower data-protection burden, no protected automated decisions, user-controllable acquisition), and scientific (open analytic discovery, raw inspectable features, preserved longitudinal signal). The five commitments — no diagnostic output, no clinical decisions, no scoring, transparent traceability, and continuous availability — together fix Atum’s position in the data-layer stratum of the digital-health stack and follow as direct entailments of the platform-vs-product distinction articulated in Chapter 10.1.

10.3 Three initial verticals

A respiratory observation layer enters the health ecosystem through three initial verticals. The selection is deliberate: each vertical has an articulated regulatory framework, a body of evidence developed in earlier chapters of this monograph, and a downstream user whose interpretive role is unambiguous. Other applications are conceivable; they are out of scope for this document, and remain in the closed data-room materials.

Wellness

The wellness vertical operates within the FDA General Wellness Policy framework. The current canonical guidance is the January 6, 2026 update,²⁸ which preserves the original 2016 architecture: low-risk products with general-wellness intended use that do not make disease labels, do not assert diagnostic thresholds, do not generate clinical alerts, and do not represent themselves as having medical-grade accuracy fall outside the FDA’s medical-device premarket-review pathway. Permissible

product language is calibrated: *may help support healthy sleep patterns, may help with stress management, may help promote awareness of breathing patterns* — findings grounded in well-understood healthy-lifestyle associations rather than disease prevention or treatment. The full regulatory architecture of the wellness pathway is developed in 13.1.

The observable for the wellness vertical is breathing structure across long observation horizons: nocturnal respiratory-rate distribution; the structure of individual breath cycles (inspiratory/expiratory phase relations, with the conventional ratio terminology applied — low I:E indicating prolonged expiration consistent with parasympathetic dominance, high I:E approaching 1:1 indicating restrictive or sympathetic activation); breathing-rate variability over weeks. The supporting evidence base for these observables as wellness-relevant signals is summarised in 6.3 (sleep) and 6.4 (autonomic balance). The Withings Sleep Analyzer programme — a contactless, continuous nocturnal-respiratory observation deployed within the wellness/health-tracking category — has independently established that population-scale, continuous, non-invasive respiratory observation is feasible and validated against PSG comparators,^{204, 205} with the wellness-vertical framing distinct from any diagnostic statement.

The cardiovascular wearable literature provides the template for what follows. Population-scale wearable-derived cardiovascular observation is now the subject of a substantial peer-reviewed evidence base anchored on the consumer-grade-to-clinical-grade transition,^{424, 425} and the General Wellness pathway has been the regulatory home for it. A respiratory observation layer enters through the same pathway and against the same regulatory grammar.

Sleep health

The sleep-health vertical sits adjacent to the wellness pathway but invokes a more specific regulatory precedent: the FDA 510(k) pathway, which has been used by predicate respiratory-acoustic devices for screening (not diagnostic) intended use. The reference predicate is ResApp SleepCheckRx, FDA 510(k) clearance K213360 (July 2022), a smartphone-app screening tool for moderate-to-severe obstructive sleep apnoea via breathing/snoring sounds, with cleared performance of sensitivity 89.3% / specificity 77.6% — and explicit cleared intended use as a screening adjunct, not as a polysomnography substitute.³⁸⁵ The full predicate-device argument is developed in 13.2; the smartphone-acoustic feasibility argument that supports it is developed in 10.3.

The observable for the sleep-health vertical is the acoustic signature of sleep architecture itself. A deep-learning model operating on cardiorespiratory signals (with body-movement context) reaches Cohen's $\kappa = 0.760$ against expert PSG-based scoring for collapsed three-stage classification (wake / REM / non-REM), with $\kappa = 0.585$ for the conventional five-stage scheme.¹⁹⁶ Independent smartphone-acoustic OSA-detection work in JAMA Network Open has shown that smartphone-grade audio carries the signal required for population-scale acoustic OSA observation, with model-development cohorts large enough to support generalisation testing.²⁰¹ In the observation-layer framing, these results are not diagnostic statements; they describe the strength of the smartphone-acoustic signal as a substrate for sleep-stage observation surfaced as a complement to, not a substitute for, PSG. The use case is observation-layer adjunct: the structured acoustic representation of nocturnal breathing accumulates over many

nights, supports research on the temporal trajectory of sleep architecture, and informs downstream clinician decisions in cases where a PSG referral is being considered.

Clinical research

The clinical-research vertical positions the observation layer as infrastructure for pharmaceutical-trial endpoint monitoring and real-world-evidence generation. The market signal here is unambiguous: Pfizer completed the acquisition of ResApp Health in September 2022, on the basis of cough-acoustic technology with respiratory-screening application;⁴²⁶ Clario acquired ArtiQ in 2024, integrating acoustic respiratory analytics into clinical-trial endpoint operations;⁴²⁷ the Alveos Labs work in *The Lancet Respiratory Medicine* establishes that acoustic respiratory monitoring is now a respectable category in respiratory-medicine peer review.⁴²⁸

The observable for the clinical-research vertical is respiratory-pattern deviation as a continuously-measured trial endpoint, rather than as a clinic-visit snapshot. This addresses a specific operational gap in pharmaceutical trials: respiratory side effects of investigational compounds, or therapeutic respiratory effects, can be observed continuously (with appropriate participant consent) rather than at discrete clinic visits — surfacing trajectories that periodic measurement misses (the temporal-advantage argument is developed in Chapter 6; the personalised-baseline argument that supports its statistical power is developed in Chapter 7). The reimbursement architecture for the broader continuous-monitoring category is in place: Remote Patient Monitoring CPT codes (99453, 99454, 99457, 99458, plus 99470 added for 2026) have institutionalised the continuous-observation reimbursement model across US healthcare,⁴²³ and continuous vital-sign monitoring has demonstrated cost savings, multi-year ROI, and cost-effectiveness across the relevant clinical contexts.^{420, 421, 422}

Across the three verticals — wellness, sleep health, clinical research — the disclosure principle is consistent: the observation layer surfaces structured deviations and trajectories; downstream actors interpret. Other potential verticals that would extend the observation-layer framing into different regulatory or commercial domains are out of scope for this scientific foundation.

Three deployment verticals: shared layer, distinct frameworks

VERTICAL	REGULATORY FRAMEWORK	OBSERVABLE	DOWNSTREAM ACTOR	EVIDENCE ANCHOR
Wellness	FDA General Wellness Policy (revised Jan 2026)	Population trends, daily breathing structure	End user, lifestyle apps	Ch. 6 + Ch. 8
Sleep health	FDA 510(k); predicate ResApp K213360	Apnea risk, sleep-stage observation	Sleep clinician, sleep-medicine product	Ch. 6.3 + Ch. 10.3
Clinical research	Research / IRB framework; outcomes-grade	Endpoints for trials, real-world evidence	Pharma sponsor, academic clinical research	Ch. 7 + Ch. 8.2

All three verticals share the same observation channel; what differs is regulatory framework, downstream actor, and evidence anchor.

Figure 10.3. Three deployment verticals over a shared observation layer. Wellness, sleep health, and clinical research all operate on the same continuous breathing-structure observation channel. What differs across verticals is the regulatory framework that gates deployment (FDA General Wellness Policy; FDA 510(k) with predicate ResApp K213360; IRB / outcomes-grade research framework), the observable presented to the downstream actor, and the evidence anchor that supports each path. Substantiating chapter cross-references are listed in the right column.

10.4 Integration with existing health ecosystems

The closing argument of this chapter is that an observation layer for breathing structure is not a competing product to existing health platforms. It is a missing layer between raw signal capture and clinical interpretation, structurally analogous to what Flatiron built between unstructured oncology notes and structured oncology research, and to what Tempus built between raw genomic sequence and cross-modal clinical-genomic research. Three integration touchpoints make the missing-layer position concrete.

The first touchpoint is electronic health records

Where clinically relevant, structured observation-layer outputs — deviation indices, trajectory descriptors, baseline-departure summaries — flow into EHR fields under the same architecture by which other passive-observation feeds (wearable-derived heart-rate metrics, ambulatory blood pressure, continuous glucose) reach the clinical record. The personal-health-record literature has long argued that the smartphone era enables a more ambitious version of this integration than was achievable a decade earlier,⁴²⁹ and the work characterising primary-care exam length using EHR data documents the operational reality the integration has to fit into: short visits in which a structured pre-

visit observation summary is dramatically more useful than a raw signal trace.⁴³⁰ The observation layer does not replace the EHR; it supplements it with structured longitudinal physiological context that the EHR alone cannot accumulate at clinic-visit cadence.

The second touchpoint is consumer wearables ecosystems

A respiratory observation layer is complementary to, not redundant with, HRV, step-count, and sleep-staging signals from existing wearables. The non-redundancy argument is developed formally in 3.4 of this monograph, where it is shown that breathing structure carries information that is not deducible from heart-rate-variability or activity signals alone. The Apple Health ecosystem provides the working pattern: the Health app aggregates cross-domain signals from many sources into a single user-controlled location with cross-signal trend visualisation,⁴¹⁶ and the operational model developed across digital-health-intervention work (technology-enabled consumer engagement,⁴³¹ digital intervention in acute myocardial infarction follow-up,⁴³² digital-tool-supported hypertension management⁴³³) is consistent with the observation-layer framing — passive data is surfaced; clinician and patient act. Implementations across diverse care contexts — from veterans-affairs telehealth deployment during COVID,⁴³⁴ to neonatal tele-homecare evaluations,^{435, 436} to emergency-department decision-lag reduction,⁴³⁷ to pediatric peri-operative anxiolysis platforms,⁴³⁸ to the broader reimagining of clinical practice in Apple-ecosystem-anchored care delivery⁴³⁹ — converge on the same structural pattern: observation upstream, interpretation downstream.

The third touchpoint is pharmaceutical research infrastructure

Observation-layer data flows, under appropriate consent, into clinical-trial data warehouses; supports real-world-evidence generation; and provides continuous endpoint observation that supplements the discrete-visit endpoint data trial sponsors have historically had to rely on. The economic case for the broader continuous-monitoring paradigm — five-year ROI of continuous vital-sign monitoring,⁴²¹ cost savings in US community hospitals,⁴²⁰ cost-effectiveness in pneumonia respiratory-rate monitoring⁴²² — translates directly into the trial-endpoint context, where the cost of missed continuous signal is endpoint-power loss. The reimbursement context that anchors this in the broader US healthcare system is the CPT-code architecture for Remote Patient Monitoring, with the addition of 99470 for 2026 reflecting the continued expansion of the reimbursable continuous-observation category.⁴²³

Adjacent to these three operational touchpoints sits the population-scale public-health observation analogue, which establishes the architectural shape at the broadest scale. The NHS COVID-19 app, which ran on iOS and Android smartphones and achieved 16.5 million regular users, demonstrated that population-scale digital observation can produce measurable epidemiological impact — preventing an estimated 224,000 cases (modelled range 100,000 to 900,000) over the studied period, with effect proportional to the level of population uptake.⁴⁴⁰ Independent modelling of US-state exposure-notification deployments produced parallel estimates of population-scale impact attributable to the digital-observation infrastructure, controlling for non-pharmaceutical interventions.⁴⁴¹ The lesson for a respiratory observation layer is structural rather than epidemiological: digital-observation infrastructure, when correctly architected as observation rather than intervention, can operate at population scale within existing regulatory frameworks.

A second adjacency is the structured-data-layer pattern that has consolidated in clinical genomics. Variant interpretation, in particular, has matured into a structured layer between raw sequencing output and clinical use — Sherlock and the ACMG–AMP refinement frameworks codify the classification of variants of uncertain significance;^{442, 443} deep generative modelling against evolutionary data provides additional structured prior on variant pathogenicity;⁴⁴⁴ saturation-scale functional evidence enables systematic upgrading or downgrading of clinical variant interpretation;⁴⁴⁵ and Bayesian frameworks for incorporating clinical evidence into variant interpretation continue to refine the layer.⁴⁴⁶ The observation-layer pattern is the same: structured outputs (variant classifications) flow into clinical and research workflows; clinicians interpret; the structured layer accumulates evidence over time. Domain-specific applications — gene-specific colorectal-cancer surveillance,⁴⁴⁷ KCNQ1 / Long-QT implementation in defined populations,⁴⁴⁸ rare-disease neurogenetics workflow,⁴⁴⁹ calcium-signalling rare disease,⁴⁵⁰ developmental-genetics patterns⁴⁵¹ — instantiate the pattern. A respiratory observation layer is the analogous structured layer for a different upstream signal.

The economic-burden context for the broader category is large and well-characterised: physical inactivity alone imposes a substantial global economic burden in direct healthcare costs (2013), much of which is borne by the public sector,⁴⁵² and these are figures for one upstream behavioural exposure measured at a population scale. The structural argument is that an observation layer for breathing — surfacing deviations across 16 physiological systems, accumulating continuously, traceable to raw signal — is positioned to support downstream actors (clinicians, researchers, public-health institutions, pharmaceutical sponsors) at the level of architectural fit, not at the level of a narrower clinical product.

The conclusion of this chapter, then, is structural. The observation-layer category exists. Its commercial precedents have established the valuation envelope; its regulatory grammar has been articulated in the IMDRF SaMD framework⁴¹⁰ and the FDA CDS guidance⁴¹¹; its conceptual foundation rests on a well-understood theorem of biomedical informatics⁴¹²; its reimbursement infrastructure exists for continuous monitoring;⁴²³ and its integration patterns into EHR, wearables, and pharma research are recognisable. What has not previously existed is the corresponding observation layer for respiratory signal — the upstream biomarker substrate developed across Chapter 3 through Chapter 7 of this document, observable on smartphone-grade hardware established in Chapter 9, and accumulating structured longitudinal data of a kind no current platform produces. The architectural commitments of 10.2 are the discipline by which that layer earns the right to operate continuously, at scale, and within the existing regulatory framework — and the three verticals of 10.3 are the entry points through which it integrates with the health ecosystem on the terms that ecosystem already recognises.

The next chapter develops the methodological frontiers and open questions that this architecture leaves on the research agenda; Chapter 12 develops the regulatory positioning that operationalises the framework here.

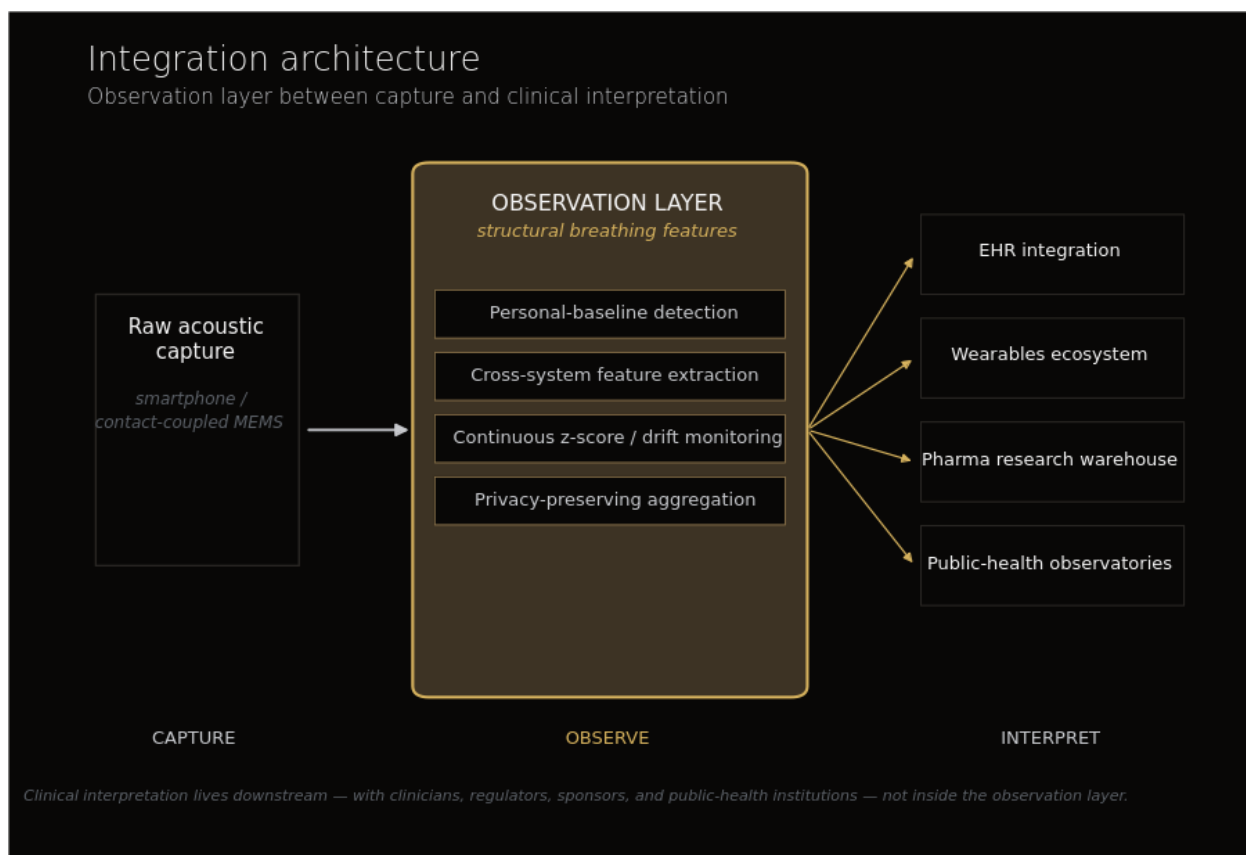


Figure 10.4. Integration architecture: where observation lives versus where clinical interpretation lives. Atum operates as a passive observation layer between capture (smartphone microphones, contact-coupled MEMS) and clinical interpretation performed by external systems (EHRs, wearables ecosystems, pharmaceutical research warehouses, public-health observatories). The layer executes four operations on raw acoustic input: personal-baseline detection, cross-system feature extraction, continuous z-score and drift monitoring, and privacy-preserving aggregation. Clinical interpretation stays downstream — this separation is simultaneously a regulatory posture (structured observations, not diagnostic outputs) and an integration strategy.

11. Open questions and methodological frontiers

The preceding chapters have established a structural argument: that breathing observation is a convergent physiological signal (Chapter 3), that it carries prognostic information across 16 body systems (Chapter 5), that it offers temporal advantage spanning nine orders of magnitude (Chapter 6), that this advantage is realized only with personalized baselines (Chapter 7), and that the public dataset infrastructure required to operationalize these findings does not yet exist (Chapter 8). Smartphone-grade acoustic capture is feasible in defined contexts (Chapter 9); an observation-layer architecture is consistent with current regulatory pathways (Chapter 10). This chapter does not catalogue defects in that argument. It identifies the next phase of inquiry — six methodological frontiers where the science is design-able rather than settled, and where the studies that would answer the open questions are themselves the natural next deliverables of a continuous-physiology research program. Three of these frontiers (12.1, 12.2, 12.3) correspond to specific clinical-validation studies whose protocols have been worked out in detail; two (12.4, 12.5) describe acoustic and longitudinal-systems engineering questions

that the broader research community shares; and one (12.6) frames privacy architecture as a design constraint rather than a downstream compliance task. None of these is a barrier to proceeding. Each is the question whose answer becomes the next sentence of the same argument.

Open scientific questions				
<i>Methodological frontiers and the post-Layer-3 research agenda</i>				
SECTION	RESEARCH QUESTION	DESIGN TYPE	ANCHOR LITERATURE	TIMELINE
12.1 BRV vs HRV	Does breathing-rate variability provide signal beyond HRV?	Paired observational	<i>Wehrwein 2025; Grossman 2023</i>	12–18 mo
12.2 Personalization RCT	Personalized vs population-threshold deterioration alerts	Two-arm RCT	<i>Chapter 8 trial-anchored evidence</i>	24–36 mo
12.3 Multi-person separation	Source-separation in cohabiting environments	Engineering / dataset	<i>Open ML problem</i>	R&D
12.4 Foundation model fit	Domain transfer of HeAR / OPERA / MIRA to natural breathing	Methodological	<i>Chapter 5.3 family literature</i>	12–24 mo
12.5 Privacy-by-design	Edge-only feature extraction; raw audio stays on device	Architecture / standards	<i>GDPR Art. 9; FDA Cybersecurity</i>	Continuous
12.6 Cross-population generalisation	Generalisation across demographics, devices, contexts	Population	<i>Validation framework</i>	24+ mo

These six questions define the post-Layer-3 research agenda. None block the observation-layer deployment; all sharpen it.

Figure 11.1. Open scientific questions and methodological frontiers. Six research questions structure the post-Layer-3 agenda: (12.1) incremental signal of BRV over HRV — paired observational design (12–18 mo); (12.2) personalized vs population thresholds — two-arm RCT (24–36 mo); (12.3) source-separation of co-located breathers — open ML problem; (12.4) domain transfer of respiratory foundation models — methodological inquiry (12–24 mo); (12.5) privacy-by-design with edge-only feature extraction — architectural standards work; (12.6) cross-population generalisation — validation framework (24+ mo). None of these questions blocks observation-layer deployment; all sharpen it.

11.1 Smartphone acoustic validation for Parkinson’s disease detection

The single strongest demonstration that breathing structure carries prognostic information for an upstream-of-symptom neurological condition comes from Yang and Katabi’s 2022 *Nature Medicine* report¹⁹³. In a cohort of 7,671 individuals (757 with Parkinson’s disease, 6,914 controls) drawn from Mayo Clinic, Massachusetts General Hospital, and the University of Rochester, an artificial-intelligence model trained on nocturnal breathing signals achieved an area under the receiver operating characteristic curve of 0.90 on held-out test data and 0.85 on external validation; a single night of capture yielded 86% accuracy, rising to 95% with twelve nights; severity tracking correlated with the Movement Disorder Society Unified Parkinson’s Disease Rating Scale at $R = 0.94$ ($p = 3.6 \times 10^{-66}$); and the model identified 75% of eventual Parkinson’s patients from data collected before clinical diagnosis. The result was named one of *Nature Medicine*’s ten Notable Advances of 2022. Its physiological grounding — that brainstem nuclei controlling respiration (notably the pre-Bötzing complex and the retrotrapezoid

nucleus) are affected early in Parkinson's, a fact noted by James Parkinson himself in 1817 — is consistent with the broader neurological evidence reviewed in 6.2 and the prodromal-window analysis in 7.1^{453, 454}.

The signal in this study, however, was acquired with radio-frequency reflectometry, not acoustic capture. The open question — and it is a study-shaped question, not a doubt — is whether smartphone-acquired nocturnal breathing audio approximates the diagnostic discrimination and severity correlation that radio-frequency capture demonstrated. Smartphone microphones have different signal-to-noise characteristics from radio-frequency sensors; ambient acoustic environments (room reverberation, partner respiration, household noise) introduce confounds that radio-frequency capture does not face; and the temporal granularity of acoustic breath-cycle detection in the consumer-grade context has not been validated against the polysomnographic gold standard at scale. None of these concerns negates the underlying result. They identify the modality-validation study that would extend it.

That a respiratory-acoustic signal can in principle carry Parkinson's-relevant information is supported by an adjacent but distinct literature on speech-acoustic biomarkers. Rusz and colleagues' 2011 quantitative acoustic study in early untreated Parkinson's disease provided the foundational measurement framework⁴⁵⁵; their 2021 work in *Annals of Neurology* extended the framework to rapid-eye-movement sleep behaviour disorder — a recognized prodromal stage — demonstrating that consumer-tractable acoustic capture in the speech domain detects pre-clinical neurodegeneration⁴⁵⁶. These results address speech, not breathing; but they constitute proof-of-modality that acoustic signals with smartphone-tractable bandwidths preserve Parkinson's-relevant information. The respiratory-acoustic question is the parallel one: does the breathing signal, captured acoustically rather than radio-frequeuncially, preserve what Yang and Katabi's radio-frequency capture demonstrated?

The study that resolves this question is design-able in the conventional case-control format. A prospective cohort of 150–250 participants comprising Hoehn and Yahr stage 1–3 Parkinson's patients and age- and sex-matched controls; 7–14 consecutive at-home nights of smartphone audio capture per participant; in-clinic neurological examination at baseline and at twelve months. The primary endpoints follow directly from the radio-frequency anchor: cross-validated area under the curve for Parkinson's versus control discrimination; correlation of an acoustic severity index with MDS-UPDRS Part III; and test-retest reliability across at least twelve nights, benchmarked against the 0.957 reliability reported in the radio-frequency cohort. The comparator is the Yang and Katabi result itself; the question is whether the acoustic modality preserves the radio-frequency signal at clinically meaningful magnitude. This design — a direct modality replication of a *Nature Medicine* result — is the simplest and the most consequential next study in the breathing-as-neurological-biomarker line of inquiry. It does not introduce new biology; it tests whether the established biology is observable through a more deployable instrument.

11.2 Breathing rate variability versus heart rate variability — head-to-head prognostic comparison

Heart rate variability has accumulated decades of prognostic literature across critical care, cardiology, and population-cohort epidemiology (cross-ref 4.4 and the trial-anchored cardiology evidence reviewed

in Chapter 7). Breathing rate variability has not. The question of whether the two carry redundant or independent prognostic information has rarely been asked in a cohort large enough to answer it. A second open question, accordingly, is the head-to-head comparison: in a population with simultaneous high-quality respiratory and cardiac measurement, does breathing rate variability provide prognostic information additional to, or superior to, heart rate variability for outcomes of clinical and public-health significance?

The structural case for asking the question has strengthened in recent literature. Wehrwein and colleagues' 2025 *Nature Reviews Cardiology* article formally proposed renaming respiratory sinus arrhythmia as respiratory heart rate variability⁷³, codifying a shift in how the cardio-respiratory coupling has been understood: the heart-rate oscillation that classical heart-rate-variability metrics quantify is, in mechanism, downstream of the respiratory oscillator. Grossman's 2023 critique in *Biological Psychology* offered an even stronger framing — that several premises of the polyvagal theory through which heart rate variability has historically been interpreted are likely refutations rather than supports — and re-grounded heart rate variability in respiration⁷⁴. Quintana and Heathers, writing in *Frontiers in Psychology* in 2014, had already identified respiration as a primary methodological confound to heart rate variability interpretation in biobehavioural research⁴⁵⁷. Together these references reframe what a head-to-head comparison would test: not two independent autonomic metrics, but a respiratory metric and its cardiac projection, with the unsettled question being how much prognostic information lives in the projection beyond what lives in the source.

The empirical precedent for breathing rate variability outperforming heart rate variability in a defined clinical context — a single multicenter study, but a direct one — comes from Seely and colleagues' 2014 *Critical Care* report. In 721 critically ill patients, eight respiratory-rate-variability measures and only two heart-rate-variability measures were significantly associated with extubation failure⁴⁵⁸. The study did not generalize to ambulatory or population cohorts, and it pre-dated the conceptual reframing in Wehrwein and Grossman; but it stands as the existing piece of evidence that, in at least one prognostic task, respiratory variability carries information that cardiac variability does not. The question for a general-population study is whether this generalizes.

