

PSYCHE

WORKING PAPER · JUNE 2026

The Emotional Layer

A research agenda for human state models: why physiology is the last missing modality, how the interpretive layer is built, and what it makes possible.

Today's AI can reason, write, and predict.

It cannot feel a racing heart.

It cannot hear the strain under a steady voice.

It cannot see the stress that never makes it into words.

AI has no body. That is the gap no model can close.

Consider a moment that happens millions of times a day. Someone is asked how they are doing. "I'm fine. Just a busy week." The voice is steady. The face is composed. Every signal the person controls says calm.

But heart rate variability has been suppressed for thirty-six hours. Electrodermal activity is elevated, with frequent spontaneous responses. Sleep has fragmented three nights running. Every signal the person does not control says otherwise.

Two channels. One moment. Opposite stories. Today's AI systems, and most of today's clinical instruments, can only read the first channel: the performative one. Self-report, facial expression, word choice. The channel people curate. The channel that can be faked, and routinely is, especially by the people who most need to be seen accurately.

The second channel is physiological. It happens inside the body, below the level of performance. It is the more honest channel, and it is almost entirely invisible to the systems that increasingly mediate human life: telehealth platforms, voice assistants, support systems, AI companions.

This is not a software gap. It is a measurement gap. And measurement gaps are not closed by better reasoning. They are closed by better instruments.

Language is solved. Vision is solved. The body is not.

The last decade of AI progress is a story of modalities falling in sequence. Text, then images, then audio, then video. Each fell when two things converged: large paired datasets and architectures that could learn from them. What remains unmodeled is the signal generated continuously by every human body: cardiac rhythm, skin conductance, vocal micro-dynamics, sleep architecture, movement.

DEFINITION

Human state models are systems that infer emotional, cognitive, and physiological state from signals generated continuously by the body: signals less vulnerable to performance, self-report, and recall bias than any other data source available to AI.

The architecture of the problem is two-sided. On one side, the stimulus: what is happening to a person. On the other, the response: what is happening inside them. Existing AI sees only the stimulus side. The interpretive layer is the mapping between the two.

STIMULUS CHANNELS

What is happening to the person

- Voice and speech
- Text and language
- Video and facial expression
- Audio environment

RESPONSE CHANNELS

What is happening inside the person

- Cardiac and heart rate variability
- Electrodermal activity
- Vocal micro-prosody
- Movement and gaze patterns

The sensors arrived before the meaning did.

Hundreds of millions of people now wear devices that record heart rate variability, skin temperature, and sleep architecture every night. Voice, the one sensor everyone already carries, is recorded in every telehealth visit, support call, and clinical encounter. The raw signal exists at population scale. What does not exist is the layer that turns it into meaning.

Wearables produce data without interpretation. A recovery score is not an emotional state. A sleep graph is not an answer to the question a clinician actually asks: how is this person, really? The infrastructure now exists to study human state at unprecedented scale. The interpretation does not. That asymmetry is the opportunity.

Three forces make this the moment. First, model capability is commoditizing. Reasoning improves every quarter, for everyone, at falling cost. But no model gains a body. Capability advances are orthogonal to physiological access, which means the value migrates to whoever holds the data that connects signal to state.

Second, the field has begun paying for emotional capability. Google DeepMind's absorption of Hume AI confirmed that frontier labs consider emotion a strategic layer. What no acquirer obtained, because it does not yet exist at scale, is a paired dataset linking non-performative physiological signal to ground-truth emotional state. Emotion AI built on actors and posed expressions inherits the performative channel's weakness. Datasets built on the response channel must be created, not scraped.

Third, consumer wellness fatigue is real, and it is evidence for this thesis rather than against it. People are tired of scores, streaks, and dashboards that describe data without meaning. The unmet demand is interpretive: systems that understand, in language a person or a clinician can act on.

The signal is there. We measured it.

Our research at Carnegie Mellon's Human-Computer Interaction Institute, published as a peer-reviewed preprint on medRxiv in 2025, asked a simple question: can next-day mood states be predicted from consumer wearable signals alone, with no self-report as input?

~70%

accuracy predicting daily mood states from consumer wearable data, using elastic net regression over passively collected physiological features. No questionnaires. No diaries. Signal alone.

The most informative features were physiological and behavioral, not performative: nightly skin temperature, NREM sleep, heart rate variability, and activity patterns. The body, measured passively by a device people already own, carries enough information to anticipate how someone will feel tomorrow.

This matters for what it implies. Consumer-grade sensors, the noisiest instruments available, already carry recoverable emotional signal. Purpose-built supervision, with validated psychophysiological ground truth instead of self-report labels, raises the ceiling. That is the dataset Psyche exists to build.

STATED HONESTLY

These are early results with real limitations. The cohort is modest in size. Labels in this study were self-reported mood states, the very instrument we argue against, used here because it is the field's current standard. The lab-to-life gap in affective computing is well documented, and accuracy on a research cohort is not a deployed product.

We publish the limitations alongside the findings because a measurement company that overstates its measurements has nothing.

Four problems, in order.

01 Ground-truth supervision

Building paired datasets where physiological signal is supervised by validated psychophysiological induction protocols rather than self-report. Label quality is the ceiling on everything downstream; this is the field's binding constraint and our first investment.

02 Physiological signal modeling

Architectures for multimodal physiological time series: cardiac, electrodermal, vocal, movement. Signals that are irregular, person-specific, and context-dependent, and that demand different inductive biases than text or images.

03 Emotional representation learning

Representations of human state that transfer across people, sensors, and settings, and that resolve states the performative channel conflates: stress from depletion, calm from suppression, recovery from absence of signal.

04 Biometric rights

Frameworks ensuring individuals retain agency, ownership, and direct benefit from the emotional data they generate. The interpretive layer reads the most intimate signal a person produces; the governance is part of the architecture, not an afterthought.

The last generation of technology helped machines understand information. The next generation will help machines understand *people*.

We believe emotional intelligence will become as foundational to AI as language is today. We believe the path runs through the body, through the channel that cannot be performed, supervised by ground truth rather than recollection. And we believe the layer between physiological signal and human meaning is too important to be built carelessly, which is why we are building it in the open, with the limitations stated and the governance designed in.

Psyche exists to build that layer.

Psyche Health Inc.

Ajmail Matin, Co-Founder & CEO · ajmail@hellopsyche.com

Mayank Goel, Co-Founder & CTO · mayank@hellopsyche.com

hellopsyche.com