

The Body as a Longitudinal Diagnostic Record

ABSURD Health Institute · Volume Two · Research Essay 51

Research series: 08 · Signals, measurement and responsible medical AI

Abstract

A single photograph or laboratory result is a snapshot of physiology; a sequence of observations reveals change. This paper advances a longitudinal phenotype model in which images, symptoms, movement patterns and biometrics are interpreted relative to a person's baseline. It argues that detecting change may sometimes be more useful than comparing an individual with one fixed population threshold—provided that noise, bias and false alarms are explicitly modeled.

The record between appointments

A patient's trajectory contains diagnostic information that one isolated laboratory result cannot preserve.

Human appearance varies with lighting, hydration, sleep, camera perspective, skin pigmentation, medications and ordinary aging. A facial color change cannot be interpreted as jaundice without controlling image capture and confirming clinical signs. Wearable signals vary by sensor placement and activity. A longitudinal model should therefore estimate the smallest reliable change for each modality before claiming that a new observation reflects disease. Preprocessing, calibration and consistent acquisition are central to causal credibility.

The conceptual architecture separates data quality, baseline estimation, deviation detection and clinical interpretation. A model might flag a new tremor pattern or sustained resting heart-rate shift, but it must distinguish benign variability from potentially concerning change. Multimodal agreement can improve confidence in some settings, yet correlated measurement errors can create false reassurance. Machine learning should support triage rather than pretend to know an underlying diagnosis from nonspecific features.

Why within-person trajectories differ from snapshots

A body generates signals at different temporal resolutions. Skin changes may evolve over weeks, resting heart rate can shift over hours, and movement patterns vary by task and medication timing. A longitudinal health record should therefore preserve the time scale and measurement reliability of each modality. Combining every feature into a single daily health score risks losing information and treating correlated artifacts as independent evidence. For example, travel can simultaneously change sleep, temperature, resting heart rate and lighting in photographs without producing a new disease. A model that attributes all changes to a systemic pathology would be sophisticated in appearance but wrong in practice.

A clinical visit captures a temporary state, whereas a longitudinal record contains a sequence of exposures, symptoms, recovery patterns and physiological measurements. Temporal relationships can be clinically useful: a person's typical resting heart rate may be more informative than a population-wide threshold when it changes abruptly during illness. Yet self-report, device calibration and seasonal or behavioral variation all create noise. The body is not a perfect recording instrument; interpretation requires a history of how the measurements were produced.

A clinically meaningful timeline distinguishes symptom onset, exposure, treatment and measurement dates. Missing a week of data during the period of greatest illness may bias an apparently smooth trend. Medication changes can shift multiple metrics simultaneously, and device replacement can create artificial discontinuities. A record that documents uncertainty, missingness and competing triggers is more useful than a graph designed to look continuous regardless of evidence.

A gradual decline in gait, sleep, heart-rate recovery or appetite may be clinically informative even when each reading remains inside a broad reference interval. Yet people also vary normally across days, and frequent measurements generate noise and false positives. The mechanistic point is about within-person change, temporal relationships and recovery after perturbation. A signal becomes useful when it is reliable, interpretable and linked to a clinical question, not simply because it is continuously collected.

Testing earlier clinical recognition

The research task is to show whether access to longitudinal signals changes management or reduces delayed diagnosis. A model that reconstructs what happened after clinicians already know the outcome has not demonstrated forecasting ability. Pre-registered alert rules and prospective blinded validation can test whether early changes truly predict future deterioration. Privacy, data portability and clear consent are essential when the record includes photographs, psychiatric symptoms and other sensitive information.

Chronology can detect otherwise invisible transitions. A resting pulse that gradually increases over several days might appear normal on any one day but abnormal relative to an individual's prior pattern. This is one reason continuous observation may contribute to early warning. Yet changes can also reflect travel, exercise, caffeine, sensor wear and medication shifts. Without recording context, a visually striking graph can encourage false causal certainty.

The baseline has to be earned. A stable reference period should include sufficient observation and an assessment of ordinary weekly variation. A single week is not necessarily representative of a person with rotating shifts or recurrent illness. Longitudinal systems need to document missing data and recalibration after sensor or software changes. Otherwise apparent physiological shifts may be changes in the measuring instrument.

The outcome study. A prospective system could test whether providing contextualized trend reports to clinicians improves appropriate intervention compared with usual records. Researchers should count both useful actions and harms from excessive alerts. The record's fundamental value may be better reconstruction of exposures and recovery rather than automated disease prediction. Both are legitimate outcomes if they are described accurately and supported by evidence.

Compare longitudinal models with visit-based assessment for prospectively adjudicated conditions. Control for baseline instability, device changes and missing data. The model should improve sensitivity at an acceptable false-alarm burden; a patient-specific threshold that alarms every week is not evidence of early detection.

A scientifically serious system should define baseline stability, thresholds for sustained change and a strategy for periods when data quality is poor. It should also state whether a flagged deviation is merely interesting, requires routine clinical follow-up or warrants urgent care. The difference among those three actions is more important than an attractive dashboard. External prospective validation remains necessary.

Change within one patient can be more informative than distance from a population mean

A person's resting heart rate may increase gradually while remaining inside a published reference interval, or their walking speed may decline long before a specialist visit. A longitudinal model can potentially recognize that departure from the person's own stable baseline. The difficulty is distinguishing a pathological trajectory from changes in sleep, travel, season, device firmware or training status. An algorithm should therefore estimate measurement reliability and quantify whether the pattern persists beyond normal variation. The paper needs a real endpoint, such as independently confirmed infection or functional decline, and a specified response to an alert. Merely providing a colorful history graph is not diagnostic progress. Comparisons should include the performance of a simple fixed threshold, a clinician reviewing the same data and a sophisticated model. If added complexity produces earlier recognition without an unacceptable surge in false alarms, the argument for continuous physiology becomes practical rather than rhetorical.

Frequency without reliability creates noise

The longitudinal-record approach can also mislead when it transforms every coincident change into a causal story. A patient's sleep may worsen after travel, heart rate may increase after a medication adjustment, and appetite may decline during infection. Several variables moving together do not reveal which one initiated the sequence. Recording exact dates of exposures and symptoms, together with missing observations, can improve inference. A useful system should separate confirmed measurements from self-reported recollections and mark uncertainty rather than filling gaps with narrative. Clinical value must be demonstrated prospectively: does the record permit earlier correct assessment, better coordination or more accurate monitoring of recovery than ordinary care? If the main benefit is improved communication, that is still worthwhile. The dangerous step is promoting a well-designed timeline into an autonomous diagnostic verdict without external validation and evidence of net benefit.

The opposite hypothesis is that longitudinal accumulation mainly amplifies artifacts and spurious correlations. That concern is real: repeated monitoring increases the number of opportunities for a false-positive flag and can worsen anxiety. Good validation therefore requires longitudinal test-retest reliability, external cohorts, clinically meaningful outcomes and evidence that the alert improves decisions.

This paper supplies a conceptual foundation for Signals as a research direction, not evidence that the current product can diagnose or predict any specific disease. Public claims should be limited to functions that have been independently validated. The promise is disciplined tracking of change, not medical clairvoyance.

Sources and research leads

1. [Quer et al. Wearable sensor data and self-reported symptoms for COVID-19 detection](#)
2. [Esteva et al. Dermatologist-level classification of skin cancer with deep neural networks](#)

Editorial status: Third-pass conceptual research manuscript. The argument has been individually strengthened and shared boilerplate removed. Original reference leads were retained; a complete independent source-by-source scholarly verification has not been performed. No new clinical trial data are presented.