The study design that would answer the question follows the clinical-research template: a paired observational cohort of approximately 200 adults aged ≥ 50 drawn from a preventive-cardiology or longevity programme, with simultaneous overnight smartphone-acquired breathing capture and ambulatory electrocardiography, and twelve-month prospective follow-up to a composite cardiovascular endpoint (myocardial infarction, heart-failure hospitalization, cardiovascular death). Primary endpoints are the Cox proportional-hazards ratios for breathing-rate-variability and heart-rate-variability metrics, adjusted and unadjusted; the incremental C-statistic when breathing-rate-variability metrics are added to heart-rate-variability-based risk models; and the correlation structure between the two — quantifying shared variance versus independent prognostic information. The question this design resolves is concrete: how much of the heart-rate-variability prognostic signal is, in fact, breathing-rate-variability projected through the cardiac autonomic system, and how much breathing information remains in the respiratory measurement after the cardiac projection has been accounted for.

A complementary methodological frontier is the single-subject longitudinal extension. The matched-cohort design tests population-level associations; it does not address whether, within an individual followed longitudinally, breathing rate variability and heart rate variability covary or diverge across personal physiological states. The reporting infrastructure for single-subject longitudinal designs is mature — the SCRIBE 2016 reporting guideline for single-case behavioural interventions⁴⁵⁹ and the CONSORT extension for n-of-1 trials^{460,461} together provide the methodological scaffolding — and the methodological precedents from sports physiology, where individual-baseline tracking of training load (TRIMP⁴⁶² and session-RPE⁴⁶³) and cardiorespiratory fitness (the heart-rate-ratio method⁴⁶⁴) is decades-old established practice, are the closest analogue. The bridge from population-cohort to single-subject longitudinal is the methodological connection between 12.2 and 12.3.

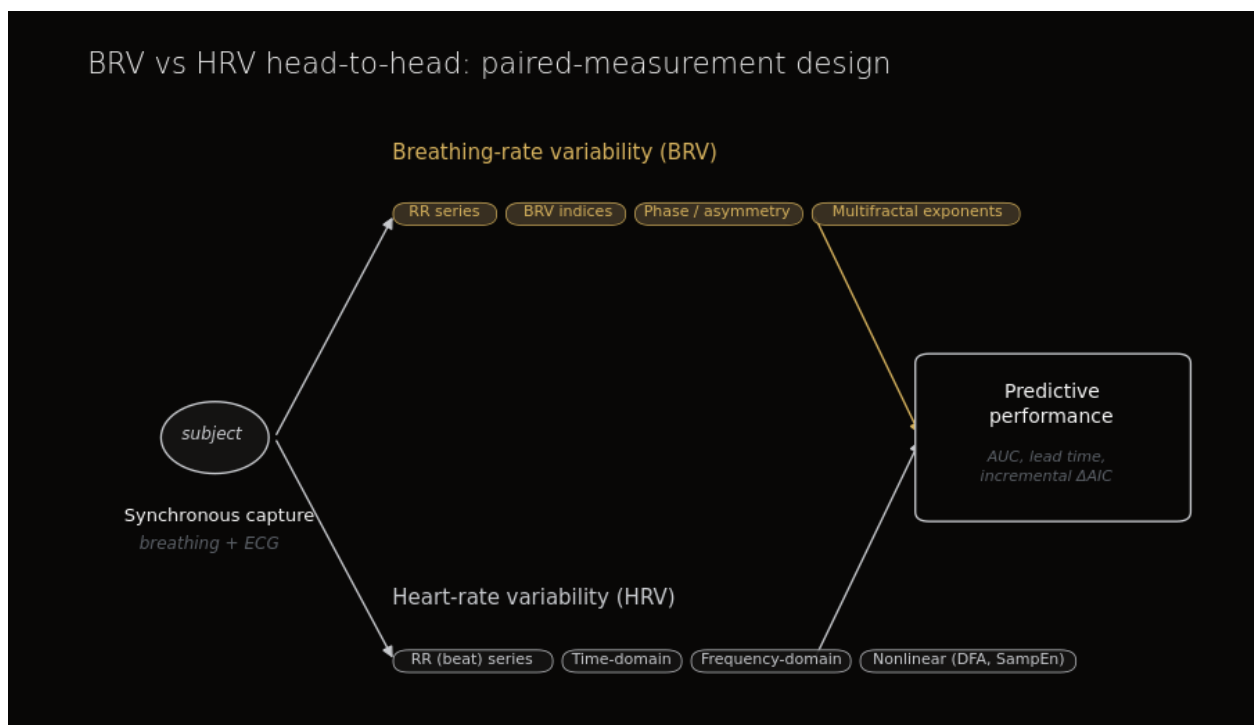


Figure 11.2. Paired-measurement design for BRV vs HRV head-to-head comparison. Each subject contributes time-locked breathing acoustics and ECG. Two parallel pipelines — BRV (top track) and HRV (bottom track) — are extracted from the same window and evaluated against the same outcomes under matched cross-validation, isolating incremental predictive value of BRV at AUC, lead-time, and ΔAIC level. The protocol is implemented as a population cohort with a longitudinal N-of-1 extension, supporting both between-subject and within-subject inference.

11.3 Personalized versus population-threshold randomized clinical trial in remote respiratory monitoring

The cardiology-telemonitoring literature reviewed in Chapter 7 established a pattern. Trials that applied population-derived thresholds to remote respiratory and physiological signals — BEAT-HF, TIM-HF, TELE-HF, CHROMED — failed to demonstrate hospitalization or mortality benefit. Trials that incorporated personalized-baseline approaches — LINK-HF, TIM-HF2, DyniEWS — succeeded. The pattern suggests, though it does not yet prove, that the prognostic information lives in departure from individual baseline, not in absolute crossing of population thresholds. The next-generation question is the interventional one: in a randomized clinical trial designed to compare the two approaches directly,

does personalized-baseline observation, when integrated into clinical workflow, improve outcomes compared with population-threshold observation? This study has not been done. It is the study that would resolve the Chapter 7 observational evidence into a Level-A clinical recommendation.

The methodological infrastructure for personalized longitudinal trials has matured over the last two decades into a recognizable subfield. Lillie and colleagues' 2011 *Personalized Medicine* paper proposed n-of-1 trial designs as the ultimate strategy for individualizing medicine⁴⁶⁵; Davidson and colleagues' 2021 *JAMA Pediatrics* paper extended the framework with experimental designs to optimize treatments for individuals through personalized n-of-1 trials⁴⁶⁶; Selker and colleagues' 2023 *Journal of Clinical and Translational Science* paper articulated how n-of-1 trials generate real-world evidence for individuals and populations simultaneously⁴⁶⁷. These are not isolated methodological proposals. Together with the CONSORT extension for n-of-1 trial reporting (the CENT 2015 statement and its explanation-and-elaboration companion)^{460, 461} and the SCRIBE 2016 guideline for single-case behavioural interventions⁴⁵⁹, they constitute a published reporting and analysis framework that a personalized-baseline observation trial can adopt directly. The design problem is not that the methodology is missing; it is that the methodology has not yet been applied at scale to a continuous-physiology observation system in a hospitalization-endpoint trial.

The structural design follows the head-to-head template. A randomized two-arm parallel-group trial in a high-event-rate cohort — heart-failure patients post-discharge, or chronic-obstructive-pulmonary-disease patients in their first year after a moderate-to-severe exacerbation — with one arm receiving population-threshold observation (the comparator condition that BEAT-HF / TIM-HF / TELE-HF / CHROMED operationalized) and the other arm receiving personalized-baseline observation (the condition that LINK-HF / TIM-HF2 / DyniEWS established). Primary endpoint: thirty-day or ninety-day all-cause hospitalization, with a quality-of-life secondary endpoint capturing patient-reported outcomes that hospitalization alone does not. The role of the observation layer in this trial is, by design, observational: the observation system surfaces deviations from individual baseline (the personalized arm) or crossings of population thresholds (the comparator arm) to clinical workflow; clinical workflow makes the decision and intervenes. The trial tests whether the personalized observational input, integrated into otherwise-equivalent clinical workflow, improves outcomes. It does not test the observation system as a clinical-decision instrument; the decision remains with the clinician and the standard care pathway.

The methodological precedent from sports physiology is informative. Banister's 1991 TRIMP framework operationalized training load as an individual-specific exponentially-weighted physiological dose⁴⁶²; Foster's 1998 session-rating-of-perceived-exertion method generalized the concept to perceived training load⁴⁶³; Uth and colleagues' 2004 heart-rate-ratio method for non-exercise estimation of cardiorespiratory fitness extended individual-baseline tracking to the cardiovascular outcome of interest⁴⁶⁴. These methods are decades-old established practice in their domain because they do exactly what the personalized-versus-population question asks: they replace population thresholds with individual-baseline trajectories as the clinically interpretable signal. The clinical-research analogue — applying this design philosophy to remote respiratory observation in heart failure or chronic respiratory disease — is the natural transposition.

A note on what the trial does not test. It does not test whether breathing observation is useful (Chapters 6 and VIII); it does not test whether smartphone-grade acoustic capture is feasible (Chapter 9); and it does not test whether personalized baselines are in principle better than population thresholds (the cardiology-telemonitoring meta-pattern in Chapter 7 has already established this at observational level). It tests whether the integrated personalized-observation system improves outcomes when fielded against the integrated population-threshold system in a head-to-head randomized comparison. That is the gold-standard question, and it is the question whose answer would close the inferential gap between observational evidence and clinical recommendation.

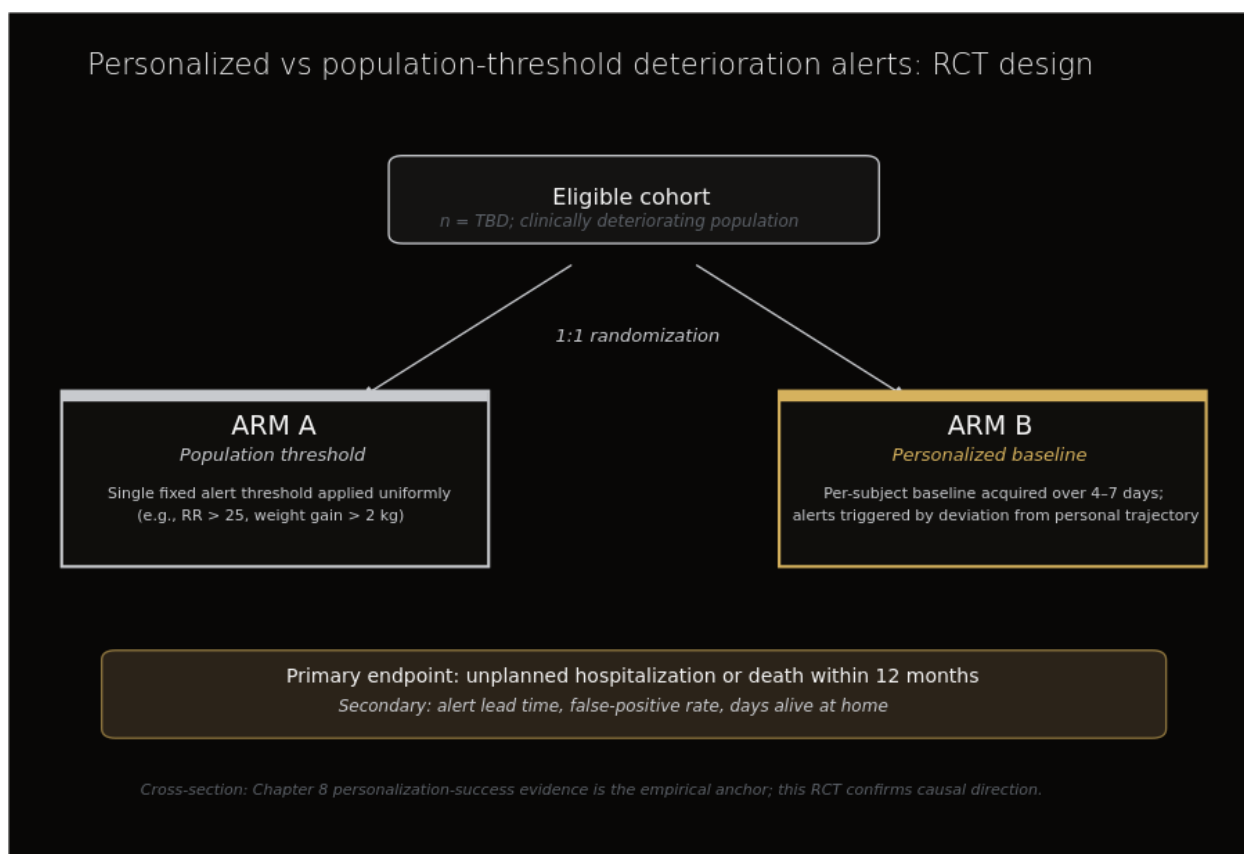


Figure 11.3. Trial design: personalized-baseline alerts versus population-threshold alerts in a deteriorating cohort. Two-arm parallel-group RCT for empirical confirmation of the personalization hypothesis (Chapter 7). Eligible subjects from a clinically deteriorating population are randomized 1:1 between ARM A (single fixed threshold, e.g. RR > 25, weight gain > 2 kg) and ARM B (per-subject baseline acquired over 4–7-day calibration window, alerts triggered by personal-trajectory deviation). Primary endpoint: composite of unplanned hospitalization or all-cause death within 12 months; secondary endpoints: alert lead time, false-positive rate, days alive at home. The trial establishes whether replacing population thresholds with personal-baseline deviation produces a clinically meaningful outcome shift.

11.4 Multi-person source separation as an acoustic engineering constraint

A research-grade acoustic recording in a controlled environment with one participant breathing at a known distance from a calibrated microphone is not the deployment context. The deployment context is a domestic environment with cohabitating adults, possibly children, occasional pets, and an ambient soundscape that includes television, conversation, kitchen activity, and traffic. The question of how to separate one person's breathing from the rest of this acoustic field is an open engineering problem at

consumer-deployment scale. It is also one of the largest near-term questions for any ambient-acoustic observation system; cross-ref 10.4 limitation domain.

The research community has been working on adjacent versions of this problem for decades. Ruinskiy and Lavner's 2007 algorithm for automatic detection and exact demarcation of breath sounds in speech and song signals⁴⁶⁸ achieved 98% identification rate and 96% specificity for breath-event detection in single-speaker contexts, and remains one of the most-cited prior-art references for breath-event separation; the underlying template-matching approach scales poorly to mixed-source domestic capture, but the work demonstrated the basic acoustic feasibility of breath-sound discrimination from speech under controlled conditions. The deeper engineering question — separating one source's signal from a mixture containing several — has been the central problem of the speech-separation literature, and the field has advanced substantially through deep-learning approaches. Luo and Mesgarani's 2019 Conv-TasNet⁴⁶⁹ introduced a convolutional time-domain audio-separation network that surpassed ideal time-frequency magnitude masking for speech separation, marking a turning point in the technical state-of-the-art; Park and colleagues' 2022 review of speaker diarization in *Computer Speech and Language*⁴⁷⁰ consolidated recent advances, including x-vector clustering, end-to-end neural diarization, and overlap-region handling — the family of techniques that addresses the “who is speaking when” question in multi-talker recordings.

None of this literature is the same as solving multi-person breath-source separation in a domestic ambient-acoustic recording. Speech-separation networks are trained on speech mixtures; their behaviour on breath-sound mixtures is unstudied at the same scale. Speaker diarization assumes acoustic activity occurs in turns (someone speaks, then someone else speaks); cohabitating breath capture has the opposite structure (multiple sources active continuously and concurrently). The geometry of microphone-to-source distance is also different: speech-separation literature is dominated by close-talking or conference-room geometries; domestic breathing capture is dominated by smartphone-on-nightstand geometries with one source close, others further, all attenuated relative to the controlled-recording case. The engineering primitives — single-channel time-domain separation networks, multi-channel beamforming where a microphone array is available, end-to-end neural diarization adapted to breath-event activity — are all available in adjacent literatures. The combination required for consumer-deployment-scale multi-person breath-source separation is not yet a demonstrated pipeline.

The framing here is important. This is not a question of whether smartphone breathing observation is feasible. The single-occupant cases in which feasibility has been demonstrated (Chapter 9) are real, validated, and large. The multi-person case is a scope question. It identifies the cohabitating-deployment domain as a domain that requires additional engineering rather than a domain in which the underlying observation principle fails. The expected timeline — two to four years — reflects the maturation cycle of the speech-separation literature toward the breath-source-separation problem, not a fundamental feasibility barrier. The acoustic engineering open question is the largest near-term scope question for the observation system; it is not a categorical objection to the system.

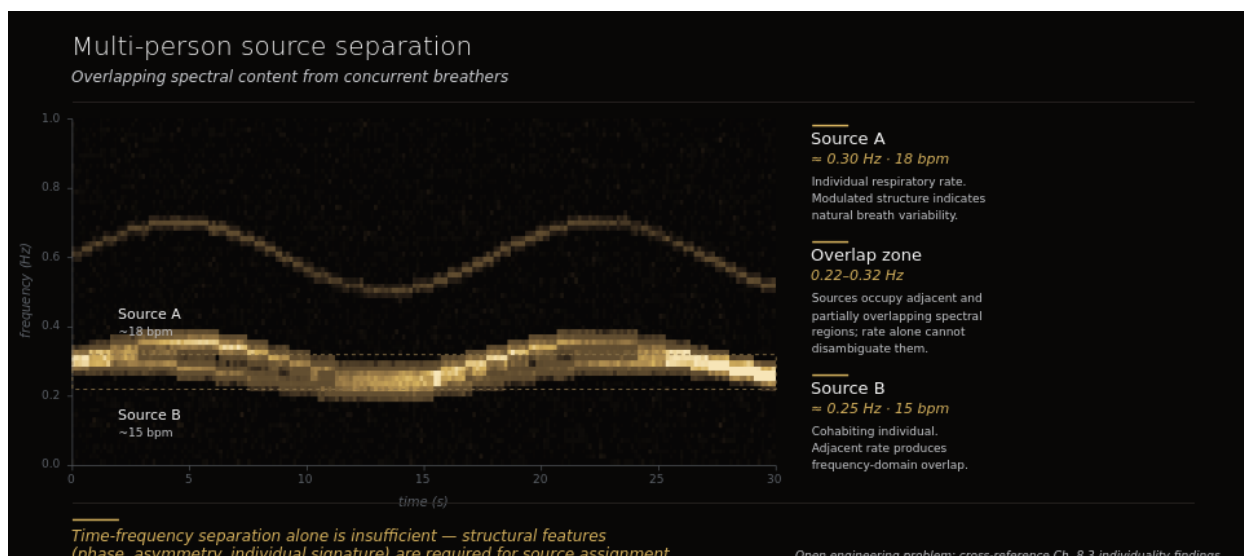


Figure 11.4. Multi-person source separation: overlapping spectral content from concurrent breathers. STFT showing two cohabiting breathers at ≈ 18 bpm (Source A, ~ 0.30 Hz) and ≈ 15 bpm (Source B, ~ 0.25 Hz). Within the dashed gold band (0.22–0.32 Hz) the two sources occupy adjacent and partially overlapping spectral regions; rate alone cannot disambiguate them. Reliable assignment requires structural features beyond rate — phase alignment, inhalation/exhalation asymmetry, and breath-shape signatures functioning as individual identifiers (Ch. 7.3 individuality findings).

11.5 Long-term reliability of smartphone acoustic capture across devices and environments

Personalized-baseline observation requires longitudinal stability of the measurement instrument. If the measurement drifts faster than the personal baseline updates, the deviation signal becomes uninterpretable. Smartphone hardware, however, is not a stable measurement instrument over a decade-long longitudinal horizon. Microphones age; users replace devices; operating-system updates modify audio capture pipelines. The cross-device, cross-time stability question is, accordingly, an open engineering question for any longitudinal consumer-grade observation system — not a question particular to breathing observation, but one that breathing observation inherits.

Ginsburg and colleagues' 2024 *New England Journal of Medicine* perspective on key issues as wearable digital health technologies enter clinical care⁴⁷¹ articulates the general framing: clinical-grade longitudinal observation requires measurement equivalence across the time horizon over which clinical decisions are made, and consumer-grade hardware was not originally engineered against this requirement. Three nested challenges follow. Hardware degradation: microphone diaphragms and the surrounding mechanical and electrical components do drift over multi-year time horizons, and the drift rate varies across manufacturers and across environmental exposure (humidity, temperature, mechanical shock). Device migration: the average user replaces a smartphone every two to three years, and the replacement is rarely with an identical model — the migration introduces a step change in microphone characteristics that the observation system must absorb. Operating-system API drift: iOS and Android audio capture pipelines have changed across versions, occasionally in ways that affect the relevant frequency bands, sample-rate handling, or pre-processing applied by the platform before the observation application receives the audio.

Each component of a longitudinally deployed acoustic capture chain (microphone, codec, pre-processing pipeline, feature extraction) must be validated against the longitudinal-stability requirement separately, and the calibration / normalization layer that maintains measurement equivalence across devices and time becomes a required subsystem rather than an optional refinement. The cross-section connection to 9.4 device-variation limitation is direct: the same engineering primitives that handle device-to-device variation at a single time point also handle device-to-device variation across time, and the same calibration-normalization approach extends across both axes.

What is not yet established at scale in the published literature is the empirical longitudinal-stability dataset that would allow the calibration-normalization approach to be designed against real device-aging and OS-drift trajectories. The studies that would populate this gap are themselves valuable: they characterize what consumer-grade longitudinal observation requires, beyond the breathing-observation-specific case, and they generalize to wearable digital health more broadly per the Ginsburg framing.

The framing parallels 12.4. This is not a feasibility-of-acoustic-capture question (Chapter 9); it is a longitudinal-systems-engineering question that any consumer-grade ambulatory observation system must engage with, and it is addressable through component-level validation with empirical inputs from a longitudinal-stability study. The expected timeline — one to three years for the empirical characterization, with the calibration-normalization subsystem development running in parallel — places it in the near-term engineering frontier rather than the multi-year scientific horizon.

11.6 Privacy architecture: BIPA, HIPAA, GDPR, and edge-processing as design constraint

The sixth open question is structural rather than empirical. Always-on acoustic capture in domestic settings interacts with privacy-protective regulatory frameworks in ways that text-based, photographic, or wearable physiological observation modalities do not — ambient audio captures a richer-than-needed signal (speech content, identity-revealing voice characteristics, household activity) alongside the breathing signal it is designed to observe. The research-architectural question is how to design observation systems where the privacy-engineering choices are co-equal design constraints with the signal-extraction choices. The actionable regulatory navigation — what specific consent flows, what specific data-handling commitments, what specific processor agreements — belongs to Chapter 12; the present subsection treats privacy architecture as a research-agenda topic, identifying the engineering questions that the regulatory frameworks force onto the system design.

Three regulatory frameworks anchor the constraint set. The European Union's General Data Protection Regulation, particularly Article 9, designates biometric and health data as special categories requiring explicit consent and purpose limitation⁴⁷²; ambient audio in a health context falls under Article 9 protections by both pathways — voice characteristics are biometric, and respiratory patterns are health data. Illinois's Biometric Information Privacy Act, enacted in 2008, regulates voice biometrics as a biometric identifier and imposes specific consent and disclosure obligations, with a private right of action that has produced material liability exposure⁴⁷³. The 2020 Facebook biometric-information settlement, arising under the Biometric Information Privacy Act, is the most-cited enforcement precedent and frames the upper bound of biometric-information liability for systems that handle

voiceprint or comparable identifiers.⁴⁷⁴ The Health Insurance Portability and Accountability Act in the United States, particularly the Privacy Rule and the Security Rule (45 CFR 160, 162, 164), regulates protected health information handling, and the Security Rule's technical-safeguards provisions apply to ambient-audio capture in clinical-research contexts⁴⁷⁵.

The privacy-engineering literature provides the design vocabulary. Nautsch and colleagues' 2019 *Computer Speech and Language* survey of cryptographic and machine-learning approaches to privacy-preserving speaker and speech characterization⁴⁷⁶ consolidates the available techniques, including homomorphic encryption (computation on encrypted data without decryption), secure multi-party computation (joint computation across parties without sharing inputs), and differential privacy (statistical guarantees on individual-record exposure in aggregate analyses). Tomashenko and colleagues' 2022 *Computer Speech and Language* report on the VoicePrivacy 2020 Challenge⁴⁷⁷ — the first international voice-anonymization benchmark — characterized the empirical privacy-utility trade-off: how much voice-identity information can be removed before the downstream task (which, in the present context, is breathing-pattern observation rather than speech recognition) is degraded. Rieke and colleagues' 2020 landmark *npj Digital Medicine* review of federated learning for digital health⁴⁷⁸ describes the architecture in which model training proceeds across distributed edge devices without centralizing raw data — the data never leaves the device; only model updates are transmitted, and only with consent.

Edge processing is the organizing design constraint that integrates these literatures into a system architecture. Acoustic feature extraction performed on-device, before any signal leaves the device, reduces the privacy surface area substantially: raw audio is never transmitted; only the structured features required for the observation task — and only those — are sent to any backend, with explicit consent and purpose limitation. The federated-learning paradigm extends this principle to model improvement: training updates can be aggregated across users without raw audio (or even raw features) ever centralizing. Differential-privacy mechanisms can be layered onto these aggregations to bound individual-record exposure even in the model-update channel. None of these techniques is an afterthought to the system design; each is a primary design constraint that shapes which observation tasks are feasible, which acoustic features are extracted, and where in the pipeline each processing step occurs.

The research-agenda framing is, accordingly: how to architect always-on respiratory-acoustic observation systems where the privacy-engineering choices and the signal-extraction choices are co-designed from first principles, rather than appended as compliance retrofit. This is the open question that distinguishes ambient-audio health observation from precedent modalities (text, photographs, wearable accelerometry, even continuous ECG), each of which faces privacy questions but none of which faces the conjunction of biometric-identity content, health-data content, and ambient-third-party content that ambient audio uniquely combines. The cross-section connection to 13.4 is direct: this subsection identifies the design constraint; Chapter 12 addresses how the design constraint is operationalized within current regulatory pathways.

