Self-Report Over Surveillance: The Case for User-Declared Cognitive State

Dominikos Pritis
Independent Researcher
Kumbatio / entro314.labs, Athens, Greece
dominikos@kumbat.io

Preprint — DOI: 10.5281/zenodo.21915163

Abstract

Workplace technology increasingly aims to infer employee cognitive and emotional states through biometric sensors, behavioral telemetry, and AI-driven analysis. This paper argues that for cognitive state—specifically, a user's current capacity for focused, complex work—self-report is not merely an acceptable fallback but the ethically and practically superior approach. We examine the reliability problems of biometric cognitive-state inference, the privacy and autonomy violations it entails, the power asymmetries it creates, and the specific harms it imposes on neurodivergent workers whose patterns deviate from normative baselines. We propose that systems seeking to adapt to cognitive capacity should trust user self-declaration, and we describe four design constraints—agency, transparency, non-punitive adaptation, and neutral framing—that prevent self-report mechanisms from reintroducing the coercion they replace. The position is grounded in an implemented system, energy-system, which models cognitive capacity as explicit, user-controlled application state.

Keywords: self-report, surveillance, cognitive state, biometrics, privacy, neurodivergent design, workplace monitoring

1 Introduction

A growing category of workplace technology promises to detect when employees are focused, distracted, stressed, fatigued, or disengaged. The mechanisms vary: eye-tracking cameras, keyboard and mouse dynamics, heart rate variability monitors, facial expression analysis, application usage patterns, and machine-learning models trained on behavioral telemetry (Ball, 2010; Ravid et al., 2020).

The stated purpose is benign: adapt the work environment to the worker's state. Reduce interruptions when someone is in flow. Flag burnout before it becomes a crisis. Match task difficulty to current capacity.

The unstated consequence is surveillance. These systems collect continuous biometric and behavioral data, process it through opaque models, and produce assessments of internal cognitive and emotional states that the person being assessed may not agree with, may not be aware of, and cannot override.

This paper argues that for cognitive capacity—the variable that determines what kind of work a person can meaningfully do at a given moment (Pritis, 2026a)—self-report is the correct design choice. Not because biometric inference is technically immature (though it is), but because the relationship between a system and its user should be one of trust, not observation.

2 The Reliability Problem

2.1 Weak Signal, Strong Claim

Physiological signals correlate weakly and inconsistently with subjective cognitive capacity at the individual level. Heart rate variability (HRV) is a commonly cited marker of cognitive load, but meta-analyses show that the relationship between HRV and task performance is moderated by task type, individual fitness, medication, time of day, and emotional state (Charles and Nixon, 2019). Skin conductance responds to arousal, not to the valence or cognitive quality of that arousal; a person solving a difficult problem and a person experiencing anxiety may produce similar galvanic skin response profiles.

Facial expression analysis has been extensively criticized for cultural bias, low accuracy in naturalistic settings, and the fundamental problem that facial expressions do not reliably map to internal states (Barrett et al., 2019). Keyboard and mouse dynamics can detect gross changes in motor behavior but cannot distinguish cognitive fatigue from physical discomfort, boredom, or a simple change in task type.

The core problem is that cognitive capacity is a subjective, multi-dimensional state whose variation across the day and across individuals is extensively documented (Pritis, 2026b). A person's ability to do complex work at a given moment depends on sleep quality, emotional context, medication timing, recent social interactions, task interest, ambient noise, physical comfort, and dozens of other factors. No external sensor captures this composite. The person experiencing it does.

2.2 The Baseline Problem

Biometric inference systems require normative baselines: what does "normal" look like for this person, so that deviations can be detected? This is already problematic for neurotypical populations, where individual variation in resting heart rate, typing patterns, and facial expressiveness is substantial.

For neurodivergent people, the baseline problem becomes acute. ADHD involves highly variable executive function across time and context (Barkley, 1997). A person with ADHD may type rapidly during hyperfocus and slowly during executive-function depletion, but both states can involve productive work of different kinds. A system trained on neurotypical baselines will misclassify ADHD patterns. A system trained on the individual's own patterns will still struggle, because the defining characteristic of ADHD is that patterns are inconsistent.

Depression introduces a different baseline problem. Depressive cognitive impairment fluctuates in ways that do not follow predictable cycles (Rock et al., 2014). Autistic individuals may have atypical facial expressiveness, atypical physiological arousal patterns, and communication styles that diverge from training data norms (Milton, 2012).