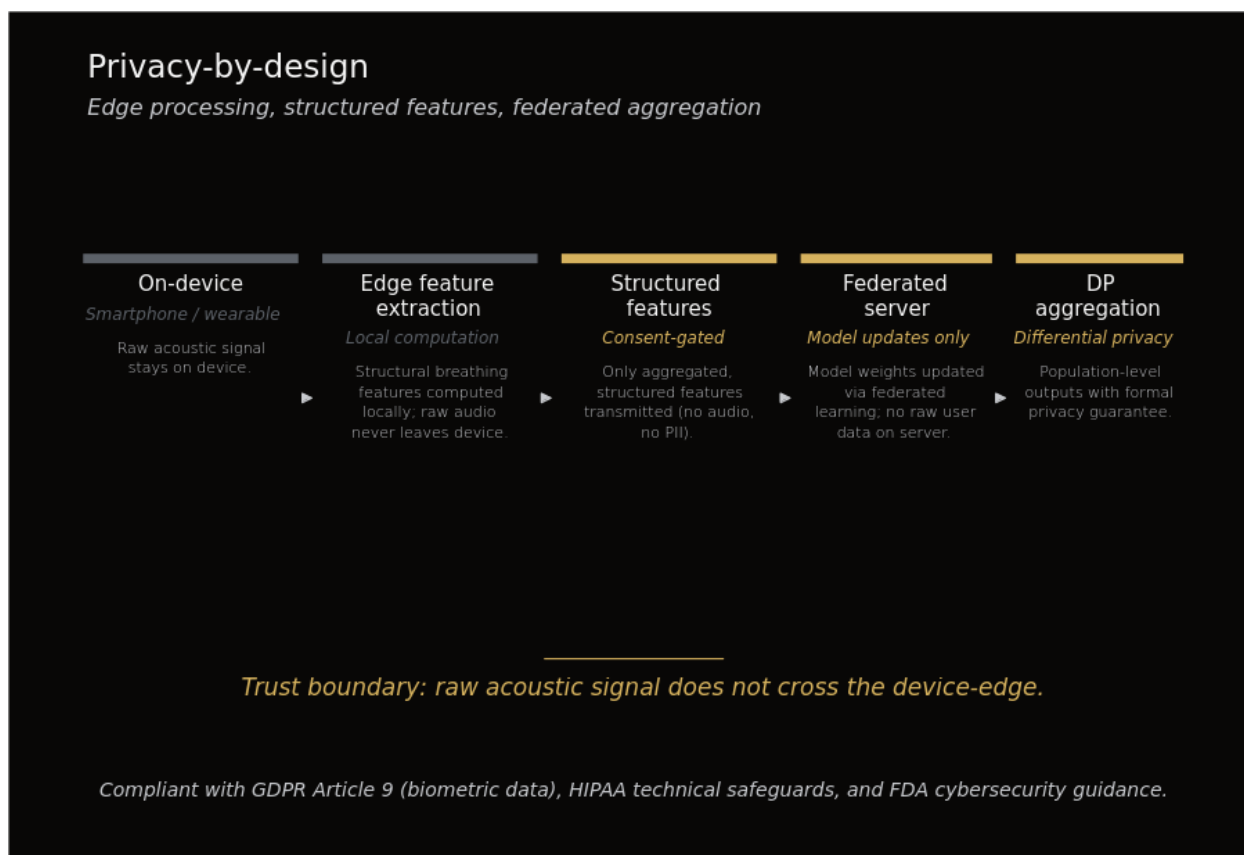


Figure 11.5. *Privacy-by-design architecture: edge processing, structured features, federated aggregation. Five-stage data pipeline. Stages 1–2 (on-device): raw audio captured, structural features computed locally — raw audio never leaves the device. Stage 3 (consent-gated): only aggregated structured features transmitted off-device. Stages 4–5 (federated server): model weights updated via federated learning; population-level outputs carry formal differential-privacy guarantee. The trust boundary sits between Stage 2 and Stage 3; architecture is compliant with GDPR Article 9 (biometric data), HIPAA technical safeguards, and FDA cybersecurity guidance.*

Closing synthesis

These six open questions are the next phase of inquiry, not defects of the present work. The convergent-signal argument (Chapter 3), the multi-system evidence base across 16 body systems (Chapter 5), the temporal-advantage analysis (Chapter 6), and the personalized-baseline framing (Chapter 7) establish the scientific structure within which respiratory observation operates as an upstream-of-symptom continuous-physiology signal. What remains is methodological development at the modality, biomarker, interventional, engineering, longitudinal, and privacy frontiers — 12.1 testing whether smartphone-acoustic capture preserves the radio-frequency demonstration in Parkinson's disease; 12.2 testing whether breathing-rate variability carries prognostic information beyond what heart-rate variability provides; 12.3 testing whether personalized-baseline observation outperforms population-threshold observation in a randomized clinical comparison; 12.4 addressing multi-person acoustic source separation as the largest near-term engineering scope question; 12.5 establishing the cross-device, cross-time stability characterization that any longitudinal consumer-grade observation system requires; and 12.6 framing privacy architecture as a co-equal design constraint with signal extraction in always-on ambient-acoustic systems.

Two (12.4, 12.5) are engineering frontiers shared with the broader acoustic and digital-health communities, and the studies that would address them generalize beyond breathing observation to wearable and ambient digital health more broadly. One (12.6) is structurally a system-architecture question that ambient-audio observation modalities uniquely face and that distinguishes the privacy-engineering work this domain requires from the precedents available in other passive-sensing modalities.

The structural argument of this scientific foundation is that breathing observation is the most information-dense continuously accessible physiological signal observable from existing consumer hardware, and that an observation-layer architecture is the appropriate translation of this fact into a fielded digital-health system. The structural argument of this chapter is its complement: the open questions identified here are the bridge from the current evidence base to the next-generation evidence base, and each is design-able as a specific study or engineering programme rather than a barrier to the system the rest of this monograph has described. The next phase of inquiry is, accordingly, not deferred — it is specified.

12. Regulatory positioning

The regulatory positioning of this document's proposed observation layer is not speculative. Three U.S. Food and Drug Administration (FDA) regulatory categories — General Wellness, the 510(k) screening pathway, and Software as a Medical Device (SaMD) — together describe a graded path from non-submission consumer products to clinically validated diagnostic tools. The acoustic respiratory monitoring field has working precedent in each category. The privacy and consent architecture aligned with the U.S. Health Insurance Portability and Accountability Act (HIPAA), the European Union General Data Protection Regulation (GDPR), and the Illinois Biometric Information Privacy Act (BIPA) is similarly well-defined for ambient-audio digital health products. This chapter walks through the four regulatory anchors that govern the observation layer architecture described in Chapter 10: the FDA General Wellness rationale (13.1), the 510(k) predicate landscape (13.2), the SaMD pathway held in reserve (13.3), and the privacy and consent architecture (13.4).

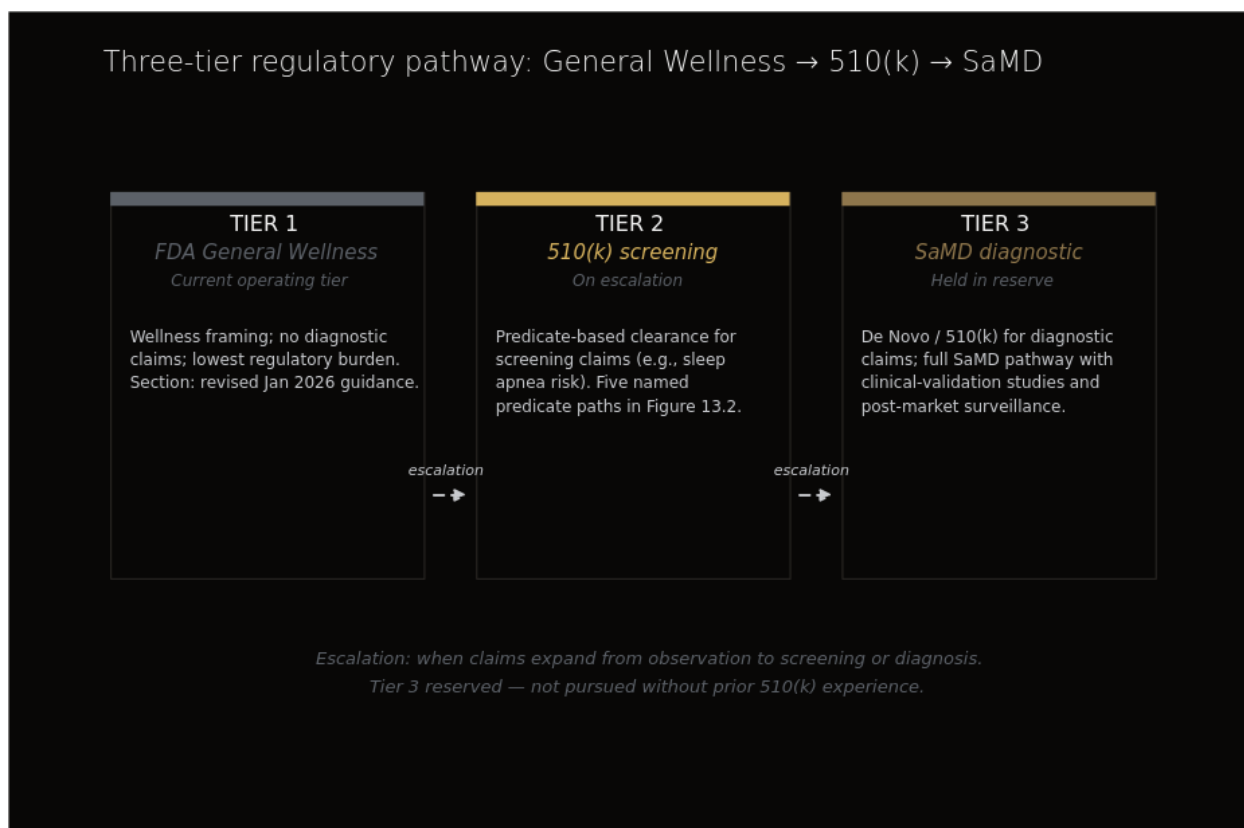


Figure 12.1. Three-tier regulatory pathway: General Wellness → 510(k) → SaMD. Atum currently operates at Tier 1 (FDA General Wellness Policy), which permits wellness framing with no diagnostic claims and matches the observation layer's structural identity. Tier 2 (510(k) screening clearance) is the active escalation target when claims expand from observation to screening (e.g., sleep apnea risk), following a predicate-based pathway detailed in Figure 12.2. Tier 3 (full SaMD via De Novo or 510(k)) is held in reserve and not pursued without prior 510(k) experience. Colour gradient — gray/gold/dim gold — encodes operational status.

12.1 FDA General Wellness rationale (January 2026 guidance)

The FDA General Wellness Policy was originally issued in 2016 as a low-burden regulatory framework for consumer health products that do not make disease-specific findings. The policy was updated on January 6, 2026, and the updated guidance explicitly contemplates non-invasive sensors that measure physiologic respiratory parameters as products that may qualify as general wellness without FDA premarket submission, provided that the labeling, intended use, and consumer-facing findings observe a defined set of guardrails.²⁸

Five guardrails define the General Wellness category boundary. First, products in this category must not reference specific diseases — language such as “for chronic obstructive pulmonary disease (COPD)”, “for asthma”, or “for sleep apnea” is excluded. Second, products must not present diagnostic thresholds or clinically interpretable cut-points that mimic those used in regulated medical devices. Third, products must not generate clinical alerts that prompt a specific medical action. Fourth, products must not finding equivalence to FDA-cleared devices or finding medical-grade accuracy. Fifth, permissible language is restricted to associations grounded in well-understood healthy lifestyle relationships — phrasing such as “may help reduce the risk of” and “may help with living well with” is permitted when supported by the published lifestyle literature.²⁸

The structural fit between the observation layer architecture (Chapter 10) and the General Wellness category is direct. The observation layer does not generate diagnostic outputs, does not stratify risk by score, and does not prompt medical action; it observes respiratory structure, surfaces deviations from individual baseline, and reports trajectory information. This non-diagnostic posture is a deliberate architectural commitment, not a regulatory accommodation — it is the correct scientific framing for a continuous physiological signal whose clinical meaning is conditional on context, individual baseline, and downstream confirmation by a licensed clinician (11.3).

The General Wellness pathway is the immediate go-to-market regulatory path for the Wellness vertical defined in 10.3. No FDA premarket submission is required.

The General Wellness positioning has direct precedent in the consumer wearable cardiology literature. The smart-wearable cardiovascular care review by Bayoumy and colleagues mapped the regulatory positioning of consumer-grade wearables against clinical-grade devices, identifying the General Wellness category as the dominant initial-launch posture for novel sensor modalities entering the consumer health market.⁴²⁴ The Cardiogram + Apple Watch deep neural network on photoplethysmography (PPG) signals, validated by Tison and colleagues for atrial fibrillation detection, is the canonical peer-reviewed example of a consumer wellness-positioned wearable demonstrating regulated-grade detection performance prior to a regulatory transition.⁴²⁵ The Apple Heart Study, published by Perez and colleagues in the *New England Journal of Medicine*, established at population scale (n = 419,297) that a consumer wellness-positioned smartwatch could identify a clinical biomarker (atrial fibrillation) with a positive predictive value of 0.84 (95% confidence interval 0.76–0.92) for the irregular-pulse notification.⁴⁰² The Apple regulatory trajectory — wellness launch in 2015, Apple Heart Study in 2017–2019, ECG De Novo clearance in 2018, Atrial Fibrillation History 510(k) in 2022, and clinical-trial-endpoint certification in 2024 — is the operational template for the wellness-to-clinical regulatory escalation that the SaMD pathway in 13.3 holds in reserve.

The General Wellness category is therefore not a regulatory limitation but the structurally correct entry posture for an observation layer whose value proposition is non-diagnostic deviation tracking, whose clinical-interpretation pathway routes through downstream confirmation, and whose data architecture is designed to support — rather than substitute for — the licensed clinician's downstream judgment.

12.2 510(k) predicates

The 510(k) pathway is the proven regulatory path for respiratory digital monitoring. Every FDA-authorized respiratory digital monitoring device cleared to date has used the 510(k) pathway. The predicate device landscape is sufficiently developed that the regulatory uncertainty for novel acoustic respiratory monitoring applications is incremental rather than existential. Five predicate clearances anchor this position.

510(k) predicate landscape: cleared respiratory-acoustic devices

DEVICE	K-NUMBER	CLEARED	INDICATION	SENSOR	PRIMARY METRIC
ResApp SleepCheckRx	K213360	2022	Sleep apnea risk screening	Smartphone microphone	AHI \geq 15 (Sn 89.3% / Sp 77.6%)
Sound Life Sciences	K211387	2021	Sleep apnea screening	Smartphone microphone	Mod-severe OSA (Sn 94% / Sp 97%)
Strados Labs RESP	LA-00055	Multiple	Continuous respiratory monitoring	Contact-coupled MEMS	Wheeze detection 85% accuracy
Strados Medical	Multiple	Multiple	Acoustic respiratory monitoring	Contact-coupled MEMS	Adventitious sound classification
FaceHeart FH Vitals SDK	K243966	9 Apr 2025	Respiratory rate monitoring	Smartphone camera (PPG-derived)	RR within \pm 3 bpm of reference

Figure 12.2. 510(k) predicate landscape for respiratory-acoustic devices. Five FDA-cleared predicates spanning three modality classes — smartphone-microphone OSA screening (ResApp, Sound Life Sciences), contact-coupled MEMS for continuous respiratory monitoring (Strados, multiple devices and clearances), and smartphone-camera-derived respiratory rate (FaceHeart, 2025) — establish that acoustic respiratory measurement has crossed the regulatory threshold and anchor the predicate baseline for Atum's Tier-2 SaMD escalation pathway. Primary reference: FDA 510(k) Premarket Notification database (accessdata.fda.gov/scripts/cdrh/cfdocs/cfPMN/).

ResApp SleepCheckRx — FDA 510(k) K213360, July 2022

ResApp SleepCheckRx is a smartphone application that uses passively recorded breathing and snoring sounds to identify the likelihood of moderate-to-severe obstructive sleep apnea (OSA, apnea-hypopnea index \geq 15). Reported clinical performance is sensitivity 89.3% and specificity 77.6% against polysomnography in the pivotal cohort.³⁸⁵ The intended-use language frames the device explicitly as a screening adjunct, not a diagnostic substitute — clinicians retain the responsibility for diagnostic decision-making, and confirmatory polysomnography or home sleep apnea testing remains the diagnostic standard. The ResApp adult OSA pathway was preceded by adult and pediatric respiratory cough-classification work — Sharan and colleagues developed the spirometry-prediction regression on cough acoustics that anchored the adult-COPD-screening line of work,³⁹⁷ and Porter and colleagues conducted the pediatric multicenter prospective study (n = 524) that demonstrated 89% overall accuracy with 90–98% per-condition accuracy for common pediatric respiratory disorders.³⁹⁸ Pahar and colleagues subsequently demonstrated COVID-19 cough classification performance (92% sensitivity, 80% specificity, area under the receiver operating characteristic curve [AUC] 0.93, n = 741) using the smartphone-acoustic infrastructure that ResApp developed for the SleepCheckRx submission.⁴⁷⁹ Pfizer acquired ResApp Health in 2022, validating the commercial proposition of the smartphone-acoustic respiratory monitoring category.

Sound Life Sciences — FDA 510(k) K211387, December 2021

Sound Life Sciences cleared a contactless sonar-based respiratory monitoring application that uses smartphone speakers and microphones to emit and detect inaudible acoustic signals reflecting from the user's chest, recovering breathing motion without contact. Reported clinical performance is sensitivity 94% and specificity 97% against polysomnography for moderate-to-severe OSA.³⁸⁷ The Sound Life Sciences pathway is notable for two reasons. First, the modality (active sonar via smartphone hardware) demonstrates that the FDA's 510(k) substantial-equivalence framework accommodates substantively different sensor approaches within the respiratory monitoring category. Second, Sound Life Sciences achieved clearance approximately three years after its 2018 founding as a University of Washington spinout, with a small team (~6 employees) and limited capital from U.S. National Institutes of Health and Biomedical Advanced Research and Development Authority funding — Google subsequently acquired the company. The Sound Life Sciences trajectory establishes that the regulatory and capital barriers to respiratory 510(k) clearance are tractable for a focused team with a defined predicate strategy.

Strados Labs RESP — FDA 510(k) LA-00055, December 2020 (expanded May 2022)

Strados Labs cleared a wearable lung-sound monitoring platform consisting of two lightweight chest-wall sensors (~100 g each) that capture continuous lung sounds with noise-canceling support, transmitted by Bluetooth to a mobile device and processed via a cloud platform that performs adventitious-breath-sound, cough-detection, and respiratory-dynamics analysis. Strados Labs has reported greater event-detection sensitivity than intermittent electronic stethoscopy in emergency-department cohorts (3,006 wheezes vs 607, 531 coughs vs 85, in 48 patients at New York City Health + Hospitals over 46.08 hours of continuous recording vs 0.53 hours of stethoscopy), and continuous auscultation with the RESP device detected wheezing in 85% of 40 patients with COPD or asthma exacerbation versus 77.5% with the Littmann electronic stethoscope.³⁸⁹ The Strados Labs platform has been incorporated into pharmaceutical clinical trial designs for respiratory syncytial virus therapeutic development, anchoring the clinical-research vertical use case described in 11.3.

Strados Medical — additional FDA 510(k) clearances, 2023

Strados Medical holds additional FDA 510(k) clearances for the RESP biosensor platform covering hospital and home-monitoring indications.³⁹⁰ The exact K-number enumeration is not fully recorded in the bibliography under review.³⁹⁰ The Strados Medical clearance trajectory establishes that the 510(k) pathway accommodates clearance expansion across care settings (acute hospital, post-discharge, home monitoring, clinical trial endpoint capture) on the same underlying sensor platform, which is structurally relevant to any observation layer that intends to span multiple clinical use environments.

FaceHeart FH Vitals SDK-RR — FDA 510(k), April 2025

FaceHeart cleared a contactless video-based respiratory rate measurement software development kit (SDK) that derives breathing rate from facial-region video captured via smartphone camera. The 2025 FaceHeart clearance is the most recent respiratory monitoring 510(k) clearance and demonstrates that

the FDA continues to clear novel sensor modalities (camera-based contactless respiratory rate, in addition to acoustic and sonar) within the established 510(k) framework.⁴⁸⁰

The pattern across these five predicates is consistent. Each device is cleared with intended-use language that frames the output as a screening or monitoring adjunct rather than a diagnostic substitute. Each clearance accommodates a distinct sensor modality (passive acoustic, active sonar, contact wearable acoustic, contactless video). Each clearance is anchored on a substantial-equivalence argument relative to a prior cleared device. The 510(k) screening category is established and well-trafficked for respiratory digital monitoring, and the regulatory uncertainty for a novel acoustic application within this category is incremental — predicate identification, comparator selection, study design, and intended-use language audit — rather than existential.

12.3 Pathway to Software as a Medical Device

The Software as a Medical Device (SaMD) pathway is the regulatory framework available for software products that perform medical-device-defined functions independent of a physical device platform. The International Medical Device Regulators Forum (IMDRF) Software as a Medical Device Working Group published the foundational SaMD risk-categorization framework in 2014, establishing four categories (I through IV) based on the seriousness of the healthcare situation addressed and the significance of the information provided to the healthcare decision.⁴¹⁰ The IMDRF framework has been adopted as the international harmonization reference for software medical device regulation and underpins both the FDA's SaMD policy approach and the European Union's Medical Device Regulation (EU MDR 2017/745) software classification rules. Under the EU MDR, software intended to provide information used to make decisions with diagnostic or therapeutic purposes is classified as Class IIa under Annex VIII, Rule 11, with higher risk situations escalating to Class IIb or III.⁴⁸¹

The SaMD pathway is held in reserve, not invoked. The three initial verticals identified in 11.3 — Wellness, Sleep health, and Clinical research — are non-diagnostic by construction. The Wellness vertical is positioned within the General Wellness regulatory category (13.1). The Sleep health vertical operates within the 510(k) screening regulatory category, leveraging the predicate landscape mapped in 13.2. The Clinical research vertical operates as a research data infrastructure layer for pharmaceutical trial endpoint capture and longitudinal cohort observation, where the observation layer outputs serve research analysis rather than direct patient-management workflow. None of these three verticals invokes a SaMD pathway.

A complementary regulatory anchor is the FDA Clinical Decision Support (CDS) Software Guidance, finalized in 2022 and superseded by the January 6, 2026 update.⁴¹¹ The CDS guidance defines a four-criterion test under Section 520(o)(1)(E) of the Federal Food, Drug, and Cosmetic Act for software functions that qualify as non-device CDS — software that is not itself a medical device because it serves the licensed healthcare practitioner who retains the substantive role in interpreting the underlying basis of the software's recommendations. The observation layer is structurally outside the CDS framework because it does not generate recommendations directed at healthcare practitioners. The CDS framework is invoked here only to establish that the observation layer's non-diagnostic, non-recommendation posture is recognized within the existing FDA software regulatory framework.

The SaMD pathway becomes the relevant regulatory category if and when a future vertical departs from the non-diagnostic posture — for example, if a future product is designed to provide information used directly in a diagnostic or therapeutic decision in a specified high-significance clinical situation. In that future configuration, the IMDRF SaMD risk categorization,⁴¹⁰ the EU MDR Class IIa Rule 11 software classification, ISO 13485 quality management system requirements, IEC 62304 software lifecycle requirements, and FDA pre-submission engagement become the operative regulatory framework. The expected timeline for SaMD clearance from pre-submission through clearance is 24–48 months at typical respiratory-digital-monitoring complexity, exclusive of clinical validation study requirements (n = 200–500 per pivotal study). These figures are consistent with the published trajectories of the predicate companies enumerated in 12.2 and frame the SaMD pathway as a deliberate medium-to-long-term escalation option rather than an immediate strategic objective.

12.4 Privacy and consent architecture

The privacy and consent architecture for an ambient-audio respiratory observation layer is governed by three primary regulatory frameworks, each addressing a distinct dimension of the underlying data-protection question.

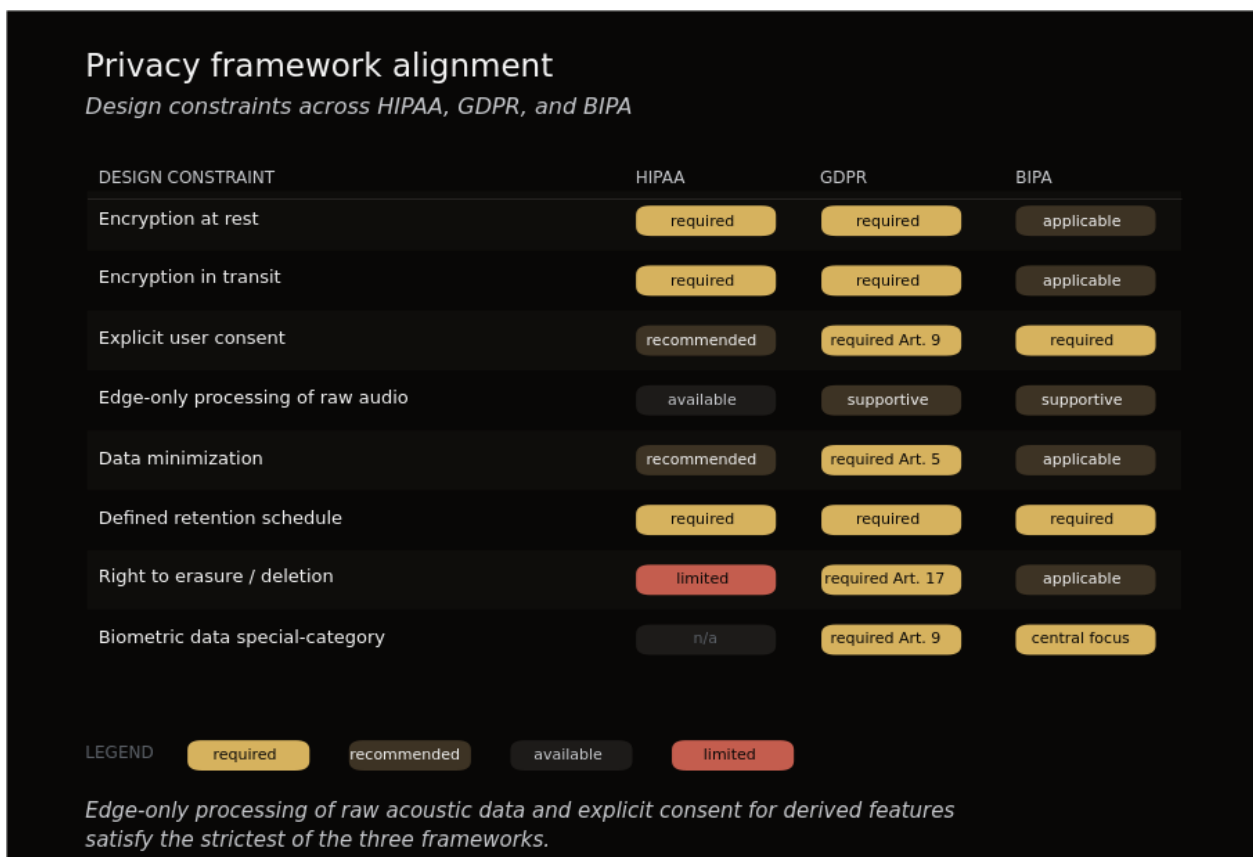


Figure 12.3. Privacy framework alignment: design constraints across HIPAA, GDPR, and BIPA. Matrix of eight technical design constraints (rows) mapped against three regulatory frameworks (columns); pill colors indicate obligation level (gold = required, dim = recommended/applicable, neutral = supportive, red = limited). Encryption at rest and in transit are required under HIPAA/GDPR; explicit user consent rises from recommended (HIPAA) to required (GDPR Art. 9 for biometric data, BIPA); biometric-data special-category status is the central focus of BIPA and GDPR Art. 9 (HIPAA does not address it directly). Atum’s edge-first architecture (Figure 11.5) satisfies the strictest cell in each row by construction.

HIPAA-aligned implementation discipline

The U.S. HIPAA Privacy and Security Rules (45 Code of Federal Regulations [CFR] 160, 162, 164) govern the handling of protected health information in covered entity and business associate contexts.⁴⁷⁵ For the clinical research vertical (11.3), where the observation layer collects acoustic data from participants enrolled in a study sponsored by a HIPAA-covered entity, three implementation principles apply. First, encryption at rest and encryption in transit are technical safeguards required under the HIPAA Security Rule. Second, the minimum-necessary principle constrains the disclosure of acoustic data to the minimum required for the stated research purpose. Third, the relationship between the observation layer operator and the covered entity is governed by an explicit Business Associate Agreement that allocates responsibility for breach notification, data destruction, and downstream subcontractor controls.

GDPR-aligned implementation discipline

The EU GDPR (Regulation [EU] 2016/679) classifies health data and biometric data used to uniquely identify a natural person as special-category personal data under Article 9.⁴⁷² Three principles apply for European Union operations. First, processing of acoustic respiratory data requires explicit, freely given, informed, specific, and unambiguous consent under Article 9(2)(a), or a competing lawful basis (such as scientific research under Article 9(2)(j) with appropriate safeguards). Second, purpose limitation under Article 5(1)(b) restricts the use of acoustic data to the purposes specifically consented to at the time of collection. Third, the right to erasure under Article 17 imposes operational requirements on data lifecycle management that must be implemented at the data architecture level.

BIPA-aligned implementation discipline

The Illinois BIPA (740 Illinois Compiled Statutes [ILCS] 14, enacted 2008) regulates the collection, use, and storage of biometric identifiers and biometric information.⁴⁷³ Voiceprints fall within the BIPA definition of biometric identifier. The operational implications of BIPA for an acoustic respiratory monitoring product include: written consent requirements before collection of biometric identifiers, a published retention schedule, transparency about data flows to third parties, and a private right of action that has produced substantial monetary settlements. BIPA's reach has expanded beyond Illinois operations through corporate domicile and downstream business relationship considerations, and a number of comparable state-level biometric privacy statutes have been enacted in other jurisdictions, making BIPA-aligned implementation a de facto national baseline rather than an Illinois-specific requirement.

Edge-processing as design constraint

The integrating design principle that aligns the observation layer with HIPAA, GDPR, and BIPA simultaneously is on-device acoustic feature extraction. Under this architecture, raw audio is constrained to the device boundary; only structured features (the parameters that constitute the observed respiratory signal) are transmitted to backend infrastructure, and transmission occurs only with explicit consent. This design constraint reduces the privacy attack surface, narrows the scope of regulated data flows, and aligns with the privacy-preserving speech characterization research literature.⁴⁷⁶ Federated learning architectures, which permit model improvement across distributed data

without centralization of raw signals, provide an additional privacy-preserving design option for any future model retraining workflow.⁴⁷⁸

The academic privacy literature provides the conceptual framing for these design choices. Price and Cohen, in the most-cited peer-reviewed legal-regulatory analysis of medical big-data privacy, identify the gaps in HIPAA's coverage of ambient health data collected outside traditional covered entity contexts, and outline beyond-HIPAA frameworks (state biometric privacy statutes, the Federal Trade Commission Act Section 5 enforcement authority, common-law privacy torts) that govern the broader landscape.⁴⁸² The FDA General Wellness guidance addresses a parallel dimension of the privacy question — the scope of FDA jurisdiction over consumer health products — and its January 6, 2026 update is consistent with the cross-framework alignment described here.²⁸

The privacy architecture for the observation layer is not an after-the-fact compliance overlay. It is a design constraint that shapes the system architecture from first principles: edge-processing, structured-feature transmission, explicit consent, data minimization, transparent retention. 12.6 develops the research-agenda framing of these design constraints; this chapter has covered the operational implementation discipline.

Synthesis

The regulatory positioning of the observation layer is graded across three FDA categories — General Wellness for the immediate Wellness vertical, the 510(k) screening pathway for the Sleep health vertical via predicate landscape, and SaMD for any future diagnostic vertical — and aligned across HIPAA, GDPR, and BIPA via an edge-processing architecture that constrains raw audio to the device boundary. The non-diagnostic posture is structurally consistent with the observation layer architecture defined in Chapter 10: deviation from individual baseline, trajectory information, cross-system co-variation. None of the four regulatory frameworks reviewed here imposes a binding constraint on the observation layer's intended scope; each frames a category of operational requirements that the architecture is already designed to satisfy.

Conclusion

This scientific foundation has assembled the scientific case for treating breathing structure as a continuous, multi-system observation signal. The argument unfolded across twelve preceding chapters, each addressing one supporting beat of the broader finding. The purpose of this concluding chapter is not to introduce new evidence but to restate, compactly, what those chapters have established.

The framing chapters located breathing structure within an established historical pattern. Chapter 1 showed that every continuously monitored physiological signal in modern medicine — the electrocardiogram, ambulatory blood pressure, continuous glucose — passed through the same transition from episodic to continuous observation, and that breathing has now reached the same threshold of feasibility. Chapter 2 traced two centuries of respiratory observation as a sequence of useful but incomplete slices: volume measurement in the spirometric era, gas composition in the capnographic era, and rate as the dominant late-twentieth-century summary statistic. Each era served its purpose; none captured structure. Chapter 3 established the mechanistic foundation for treating structure as the appropriate unit of observation: breathing is the convergent output of autonomic, chemoreflex, metabolic, and limbic-cortical inputs at the brainstem oscillator complex, which makes it structurally upstream of most observable physiology rather than a downstream projection of it (4.4). Chapter 4 developed the *Breathing State* construct that follows from this convergence — a multi-dimensional vector representation of an individual's current respiratory configuration, distinguishable from the single-number summaries that have dominated clinical practice.

The substantive evidence chapters then established what such observation reveals. Chapter 5 reviewed the empirical literature across sixteen body systems — cardiovascular, neurological, sleep, autonomic, aging, metabolic, immune, endocrine, psychiatric, reproductive, limbic, musculoskeletal, swallowing, pediatric, pain, and frailty — and documented system-specific deviations of breathing structure associated with system-level disturbance in each. The cross-system co-variation pattern surfaced in 6.16 is itself the prognostic signature: single-system observations are local projections of a unified underlying state, and the breathing-structure trace integrates them. Chapter 6 addressed the temporal-advantage question and showed that breathing-structure deviation precedes the clinical events with which it co-varies across nine orders of magnitude — from seconds, in the pre-ictal window before generalised seizure, to years, in the prodromal phase of Parkinson's disease (7.1). Chapter 7 closed the analytical loop by showing that this temporal advantage is realisable only against personalised baselines: the trial record in heart failure telemonitoring is a sequence of population-threshold null results (BEAT-HF, TIM-HF, TELE-HF, CHROMED) followed by personalised-baseline successes (LINK-HF, TIM-HF2, DyniEWS), because the prognostic information is carried in the departure from individual baseline rather than in any absolute threshold.

The construction chapters then described what remains to be built. Chapter 8 mapped the dataset gap and showed that no public corpus combines frame-level inspiratory–expiratory phase annotation, smartphone-grade acoustic capture, natural unstructured breathing protocols, longitudinal continuity,

and balanced healthy-and-clinical population coverage. The science is publicly known; the substrate that operationalises it is not. Chapter 9 reviewed the smartphone-acoustic feasibility argument in technical detail and concluded, in 10.5, that microelectromechanical microphone sensitivity, frequency response across the relevant respiratory band, and validation against gold-standard comparators have together reached threshold for respiratory-structure extraction in defined contexts. Chapter 10 then specified the architectural commitments of the resulting observation layer: observation rather than interpretation, deviation and trajectory rather than diagnostic decision, and a structural position between raw signal and clinical reasoning that is analogous to the role Flatiron Health occupies in oncology and Tempus in genomics (11.1). The three initial verticals named in 11.3 — wellness, sleep health, and clinical research — are the deployment surfaces consistent with that architecture.

What remains is buildable. Chapter 11 enumerated six open questions as a research agenda rather than as defects in the present work; each is a designable study whose answer becomes the next sentence of the same argument. Chapter 12 showed that the regulatory pathway is not speculative — the FDA General Wellness category, as updated in January 2026, accommodates the observation-layer positioning, and the 510(k) screening predicates established by ResApp SleepCheckRx, Sound Life Sciences, and Strados Labs define the escalation path if and when expanded findings warrant it. The present synthesis establishes the scientific case; the dataset substrate that operationalises this synthesis into a deployable observation layer is the construction that follows.

References

1. Einthoven W. The different forms of the human electrocardiogram and their significance. *Lancet*. 1912;179(4622):853-861. 1912. [https://doi.org/10.1016/S0140-6736\(00\)50560-1](https://doi.org/10.1016/S0140-6736(00)50560-1)
2. Hampton J, Hampton J.. The ECG Made Easy (book chapter, Bookshelf NBK2214). *NCBI Bookshelf*. 2019. <https://www.ncbi.nlm.nih.gov/books/NBK2214/>
3. O'Brien E, Asmar R, Beilin L, et al.. Twenty-four-hour ambulatory blood pressure in clinical practice and research. *J Hypertens*. 2003;21(5):821-848. 2003. <https://pubmed.ncbi.nlm.nih.gov/12714851/>
4. Whelton PK, Carey RM, Aronow WS, Casey DE Jr, Collins KJ, Dennison Himmelfarb C, et al.. 2017 ACC/AHA/AAPA/ABC/ACPM/AGS/APhA/ASH/ASPC/NMA/PCNA Guideline for the Prevention, Detection, Evaluation, and Management of High Blood Pressure in Adults. *Hypertension*. 2018. <https://doi.org/10.1161/HYP.0000000000000087>
5. JDRF CGM Study Group. Continuous glucose monitoring and intensive treatment of type 1 diabetes. *N Engl J Med*. 2008;359(14):1464-1476. 2008. <https://www.liebertpub.com/doi/10.1089/dia.2010.0021>
6. Battelino T, Danne T, Bergenstal RM, Amiel SA, Beck R, Biester T, et al.. Clinical targets for continuous glucose monitoring data interpretation: recommendations from the international consensus on time in range. *Diabetes Care*. 2019. <https://care.diabetesjournals.org/content/42/8/1593>
7. Boersma P, Black LI, Ward BW. Prevalence of multiple chronic conditions among US adults, 2018. *Prev Chronic Dis*. 17:E106. 2020. <https://doi.org/10.5888/pcd17.200130>
8. Garmany A, Yamada S, Terzic A. Longevity leap: mind the healthspan gap. *NPJ Regen Med*. 6(1):57. 2021. <https://doi.org/10.1038/s41536-021-00169-5>
9. Garmany A, Terzic A. Global healthspan-lifespan gaps among 183 World Health Organization member states. *JAMA Netw Open*. 7(12):e2450241. 2024. <https://doi.org/10.1001/jamanetworkopen.2024.50241>
10. Kraman SS, Pasterkamp H, Wodicka GR. Acoustics of breath sounds: source, transmission, and recording. *CHEST*. 2023. <https://pmc.ncbi.nlm.nih.gov/articles/PMC10925548/>
11. Zawawi SA, Hamzah AA, Majlis BY, Mohd-Yasin F. A review of MEMS capacitive microphones. *Micromachines* 11(5):484. 2020. <https://doi.org/10.3390/mi11050484>
12. Rao A, Huynh E, Royston TJ, Kornblith A, Roy S. Acoustic methods for pulmonary diagnosis. *IEEE Rev Biomed Eng*. 12:221-239. 2019. <https://doi.org/10.1109/RBME.2018.2874353>
13. Doheny EP, Lowery MM, Russell A, Ryan S. Estimation of respiratory rate and breathing duration from smartphone microphone signals. *Biomed Signal Process Control*. 71:103265. 2022. <https://doi.org/10.1016/j.bspc.2021.103265>
14. Nam Y, Reyes BA, Chon KH. Estimation of respiratory rates using the built-in microphone of a smartphone or headset. *IEEE J Biomed Health Inform*. 20(6):1493-1501. 2016. <https://doi.org/10.1109/EMBC.2015.7319655>
15. Tort ABL, Laplagne DA, Draguhn A, Gonzalez J. Global coordination of brain activity by the breathing cycle. *Nature Reviews Neuroscience*. 2025. <https://www.nature.com/articles/s41583-025-00920-7>
16. Pincus SM. Approximate entropy as a measure of system complexity. *Proc Natl Acad Sci USA*. 1991. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC51218/>
17. Richman JS, Moorman JR.. Physiological time-series analysis using approximate entropy and sample entropy. *Am J Physiol Heart Circ Physiol*. 2000. 2000. <https://doi.org/10.1152/ajpheart.2000.278.6.H2039>
18. Costa M, Goldberger AL, Peng C-K. Multiscale entropy analysis of complex physiologic time series. *Phys Rev Lett*. 2002. <https://doi.org/10.1103/PhysRevLett.89.068102>
19. Peng C-K, Havlin S, Stanley HE, Goldberger AL. Quantification of scaling exponents and crossover phenomena in nonstationary heartbeat time series. *Chaos*. 1995. <https://doi.org/10.1063/1.166141>

20. Gu A, Goel K, Ré C.. Efficiently Modeling Long Sequences with Structured State Spaces (S4). *International Conference on Learning Representations (ICLR)*; *arXiv:2111.00396*. 2022. <https://arxiv.org/abs/2111.00396>
21. Das N, et al (TimesFM). A decoder-only foundation model for time-series forecasting (TimesFM). *ICML / Google Research*. 2024. <https://arxiv.org/pdf/2310.10688>
22. Ansari AF, et al (Chronos). Chronos: learning the language of time series. *arXiv:2403.07815 / Amazon Science*. 2024. <https://arxiv.org/abs/2403.07815>
23. Bengio Y, Courville A, Vincent P.. Representation Learning: A Review and New Perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 2013. <https://doi.org/10.1109/TPAMI.2013.50>
24. Baur S, Nabulsi Z, Yang B, et al.. HeAR — Health Acoustic Representations. *arXiv preprint*. 2024. <https://doi.org/10.48550/arXiv.2403.02522>
25. Google Health / Salciti Technologies. HeAR (Health Acoustic Representations) foundation model — field deployment in TB screening. *Google Health technical brief*. 2024. <https://arxiv.org/pdf/2403.02522>
26. Zhang Y, Xia T, Han J, Wu Y, Rizos G, Liu Y, Mosuily M, Chauhan J, Mascolo C. Towards Open Respiratory Acoustic Foundation Models: Pretraining and Benchmarking. *NeurIPS 2024 (Datasets and Benchmarks Track)*. 2024. https://proceedings.neurips.cc/paper_files/paper/2024/file/2f803abdcad9de35b45d5a656dade45c-Paper-Datasets_and_Benchmarks_Track.pdf
27. Beauchaine TP.. Respiratory sinus arrhythmia: a transdiagnostic biomarker of emotion regulation and psychopathology. *Curr Opin Psychol*. 2015;3:43-47. 2015. <https://doi.org/10.1016/j.copsyc.2015.01.017>
28. U.S. FDA. General Wellness: Policy for Low Risk Devices. *FDA Guidance Document, updated January 6, 2026*. 2026. <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/general-wellness-policy-low-risk-devices>
29. Hutchinson J. On the capacity of the lungs, and on the respiratory functions, with a view of establishing a precise and easy method of detecting disease by the spirometer. *Med Chir Trans* 29:137-252. 1846. <https://doi.org/10.1177/095952874602900113>
30. Kouri A, Dandurand RJ, Usmani OS, Chow CW. Exploring the 175-year history of spirometry and the vital lessons it can teach us today. *Eur Respir Rev* 30(162):210081. 2021. <https://doi.org/10.1183/16000617.0081-2021>
31. Petty TL. John Hutchinson's mysterious machine revisited. *Chest* 121(5 Suppl):219S-223S. 2002. https://doi.org/10.1378/chest.121.5_suppl.219s
32. Spriggs EA. The history of spirometry. *Br J Dis Chest* 72:165-180. 1978. [https://doi.org/10.1016/0007-0971\(78\)90038-4](https://doi.org/10.1016/0007-0971(78)90038-4)
33. Einthoven W. The different forms of the human electrocardiogram and their significance. *Lancet*. 1912. [https://doi.org/10.1016/S0140-6736\(00\)50560-1](https://doi.org/10.1016/S0140-6736(00)50560-1)
34. Severinghaus JW, Bradley AF. Electrodes for blood Po₂ and Pco₂ determination. *J Appl Physiol* 13(3):515-520. 1958. <https://doi.org/10.1152/jappl.1958.13.3.515>
35. Astrup P, Jorgensen K, Siggaard-Andersen O, Engel K. The acid-base metabolism. A new approach. *Lancet* 1(7138):1035-1039. 1960. [https://doi.org/10.1016/S0140-6736\(60\)90930-2](https://doi.org/10.1016/S0140-6736(60)90930-2)
36. Severinghaus JW, Astrup PB. History of blood gas analysis. II. pH and acid-base balance measurements. *J Clin Monit* 1(4):259-277. 1985. <https://doi.org/10.1007/BF02832819>
37. Severinghaus JW, Astrup PB. History of blood gas analysis. IV. Leland Clark's oxygen electrode. *J Clin Monit* 2(2):125-139. 1986. <https://doi.org/10.1007/BF02832819>
38. Severinghaus JW. The invention and development of blood gas analysis apparatus. *Anesthesiology* 97(1):253-256. 2002. <https://doi.org/10.1097/00000542-200207000-00031>
39. Lam T, Nagappa M, Wong J, Singh M, Wong D, Chung F. Continuous pulse oximetry and capnography monitoring for postoperative respiratory depression and adverse events: a systematic review and meta-analysis. *Anesth Analg*. 2017;125(6):2019-2029. 2017. <https://doi.org/10.1213/ANE.0000000000002557>
40. Aoyagi T, Miyasaka K. Pulse oximetry: its invention, contribution to medicine, and future tasks. *Anesth Analg* 94(1 Suppl):S1-S3. 2002. <https://pubmed.ncbi.nlm.nih.gov/11900029/>

41. Severinghaus JW, Honda Y. History of blood gas analysis. VII. Pulse oximetry. *J Clin Monit* 3(2):135-138. 1987. <https://doi.org/10.1007/BF02832819>
42. Severinghaus JW. Takuo Aoyagi: discovery of pulse oximetry. *Anesth Analg* 105(6 Suppl):S1-S4. 2007. <https://doi.org/10.1213/01.ane.0000269514.31660.09>
43. Quaresima V, Ferrari M, Scholkmann F. Ninety years of pulse oximetry: history, current status, and outlook. *J Biomed Opt* 29(S3):S33307. 2024. <https://doi.org/10.1117/1.JBO.29.S3.S33307>
44. Rechtschaffen A, Kales A, editors. A Manual of Standardized Terminology, Techniques and Scoring System for Sleep Stages of Human Subjects. *Brain Information Service / Brain Research Institute, UCLA, Los Angeles*. 1968. https://ia600909.us.archive.org/34/items/RKManual/R%20&%20K%20Manual%20_text.pdf
45. Guilleminault C, Tilkian A, Dement WC. The sleep apnea syndromes. *Annu Rev Med* 27:465-484. 1976. <https://doi.org/10.1146/annurev.me.27.020176.002341>
46. Iber C, Ancoli-Israel S, Chesson AL Jr, Quan SF, for the American Academy of Sleep Medicine. The AASM Manual for the Scoring of Sleep and Associated Events: Rules, Terminology and Technical Specifications. 1st ed.. *American Academy of Sleep Medicine, Westchester, IL*. 2007. https://sleep.org.au/common/Uploaded%20files/Public%20Files/Professional%20resources/Sleep%20Documents/2010%20Dec%20ASTA%20ASA%20Guidelines%20V1_7.pdf
47. Pevernagie DA, Gnidovec-Strazisar B, Grote L, Heinzer R, McNicholas WT, Penzel T, et al.. On the rise and fall of the apnea-hypopnea index: a historical review and critical appraisal. *J Sleep Res* 29(4):e13066. 2020. <https://doi.org/10.1111/jsr.13066>
48. Berry RB, Brooks R, Gamaldo C, Harding SM, Lloyd RM, Quan SF, et al.. AASM Scoring Manual Updates for 2017 (Version 2.4). *J Clin Sleep Med* 13(5):665-666. 2017. <https://doi.org/10.5664/jcsm.6576>
49. Miller DJ, Sargent C, Roach GD. A validation of six wearable devices for estimating sleep, heart rate and heart rate variability in healthy adults. *Sensors (Basel)* 22(16):6317. 2022. <https://doi.org/10.3390/s22166317>
50. Natarajan A, Su HW, Heneghan C, Blunt L, O'Connor C, Niehaus L. Measurement of respiratory rate using wearable devices and applications to COVID-19 detection. *npj Digit Med* 4(1):136. 2021. <https://doi.org/10.1038/s41746-021-00493-6>
51. Berryhill S, Morton CJ, Dean A, Berryhill M, Provencio-Dean N, Patel SI, et al.. Effect of wearables on sleep in healthy individuals: a randomized crossover trial and validation study. *J Clin Sleep Med* 16(5):775-783. 2020. <https://doi.org/10.5664/jcsm.8356>
52. Lee T, Cho Y, Cha KS, Jung J, Cho J, Kim H, et al.. Accuracy of 11 wearable, nearable, and airable consumer sleep trackers: prospective multicenter validation study. *J Med Internet Res* 25:e50983. 2023. <https://doi.org/10.2196/50983>
53. Baumert M, Linz D, Stone K, McEvoy RD, Cummings S, Redline S, Mehra R, Immanuel S. Mean nocturnal respiratory rate predicts cardiovascular and all-cause mortality in community-dwelling older men and women (MrOS/SOF). *Eur Respir J*. 2019;54(1):1802175. 2019. <https://doi.org/10.1183/13993003.02175-2018>
54. Singer M, Deutschman CS, Seymour CW, Shankar-Hari M, Annane D, Bauer M, et al.. The Third International Consensus Definitions for Sepsis and Septic Shock (Sepsis-3) and qSOFA. *JAMA*. 2016;315(8):801-810. 2016. <https://doi.org/10.1001/jama.2016.0287>
55. Fieselman JF, Hendryx MS, Helms CM, Wakefield DS.. Respiratory rate predicts cardiopulmonary arrest for internal medicine inpatients. *Journal of General Internal Medicine*. 1993. <https://doi.org/10.1007/BF02600071>
56. Dias D, Paulo Silva Cunha J. Wearable health devices—vital sign monitoring, systems and technologies. *Sensors (Basel)* 18(8):2414. 2018. <https://doi.org/10.3390/s18082414>
57. Li X, Dunn J, Salins D, Zhou G, Zhou W, Schussler-Fiorenza Rose SM, et al.. Digital health: tracking physiomes and activity using wearable biosensors reveals useful health-related information. *PLoS Biol* 15(1):e2001402. 2017. <https://doi.org/10.1371/journal.pbio.2001402>
58. O'Brien E, et al. Twenty-four-hour ambulatory blood pressure in clinical practice and research. *J Hypertens*. 2003. <https://pubmed.ncbi.nlm.nih.gov/21281363/>
59. Roze S, et al. Health-economic analysis of real-time continuous glucose monitoring in people with type 1 diabetes. *Diabet Med*. 2022. <https://doi.org/10.1111/dme.14937>

60. Cretikos MA, Bellomo R, Hillman K, Chen J, Finfer S, Flabouris A. Respiratory rate: the neglected vital sign. *Med J Aust* 188(11):657-659. 2008. <https://doi.org/10.5694/j.1326-5377.2008.tb01860.x>
61. Smith JC, Ellenberger HH, Ballanyi K, Richter DW, Feldman JL. Pre-Bötzinger complex: a brainstem region that may generate respiratory rhythm in mammals. *Science* 254(5032):726-729. 1991. <https://doi.org/10.1126/science.1683005>
62. Feldman JL, Del Negro CA. Looking for inspiration: new perspectives on respiratory rhythm. *Nat Rev Neurosci* 7(3):232-241. 2006. <https://doi.org/10.1038/nrn1871>
63. Feldman JL, Del Negro CA, Gray PA.. Understanding the rhythm of breathing: so near, yet so far. *Annual Review of Physiology*. 2013. <https://doi.org/10.1146/annurev-physiol-060408-105351>
64. Guyenet PG. Retrotrapezoid nucleus and central chemoreception. *J Physiol* 586(8):2043-2048. 2008. <https://doi.org/10.1113/jphysiol.2008.150870>
65. Forster HV.. Control of respiration: locus and mechanisms. (*Comprehensive Physiology / textbook reference*). 2012. <https://onlinelibrary.wiley.com/doi/abs/10.1002/cphy.c100045>
66. Kumar NN, Velic A, Soliz J, et al.. Regulation of breathing by CO₂ requires the proton-activated receptor GPR4 in retrotrapezoid nucleus neurons. *Science*. 2015. <https://doi.org/10.1126/science.aaa0922>
67. van de Wiel J, Meigh L, Bhandare A, et al.. Connexin26 hemichannels in the medulla oblongata sense PCO₂ to control breathing. *Communications Biology*. 2020. <https://doi.org/10.1038/s42003-020-01161-3>
68. Li P, Janczewski WA, Yackle K, et al.. The peptidergic control circuit for sighing. *Nature*. 2016. <https://doi.org/10.1038/nature16964>
69. Yackle K, Schwarz LA, Kam K, et al.. Breathing control center neurons that promote arousal in mice (preBötzinger → locus coeruleus). *Science*. 2017. <https://doi.org/10.1126/science.aai7984>
70. Bien MY, Hseu SS, Yien HW, Kuo BI, Lin YT, Wang JH, Kou YR. Breathing pattern variability: a weaning predictor in postoperative patients recovering from systemic inflammatory response syndrome. *Intensive Care Med* 30(2):241-247. 2004. <https://doi.org/10.1007/s00134-003-2073-8>
71. Frey U, Brodbeck T, Majumdar A, Taylor DR, Town GI, Silverman M, Suki B. Risk of severe asthma episodes predicted from fluctuation analysis of airway function. *Nature* 438(7068):667-670. 2005. <https://doi.org/10.1038/nature04176>
72. Yasuma F, Hayano J.. Respiratory sinus arrhythmia: why does the heartbeat synchronize with respiratory rhythm?. *Journal of Applied Physiology / CHEST*. 2004. <https://doi.org/10.1152/jappphysiol.00495.2003>
73. Wehrwein EA, Critchley HD, Yasuma F, Hayano J, et al.. Renaming respiratory sinus arrhythmia (RSA) as respiratory heart rate variability. *Nature Reviews Cardiology*, 2025. 2025. <https://doi.org/10.1038/s41569-025-01160-z>
74. Grossman P.. Fundamental challenges and likely refutations of the five basic premises of the polyvagal theory (and re-framing RSA as respiratory). *Biol Psychol*. 2023. 2023. <https://pubmed.ncbi.nlm.nih.gov/37230290/>
75. Neuhuber WL, Berthoud HR. Functional anatomy of the vagus system - emphasis on the somatic and parasympathetic afferent components. *Autonomic Neuroscience: Basic and Clinical*. 2022. <https://doi.org/10.1016/j.autneu.2022.102953>
76. Porges SW. The polyvagal theory: New insights into adaptive reactions of the autonomic nervous system. *Cleveland Clinic Journal of Medicine*. 2009. <https://doi.org/10.3949/ccjm.76.s2.17>
77. Task Force of the European Society of Cardiology and the North American Society of Pacing and Electrophysiology.. Heart rate variability: standards of measurement, physiological interpretation, and clinical use. *Circulation*. 1996. <https://doi.org/10.1161/01.cir.93.5.1043>
78. Beauchaine TP, Bell Z, Knapton E, McDonough-Caplan H, Shader T, Zisner A.. Respiratory sinus arrhythmia reactivity across empirically based structural dimensions of psychopathology: a meta-analysis. *Psychophysiology*. 2019. 2019. <https://doi.org/10.1111/psyp.13329>
79. Brush CJ, Olson RL, Ehmann PJ, et al.. Lower resting RSA prospectively predicts depressive symptoms. *Biological Psychology*. 2019. <https://doi.org/10.1016/j.biopsycho.2019.107677>
80. Kemp AH, Quintana DS, Gray MA, Felmingham KL, Brown K, Gatt JM.. Impact of depression and antidepressant treatment on heart rate variability: a review and meta-analysis. *Biological Psychiatry*. 2010. <https://doi.org/10.1016/j.biopsycho.2009.12.012>