In each case, the person most likely to be misread by biometric inference is the person who would benefit most from accurate capacity-aware adaptation.

3 The Autonomy Problem

3.1 Who Owns the Assessment?

When a system infers your cognitive state, the assessment belongs to the system, not to you. You may not know what state the system has assigned you. You may not agree with it. You may not be able to change it.

This is a qualitative difference from self-report. When you declare "I am at low capacity right now," you retain authority over the description of your own experience. When a sensor array and a machine-learning model declare "this person is at low capacity right now," that authority has been transferred to the system.

The transfer matters because cognitive state is not like room temperature. It is not an objective quantity that a sufficiently accurate thermometer can measure correctly. It is a subjective experience that the person is the foremost authority on. Overriding that authority is not a measurement improvement. It is a power shift.

3.2 Coercion Through Inference

Even when biometric monitoring is described as voluntary, the power dynamics of employment make genuine voluntariness difficult. An employee who opts out of "wellness monitoring" signals that they have something to hide. An employee who opts in consents to continuous assessment of their internal states by systems they do not control and cannot meaningfully audit.

The problem intensifies when inferred states drive consequential decisions. If a system determines that an employee is "unfocused" and routes complex tasks away from them, the employee has been evaluated and acted upon without due process, without transparency, and without recourse.

3.3 The Neurodivergent Trap

For neurodivergent workers, biometric inference creates a specific trap. Their patterns deviate from normative baselines. A system trained on those baselines will systematically misread them. If misreadings drive adaptation (reduced task complexity, additional check-ins, flagged performance), the system amplifies the very marginalization it claims to address.

A self-report system avoids this entirely. The neurodivergent worker declares their own state. The system adapts to that declared state. No normative baseline is required. No inference model needs to accommodate atypical patterns. The user is the authority.

4 The Privacy Problem

Continuous biometric and behavioral monitoring collects intimate data about physiological and cognitive states. This data, once collected, creates multiple privacy risks:

Scope creep: Data collected for "cognitive state adaptation" can be repurposed for performance evaluation, attendance monitoring, or insurance risk assessment. The history of workplace data collection shows that beneficial framing at introduction does not prevent extractive use over time (Ball, 2010).

Inference depth: Machine learning models trained on behavioral telemetry can infer information beyond their stated purpose. Keystroke dynamics collected for one purpose can be used to identify individuals and track behavioral patterns over time (Monaco and Tappert, 2018). The data collected for one purpose becomes a surveillance substrate for unlimited inference.

Data persistence: Biometric data, once collected, is difficult to meaningfully delete. Energy patterns, stress profiles, and cognitive-state histories constitute a detailed record of a person's internal life over time.

Self-report systems minimize these risks by design. The only data stored is the user's declared state—a level, a timestamp, and a source—plus two mechanical fields (a revision counter and a random producer identifier) that order concurrent writes across tabs or devices. No biometric data is collected. No behavioral telemetry is logged.

5 The Design of Trustworthy Self-Report

Rejecting surveillance does not mean accepting any self-report mechanism. Self-report systems can themselves become coercive if poorly designed. We identify four constraints:

5.1 Agency

The user must be able to override the system at any time, with no friction and no negative consequences. The system is advisory. The user is authoritative.

5.2 Transparency

When the application adapts behavior based on declared energy state, the reason for the adaptation must be clear. An unexplained change creates anxiety and confusion. A labeled adaptation preserves trust.

5.3 Non-Punitive Adaptation

Reduced capacity must not result in reduced capability. If a low-energy declaration causes the system to hide controls or lock out features, the system punishes the user for honesty. The correct behavior is to reduce friction while preserving access to all critical functions.

This constraint addresses the "locked out when I need help most" problem. A person at low energy may need navigation tools more than a person at peak energy, precisely because they have fewer cognitive resources to find alternatives.

5.4 Neutral Framing

The language and visual design of the system must not associate low energy with failure. Energy level is context, not character. This matters especially for people with depression, for whom shame is a clinical symptom.

6 Implementation

We have implemented these principles in energy-system, a TypeScript library that treats cognitive capacity as first-class, user-controlled application state (Pritis, 2026a). State is set by explicit user interaction. The library supports a source field on each state change (manual, scheduled, inferred), making provenance transparent. No biometric data is collected, processed, or stored. Persistence stores only the declared state. All adaptation is mediated through strategies—pure functions that map declared state to behavior changes.