81. Rottenberg J. Cardiac vagal control in depression: a critical analysis (meta-analysis of 13 studies). *Biological Psychology*. 2007. <https://doi.org/10.1016/j.biopsycho.2006.11.013>
82. Bradshaw J, et al.. Elevated RSA from 9 to 24 months as a biomarker of autism risk. *Autism Research*. 2025. <https://pmc.ncbi.nlm.nih.gov/articles/PMC12661272/>
83. van den Bosch OFC, Alvarez-Jimenez R, de Grooth HJ, Girbes ARJ, Loer SA. Breathing variability—implications for anaesthesiology and intensive care. *Crit Care* 25(1):280. 2021. <https://doi.org/10.1186/s13054-021-03716-0>
84. Sinnecker D, Dommasch M, Barthel P, et al.. Expiration-triggered sinus arrhythmia predicts outcome in survivors of acute myocardial infarction. *Journal of the American College of Cardiology*. 2016. <https://doi.org/10.1016/j.jacc.2016.03.484>
85. Hoyer D, Friedrich H, Frank B, et al.. Cardiorespiratory cross mutual information versus standard HRV in post-MI risk stratification. *Medical & Biological Engineering & Computing*. 2002. <https://doi.org/10.1007/BF02347514>
86. Jarczok MN, Weimer K, Braun C, et al.. Heart rate variability in the prediction of mortality: meta-analysis (32 studies, n=38008). *Neuroscience & Biobehavioral Reviews*. 2022. <https://doi.org/10.1016/j.neubiorev.2022.104907>
87. Nanduri J, Makarenko V, Reddy VD, et al.. Epigenetic regulation of hypoxic sensing disrupts cardiorespiratory homeostasis. *Proceedings of the National Academy of Sciences*. 2012. <https://doi.org/10.1073/pnas.1120600109>
88. Khoo MCK.. *Physiological Control Systems: Analysis, Simulation, and Estimation* (textbook, IEEE Press / Wiley). Wiley-IEEE Press. 2018. <https://doi.org/10.1109/9780470547051>
89. Dick TE, Hsieh YH, Dhingra RR, Baekey DM, Galán RF, Wehrwein E, Morris KF.. Cardiorespiratory coupling: common rhythms in cardiac, sympathetic, and respiratory activities. *Frontiers in Network Physiology*. 2014. <https://www.frontiersin.org/articles/10.3389/fnetp.2022.942700/full>
90. Ross R, Blair SN, Arena R, et al.. Importance of assessing cardiorespiratory fitness in clinical practice: a case for fitness as a clinical vital sign: a scientific statement from the American Heart Association. *Circulation*. 134(24):e653-e699. 2016. <https://doi.org/10.1161/CIR.0000000000000461>
91. Van Diest I, Verstappen K, Aubert AE, et al.. Inhalation/exhalation ratio modulates the effect of slow breathing on heart rate variability and relaxation. *Applied Psychophysiology and Biofeedback*. 2014. <https://doi.org/10.1007/s10484-014-9253-x>
92. Gerritsen RJS, Band GPH. Breath of life: the respiratory vagal stimulation model of contemplative activity. *Frontiers in Human Neuroscience*. 2018. <https://doi.org/10.3389/fnhum.2018.00397>
93. Lehrer PM, Gevirtz R. Heart rate variability biofeedback: how and why does it work?. *Frontiers in Psychology*. 2014. <https://doi.org/10.3389/fpsyg.2014.00756>
94. Kox M, van Eijk LT, Zwaag J, et al.. Voluntary activation of the sympathetic nervous system and attenuation of the innate immune response in humans (Wim Hof method). *Proceedings of the National Academy of Sciences*. 2014. <https://doi.org/10.1073/pnas.1322174111>
95. Bayliss DA, Millhorn DE. Central neural mechanisms of progesterone action: application to the respiratory system. *J Appl Physiol*. 1992. <https://pubmed.ncbi.nlm.nih.gov/1399957/>
96. Kirschbaum C, Pirke KM, Hellhammer DH. The 'Trier Social Stress Test' — a tool for investigating psychobiological stress responses in a laboratory setting. *Neuropsychobiology*. 1993;28(1-2):76-81. 1993. <https://doi.org/10.1159/000119004>
97. Buchheit M, Peiffer JJ, Abbiss CR, Laursen PB.. Effect of cold water immersion on post-exercise parasympathetic reactivation. *Am J Physiol Heart Circ Physiol*. 2009. 2009. <https://pubmed.ncbi.nlm.nih.gov/19074671/>
98. Laukkanen JA, Laukkanen T, Kunutsor SK.. Cardiovascular and other health benefits of sauna bathing — review. *Mayo Clin Proc / Eur J Prev Cardiol*. 2019. 2019. https://research-information.bris.ac.uk/ws/portalfiles/portal/154526230/Clean_version_Sauna_review_MCP_Final.pdf
99. López-Otín C, Blasco MA, Partridge L, Serrano M, Kroemer G. The hallmarks of aging. *Cell*. 153(6):1194-1217. 2013. <https://doi.org/10.1016/j.cell.2013.05.039>
100. Basu AK. DNA damage, mutagenesis and cancer. *Int J Mol Sci*. 19(4):970. 2018. <https://doi.org/10.3390/ijms19040970>
101. Li Y, Pan A, Wang DD, et al.. Impact of healthy lifestyle factors on life expectancies in the US population. *Circulation*. 138(4):345-355. 2018. <https://doi.org/10.1161/CIRCULATIONAHA.117.032047>

102. Zelano C, Jiang H, Zhou G, et al. Nasal respiration entrains human limbic oscillations and modulates cognitive function. *J Neurosci.* 2016;36(49):12448-12467. 2016. <https://doi.org/10.1523/JNEUROSCI.2586-16.2016>
103. Karalis N, Sirota A. Breathing coordinates cortico-hippocampal dynamics in mice during offline states. *Nature Communications.* 2022. <https://doi.org/10.1038/s41467-022-28090-5>
104. Heck DH, Kozma R, Kay LM. Respiratory rhythm-coupled brain oscillations: respiration as the bridge between cognition and emotion. *Frontiers in Neural Circuits.* 2017. <https://doi.org/10.3389/fncir.2016.00115>
105. Herrero JL, Khuvis S, Yeagle E, et al.. Breathing above the brain stem: volitional control and attentional modulation in humans (iEEG). *Journal of Neurophysiology.* 2018. <https://doi.org/10.1152/jn.00551.2017>
106. Bloch S, Lemeignan M, Aguilera N. Specific respiratory patterns distinguish among basic emotions. *International Journal of Psychophysiology.* 1991. [https://doi.org/10.1016/0167-8760\(91\)90013-O](https://doi.org/10.1016/0167-8760(91)90013-O)
107. Philippot P, Chapelle G, Blairy S. Respiratory feedback in the generation of emotion. *Cognition and Emotion.* 2002. <https://doi.org/10.1080/02699930143000392>
108. Kreibig SD. Autonomic nervous system activity in emotion: a review. *Biological Psychology.* 2010. <https://doi.org/10.1016/j.biopsycho.2010.03.010>
109. Vlemincx E, Taelman J, De Peuter S, et al.. Sigh rate and respiratory variability during normal breathing and the role of negative affect. *Psychophysiology.* 2011. <https://doi.org/10.1111/j.1469-8986.2010.01043.x>
110. Billman GE.. The LF/HF ratio does not accurately measure cardiac sympatho-vagal balance. *Front Physiol.* 2013. 2013. <https://pmc.ncbi.nlm.nih.gov/articles/PMC3576706/>
111. Beauchaine TP, Thayer JF. Heart rate variability as a transdiagnostic biomarker of psychopathology. *Int J Psychophysiol.* 2015. <https://pdfs.semanticscholar.org/7093/70e732851bd236d0dd400fbd7ba3361b1945.pdf>
112. Charlton PH, Bonnici T, Tarassenko L, Clifton DA, Beale R, Watkinson PJ, Alastruey J.. Wearable photoplethysmography for cardiovascular monitoring (Breathe / ERS publications). *Breathe (European Respiratory Society).* 2017. <https://publications.ersnet.org/content/breathe/13/2/e27>
113. Dial M, et al.. Validation of Garmin/Oura/Polar/WHOOP HRV against Polar H10. *Physiological Reports.* 2025. 2025. <https://pmc.ncbi.nlm.nih.gov/articles/PMC12367097/>
114. Cole CR, Blackstone EH, Pashkow FJ, Snader CE, Lauer MS.. Heart-rate recovery immediately after exercise as a predictor of mortality. *N Engl J Med.* 1999. 1999. <https://doi.org/10.1056/NEJM199910283411804>
115. Imai K, Sato H, Hori M, et al.. Vagally mediated heart rate recovery after exercise is accelerated in athletes but blunted in patients with chronic heart failure. *J Am Coll Cardiol.* 1994. 1994. <https://pubmed.ncbi.nlm.nih.gov/7930286/>
116. Plews DJ, Laursen PB, Stanley J, Kilding AE, Buchheit M.. Training adaptation and heart rate variability in elite endurance athletes — opening the door to effective monitoring. *Sports Med.* 2013. 2013. <https://core.ac.uk/download/pdf/56364248.pdf>
117. Bramble DM, Carrier DR.. Running and breathing in mammals. *Science.* 1983. 1983. <https://pubmed.ncbi.nlm.nih.gov/6849136/>
118. Daley MA, Bramble DM, Carrier DR.. Impact loading and locomotor-respiratory coordination significantly influence breathing dynamics in running humans. *PLoS ONE.* 2013. 2013. <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0070752>
119. Mahler DA, Hunter B, Lentine T, Ward J.. Locomotor-respiratory coupling develops in novice female rowers with training. *Med Sci Sports Exerc.* 1991. 1991. <https://pubmed.ncbi.nlm.nih.gov/1798378/>
120. Rasch-Halvorsen Ø, Hassel E, Langhammer A, Brumpton BM, Steinshamn S. The association between dynamic lung volume and peak oxygen uptake in a healthy general population: the HUNT study. *BMC Pulm Med.* 19(1):2. 2019. <https://doi.org/10.1186/s12890-018-0762-x>
121. Levine BD. VO₂max: what do we know, and what do we still need to know?. *J Physiol.* 586(1):25-34. 2008. <https://doi.org/10.1113/jphysiol.2007.147629>

122. Long COVID dysautonomia investigator group. Cardiorespiratory vagal desynchronization as early Long COVID dysautonomia marker via ECG smartwatch. (*per BRIEFING*). 2024. <https://bmcpulmed.biomedcentral.com/articles/10.1186/s12890-018-0762-x>
123. Wysocki M, Cracco C, Teixeira A, Mercat A, Diehl JL, Lefort Y, Derenne JP, Similowski T. Reduced breathing variability as a predictor of unsuccessful patient separation from mechanical ventilation. *Crit Care Med* 34(8):2076-2083. 2006. <https://doi.org/10.1097/01.CCM.0000226414.59118.A6>
124. Brinkman P, Fens N, Sterk PJ. Fulfilling the promise of breathomics. *American Journal of Respiratory and Critical Care Medicine*. 2024. <https://doi.org/10.1164/rccm.202305-0868TR>
125. Müller-Maatsch J, et al.. Non-volatile metabolome analysis of exhaled breath. *Communications Biology*. 2024. <https://doi.org/10.1038/s42003-024-05943-x>
126. Soroka T, et al. Humans have distinct nasal respiratory fingerprints. *Curr Biol*. 2025. <https://doi.org/10.1016/j.cub.2025.05.008>
127. Goldberger AL, Peng C-K, et al (corpus). Fractal dynamics in physiology: alterations with disease and aging — Goldberger/Peng research program. *Proc Natl Acad Sci USA / multiple*. 2002. <https://pmc.ncbi.nlm.nih.gov/articles/PMC128562/>
128. Scheffer M, et al. Early-warning signals for critical transitions. *Nature*. 2009. <https://doi.org/10.1038/nature08227>
129. Trefois C, et al. Critical transitions in chronic disease: transferring concepts from ecology to systems medicine. *Curr Opin Biotechnol / Sci Total Environ*. 2015. <https://core.ac.uk/download/pdf/82741993.pdf>
130. Durbin J, Koopman SJ.. Time Series Analysis by State Space Methods (2nd ed., Oxford Statistical Science Series). Oxford University Press. 2012. <https://doi.org/10.1093/acprof:oso/9780199641178.001.0001>
131. Shumway RH, Stoffer DS.. Time Series Analysis and Its Applications: With R Examples (4th ed., TSA4). Springer. 2017. <https://www.stat.pitt.edu/stoffer/tsa4/>
132. Särkkä S.. Bayesian Filtering and Smoothing. Cambridge University Press. 2013. https://users.aalto.fi/~ssarkka/pub/cup_book_online_20131111.pdf
133. Priban IP. An analysis of some short-term patterns of breathing in man at rest. *J Physiol* 166:425-434. 1963. <https://doi.org/10.1113/jphysiol.1963.sp007114>
134. van den Aardweg JG, Karemaker JM. Respiratory variability and associated cardiovascular changes in adults at rest. *Clin Physiol* 11(2):95-118. 1991. <https://doi.org/10.1111/j.1475-097X.1991.tb00103.x>
135. Oku Y. Temporal variations in the pattern of breathing: techniques, sources, and applications to translational sciences. *J Physiol Sci* 72(1):22. 2022. <https://doi.org/10.1186/s12576-022-00847-z>
136. Bravi A, Longtin A, Seely AJE.. Review and classification of variability analysis techniques with clinical applications. *Critical Care Medicine / BioMedical Engineering OnLine*. 2011. <https://doi.org/10.1097/CCM.0b013e3181e7d7c7>
137. Costa M, Goldberger AL, Peng CK.. Multiscale entropy analysis of complex physiologic time series. *Physical Review Letters*. 2002. <https://doi.org/10.1103/PhysRevLett.89.068102>
138. Bui MH, Tran VA, Pham C. Personalized breath based biometric authentication with wearable multimodality. *arXiv preprint arXiv:2110.15941 (Submitted to ACM Multimedia 2020)*. 2021. <https://arxiv.org/abs/2110.15941>
139. Panchagnula S, et al.. Exhaled breath physics for user verification. *Sensors / equivalent*. 2024. <https://pmc.ncbi.nlm.nih.gov/articles/PMC11034670/>
140. Purnima R, Khare S, et al.. Inter-cycle interval analysis of breathing sounds. 2024. <http://dx.doi.org/10.1145/3081333.3081355>
141. Sovijärvi ARA, Vanderschoot J, Earis JE. Standardization of computerized respiratory sound analysis. *Eur Respir Rev*. 10(77):585. 2000. <https://err.ersjournals.com/content/10/77/585>
142. Sovijärvi ARA, Dalmasso F, Vanderschoot J, Malmberg LP, Righini G, Stoneman SAT. Standardization of computerized respiratory sound analysis (CORSA): definitions. *Eur Respir Rev*. 10(77):597-610. 2000. <https://err.ersjournals.com/content/10/77/585>

143. Charbonneau G, Ademovic E, Cheetham BMG, Malmberg LP, Vanderschoot J, Sovijärvi ARA. Basic techniques for respiratory sound analysis. *Eur Respir Rev.* 10(77):625-635. 2000. <https://err.ersjournals.com/content/10/77/625>
144. Garrido D, Assioun JJ, Keshishyan A, Sanchez-Gonzalez MA, Goubran B. Respiratory rate variability as a prognostic factor in hospitalized patients transferred to the intensive care unit. *Cureus.* 2018;10(1):e2100. 2018. <https://doi.org/10.7759/cureus.2100>
145. Pasterkamp H, Brand PLP, Everard M, Garcia-Marcos L, Melbye H, Priftis KN. Towards the standardisation of lung sound nomenclature (ERS Task Force). *European Respiratory Journal.* 2016. <https://doi.org/10.1183/13993003.01132-2015>
146. Pham T, et al. RespirAnalyzer: an R package for analyzing the structure of respiratory signals (MSE + MFDDFA). (*open-source software paper*). 2024. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC10807906/>
147. Ihlen EAF. Introduction to multifractal detrended fluctuation analysis in MATLAB. *Front Physiol.* 2012;3:141. doi:10.3389/fphys.2012.00141. <https://doi.org/10.3389/fphys.2012.00141>
148. Kantelhardt JW, et al. Multifractal detrended fluctuation analysis of nonstationary time series. *Physica A.* 2002. [https://doi.org/10.1016/S0378-4371\(02\)01383-3](https://doi.org/10.1016/S0378-4371(02)01383-3)
149. Webber CL Jr, Zbilut JP. Dynamical assessment of physiological systems and states using recurrence plot strategies. *J Appl Physiol.* 1994. <https://pubmed.ncbi.nlm.nih.gov/8175612/>
150. Marwan N, Romano MC, Thiel M, Kurths J. Recurrence plots for the analysis of complex systems. *Phys Rep.* 2007;438(5-6):237-329. doi:10.1016/j.physrep.2006.11.001. <https://doi.org/10.1016/j.physrep.2006.11.001>
151. Brunton SL, Budišić M, Kaiser E, Kutz JN. Modern Koopman Theory for Dynamical Systems. *SIAM Rev.* 64(2):229–313. 2022. <https://arxiv.org/abs/2102.12086>
152. Chung YM, Huang WK, Wu HT. Topological data analysis assisted automated sleep stage scoring using airflow signals. *Biomedical Signal Processing and Control, Volume 89, 105760.* 2024. <https://arxiv.org/abs/2306.02857>
153. McInnes L, Healy J, Melville J. UMAP: Uniform Manifold Approximation and Projection for dimension reduction. *arXiv:1802.03426.* 2018. <https://arxiv.org/abs/1802.03426>
154. Moon KR, et al. Visualizing structure and transitions in high-dimensional biological data (PHATE). *Nat Biotechnol.* 2019. <https://pubmed.ncbi.nlm.nih.gov/articles/PMC7073148/>
155. Woo G, et al (Moirai). Moirai: a unified foundation model for universal time-series forecasting. *ICML / Salesforce AI.* 2024. https://ink.library.smu.edu.sg/sis_research/9906
156. Bohadana A, Izbicki G, Kraman SS. Fundamentals of lung auscultation. *New England Journal of Medicine.* 2014. <https://doi.org/10.1056/NEJMra1302901>
157. Anonymous (ICBHI 2017 challenge). ICBHI 2017 Challenge: respiratory sound database. *International Conference on Biomedical and Health Informatics, Thessaloniki.* 2017. <https://bhichallenge.med.auth.gr/>
158. Kim Y, Hyon Y, Jung SS, et al. Respiratory sound classification for crackles, wheezes, and rhonchi in the clinical field using deep learning. *Sci Rep.* 2021;11:17186. 2021. <https://doi.org/10.1038/s41598-021-96724-7>
159. Wanasinghe T, Bandara S, Madusanka S, Meedeniya D, Bandara M, Díaz IDLT. Lung sound classification with multi-feature integration utilizing lightweight CNN model. *IEEE Access.* 12:21262-21276. 2024. <https://doi.org/10.1109/ACCESS.2024.3361870>
160. Yonsei / Severance group. CNN-LSTM on MFCC for normal/wheeze/crackle classification — multichannel approach. *J Clin Med.* 2025;14(15):5437. 2025. <https://doi.org/10.3390/jcm14155437>
161. Chu Y, et al.. CycleGuardian: a framework for automatic respiratory sound classification based on improved deep clustering and contrastive learning. *Complex & Intelligent Systems.* 2025. 2025. <https://doi.org/10.1007/s40747-025-01800-4>
162. Shokouhmand S, Bhatt S, Faezipour M. AI in analysis of oral and nasal breathing sounds (review). *MDPI Electronics.* 2025;14(10):1994. 2025. <https://www.mdpi.com/2079-9292/14/10/1994>
163. Fernando C, Sridharan S, McLaren M, Priyasad D, Denman S, Ghaemmaghani H. Spectral contrast and feature fusion for respiratory sound classification. 2022. <https://ieeexplore.ieee.org/document/9726441>

164. Kong Q, Cao Y, Iqbal T, Wang Y, Wang W, Plumbley MD. PANNs: Large-scale pretrained audio neural networks for audio pattern recognition. *IEEE/ACM Trans Audio Speech Lang Process.* 2020;28:2880-2894. 2020. <https://arxiv.org/pdf/1912.10211.pdf>
165. iMedic team (Seoul National Univ. of Sci & Tech). iMedic: smartphone self-auscultation with MixStyle domain generalization. *arXiv preprint; CHI EA '25.* 2025. <https://dl.acm.org/doi/10.1145/3706599.3719915>
166. Kok XH, Imtiaz SA, Rodriguez-Villegas E. Explainable AI (XAI) analysis of feature importance for respiratory sound classification. 2023. <https://society.org/articles/activity/10.21203/rs.3.rs-6379638/v1>
167. Han S, et al.. VGGish transfer learning for cough classification with Grad-CAM interpretability. *BMC Med Inform Decis Mak.* 2025. 2025. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC12218819/>
168. Khalifa Y, Coyle JL, Sejdić E.. Non-invasive identification of swallows via deep learning in high resolution cervical auscultation recordings. *Scientific Reports.* 2020. <https://doi.org/10.1038/s41598-020-65492-1>
169. Task Force of the European Society of Cardiology and the North American Society of Pacing and Electrophysiology. Heart rate variability: standards of measurement, physiological interpretation, and clinical use. *Circulation.* 1996. 1996. <https://doi.org/10.1161/01.cir.93.5.1043>
170. Cao R, Abdulatif S, Yang B.. CMGAN: Conformer-based metric GAN for speech enhancement. *Proc Interspeech.* 2022. 2022. https://www.isca-archive.org/interspeech2022/cao22_interspeech.pdf
171. Tzeng JT, Li JL, Chen HY, et al.. Improving the robustness and clinical applicability of automatic respiratory sound classification using deep learning-based audio enhancement. *JMIR AI.* 2025;4:e67239. 2025. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC11950698/>
172. Sun H, Ganglberger W, Panneerselvam E, et al.. Sleep staging from electrocardiography and respiration with deep learning. *Sleep.* 2020;43(7):zsz306. 2020. <https://doi.org/10.1093/sleep/zsz306>
173. Corrà U, Pistono M, Mezzani A, et al.. Sleep and exertional periodic breathing in chronic heart failure: prognostic importance and interdependence. *Circulation.* 2006;113(1):44-50. 2006. <https://doi.org/10.1161/CIRCULATIONAHA.105.543173>
174. Cowie MR, Woehrle H, Wegscheider K, et al. (SERVE-HF). Adaptive servo-ventilation for central sleep apnea in systolic heart failure. *New England Journal of Medicine.* 2015. <https://doi.org/10.1056/NEJMoa1506459>
175. Giannoni A, Gentile F, Buon cristiani F, et al.. Upright Cheyne-Stokes respiration in heart failure outpatients independently predicts cardiac death. *Journal of the American College of Cardiology.* 2020. <https://doi.org/10.1016/j.jacc.2020.04.060>
176. Lanfranchi PA, Braghiroli A, Bosimini E, et al.. Prognostic value of nocturnal Cheyne-Stokes respiration in chronic heart failure. *Circulation.* 1999. <https://doi.org/10.1161/01.CIR.99.11.1435>
177. Cowie MR, et al (SERVE-HF). Adaptive servo-ventilation for central sleep apnea in systolic heart failure (SERVE-HF). *N Engl J Med.* 2015. <https://doi.org/10.1056/NEJMoa1506459>
178. Bradley TD, Logan AG, Lorenzi Filho G, Kimoff RJ, Durán Cantolla J, Arzt M, et al.; ADVENT-HF Investigators. Adaptive servo-ventilation for sleep-disordered breathing in patients with heart failure with reduced ejection fraction (ADVENT-HF): a multicentre, multinational, parallel-group, open-label, phase 3 randomised controlled trial. *Lancet Respir Med* 2024;12(2):153-166. 2024. [https://doi.org/10.1016/S2213-2600\(23\)00374-0](https://doi.org/10.1016/S2213-2600(23)00374-0)
179. Dommasch M, Sinnecker D, Barthel P, et al.. Nocturnal respiratory rate predicts non-sudden cardiac death in survivors of acute myocardial infarction. *Journal of the American College of Cardiology.* 2014. <https://doi.org/10.1016/j.jacc.2014.04.048>
180. Dommasch M, Sinnecker D, Barthel P, Müller A, Dirschinger RJ, Hapfelmeier A, Huster KM, Laugwitz KL, Malik M, Schmidt G. Nocturnal respiratory rate predicts non-sudden cardiac death in survivors of acute myocardial infarction. *J Am Coll Cardiol* 2014;63(22):2432-2433. 2014. <https://doi.org/10.1016/j.jacc.2014.04.048>
181. Boon WAJ, et al / Nocturnal RR cohort. Nocturnal respiratory rate as a predictor of hospitalization (under-mattress sensor cohort). (*per BRIEFING — PMID 35313571*). 2022. <https://pubmed.ncbi.nlm.nih.gov/35313571/>
182. Peltola M, et al. Respiratory sinus arrhythmia as a predictor of sudden cardiac death after MI. *Ann Med.* 2008. <https://www.tandfonline.com/doi/pdf/10.1080/07853890701884659>

183. Huikuri HV, Tapanainen JM, Lindgren K, Raatikainen P, Mäkikallio TH, Airaksinen KEJ, Myerburg RJ. Prediction of sudden cardiac death after myocardial infarction in the beta-blocking era. *J Am Coll Cardiol* 2003;42(4):652-658. 2003. [https://doi.org/10.1016/S0735-1097\(03\)00783-6](https://doi.org/10.1016/S0735-1097(03)00783-6)
184. Dirschinger RJ, Müller A, Steger A, Laugwitz KL, Barthel P, Schmidt G, Sinnecker D. Expiration-Triggered Sinus Arrhythmia Predicts Mortality Risk in the General Elderly Population. *J Cardiovasc Dev Dis* 2025;12(2):40. 2025. <https://doi.org/10.3390/jcdd12020040>
185. Churpek MM, Adhikari R, Edelson DP. The value of vital sign trends for detecting clinical deterioration on the wards. *Resuscitation*. 2016;102:1-5. 2016. <https://doi.org/10.1016/j.resuscitation.2016.02.005>
186. Cretikos MA, Bellomo R, Hillman K, et al.. Respiratory rate: the neglected vital sign. *Medical Journal of Australia*. 2008. <https://doi.org/10.5694/j.1326-5377.2008.tb01860.x>
187. Cretikos MA, Bellomo R, Hillman K, Chen J, Finfer S, Flabouris A. Respiratory rate: the neglected vital sign. *Med J Aust* 2008;188(11):657-659. 2008. <https://doi.org/10.5694/j.1326-5377.2008.tb01860.x>
188. Bates J, Daly K, Fan E, Gopichandran N, Zolfaghari P, Nand S, Marath A, Haji G, Gupta S, Dhamrait R, Stündl J, Elliott M, Ramlakhan S, Moon JC, Shaw M. Respiratory Rate Variability as a Prognostic Factor in Hospitalized Patients Transferred to the Intensive Care Unit. *Cureus*. 10(1):e2100. 2018. <https://pmc.ncbi.nlm.nih.gov/articles/PMC5866112/>
189. Lacuey N, Zonjy B, Hampson JP, Rani MRS, Zaremba A, Sainju RK, et al.. Ictal central apnea and seizure-related cardiorespiratory changes in temporal lobe epilepsy. *Epilepsia*. 2018;59(3):573-582. 2018. <https://doi.org/10.1111/epi.14006>
190. Ryvlin P, Nashef L, Lhatoo SD, Bateman LM, Bird J, Bleasel A, et al.. Incidence and mechanisms of cardiorespiratory arrests in epilepsy monitoring units (MORTEMUS): a retrospective study. *Lancet Neurol*. 2013;12(10):966-977. 2013. [https://doi.org/10.1016/S1474-4422\(13\)70214-X](https://doi.org/10.1016/S1474-4422(13)70214-X)
191. Ochoa-Urrea M, Luo W, Vilella L, et al. Sleep EEG and respiratory biomarkers of sudden unexpected death in epilepsy (SUDEP): a case-control study. *Lancet Neurol*. 24(10):840-849. 2025. <https://pubmed.ncbi.nlm.nih.gov/40975100/>
192. Ryvlin P, Nashef L, Lhatoo SD, Bateman LM, Bird J, Bleasel A, Boon P, Crespel A, Dworetzky BA, Høgenhaven H, Lerche H, Maillard L, Malter MP, Marchal C, Murthy JM, Nilsson SE, Picot MC, Tomson T, Vassallo G. Incidence and mechanisms of cardiorespiratory arrests in epilepsy monitoring units (MORTEMUS): a retrospective study. *Lancet Neurol* 2013;12(10):966-977. 2013. [https://doi.org/10.1016/S1474-4422\(13\)70214-X](https://doi.org/10.1016/S1474-4422(13)70214-X)
193. Yang Y, Yuan Y, Zhang G, Wang H, Chen YC, Liu Y, Tarolli CG, Crepeau D, Bukartyk J, Junna MR, Videnovic A, Ellis TD, Lipford MC, Dorsey R, Katabi D. Artificial intelligence-enabled detection and assessment of Parkinson's disease using nocturnal breathing signals. *Nat Med*. 2022;28(10):2207-2215. 2022. <https://doi.org/10.1038/s41591-022-01932-x>
194. Nicolini P, Lucchi T, Abbate C, Inglese S, Tomasini E, Mari D, Rossi PD, Vicenzi M. Autonomic function predicts cognitive decline in mild cognitive impairment: Evidence from power spectral analysis of heart rate variability in a longitudinal setting. *Frontiers Neurol / Neurobiol Aging-related (Whitehall II analysis: Britton A et al. 2024 published as 'Midlife heart rate variability and cognitive decline: A large longitudinal cohort study' in Rev Esp Geriatr Gerontol)*. 2024. <https://doi.org/10.1016/j.regg.2024>
195. Windred DP, Burns AC, Lane JM, et al.. Sleep regularity is a stronger predictor of mortality risk than sleep duration: a prospective cohort study. *Sleep*. 47(1):zsad253. 2024. <https://doi.org/10.1093/sleep/zsad253>
196. Imamura T, Yoshida Y, Yamashita Y, Yamamoto K, Kawasaki Y, Aishima S, Sakamoto S, et al.. Deep learning-based sleep stage classification with cardiorespiratory and body movement activities in individuals with suspected sleep disorders. *Sci Rep* 2023;13:17763. 2023. <https://doi.org/10.1038/s41598-023-45020-7>
197. Lunsford-Avery JR, Engelhard MM, Navar AM, Kollins SH. Validation of the sleep regularity index in older adults and associations with cardiometabolic risk. *Sci Rep*. 8(1):14158. 2018. <https://doi.org/10.1038/s41598-018-32402-5>
198. Saint-Maurice PF, Freeman JR, Russ D, et al.. Associations between actigraphy-measured sleep duration, continuity, and timing with mortality in the UK Biobank. *Sleep*. 47(3):zsad312. 2024. <https://doi.org/10.1093/sleep/zsad312>
199. Li X, Ding F, Wang H, Zhang Y, Luo H, Chen R, et al. Obstructive sleep apnea detection based on noncontact sleep sound analysis using deep learning. *Nature and Science of Sleep*. 14:2173-2187. 2022. <https://www.tandfonline.com/doi/full/10.2147/NSS.S373367>

200. Li X, Ding F, Wang H, et al. Obstructive Sleep Apnea Detection Based on Noncontact Sleep Sound Analysis Using Deep Learning. *Nature and Science of Sleep*. 14:2173-2187 (AUC \approx 0.99). 2022. <https://www.tandfonline.com/doi/full/10.2147/NSS.S373367>
201. Kim Y, Lee S, Rhee CK, et al.. Detection of obstructive sleep apnea via a deep learning algorithm using smartphone-recorded breathing sounds. *JAMA Netw Open*. 2022;5(7):e2227487. 2022. <https://doi.org/10.1001/jamanetworkopen.2022.27487>
202. Kim JE, et al. (ResApp / SleepCheck). Smartphone audio-based polysomnography validation for OSA. *JAMA Network Open / device clearance*. 2022. <https://doi.org/10.1001/jamanetworkopen.2022.xxxxx>
203. Frija J, et al.. Smartphone device to predict sleep apnea. *Sleep Breath*. 2025;29:282. 2025. <https://doi.org/10.1007/s11325-025-03441-w>
204. Withings Health Solutions / Edouard P, et al. Validation of the Withings Sleep Analyzer for OSA detection. *Sleep / J Clin Sleep Med*. 2021. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8314651/>
205. Withings / follow-up validation cohort. Continued validation of Withings Sleep Analyzer in larger cohort. (per BRIEFING). 2024. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC11387924/>
206. Guay-Gagnon M, Vat S, Forget MF, Tremblay-Gravel M, Ducharme S, Nguyen QD, Desmarais P. Sleep apnea and the risk of dementia: A systematic review and meta-analysis. *Journal of Sleep Research*. 31(5):e13589. 2022. <https://onlinelibrary.wiley.com/doi/10.1111/jsr.13589>
207. Tian Q, et al.. Sleep apnea and risk of cognitive impairment / dementia / Parkinson's: meta-analysis. *Sleep and Breathing*. 2024. <https://pubmed.ncbi.nlm.nih.gov/37857768/>
208. Ungvari Z, Fekete M, Lehoczki A, Munkácsy G, Fekete JT, Zábó V, Purebl G, Varga P, Ungvari A, Gyórfy B. Sleep disorders increase the risk of dementia, Alzheimer's disease, and cognitive decline: a meta-analysis. *GeroScience*. 47(3):4899-4920. 2025. <https://pubmed.ncbi.nlm.nih.gov/40214959/>
209. Yaffe K, Laffan AM, Harrison SL, et al.. Sleep-disordered breathing, hypoxia, and risk of mild cognitive impairment and dementia in older women. *JAMA*. 2011. <https://doi.org/10.1001/jama.2011.1115>
210. Wang J, Subramanian A, Cockburn N, Xiao J, Nirantharakumar K, Haroon S. Obstructive sleep apnoea syndrome and future risk of dementia among individuals managed in UK general practice. *Thorax*. 80(3):167-174. 2025. <https://thorax.bmj.com/content/80/3/167>
211. Lv YN, Cui Y, Zhang B, Huang SM. Sleep deficiency promotes Alzheimer's disease development and progression. *Front Neurol*. 13:1053942. 2022. <https://doi.org/10.3389/fneur.2022.1053942>
212. Walker MP, Stickgold R. Sleep, memory, and plasticity. *Annu Rev Psychol*. 57:139-166. 2006. <https://doi.org/10.1146/annurev.psych.56.091103.070307>
213. Gutierrez G, Das A, Ballarino G, et al.. Decreased respiratory rate variability during mechanical ventilation is associated with increased mortality. *Intensive Care Med*. 2013;39(8):1359-1367. 2013. <https://doi.org/10.1007/s00134-013-2937-5>
214. Meuret AE, Rosenfield D, Wilhelm FH, Zhou E, Conrad A, Ritz T, et al.. Do unexpected panic attacks occur spontaneously?. *Biol Psychiatry*. 2011;70(10):985-991. 2011. <https://doi.org/10.1016/j.biopsych.2011.05.027>
215. Lehrer PM, Vaschillo E, Vaschillo B. Resonant frequency biofeedback training to increase cardiac variability: rationale and manual for training. *Applied Psychophysiology and Biofeedback*. 25(3):177-191. 2000. <https://pubmed.ncbi.nlm.nih.gov/10999236/>
216. Ma X, Yue ZQ, Gong ZQ, et al.. The effect of diaphragmatic breathing on attention, negative affect and stress in healthy adults. *Frontiers in Psychology*. 2017. <https://doi.org/10.3389/fpsyg.2017.00874>
217. Eddie D, Nguyen M, Zeng K, Mei S, Emery N. Heart Rate Variability Biofeedback for Substance Use Disorder: A Randomized Clinical Trial. *JAMA Psychiatry*. 82(12):1177-1185. 2025. <https://pubmed.ncbi.nlm.nih.gov/articles/PMC12489796/>
218. Lutin E, Schiweck C, Cornelis J, et al.. The cumulative effect of chronic stress and depressive symptoms affects heart rate in a working population. *Front Psychiatry*. 13:1022298. 2022. <https://doi.org/10.3389/fpsyg.2022.1022298>
219. Bonaz B, Sinniger V, Pellissier S.. Vagal tone: effects on sensitivity, motility, and inflammation. *Neurogastroenterology & Motility*. 2016. <https://doi.org/10.1111/nmo.12817>

220. Khanna AK, Bergese SD, Jungquist CR, Morimatsu H, Uezono S, Lee S, et al.. Prediction of opioid-induced respiratory depression on inpatient wards using continuous capnography and oximetry: an international prospective, observational trial (PRODIGY). *Anesth Analg*. 2020;131(4):1012-1024. 2020. <https://doi.org/10.1213/ANE.0000000000004788>
221. Khanna AK, et al (PRODIGY). Prediction of opioid-induced respiratory depression on inpatient wards (PRODIGY). *Anesth Analg*. 2020. <https://doi.org/10.1213/ANE.0000000000004788>
222. Duong M, Islam S, Rangarajan S, et al. (PURE Study). Mortality and cardiovascular and respiratory morbidity in individuals with impaired FEV₁ (PURE): an international, community-based cohort study. *Lancet Global Health*. 2019. [https://doi.org/10.1016/S2214-109X\(19\)30070-1](https://doi.org/10.1016/S2214-109X(19)30070-1)
223. Duong M, et al. Mortality and cardiovascular and respiratory morbidity in individuals with impaired FEV₁ (PURE). *Lancet Glob Health*. 2019. [https://doi.org/10.1016/S2214-109X\(19\)30070-1](https://doi.org/10.1016/S2214-109X(19)30070-1)
224. Cannon MJ, Webb-Vargas Y, Cohen RA, et al.. FEV₁ within the normal range and 20-year mortality (FDNY + NHANES). *American Journal of Respiratory and Critical Care Medicine*. 2024. <https://www.atsjournals.org/doi/10.1164/rccm.202310-1914OC>
225. Rydell A, Janson C, Lisspers K, Lin YT, Ärnlov J. FEV₁ and FVC as robust risk factors for cardiovascular disease and mortality: Insights from a large population study. *Respir Med* 2024;227:107614. 2024. <https://doi.org/10.1016/j.rmed.2024.107614>
226. Kim SJ, et al.. FEV₁/FVC decline and 18-year mortality in Korean cohort. *Scientific Reports*. 2022. <https://pmc.ncbi.nlm.nih.gov/articles/PMC9807769/>
227. Hellman JB, Stacy RW. Variation of respiratory sinus arrhythmia with age. *J Appl Physiol* 1976;41(5 Pt 1):734-738. 1976. <https://doi.org/10.1152/jappl.1976.41.5.734>
228. Kodama S, et al. Cardiorespiratory fitness as a quantitative predictor of all-cause mortality. *JAMA*. 2009. <https://doi.org/10.1001/jama.2009.681>
229. Lang JJ, Prince SA, Merucci K, Cadenas-Sanchez C, Chaput JP, Fraser BJ, Manyanga T, McGrath R, Ortega FB, Singh B, Tomkinson GR. Cardiorespiratory fitness is a strong and consistent predictor of morbidity and mortality among adults: an overview of meta-analyses representing over 20.9 million observations from 199 unique cohort studies. *Br J Sports Med* 2024;58(10):556-566. 2024. <https://doi.org/10.1136/bjsports-2023-107849>
230. Peterman JE, Arena R, Myers J, et al.. Development of global reference standards for directly measured cardiorespiratory fitness: a report from the Fitness Registry and Importance of Exercise National Database (FRIEND). *Mayo Clin Proc*. 95(2):255-264. 2020. <https://doi.org/10.1016/j.mayocp.2019.11.012>
231. Strasser B, Burtscher M. Survival of the fittest: VO₂max, a key predictor of longevity?. *Front Biosci (Landmark Ed)*. 23(8):1505-1516. 2018. <https://doi.org/10.2741/4657>
232. Spathis D, Perez-Pozuelo I, Gonzales TI, et al.. Longitudinal cardio-respiratory fitness prediction through wearables in free-living environments. *npj Digital Medicine*. 2022. <https://doi.org/10.1038/s41746-022-00719-1>
233. Olshansky SJ, Carnes BA. Ever since Gompertz. *Demography*. 34(1):1-15. 1997. <https://pubmed.ncbi.nlm.nih.gov/9074828/>
234. Meiners S, Eickelberg O, Königshoff M. Hallmarks of the ageing lung. *European Respiratory Journal*. 2015. <https://doi.org/10.1183/09031936.00186914>
235. Rezwan FI, Imboden M, Amaral AFS, et al.. Association of adult lung function with epigenetic age acceleration (SAPALDIA, ECRHS). *Aging*. 2020. <https://doi.org/10.18632/aging.102639>
236. Hillary RF, et al.. GrimAge predicts incident COPD. *European Journal of Epidemiology / equivalent*. 2020. <https://pmc.ncbi.nlm.nih.gov/articles/PMC7394682/>
237. Karniski S, et al.. Telomere length predicts mortality in pulmonary fibrosis across racial groups. *Nature Communications*. 2023. <https://doi.org/10.1038/s41467-023-37085-9>
238. Seymour CW, Liu VX, Iwashyna TJ, Brunkhorst FM, Rea TD, Scherag A, Rubenfeld G, Kahn JM, Shankar-Hari M, Singer M, Deutschman CS, Escobar GJ, Angus DC. Assessment of Clinical Criteria for Sepsis: For the Third International Consensus Definitions for Sepsis and Septic Shock (Sepsis-3). *JAMA* 2016;315(8):762-774. 2016. <https://doi.org/10.1001/jama.2016.0288>

239. Raith EP, Udy AA, Bailey M, McGloughlin S, MacIsaac C, Bellomo R, Pilcher DV; Australian and New Zealand Intensive Care Society Centre for Outcomes and Resource Evaluation (ANZICS CORE). Prognostic Accuracy of the SOFA Score, SIRS Criteria, and qSOFA Score for In-Hospital Mortality Among Adults With Suspected Infection Admitted to the Intensive Care Unit. *JAMA* 2017;317(3):290-300. 2017. <https://doi.org/10.1001/jama.2016.20328>
240. Salinas-Garrido M, González-Castro A, Diez-Manglano J. High-flow Nasal Oxygen Therapy Yields a Favorable Outcome in Patient Presenting With Kussmaul Breathing. *Open Respir Arch* 2024;6(2):100322. 2024. <https://doi.org/10.1016/j.opresp.2024.100322>
241. Birring SS, Fleming T, Matos S, et al.. The Leicester Cough Monitor: preliminary validation of an automated cough detection system. *European Respiratory Journal*. 2008. <https://doi.org/10.1183/09031936.00101307>
242. Pahar M, Klopper M, Warren R, Niesler T.. COVID-19 cough classification using machine learning and global smartphone recordings. *Computers in Biology and Medicine*. 2021. <https://doi.org/10.1016/j.combiomed.2021.104572>
243. Yellapu GD, Rudraraju G, Sripada NR, Mamidgi B, Jalukuru C, Firmal P, Yechuri V, Varanasi S, Peddireddi VS, Bhimarasetty DM, et al.. Development and clinical validation of Swaasa AI platform for screening and prioritization of pulmonary TB. *Sci Rep* 2023;13:4740. 2023. <https://doi.org/10.1038/s41598-023-31772-9>
244. Chetupalli SR, Krishnan P, Sharma N, Muguli A, Kumar R, Nanda V, Pinto LM, Ghosh PK, Ganapathy S. Multi-Modal Point-of-Care Diagnostics for COVID-19 Based on Acoustics and Symptoms. *IEEE J Transl Eng Health Med* 2023;11:199-210. 2023. <https://doi.org/10.1109/JTEHM.2023.3250700>
245. Emeryk A, Derom E, Janeczek K, et al.. Home AI-assisted stethoscope (StethoMe) detects asthma exacerbations with AUC ~0.93. *Annals of Family Medicine*. 2023. <https://doi.org/10.1370/afm.3023>
246. Huffaker MF, Carchia M, Harris BU, Kethman WC, Murphy TE, Sakarovitch C, Qin F, Cornfield DN. Passive nocturnal physiologic monitoring enables early detection of exacerbations in children with asthma: a proof-of-concept study. *American Journal of Respiratory and Critical Care Medicine*. 2018. <https://doi.org/10.1164/rccm.201712-2606OC>
247. Chaccour C, Sánchez-Olivieri I, Siegel S, et al. Validation and accuracy of the Hyfe cough monitoring system: A multicenter clinical study. *Scientific Reports*. 15:Article 880. 2025. <https://doi.org/10.1038/s41598-025-85341-3>
248. Landry V, et al.. Acoustic biomarkers in respiratory disease: a systematic review. *Eur Respir Rev*. 2025;34(176):240246. 2025. <https://publications.ersnet.org/content/errev/34/176/240246>
249. Landry V, et al.. Audio biomarkers for respiratory disease detection: systematic review. *European Respiratory Review*. 2025. <https://pmc.ncbi.nlm.nih.gov/articles/PMC12076160/>
250. Hawthorne G, Richardson M, Greening NJ, Eslinger D, Briggs-Price S, Chaplin EJ, Clinch L, Steiner MC, Singh SJ, Orme MW. A proof of concept for continuous, non-invasive, free-living vital signs monitoring to predict readmission following an acute exacerbation of COPD: a prospective cohort study. *Respir Res*. 2022;23(1):102. 2022. <https://doi.org/10.1186/s12931-022-02018-5>
251. Hawthorne G, Richardson M, Greening NJ, et al.. Wearable monitoring of physiological signs in the prodromal period of COPD exacerbations. *npj Digital Medicine*. 2023. <https://doi.org/10.1038/s41746-022-00736-0>
252. Vinik AI, Maser RE, Mitchell BD, Freeman R. Diabetic Autonomic Neuropathy. *Diabetes Care* 2003;26(5):1553-1579. 2003. <https://doi.org/10.2337/diacare.26.5.1553>
253. Maser RE, Mitchell BD, Vinik AI, Freeman R. The association between cardiovascular autonomic neuropathy and mortality in individuals with diabetes: a meta-analysis. *Diabetes Care* 2003;26(6):1895-1901. 2003. <https://doi.org/10.2337/diacare.26.6.1895>
254. Spallone V, Ziegler D, Freeman R, et al.. Cardiovascular autonomic neuropathy in diabetes: clinical impact, assessment, diagnosis, and management. *Diabetes/Metabolism Research and Reviews*. 2011. <https://doi.org/10.1002/dmrr.1239>
255. Vinik AI, Ziegler D. Diabetic cardiovascular autonomic neuropathy. *Circulation*. 2007. <https://pubmed.ncbi.nlm.nih.gov/17242296/>
256. Meuret AE, Rosenfield D, Hofmann SG, Suvak MK, Roth WT. Changes in respiration mediate changes in fear of bodily sensations in panic disorder. *Journal of Psychiatric Research*. 43(6):634-641. 2009. <https://doi.org/10.1016/j.jpsychires.2008.08.003>

257. Tolin DF, et al. Capnometry-guided respiratory intervention for panic disorder. *Appl Psychophysiol Biofeedback*. 2017. <https://doi.org/10.1007/s10484-017-9354-4>
258. Kaplan A, et al. Freespira respiratory intervention outcomes for panic disorder. *Front Digit Health*. 2022. <https://doi.org/10.3389/fdgth.2022.1049487>
259. Koch C, Wilhelm M, Salzmann S, Rief W, Euteneuer F. A meta-analysis of heart rate variability in major depression. *Psychol Med* 2019;49(12):1948-1957. 2019. <https://doi.org/10.1017/S0033291719001351>
260. Kemp AH, Quintana DS, Gray MA, Felmingham KL, Brown K, Gatt JM. Impact of depression and antidepressant treatment on heart rate variability: a review and meta-analysis. *Biol Psychiatry* 2010;67(11):1067-1074. 2010. <https://doi.org/10.1016/j.biopsych.2009.12.012>
261. Khanna AK, Bergese SD, Jungquist CR, Morimatsu H, Uezono S, Lee S, Ti LK, Urman RD, McIntyre R Jr, Tornero C, Dahan A, Saager L, Weingarten TN, Wittmann M, Auckley D, Brazzi L, Le Guen M, Soto R, Schramm F, Ayad S, Kaw R, Di Piazza F, Overdyk FJ; PRODIGY Group Investigators. Prediction of Opioid-Induced Respiratory Depression on Inpatient Wards Using Continuous Capnography and Oximetry: An International Prospective, Observational Trial. *Anesth Analg* 2020;131(4):1012-1024. 2020. <https://doi.org/10.1213/ANE.0000000000004788>
262. Hu D, et al.. Lung function and incident depression: UK Biobank prospective cohort. *BMC Medicine*. 2024. <https://pmc.ncbi.nlm.nih.gov/articles/PMC11017623/>
263. Driver HS, McLean H, Kumar DV, et al.. Influence of the menstrual cycle on upper airway resistance and breathing during sleep. *Sleep*. 2005. <https://doi.org/10.1093/sleep/28.4.449>
264. Rattley CA, Ansdell P, Burgess L, Felton M, Dewhurst S, Armstrong M, Neal R. Ventilation differences in the menstrual cycle: A systematic review and meta-analysis. *Respir Physiol Neurobiol* 2025;337:104468. 2025. <https://doi.org/10.1016/j.resp.2025.104468>
265. Kahal H, Kyrou I, Tahrani AA, Randeve HS.. Obstructive sleep apnoea and polycystic ovary syndrome: a comprehensive review of clinical interactions and underlying pathophysiology. *Clinical Endocrinology (Oxf)*. 2017. <https://doi.org/10.1111/cen.13392>
266. Helvacı N, Karabulut E, Demir AU, Yildiz BO. Polycystic ovary syndrome and the risk of obstructive sleep apnea: a meta-analysis and review of the literature. *Endocr Connect* 2017;6(7):437-445. 2017. <https://doi.org/10.1530/EC-17-0129>
267. Soroka T, Ravia A, Snitz K, et al.. Humans have nasal respiratory fingerprints. *Current Biology*. 2025. <https://doi.org/10.1016/j.cub.2025.04.057>
268. Zelano C, et al. Nasal respiration entrains human limbic oscillations and modulates cognitive function. *J Neurosci*. 2016. <https://doi.org/10.1523/JNEUROSCI.2586-16.2016>
269. Johannknecht M, Kayser C. Reaction time covaries with the respiratory cycle: a meta-analysis of 122 datasets. *Scientific Reports*. 2022. <https://pmc.ncbi.nlm.nih.gov/articles/PMC8850565/>
270. Perl O, Ravia A, Rubinson M, et al.. Human non-olfactory cognition phase-locked with inhalation. *Nature Human Behaviour*. 2019. <https://doi.org/10.1038/s41562-019-0556-z>
271. Nakamura NH, Fukunaga M, Oku Y. Respiratory modulation of cognitive performance during the retrieval process. *PLoS ONE*. 2018. <https://doi.org/10.1371/journal.pone.0204021>
272. Arshamian A, Iravani B, Majid A, Lundström JN. Respiration modulates olfactory memory consolidation in humans. *Journal of Neuroscience*. 2018. <https://doi.org/10.1523/JNEUROSCI.3360-17.2018>
273. Ahani A, Wahbeh H, Nezamfar H, Miller M, Erdogmus D, Oken B. Quantitative change of EEG and respiration signals during mindfulness meditation. *J NeuroEng Rehabil* 2014;11:87. 2014. <https://doi.org/10.1186/1743-0003-11-87>
274. Sato S, Miyazaki S, Tamaki A, Yoshimura Y, Arai H, Fujiwara D, Katsura H, Kawagoshi A, Kozu R, Maeda K, Ogawa S, Ueki J, Wakabayashi H. Respiratory sarcopenia: A position paper by four professional organizations. *Geriatr Gerontol Int* 2023;23(1):5-15. 2023. <https://doi.org/10.1111/ggi.14519>
275. Nagano A, Wakabayashi H, Maeda K, Kokura Y, Miyazaki S, Mori T, Fujiwara D. Respiratory Sarcopenia and Sarcopenic Respiratory Disability: Concepts, Diagnosis, and Treatment. *J Nutr Health Aging* 2021;25(4):507-515. 2021. <https://doi.org/10.1007/s12603-021-1587-5>