The scheduled and inferred sources exist for future use cases (pre-set energy curves, time-of-day suggestions) but are designed as suggestions that the user can override, not as authoritative assessments.

7 Objections

7.1 "People Will Game Self-Report"

This concern reveals the underlying assumption: that the purpose of cognitive state tracking is employer oversight, not user support. In a user-facing tool designed to help individuals work more effectively, gaming is not a meaningful risk. The person is the primary beneficiary of accurate self-report. Declaring peak energy when exhausted results in the wrong experience for the person's actual state.

7.2 "Biometric Data Is More Objective"

Objectivity is not the relevant criterion. Accuracy and authority are. An "objective" measurement that consistently misreads neurodivergent patterns is not more accurate than a self-report that captures the person's actual experience. And even if biometric inference were perfectly accurate, the question of who holds authority over the description of a person's internal state is an ethical question, not a technical one.

7.3 "Self-Report Adds Cognitive Load"

Setting a five-level energy state takes one click or keystroke. The cognitive cost is trivially low. By contrast, the cognitive cost of being continuously monitored, potentially misread, and subject to opaque automated decisions is substantial and ongoing.

8 Related Work

Experience Sampling Method (ESM) and Ecological Momentary Assessment (EMA) rely on self-report precisely because subjective experience is their object of study (Csikszentmihalyi and Larson, 2014). Affective computing has pursued automated detection since its founding

(Picard, 1997), though recent work has increasingly acknowledged the limits of inference-based approaches (Barrett et al., 2019).

The workplace surveillance literature documents the harms of monitoring in detail (Ball, 2010; Ravid et al., 2020). Our contribution is to connect these critiques to a specific design alternative: energy-state self-report with principled adaptation (Pritis, 2026a).

9 Conclusion

The case for self-report over surveillance is not a concession to technical immaturity. It is a principled position about the relationship between tools and the people who use them.

People know how they feel. Systems should trust that knowledge. When a tool needs to adapt to cognitive capacity, the right mechanism is to ask, not to observe. The right architecture is explicit state, not inferred assessment. The right power relationship is one where the user holds authority over their own experience.

This is not a weaker design. It is a more honest one.

References

Ball, K. (2010). Workplace surveillance: An overview. Labor History, 51(1):87–106.

Barkley, R. A. (1997). Behavioral inhibition, sustained attention, and executive functions: Constructing a unifying theory of ADHD. Psychological Bulletin, 121(1):65–94.

Barrett, L. F., Adolphs, R., Marsella, S., Martinez, A. M., and Pollak, S. D. (2019). Emotional expressions reconsidered: Challenges to inferring emotion from human facial movements. Psychological Science in the Public Interest, 20(1):1–68.

Charles, R. L. and Nixon, J. (2019). Measuring mental workload using physiological measures: A systematic review. Applied Ergonomics, 74:221–232.

Csikszentmihalyi, M. and Larson, R. (2014). Validity and reliability of the experience-sampling method. In Flow and the Foundations of Positive Psychology, pages 35–54. Springer.

Milton, D. E. M. (2012). On the ontological status of autism: The 'double empathy problem'. Disability & Society, 27(6):883–887.

Monaco, J. V. and Tappert, C. C. (2018). The partially observable hidden Markov model and its application to keystroke dynamics. Pattern Recognition, 76:449–462.

Picard, R. W. (1997). Affective Computing. MIT Press.

Pritis, D. (2026a). Energy as state, not time: Rethinking productivity models for variable-capacity work. Kumbatio / entro314.labs. https://doi.org/10.5281/zenodo.21915159.

Pritis, D. (2026b). The myth of the flat workday: What cognitive science already knows about variable capacity. Kumbatio / entro314.labs. https://doi.org/10.5281/zenodo.21915161.

Ravid, D. M., Tomczak, D. L., White, J. C., and Behrend, T. S. (2020). EPM 20/20: A review, framework, and research agenda for electronic performance monitoring. Journal of Management, 46(1):100–126.

Rock, P. L., Roiser, J. P., Riedel, W. J., and Blackwell, A. D. (2014). Cognitive impairment in depression: A systematic review and meta-analysis. Psychological Medicine, 44(10):2029–2040.