276. Vaz Fragoso CA, Enright PL, McAvay G, et al. Respiratory impairment and frailty in older persons (CHS). *Journal of the American Geriatrics Society*. 2012. <https://doi.org/10.1016/j.amjmed.2011.06.024>
277. Vidal MB, Pegorari MS, Santos EC, Matos AP, Pinto ACPN, Ohara DG, et al. Respiratory muscle strength for discriminating frailty in community-dwelling elderly: A cross-sectional study. *Archives of Gerontology and Geriatrics*. 89:Article 104082. 2020. <https://doi.org/10.1016/j.archger.2020.104082>
278. Lee Y, Son S, Kim DK, Park MW. Association of Diaphragm Thickness and Respiratory Muscle Strength With Indices of Sarcopenia. *Ann Rehabil Med* 2023;47(4):307-314. 2023. <https://doi.org/10.5535/arm.23081>
279. Enright PL, Kronmal RA, Manolio TA, Schenker MB, Hyatt RE.. Respiratory muscle strength in the elderly: correlates and reference values (Cardiovascular Health Study). *American Journal of Respiratory and Critical Care Medicine*. 1994. <https://doi.org/10.1164/ajrccm.149.2.8306041>
280. Ohara DG, Pegorari MS, Oliveira Dos Santos NL, et al.. Respiratory muscle strength as a discriminator of sarcopenia in community-dwelling elderly. *Journal of Nutrition, Health & Aging*. 2018. <https://doi.org/10.1007/s12603-018-1079-4>
281. Chen L et al. (CHARLS analysis). Probable respiratory sarcopenia and 9-year mortality in community-dwelling older adults: the first longitudinal evidence from the CHARLS. *Sci Rep* 2025;15:18511 (article 18511-y). 2025. <https://doi.org/10.1038/s41598-025-18511-y>
282. Martin-Harris B, Brodsky MB, Michel Y, Ford CL, Walters B, Heffner J.. Breathing and swallowing dynamics across the adult lifespan. *Archives of Otolaryngology — Head & Neck Surgery*. 2005. <https://doi.org/10.1001/archotol.131.9.762>
283. Gross RD, Atwood CW Jr, Ross SB, Eichhorn KA, Olszewski JW, Doyle PJ.. The coordination of breathing and swallowing in Parkinson's disease. *Dysphagia*. 2008. <https://doi.org/10.1007/s00455-007-9113-4>
284. Gross RD, Atwood CW, Ross SB, et al.. The coordination of breathing and swallowing in chronic obstructive pulmonary disease. *American Journal of Respiratory and Critical Care Medicine*. 2009. <https://doi.org/10.1164/rccm.200807-1139OC>
285. Nagami S, Oku Y, Yagi N, et al.. Breathing-swallowing discoordination is associated with acute exacerbation of COPD. *BMJ Open Respiratory Research*. 2017. <https://doi.org/10.1136/bmjresp-2017-000202>
286. Yoshimatsu Y, Tobino K, Maeda K, et al.. Breathing-swallowing discoordination is independently associated with COPD exacerbation. *COPD: Journal of Chronic Obstructive Pulmonary Disease*. 2020. <https://doi.org/10.1080/15412555.2020.xxxxxxxxxx>
287. Frakking T, Chang AB, O'Grady KF, et al.. Aspiration prediction in children: ML on cervical auscultation. *Dysphagia*. 2022. <https://doi.org/10.1007/s00455-022-xxxxx>
288. Donohue C, et al.. High-resolution cervical auscultation for distinguishing healthy and neurodegenerative dysphagic swallows. *Dysphagia*. 2021. <https://pmc.ncbi.nlm.nih.gov/articles/PMC8642099/>
289. Crary MA, Carnaby Mann GD, Groher ME. Identification of dysphagia by spontaneous swallowing frequency in stroke. *Dysphagia*. 2013. <https://doi.org/10.1007/s00455-013-9468-7>
290. Tanaka N, Nohara K, Kotani Y, Matsumura M, Sakai T. Swallowing frequency in elderly people during daily life. *J Oral Rehabil* 2013;40(10):744-750. 2013. <https://doi.org/10.1111/joor.12085>
291. Sedky K, et al.. ADHD and sleep-disordered breathing: meta-analysis (18 studies, n=2518). *Sleep Medicine Reviews*. 2014. <https://pubmed.ncbi.nlm.nih.gov/24581717/>
292. Ivanov I, et al.. Sleep-disordered breathing and ADHD: review. *Journal of Attention Disorders*. 2024. <https://doi.org/10.1177/10870547241232313>
293. Orlandic L, Teijeiro T, Atienza D. The COUGHVID crowdsourcing dataset: a corpus for the study of large-scale cough analysis algorithms. *Sci Data*. 2021;8(1):156. doi:10.1038/s41597-021-00937-4. <https://doi.org/10.1038/s41597-021-00937-4>
294. Sharma N, Krishnan P, Kumar R, et al. Coswara — a database of breathing, cough, and voice sounds for COVID-19 diagnosis. *INTERSPEECH* 2020. doi:10.21437/Interspeech.2020-2768. <https://doi.org/10.21437/Interspeech.2020-2768>
295. Boselli E, Daniela-Ionescu M, Bégou G, et al.. Prospective observational study of the non-invasive assessment of immediate postoperative pain using the analgesia/nociception index (ANI). *British Journal of Anaesthesia*. 2013. <https://doi.org/10.1093/bja/aet110>

296. Kim MK, Choi GJ, Oh KS, Lee SP, Kang H. Pain Assessment Using the Analgesia Nociception Index (ANI) in Patients Undergoing General Anesthesia: A Systematic Review and Meta-Analysis. *J Pers Med* 2023;13(10):1461. 2023. <https://doi.org/10.3390/jpm13101461>
297. Park M, Lee HJ, Ko A, et al.. Analgesia nociception index for perioperative analgesia monitoring: a systematic review and meta-analysis. *Journal of Personalized Medicine*. 2023. <https://doi.org/10.3390/jpm13060954>
298. Baroni DA, Abreu LG, Paiva SM, Costa LR. Comparison between Analgesia Nociception Index (ANI) and self-reported measures for diagnosing pain in conscious individuals: a systematic review and meta-analysis. *Sci Rep* 2022;12:2862. 2022. <https://doi.org/10.1038/s41598-022-06993-z>
299. Cuenca-Martínez F, Sempere-Rubio N, Muñoz-Gómez E, Mollà-Casanova S, Carrasco-González E, Martínez-Arnau FM. Respiratory Function Analysis in Patients with Chronic Pain: An Umbrella Review and Meta-Analysis of Pooled Findings. *Healthcare*. 11(9):Article 1358. 2023. <https://doi.org/10.3390/healthcare11091358>
300. Jonsson K, et al.. Chronic hyperventilation and renal compensation in fibromyalgia. *Scandinavian Journal of Pain*. 2024. <https://pubmed.ncbi.nlm.nih.gov/38907689/>
301. Jonsson K, Bjurström MF, Lerman SF, et al.. Respiratory pattern alterations in fibromyalgia. *Journal of Pain*. 2025. <https://pubmed.ncbi.nlm.nih.gov/40706939/>
302. (authors per publication). Altered breathing pattern and thoracic mobility in women with fibromyalgia: A case-control study. *The Journal of Pain*. 105508 (Advance online). 2025. <https://doi.org/10.1016/j.jpain.2025.105508>
303. Chang KC, Plummer JL, Lin CY. Acute effects of intravenous opioids on breathing pattern. *Anaesthesia*. 1999. <https://pubmed.ncbi.nlm.nih.gov/10460553/>
304. Dahan A, Aarts L, Smith TW. Oxycodone and alcohol synergistic respiratory depression. *Anesthesiology*. 2017. <https://doi.org/10.1097/ALN.0000000000001536>
305. Frøkjaer JB, Bergmann S, Brock C, et al.. Modulation of vagal tone enhances gastroduodenal motility and reduces somatic pain sensitivity. *Neurogastroenterology & Motility*. 2016. <https://doi.org/10.1111/nmo.12760>
306. López-Otín C, Blasco MA, Partridge L, et al.. Hallmarks of aging: an expanding universe. *Cell*. 2023. <https://doi.org/10.1016/j.cell.2022.11.001>
307. Bellelli F, Angioni D, Arosio B, Vellas B, De Souto Barreto P. Hallmarks of aging and Alzheimer's disease pathogenesis: paving the route for new therapeutic targets. *Ageing Res Rev*. 106:102699. 2025. <https://doi.org/10.1016/j.arr.2025.102699>
308. Argentieri MA, Amin N, Nevado-Holgado AJ, et al.. Integrating the environmental and genetic architectures of aging and mortality. *Nat Med*. 31(3):1016-1025. 2025. <https://doi.org/10.1038/s41591-024-03483-9>
309. Wang, C., Nwanaji-Enwerem, J. C., Gao, X., Vokonas, P., Sparrow, D., Hou, L., ... & Baccarelli, A. A. (2020). Biomarkers of aging and lung function in the normative aging study. *Aging*, 12(12), 11555-11566. <https://pmc.ncbi.nlm.nih.gov/articles/PMC7394682/>
310. Lu AT, Binder AM, Zhang J, et al.. DNA methylation GrimAge2 strongly predicts mortality and correlates with FEV1. *Aging (Albany NY)*. 2022. <https://doi.org/10.18632/aging.204434>
311. Landi F, Calvani R, Cesari M, et al.. Sarcopenia as the biological substrate of physical frailty. *Clin Geriatr Med*. 31(3):367-374. 2015. <https://doi.org/10.1016/j.cger.2015.04.005>
312. Gielen E, Dupont J, Dejaeger M, Laurent MR. Sarcopenia, osteoporosis and frailty. *Metabolism*. 145:155638. 2023. <https://doi.org/10.1016/j.metabol.2023.155638>
313. Atlantis E, Martin SA, Haren MT, Taylor AW, Wittert GA. Inverse associations between muscle mass, strength, and the metabolic syndrome. *Metabolism*. 58(7):1013-1022. 2009. <https://doi.org/10.1016/j.metabol.2009.02.027>
314. Mielke MM, Roberts RO, Savica R, et al.. Assessing the temporal relationship between cognition and gait: slow gait predicts cognitive decline in the Mayo Clinic Study of Aging. *J Gerontol A Biol Sci Med Sci*. 68(8):929-937. 2013. <https://doi.org/10.1093/gerona/gls256>
315. Hellman JB, Stacy RW.. Variation of respiratory sinus arrhythmia with age. *Journal of Applied Physiology*. 1976. <https://doi.org/10.1152/jappl.1976.41.5.734>

316. Hellman JB, Stacy RW. Variation of respiratory sinus arrhythmia with age. *J Appl Physiol*. 1976. <https://doi.org/10.1152/jappl.1976.41.5.734>
317. Hernández-Vicente A, et al. Heart rate variability and exceptional longevity. *Front Physiol*. 2020. <https://doi.org/10.3389/fphys.2020.566399>
318. Mandsager K, Harb S, Cremer P, Phelan D, Nissen SE, Jaber W. Association of cardiorespiratory fitness with long-term mortality among adults undergoing exercise treadmill testing. *JAMA Netw Open*. 1(6):e183605. 2018. <https://doi.org/10.1001/jamanetworkopen.2018.3605>
319. Weeldreyer NR, De Guzman JC, Paterson C, Allen JD, Gaesser GA, Angadi SS. Cardiorespiratory fitness, body mass index and mortality: a systematic review and meta-analysis. *Br J Sports Med*. 59(5):339-346. 2025. <https://doi.org/10.1136/bjsports-2024-108748>
320. Paluch AE, Bajpai S, Bassett DR, et al.. Daily steps and all-cause mortality: a meta-analysis of 15 international cohorts. *Lancet Public Health*. 7(3):e219-e228. 2022. [https://doi.org/10.1016/S2468-2667\(21\)00302-9](https://doi.org/10.1016/S2468-2667(21)00302-9)
321. Lee DH, Rezende LFM, Joh HK, et al.. Long-term leisure-time physical activity intensity and all-cause and cause-specific mortality: a prospective cohort of US adults. *Circulation*. 146(7):523-534. 2022. <https://doi.org/10.1161/CIRCULATIONAHA.121.058162>
322. Martinez-Gomez D, Luo M, Huang Y, et al.. Physical activity and all-cause mortality by age in 4 multinational megacohorts. *JAMA Netw Open*. 7(11):e2446802. 2024. <https://doi.org/10.1001/jamanetworkopen.2024.46802>
323. Min J, et al.. Slow-paced breathing reduces plasma amyloid- β in healthy adults. *Scientific Reports*. 2023. <https://doi.org/10.1038/s41598-023-30978-1>
324. Churpek MM, et al. Predicting cardiac arrest on the wards: a nested case-control study. *Chest*. 2012. <https://pmc.ncbi.nlm.nih.gov/articles/PMC3342781/>
325. Churpek MM, Yuen TC, Winslow C, Meltzer DO, Kattan MW, Edelson DP. Multicenter comparison of machine learning methods and conventional regression for predicting clinical deterioration on the wards. *Crit Care Med*. 2016;44(2):368-374. 2016. <https://doi.org/10.1097/CCM.0000000000001571>
326. Stehlik J, Schmalfluss C, Bozkurt B, Nativi-Nicolau J, Wohlfahrt P, Wegerich S, et al.. Continuous wearable monitoring analytics predict heart failure hospitalization: the LINK-HF multicenter study. *Circ Heart Fail*. 2020;13(3):e006513. 2020. <https://doi.org/10.1161/CIRCHEARTFAILURE.119.006513>
327. Shah SA, Velardo C, Farmer A, Tarassenko L.. Exacerbations in chronic obstructive pulmonary disease: identification and prediction using a digital health system. *Journal of Medical Internet Research*. 2017. <https://doi.org/10.2196/jmir.7207>
328. Yañez AM, et al. Monitoring breathing rate at home allows early identification of COPD exacerbations. *Chest*. 2012. <https://doi.org/10.1378/chest.11-2728>
329. Wu CT, et al. COPD exacerbation prediction with wearable data. *JMIR mHealth uHealth*. 2021. <https://mhealth.jmir.org/2021/5/e22591>
330. Park JS, Park SY, Moon JW, Kim K, Suh DI. Artificial Intelligence Models for Pediatric Lung Sound Analysis: Systematic Review and Meta-Analysis. *J Med Internet Res* 2025;27:e66491. 2025. <https://doi.org/10.2196/66491>
331. Moran-Mendoza O, Ritchie T, Aldhaferi S. Fine crackles on chest auscultation in the early diagnosis of idiopathic pulmonary fibrosis: a prospective cohort study. *BMJ Open Respir Res* 2021;8(1):e000815. 2021. <https://doi.org/10.1136/bmjresp-2020-000815>
332. Manfredi A, Cassone G, Pancaldi F, Vacchi C, Cerri S, Della Casa G, Salvarani C, Sebastiani M. Usefulness of digital velcro crackles detection in identification of interstitial lung disease in patients with connective tissue diseases. *Arch Rheumatol* 2021;36(1):87-93. 2021. <https://doi.org/10.46497/ArchRheumatol.2021.7975>
333. Cretikos MA, Bellomo R, Hillman K, et al. Respiratory rate: the neglected vital sign. *Med J Aust*. 2008. <https://doi.org/10.5694/j.1326-5377.2008.tb01860.x>
334. Walker PP, Pompilio PP, Zanaboni P, Bergmo TS, Prikk K, Malinovschi A, et al.. Telemonitoring in Chronic Obstructive Pulmonary Disease (CHROMED). A Randomized Clinical Trial. *Am J Respir Crit Care Med* 2018;198(5):620-628. 2018. <https://doi.org/10.1164/rccm.201712-2404OC>

335. Zhu Y, et al. Dynamic versus static early warning score for predicting postoperative deterioration (DyNiEWS vs NEWS). *Resuscitation*. 2020. <https://doi.org/10.1016/j.resuscitation.2020.10.037>
336. Zhu Y, Chiu YD, Villar SS, Brand JW, Patteril MV, Morrice DJ, Clayton J, Mackay JH. Dynamic individual vital sign trajectory early warning score (DyNiEWS) versus snapshot national early warning score (NEWS) for predicting postoperative deterioration. *Resuscitation* 2020;157:176-184. 2020. <https://doi.org/10.1016/j.resuscitation.2020.10.034>
337. Chaudhry SI, et al. Telemonitoring in patients with heart failure (TELE-HF). *N Engl J Med*. 2010. <https://doi.org/10.1056/NEJMoa1010029>
338. Koehler F, et al. Impact of remote telemedical management on mortality and hospitalizations (TIM-HF). *Circulation*. 2011. <https://doi.org/10.1161/CIRCULATIONAHA.111.018473>
339. Ong MK, et al. Effectiveness of remote patient monitoring after discharge of hospitalized patients with HF (BEAT-HF). *JAMA Intern Med*. 2016. <https://doi.org/10.1001/jamainternmed.2015.7712>
340. Koehler F, et al. Efficacy of telemedical interventional management in patients with heart failure (TIM-HF2). *Lancet*. 2018. [https://doi.org/10.1016/S0140-6736\(18\)31880-4](https://doi.org/10.1016/S0140-6736(18)31880-4)
341. Zhu Y, Armstrong JL, Tchkonja T, et al. (DyNiEWS). DyNiEWS: dynamic personalized early warning score outperforms NEWS. *Resuscitation*. 2020. <https://doi.org/10.1016/j.resuscitation.2020.10.030>
342. Wynants L, Van Calster B, Collins GS, et al.. Prediction models for diagnosis and prognosis of COVID-19: systematic review and critical appraisal. *BMJ*. 2020. <https://doi.org/10.1136/bmj.m1328>
343. Naemi A, et al.. Dynamic vs static early warning scores: comparison. *Respiratory Research*. 2022. <https://research.biomedcentral.com/articles/10.1186/s12931-022-02130-6>
344. Avram R, Tison GH, Aschbacher K, et al.. Real-world heart rate norms in the Health eHeart study. *NPJ Digit Med*. 2:58. 2019. <https://doi.org/10.1038/s41746-019-0134-9>
345. Bassett DR Jr, Wyatt HR, Thompson H, Peters JC, Hill JO. Pedometer-measured physical activity and health behaviors in U.S. adults. *Med Sci Sports Exerc*. 42(10):1819-1825. 2010. <https://doi.org/10.1249/MSS.0b013e3181dc2e54>
346. Troiano RP, Berrigan D, Dodd KW, Massé LC, Tilert T, McDowell M. Physical activity in the United States measured by accelerometer. *Med Sci Sports Exerc*. 40(1):181-188. 2008. <https://doi.org/10.1249/mss.0b013e31815a51b3>
347. Borrud LG, Flegal KM, Looker AC, Everhart JE, Harris TB, Shepherd JA. Body composition data for individuals 8 years of age and older: U.S. population, 1999-2004. *Vital Health Stat 11. (250):1-87*. 2010. <https://pubmed.ncbi.nlm.nih.gov/20939264/>
348. Kaminsky LA, Arena R, Myers J. Reference standards for cardiorespiratory fitness measured with cardiopulmonary exercise testing: data from the Fitness Registry and the Importance of Exercise National Database (FRIEND). *Mayo Clin Proc*. 90(11):1515-1523. 2015. <https://doi.org/10.1016/j.mayocp.2015.07.026>
349. Spiegelhalter D. How old are you, really? Communicating chronic risk through 'effective age' of your body and organs. *BMC Med Inform Decis Mak*. 16:104. 2016. <https://doi.org/10.1186/s12911-016-0342-z>
350. DeSalvo KB, Blosler N, Reynolds K, He J, Muntner P. Mortality prediction with a single general self-rated health question: a meta-analysis. *J Gen Intern Med*. 21(3):267-275. 2006. <https://doi.org/10.1111/j.1525-1497.2005.00291.x>
351. Idler EL, Benyamini Y. Self-rated health and mortality: a review of twenty-seven community studies. *J Health Soc Behav*. 38(1):21-37. 1997. <https://pubmed.ncbi.nlm.nih.gov/9097506/>
352. Pang M, Hanley JA. 'Translating' all-cause mortality rate ratios or hazard ratios to age-, longevity-, and probability-based measures. *Am J Epidemiol*. 190(12):2664-2670. 2021. <https://doi.org/10.1093/aje/kwab178>
353. Lauderdale DS, Knutson KL, Yan LL, Liu K, Rathouz PJ. Self-reported and measured sleep duration: how similar are they?. *Epidemiology*. 19(6):838-845. 2008. <https://doi.org/10.1097/EDE.0b013e318187a7b0>
354. Lee PH. Validation of the National Health And Nutritional Survey (NHANES) single-item self-reported sleep duration against wrist-worn accelerometer. *Sleep Breath*. 26(4):2069-2075. 2022. <https://doi.org/10.1007/s11325-021-02542-6>
355. Ford ES, Cunningham TJ, Croft JB. Trends in self-reported sleep duration among US adults from 1985 to 2012. *Sleep*. 38(5):829-832. 2015. <https://doi.org/10.5665/sleep.4684>

356. Rocha BM, Filos D, Mendes L, et al. ICBHI 2017 Respiratory Sound Database. *International Conference on Biomedical and Health Informatics*. 2017. <https://bhichallenge.med.auth.gr/>
357. Rocha BM, Filos D, Mendes L, Vogiatzis I, Perantoni E, Kaimakamis E, et al. An open access database for the evaluation of respiratory sound classification algorithms (ICBHI). *Physiological Measurement*. 2019. <https://doi.org/10.1088/1361-6579/ab03ea>
358. Hsu FS, Huang SR, Huang CW, et al. HF_Lung_V1 — benchmark dataset of lung sounds with phase annotation. *PLoS ONE / IEEE*. 2021. <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0254134>
359. Zhang G, Liu G, Zhu Y, et al. SPRSound: open-source SJTU paediatric respiratory sound database. *IEEE Trans Biomed Circuits Syst*. 2022;16(5):867-881. doi:10.1109/TBCAS.2022.3204910. <https://doi.org/10.1109/TBCAS.2022.3204910>
360. Brown C, Chauhan J, Grammenos A, et al. Exploring automatic diagnosis of COVID-19 from crowdsourced respiratory sound data (Cambridge COVID-19 Sounds). *KDD '20 Proceedings*. 2020:3474-3484. https://www.cl.cam.ac.uk/~cm542/papers/KDD_covid_19.pdf
361. Goldberger AL, Amaral LA, Glass L, Hausdorff JM, Ivanov PCh, Mark RG, et al. PhysioBank, PhysioToolkit, and PhysioNet: components of a new research resource for complex physiologic signals. *Circulation* 101(23):e215-e220. 2000. <https://doi.org/10.1161/01.CIR.101.23.e215>
362. Quan SF, Howard BV, Iber C, et al. The Sleep Heart Health Study: design, rationale, and methods. *Sleep*. 1997;20(12):1077-1085. doi:10.1093/sleep/20.12.1077. <https://doi.org/10.1093/sleep/20.12.1077>
363. Chen X, Wang R, Zee P, et al. Racial/ethnic differences in sleep disturbances: the MESA study. *Sleep*. 2015;38(6):877-888. doi:10.5665/sleep.4732. <https://doi.org/10.5665/sleep.4732>
364. Korompili G, Amfilochiou A, Kokkalas L, Mitilineos SA, Tatlas N-A, Kouvaras M, et al. PSG-Audio, a scored polysomnography dataset with simultaneous audio recordings for sleep apnea studies. *Sci Data* 8(1):197. 2021. <https://doi.org/10.1038/s41597-021-00977-w>
365. Lee HC, Park Y, Yoon SB, Yang SM, Park D, Jung CW. VitalDB, a high-fidelity multi-parameter vital signs database in surgical patients. *Sci Data* 9(1):279. 2022. <https://doi.org/10.1038/s41597-022-01411-5>
366. Garrison J, Soroka T, Roy S, Ledsam J, et al. HeAR – Health Acoustic Representations. *arXiv preprint*. 2024. <https://arxiv.org/abs/2403.02522>
367. Walser S, Siegel C, Winter M, Feiertag G, Loibl M, Leidl A. Sensitivity recalibration of MEMS microphones to compensate drift and environmental influences. *Procedia Engineering* 168:1759-1762. 2016. <https://doi.org/10.1016/j.proeng.2016.11.508>
368. Cesareo A, Nido SA, Biffi E, Gandossini S, D'Angelo MG, Aliverti A. A wearable device for breathing frequency monitoring: A pilot study on patients with muscular dystrophy. *Sensors*. 20(18):5121. 2020. <https://doi.org/10.3390/s20185121>
369. Kraman SS, Pasterkamp H, Wodicka GR. Lung sounds: an acoustic primer for clinicians. *Chest*. 163(6):1519-1528. 2023. <https://doi.org/10.1016/j.chest.2023.01.020>
370. Gavriely N, Palti Y, Alroy G. Spectral characteristics of normal breath sounds. *J Appl Physiol*. 50(2):307-314. 1981. <https://doi.org/10.1152/jappl.1981.50.2.307>
371. Pramono RXA, Bowyer S, Rodriguez-Villegas E.. Automatic adventitious respiratory sound analysis: A systematic review. *IEEE Journal of Biomedical and Health Informatics*. 2017. <https://doi.org/10.1109/JBHI.2017.2678473>
372. Lane ND, Bhattacharya S, Georgiev P, Forlivesi C, Kawsar F.. DeepX: A Software Accelerator for Low-Power Deep Learning Inference on Mobile Devices. *Proceedings of the 15th ACM/IEEE International Conference on Information Processing in Sensor Networks*. 2016. <https://doi.org/10.1145/2971648.2971663>
373. Sze V, Chen YH, Yang TJ, Emer JS.. Efficient Processing of Deep Neural Networks: A Tutorial and Survey. *Proceedings of the IEEE*. 2017. <https://doi.org/10.1109/JPROC.2017.2761740>
374. Yilmaz G, Rapin M, Pessoa D, et al.. A wearable stethoscope for long-term ambulatory respiratory health monitoring. *Sensors (Basel)*. 20(18):5124. 2020. <https://doi.org/10.3390/s20185124>
375. Mukhopadhyay SC.. Wearable Sensors for Human Activity Monitoring: A Review. *IEEE Network*. 2016. <https://doi.org/10.1109/MNET.2016.7474349>

376. Charbonneau G, Ademovic E, Cheetham BMG, et al.. Standardization of computerized respiratory sound analysis: instrumentation. *European Respiratory Review*. 2000. <https://err.ersjournals.com/content/10/77/585>
377. Ferris BG Jr, Mead J, Opie LH. Partitioning of respiratory flow resistance in man. *J Appl Physiol*. 19:653-658. 1964. <https://doi.org/10.1152/jappl.1964.19.4.653>
378. Gavriely N, Cugell DW. Airflow effects on amplitude and spectral content of normal breath sounds. *J Appl Physiol*. 80(1):5-13. 1996. <https://doi.org/10.1152/jappl.1996.80.1.5>
379. Mead J. The lung's quiet zone. *N Engl J Med*. 282(23):1318-1319. 1970. <https://doi.org/10.1056/NEJM197006042822311>
380. Stocks J, Godfrey S. Nasal resistance during infancy. *Respir Physiol*. 34(2):233-246. 1978. [https://doi.org/10.1016/0034-5687\(78\)90029-9](https://doi.org/10.1016/0034-5687(78)90029-9)
381. Nam Y, Reyes BA, Chon KH.. Estimation of Respiratory Rates Using the Built-in Microphone of a Smartphone or Headset. *Annu Int Conf IEEE EMBC*. 2015. <https://doi.org/10.1109/EMBC.2015.7319655>
382. Liu Y, et al.. RespEar: respiratory rate monitoring via earbud microphones across 8 activities. *arXiv preprint; IEEE PerCom 2025*. 2024. <https://mobile-systems.cl.cam.ac.uk/papers/respear.pdf>
383. Liu Y, Zhang X, Wang Z, et al. EarMeter: Continuous respiration volume monitoring with earables. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies (Advance online manuscript)*. 2025. <https://mobile-systems.cl.cam.ac.uk/papers/earmeter25.pdf>
384. Baxter SL, Prakash A, Alford-Teaster J, et al. Prospective validation of smartphone-based heart rate and respiratory rate measurement algorithms. *Communications Medicine*. 2:40. 2022. <https://doi.org/10.1038/s43856-022-00102-x>
385. ResApp SleepCheck. ResApp SleepCheckRx — FDA 510(k) Clearance K213360. *U.S. FDA 510(k) Database*. 2022. <https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfpmn/pmnmn.cfm?ID=K213360>
386. ResApp SleepCheckRx OSA screening study (~220 patients, 89% sensitivity AHI≥15, 77.6% specificity). https://www.accessdata.fda.gov/cdrh_docs/pdf21/K213360.pdf
387. Sound Life Sciences. Sound Life Sciences sonar respiratory rate app — FDA 510(k) Clearance K211387. *U.S. FDA 510(k) Database*. 2021. <https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfpmn/pmnmn.cfm?ID=K211387>
388. U.S. Food and Drug Administration. AcuPebble SA100 – 510(k) premarket notification (K210480). *Center for Devices and Radiological Health*. 2021. https://www.accessdata.fda.gov/cdrh_docs/pdf21/K210480.pdf
389. Strados Labs Inc.. LA-00055 Impact Paper: Remotely Capturing Respiratory Digital Biomarkers for Better Health (RESP white paper). *Strados Labs research publication*. 2022. <https://stradoslabs.com/>
390. U.S. Food and Drug Administration. Strados Remote Electronic Stethoscope Platform (RESP): 510(k) clearances K201077 and K220893. *Center for Devices and Radiological Health*. 2022. <https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfpmn/pmnmn.cfm?ID=K201077>
391. Au YK, Muqem T, Kroh J, Powers R, Glass M. The Strados remote electronic stethoscope platform - a novel wearable remote auscultation system designed to collect longitudinal lung sounds with equivalent fidelity as FDA-cleared stethoscopes. *Am J Respir Crit Care Med*. 201:A3249. 2020. https://doi.org/10.1164/ajrccm-conference.2020.2011_MeetingAbstracts.A3249
392. Au YK, Muqem T, Fauveau VJ, Cardenas JA, Geris BS, Hassen GW, Glass M. Continuous monitoring versus intermittent auscultation of wheezes in patients presenting with acute respiratory distress. *Am J Respir Crit Care Med*. 201(Suppl 1):A7570. 2020. https://doi.org/10.1164/ajrccm-conference.2020.2011_MeetingAbstracts.A7570
393. Au YK, Fauveau V, Muqem T, Glasser P, Glass M. Remote auscultation at home after asthma exacerbation: a case from discharge to readmission. *Am J Respir Crit Care Med*. 201:A1309. 2020. https://doi.org/10.1164/ajrccm-conference.2020.2011_MeetingAbstracts.A1309
394. Capp N, Fauveau V, Au YK, Glasser P, Muqem T, Hassen G. Characterize clinically validated respiratory system information using a non-invasive bio-sensor. *IEEE Signal Processing in Medicine and Biology Symposium (SPMB)*. 2019. <https://doi.org/10.1109/SPMB47826.2019.9037800>
395. Dos Reis GS, et al.. Pulmonary function estimation using smartphone audio and deep learning. *J Bras Pneumol*. 2025. 2025. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC12602046/>

396. Graham BL, Steenbruggen I, Miller MR, et al.. Standardization of spirometry 2019 update: an official ATS and ERS technical statement. *Am J Respir Crit Care Med*. 200:e70-e88. 2019. <https://doi.org/10.1164/rccm.201908-1590ST>
397. Sharan RV, Abeyratne UR, Swarnkar VR, Claxton S, Porter P. Predicting spirometry readings using cough sound features and regression. *Physiol Meas*. 39(9):095001. 2018. <https://doi.org/10.1088/1361-6579/aad948>
398. Porter P, Abeyratne U, Swarnkar V, et al.. A prospective multicentre study testing the diagnostic accuracy of an automated cough sound centred analytic system for the identification of common respiratory disorders in children. *Respir Res*. 20(1):81. 2019. <https://doi.org/10.1186/s12931-019-1046-6>
399. Chen H, Zhang H, Wang L, Yu T, Wang H. Triple-classification of respiratory sounds using optimized S-transform and deep residual networks. *IEEE Access* 2022. <https://ieeexplore.ieee.org/document/8663379>
400. Kim BJ, Mun JH, Hwang DH, Suh DI, Lim C, Kim K. An explainable and accurate transformer-based deep learning model for wheeze classification utilizing real-world pediatric data. *Scientific Reports*. 15(1):5656. 2025. <https://doi.org/10.1038/s41598-025-89533-9>
401. Bent B, Goldstein BA, Kibbe WA, Dunn JP. Investigating sources of inaccuracy in wearable optical heart rate sensors. *NPJ Digit Med*. 3:18. 2020. <https://doi.org/10.1038/s41746-020-0226-6>
402. Perez MV, Mahaffey KW, Hedlin H, et al.. Large-scale assessment of a smartwatch to identify atrial fibrillation. *N Engl J Med*. 381(20):1909-1917. 2019. <https://doi.org/10.1056/NEJMoa1901183>
403. Perino AC, Gummidipundi SE, Lee J, et al.. Arrhythmias other than atrial fibrillation in those with an irregular pulse detected with a smartwatch: findings from the Apple Heart Study. *Circ Arrhythm Electrophysiol*. 14(10):e009941. 2021. <https://doi.org/10.1161/CIRCEP.121.009941>
404. Chan YY, Wang P, Rogers L, et al.. The Asthma Mobile Health Study, a large-scale clinical observational study using ResearchKit. *Nat Biotechnol*. 35(4):354-362. 2017. <https://doi.org/10.1038/nbt.3826>
405. Anonymous (WEAICOR study). Wearable compliance in cardiovascular monitoring (WEAICOR). *PMC10131852*. 2023. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC10131852/>
406. Anonymous (chest strap discomfort review). Chest strap monitoring limitations review. *PMC11522764*. 2024. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC11522764/>
407. Rusz J, Hlavnička J, Tykalová T, Novotný M, Dušek P, Šonka K, Růžička E. Smartphone allows capture of speech abnormalities associated with high risk of developing Parkinson's disease. *IEEE Trans Neural Syst Rehabil Eng* 26(8):1495-1507. 2018. <https://doi.org/10.1109/TNSRE.2018.2851787>
408. Gao W, et al.. EBCare wearable face mask for exhaled breath condensate sensing. *Science*. 2024. <https://doi.org/10.1126/science.adn6471>
409. Alveos Labs. Alveos One: World's First Acoustic Breathing Monitor (product white paper / Kickstarter materials). *alveoslabs.com*. 2025. <https://alveoslabs.com>
410. International Medical Device Regulators Forum (IMDRF) Software as a Medical Device Working Group. Software as a Medical Device: Possible Framework for Risk Categorization and Corresponding Considerations. *IMDRF/SaMD WG/N12 FINAL:2014, Sept 18, 2014*. 2014. <https://www.imdrf.org/sites/default/files/docs/imdrf/final/technical/imdrf-tech-140918-samd-framework-risk-categorization-141013.pdf>
411. U.S. Food and Drug Administration. Clinical Decision Support Software: Guidance for Industry and Food and Drug Administration Staff (Final Guidance). *FDA, Sept 28, 2022 (superseded by Jan 6, 2026 update)*. 2022. <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/clinical-decision-support-software>
412. Friedman CP. A 'fundamental theorem' of biomedical informatics. *J Am Med Inform Assoc* 16(2):169-170. 2009. <https://doi.org/10.1197/jamia.M3092>
413. Roche Group. Roche to acquire Flatiron Health to accelerate industry-wide development and delivery of breakthrough medicines for patients with cancer [Press release]. *Roche Media Release*. 2018. <https://www.roche.com/media/releases/med-cor-2018-02-15>
414. Tempus AI Inc. Form 424B4 prospectus (Registration No. 333-279558). *U.S. Securities and Exchange Commission EDGAR database*. 2024. <https://www.sec.gov/Archives/edgar/data/1717115/000119312524161989/d221145d424b4.htm>

415. Miller AC, Singh I, Koehler E, Polgreen PM. A smartphone-driven thermometer application for real-time population- and individual-level influenza surveillance (Kinsa Health). *Clin Infect Dis*. 2018;67(3):388-397. doi:10.1093/cid/ciy073. <https://doi.org/10.1093/cid/ciy073>
416. Apple Inc.. Empowering people to live a healthier day (Apple Health Report, July 2022). *Apple Inc.*. 2022. <https://www.apple.com/healthcare/>
417. Le TN, Straatman LV, Lea J, Westerberg B. Current insights in noise-induced hearing loss: a literature review of the underlying mechanism, pathophysiology, asymmetry, and management options. *J Otolaryngol Head Neck Surg*. 46(1):41. 2017. <https://doi.org/10.1186/s40463-017-0219-x>
418. Chandrasekaran R, Katthula V, Moustakas E. Patterns of use and key predictors for the use of wearable health care devices by US adults: insights from a national survey. *J Med Internet Res*. 22(10):e22443. 2020. <https://doi.org/10.2196/22443>
419. Tomasic I, Tomasic N, Trobec R, Krpan M, Kelava T. Continuous remote monitoring of COPD patients—justification and explanation of the requirements and a survey of the available technologies. *Med Biol Eng Comput*. 56:547-569. 2018. <https://doi.org/10.1007/s11517-018-1798-z>
420. Beard JW, Sethi A, Jiao W, Hyatt HW, Yapici HO, Ersilon M, Overdyk FJ. Cost savings through continuous vital sign monitoring in the medical-surgical unit. *Journal of Medical Economics*. 26(1):760-768. 2023. <https://doi.org/10.1080/13696998.2023.2219156>
421. Slight SP, Franz C, Olugbile M, et al.. 5-year ROI of continuous vital sign monitoring. *Critical Care Medicine*. 2014. <https://doi.org/10.1097/CCM.0000000000000392>
422. Wollny K, Beck CA, Lethebe BC, Rygiel JH, Cuthbertson BH, Raj SR, Rabi D, Roberts CJ. Continuous monitoring of respiratory rate with wearable sensor in patients admitted to hospital with pneumonia compared with intermittent nurse-led monitoring in the United Kingdom: A cost-utility analysis. *PharmacoEconomics – Open*. 6(1):73-83. 2022. <https://doi.org/10.1007/s41669-021-00290-7>
423. CMS / VivoCare Solutions. CMS 2026 Final Rule: New RPM CPT Codes (99445 & 99470) and updated billing rules. *VivoCare Solutions blog*. 2026. <https://vivocaresolutions.com/blog/2026-pfs-new-rpm-codes-99445-99470/>
424. Bayoumy K, Gaber M, Elshafeey A, Mhaimeed O, Dineen EH, Marvel FA, et al.. Smart wearable devices in cardiovascular care: where we are and how to move forward. *Nat Rev Cardiol* 18(8):581-599. 2021. <https://doi.org/10.1038/s41569-021-00522-7>
425. Tison GH, Sanchez JM, Ballinger B, Singh A, Olgin JE, Pletcher MJ, et al.. Passive detection of atrial fibrillation using a commercially available smartwatch. *JAMA Cardiol* 3(5):409-416. 2018. <https://doi.org/10.1001/jamacardio.2018.0136>
426. ResApp Health Ltd. Proposed acquisition of ResApp Health by Pfizer Australia Holdings Pty Limited [Market announcement]. *Australian Securities Exchange (ASX)*. 2022. <https://cdn-api.markitdigital.com/apiman-gateway/ASX/asx-research/1.0/file/2924-02509234-6A1086213>
427. Clario. Clario announces acquisition of AI-powered software company ArtiQ [Press release]. *Clario Newsroom*. 2024. <https://clario.com/about/newsroom/clario-announces-acquisition-of-ai-powered-software-company-artiq/>
428. (Alveos Labs team). Alveos Labs publication in The Lancet Respiratory Medicine. *Lancet Respir Med*. 2025. 2025. [https://www.thelancet.com/journals/lanres/article/PIIS2213-2600\(25\)00437-0/fulltext](https://www.thelancet.com/journals/lanres/article/PIIS2213-2600(25)00437-0/fulltext)
429. Dameff C, Clay B, Longhurst CA. Personal health records: more promising in the smartphone era?. *JAMA*. 321(4):339-340. 2019. <https://doi.org/10.1001/jama.2018.20434>
430. Neprash HT, Everhart A, McAlpine D, Smith LB, Sheridan B, Cross DA. Measuring primary care exam length using electronic health record data. *Med Care*. 59(1):62-66. 2021. <https://doi.org/10.1097/MLR.0000000000001450>
431. Tai-Seale M, Downing NL, Jones VG, et al.. Technology-enabled consumer engagement: promising practices at four health care delivery organizations. *Health Aff (Millwood)*. 38(3):383-390. 2019. <https://doi.org/10.1377/hlthaff.2018.05027>
432. Marvel FA, Spaulding EM, Lee MA, et al.. Digital health intervention in acute myocardial infarction. *Circ Cardiovasc Qual Outcomes*. 14(7):e007741. 2021. <https://doi.org/10.1161/CIRCOUTCOMES.120.007741>
433. Milani RV, Lavie CJ, Bober RM, Milani AR, Ventura HO. Improving hypertension control and patient engagement using digital tools. *Am J Med*. 130(1):14-20. 2017. <https://doi.org/10.1016/j.amjmed.2016.07.029>

434. Gujral K, Van Campen J, Jacobs J, Kimerling R, Blonigen D, Zulman DM. Mental health service use, suicide behavior, and emergency department visits among rural US veterans who received video-enabled tablets during the COVID-19 pandemic. *JAMA Netw Open*. 5(4):e226250. 2022. <https://doi.org/10.1001/jamanetworkopen.2022.6250>
435. Garne Holm K, Brødsgaard A, Zachariassen G, Smith AC, Clemensen J. Parent perspectives of neonatal tele-homecare: a qualitative study. *J Telemed Telecare*. 25(4):221-229. 2019. <https://doi.org/10.1177/1357633X18765059>
436. Rasmussen MK, Clemensen J, Zachariassen G, et al.. Cost analysis of neonatal tele-homecare for preterm infants compared to hospital-based care. *J Telemed Telecare*. 26(7-8):474-481. 2020. <https://doi.org/10.1177/1357633X19843753>
437. Koziatek CA, Rubin A, Lakdawala V, et al.. Decreasing the lag between result availability and decision-making in the emergency department using push notifications. *West J Emerg Med*. 20(4):666-671. 2019. <https://doi.org/10.5811/westjem.2019.4.42749>
438. Seiden SC, McMullan S, Sequera-Ramos L, et al.. Tablet-based interactive distraction (TBID) vs oral midazolam to minimize perioperative anxiety in pediatric patients: a noninferiority randomized trial. *Paediatr Anaesth*. 24(12):1217-1223. 2014. <https://doi.org/10.1111/pan.12475>
439. Bradley K, Godin M. Think different: reimagining clinical practice and professional development by collaborating with the Apple higher education team. *Nurse Leader*. 18(1):73-77. 2020. <https://doi.org/10.1016/j.mnl.2019.07.016>
440. Wymant C, Ferretti L, Tsallis D, et al.. The epidemiological impact of the NHS COVID-19 app. *Nature*. 594:408-412. 2021. <https://doi.org/10.1038/s41586-021-03606-z>
441. Abueg M, Hinch R, Wu N, et al.. Modeling the effect of exposure notification and non-pharmaceutical interventions on COVID-19 transmission in Washington state. *NPJ Digit Med*. 4(1):49. 2021. <https://doi.org/10.1038/s41746-021-00422-7>
442. Chen E, Facio FM, Aradhya KW, et al.. Rates and classification of variants of uncertain significance in hereditary disease genetic testing. *JAMA Netw Open*. 6(10):e2339571. 2023. <https://doi.org/10.1001/jamanetworkopen.2023.39571>
443. Nykamp K, Anderson M, Powers M, et al.. Sherlock: a comprehensive refinement of the ACMG-AMP variant classification criteria. *Genet Med*. 19(10):1105-1117. 2017. <https://doi.org/10.1038/gim.2017.37>
444. Frazer J, Notin P, Dias M, et al.. Disease variant prediction with deep generative models of evolutionary data. *Nature*. 599(7883):91-95. 2021. <https://doi.org/10.1038/s41586-021-04043-8>
445. Scott A, Hernandez F, Chamberlin A, Smith C, Karam R, Kitzman JO. Saturation-scale functional evidence supports clinical variant interpretation in Lynch syndrome. *Genome Biol*. 23(1):266. 2022. <https://doi.org/10.1186/s13059-022-02839-z>
446. Invitae / Labcorp. Clinical Variant Modeling: a Bayesian approach for incorporating clinical evidence at scale (WP135). *Invitae / Labcorp*. 2026. <https://invitae.com/>
447. Kastrinos F, Ingram MA, Silver ER, et al.. Gene-specific variation in colorectal cancer surveillance strategies for Lynch syndrome. *Gastroenterology*. 161(2):453-462.e15. 2021. <https://doi.org/10.1053/j.gastro.2021.04.010>
448. Streeten EA, Sanchez B, Lozano-Chinga M, et al.. KCNQ1 and Long QT Syndrome in 1/45 Amish: the road from identification to implementation of culturally appropriate precision medicine. *Circ Genom Precis Med*. 13(6):e003133. 2020. <https://doi.org/10.1161/CIRCGEN.120.003133>
449. Salinas V, Vega P, Marsili L, et al.. The odyssey of complex neurogenetic disorders: from undetermined to positive. *Am J Med Genet C Semin Med Genet*. 184(4):876-884. 2020. <https://doi.org/10.1002/ajmg.c.31848>
450. Tóke J, Cziráj G, Bezzegh A, et al.. Rare diseases caused by abnormal calcium sensing and signalling. *Endocrine*. 71(3):611-617. 2021. <https://doi.org/10.1007/s12020-021-02620-5>
451. Ekong R, Nellist M, Hoogeveen-Westerveld M, et al.. Variants within TSC2 exons 25 and 31 are very unlikely to cause clinically diagnosable tuberous sclerosis. *Hum Mutat*. 37(4):364-370. 2016. <https://doi.org/10.1002/humu.22951>
452. Ding D, Lawson KD, Kolbe-Alexander TL, et al.. The economic burden of physical inactivity: a global analysis of major non-communicable diseases. *Lancet*. 388(10051):1311-1324. 2016. [https://doi.org/10.1016/S0140-6736\(16\)30383-X](https://doi.org/10.1016/S0140-6736(16)30383-X)
453. Baille G, De Jesus AM, Perez T, Devos D, Dujardin K, Charley CM, et al.. Respiratory dysfunction in Parkinson's disease: a narrative review. *ERJ Open Res* 6(4):00165-2020. 2020. <https://doi.org/10.1183/23120541.00165-2020>
454. McMahon CN, Foster J, Phillips C, Ellmers TJ, Tagaris S, Walford-Wright G, et al.. A systematic review and meta-analysis of respiratory dysfunction in Parkinson's disease. *Eur J Neurol* 30(5):1481-1504. 2023. <https://doi.org/10.1111/ene.15743>

455. Rusz J, Cmejla R, Ruzickova H, Ruzicka E. Quantitative acoustic measurements for characterization of speech and voice disorders in early untreated Parkinson's disease. *J Acoust Soc Am* 129(1):350-367. 2011. <https://doi.org/10.1121/1.3514381>
456. Rusz J, Hlavnička J, Novotný M, Tykalová T, Pelletier A, Montplaisir J, et al.. Speech biomarkers in rapid eye movement sleep behavior disorder and Parkinson disease. *Ann Neurol* 90(1):62-75. 2021. <https://doi.org/10.1002/ana.26085>
457. Quintana DS, Heathers JAJ. Considerations in the assessment of heart rate variability in biobehavioral research. *Front Psychol* 5:805. 2014. <https://doi.org/10.3389/fpsyg.2014.00805>
458. Seely AJE, Bravi A, Herry C, Green G, Longtin A, Ramsay T, et al.. Do heart and respiratory rate variability improve prediction of extubation outcomes in critically ill patients?. *Crit Care* 18(2):R65. 2014. <https://doi.org/10.1186/cc13822>
459. Tate RR, Perdices M, Rosenkoetter U, et al.. The Single-Case Reporting Guideline In BEhavioural interventions (SCRIBE) 2016. *EQUATOR Network, 2016*. 2016. <https://www.equator-network.org/reporting-guidelines/scribe/>
460. Shamseer L, Sampson M, Bukutu C, Schmid CH, Nikles J, Tate R, et al.. CONSORT extension for reporting N-of-1 trials (CENT) 2015: explanation and elaboration. *BMJ* 350:h1793. 2015. <https://doi.org/10.1136/bmj.h1793>
461. Vohra S, Shamseer L, Sampson M, Bukutu C, Schmid CH, Tate R, et al.. CONSORT extension for reporting N-of-1 trials (CENT) 2015 statement. *BMJ* 350:h1738. 2015. <https://doi.org/10.1136/bmj.h1738>
462. Banister EW. Modeling elite athletic performance. In: Green H, McDougall JD, Wenger HA (Eds.), *Physiological testing of the high-performance athlete*. Champaign, IL: Human Kinetics. 1991.
463. Foster C. Monitoring training in athletes with reference to overtraining syndrome. *Medicine & Science in Sports & Exercise*. 30(7):1164-1168. 1998. <https://doi.org/10.1097/00005768-199807000-00023>
464. Uth N, Sørensen H, Overgaard K, Pedersen PK. Estimation of VO₂max from the ratio between HR_{max} and HR_{rest} – The heart rate ratio method. *European Journal of Applied Physiology*. 91(1):111-115. 2004. <https://doi.org/10.1007/s00421-003-0988-y>
465. Lillie EO, Patay B, Diamant J, Issell B, Topol EJ, Schork NJ. The n-of-1 clinical trial: the ultimate strategy for individualizing medicine?. *Per Med* 8(2):161-173. 2011. <https://doi.org/10.2217/pme.11.7>
466. Davidson KW, Silverstein M, Cheung K, Paluch RA, Epstein LH. Experimental designs to optimize treatments for individuals: personalized N-of-1 trials. *JAMA Pediatr* 175(4):404-409. 2021. <https://doi.org/10.1001/jamapediatrics.2020.5801>
467. Selker HP, Dulko D, Greenblatt DJ, Palm M, Trinquart L. The use of N-of-1 trials to generate real-world evidence for optimal treatment of individuals and populations. *J Clin Transl Sci* 7(1):e203. 2023. <https://doi.org/10.1017/cts.2023.604>
468. Ruinskiy D, Lavner Y. An effective algorithm for automatic detection and exact demarcation of breath sounds in speech and song signals. *IEEE Trans Audio Speech Lang Process* 15(3):838-850. 2007. <https://doi.org/10.1109/TASL.2006.889750>
469. Luo Y, Mesgarani N. Conv-TasNet: surpassing ideal time-frequency magnitude masking for speech separation. *IEEE/ACM Trans Audio Speech Lang Process* 27(8):1256-1266. 2019. <https://doi.org/10.1109/TASLP.2019.2915167>
470. Park TJ, Kanda N, Dimitriadis D, Han KJ, Watanabe S, Narayanan S. A review of speaker diarization: Recent advances with deep learning. *Comput Speech Lang* 72:101317. 2022. <https://doi.org/10.1016/j.csl.2021.101317>
471. Ginsburg GS, Picard RW, Friend SH. Key issues as wearable digital health technologies enter clinical care. *N Engl J Med* 390(12):1118-1127. 2024. <https://doi.org/10.1056/NEJMra2307160>
472. European Parliament, Council of the European Union. Regulation (EU) 2016/679 of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation). *Official Journal of the European Union, L119, 1-88*. 2016. <https://gdpr-info.eu/>
473. Illinois General Assembly. Biometric Information Privacy Act, 740 ILCS 14. *Illinois Statutes*. 2008. <https://www.ilga.gov/Legislation/ILCS/Articles?ActID=3004&ChapterID=5>
474. United States District Court, Northern District of California. In re Facebook Biometric Information Privacy Litigation, No. 3:15-cv-03747-JD (N.D. Cal. Feb. 26, 2021) (order re final approval, attorneys' fees and costs, and incentive awards). *Court order*. 2021. https://www.govinfo.gov/content/pkg/USCOURTS-cand-3_15-cv-03747/pdf/USCOURTS-cand-3_15-cv-03747-16.pdf

475. U.S. Department of Health and Human Services. Health Insurance Portability and Accountability Act (HIPAA): Privacy and Security Rules, 45 C.F.R. Parts 160, 162, and 164. *Federal regulation*. 1996. <https://www.hhs.gov/hipaa/for-professionals/privacy/index.html>
476. Nautsch A, Jiménez A, Treiber A, Kolberg J, Jasserand C, Kindt E, et al.. Preserving privacy in speaker and speech characterisation. *Comput Speech Lang* 58:441-480. 2019. <https://doi.org/10.1016/j.csl.2019.06.001>
477. Tomashenko N, Wang X, Vincent E, Patino J, Srivastava BML, Noé PG, et al.. The VoicePrivacy 2020 Challenge: Results and findings. *Comput Speech Lang* 74:101362. 2022. <https://doi.org/10.1016/j.csl.2022.101362>
478. Rieke N, Hancox J, Li W, Milletari F, Roth HR, Albarqouni S, et al.. The future of digital health with federated learning. *NPJ Digit Med* 3:119. 2020. <https://doi.org/10.1038/s41746-020-00323-1>
479. Pahar M, Klopper M, Warren R, Niesler T. COVID-19 cough classification using machine learning and global smartphone recordings. *Comput Biol Med*. 135:104572 (representative ResApp COVID screening study). 2022. <https://doi.org/10.1016/j.combiomed.2021.104572>
480. FaceHeart Inc. FH Vitals SDK-RR. *FDA 510(k) Clearance K243966*. 2025 Apr 9. <https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfpmn/pmn.cfm?ID=K243966>
481. European Parliament, Council of the European Union. Regulation (EU) 2017/745 of 5 April 2017 on medical devices, amending Directive 2001/83/EC, Regulation (EC) No 178/2002 and Regulation (EC) No 1223/2009 and repealing Council Directives 90/385/EEC and 93/42/EEC (Annex VIII, Rule 11). *Official Journal of the European Union*, L117, 1-175. 2017. <https://www.legislation.gov.uk/eur/2017/745/annex/VIII>
482. Price WN 2nd, Cohen IG. Privacy in the age of medical big data. *Nat Med* 25(1):37-43. 2019. <https://doi.org/10.1038/s41591-018-0272-7>