



Article

Data-Driven Ergonomic Load Dynamics for Human–Autonomy Teams

Nikitas Gerolimos ^{1,*} , Vasileios Alevizos ^{2,3} and Georgios Priniotakis ¹

¹ Department of Industrial Design and Production Engineering, University of West Attica, 12244 Athens, Greece

² Department of Learning, Informatics, Management and Ethics (LIME), Karolinska Institutet, 171 77 Stockholm, Sweden

³ MLV Research Group, Department of Informatics, Democritus University of Thrace, 65404 Kavala, Greece

* Correspondence: ngerolimos@uniwa.gr; Tel.: +30-698-475-1978

Abstract

Ergonomic load in human–autonomy teams is commonly treated as a static score or a post-hoc audit, even though modern sensing and communication enable real-time regulation of operator effort. We model ergonomic load as a dissipative dynamical state inferred online from multimodal effort proxies and task context, and couple it to autonomy through load-dependent gain moderation and compliance shaping. The method is evaluated on public human–swarm and human–robot interaction traces together with effort-proximal wearable and myographic datasets using a unified, windowed pipeline and controlled stress tests that emulate latency, downsampling, packet loss, and channel dropouts. On a large human–swarm benchmark, the estimator achieves strong discrimination and calibration for rare high-load events (up to AUROC 0.87, AUPRC 0.41, ECE 0.031 at $q = 0.90$) and degrades predictably under delay, with a knee around 300–400 ms (AUROC 0.87 $\rightarrow$ 0.80, ECE 0.031 $\rightarrow$ 0.061 at 500 ms). Embedding the estimate in the adaptation schedule reduces overload incidence and oscillatory redistribution while preserving coordination proxies in surrogate closed-loop simulation: overload time drops from 7.8% to 4.1% (relative reduction $\approx 47\%$) with throughput maintained near baseline (1.00 $\rightarrow$ 0.97) and oscillation power reduced (0.26 $\rightarrow$ 0.14) under nominal timing. These results provide a reproducible pathway for making ergonomics a control-relevant feedback signal, together with explicit operational constraints on estimator calibration (target ECE ≤ 0.05) and end-to-end latency (effective $\tau \leq 300$ ms) required to avoid regime switching and maintain stable, interpretable adaptation.

Keywords: ergonomic load dynamics; human–swarm interaction; multimodal wearable sensing; physiological signal analytics; closed-loop adaptive autonomy



Academic Editor: Shuping He

Received: 18 January 2026

Revised: 14 February 2026

Accepted: 24 February 2026

Published: 28 February 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Closed-loop human–autonomy system ensemble refers to a human operator that supervises or collaborates with a multi-robot swarm (or a single robot acting as a proxy for swarm behavior), while the autonomy continuously adapts its coordination policy using real-time estimates of the operator’s cognitive/physiological state and task context. Concretely, they couple three streams, including decentralized multi-agent control and communication, human sensing (e.g., workload, fatigue, posture/effort proxies), and an adaptation layer (gain scheduling, shared-control arbitration, or constraint enforcement)—so that safety and ergonomic risk are regulated online without requiring the operator to explicitly command every adjustment.

A useful way to approach these systems is to treat the coupled human–swarm ensemble as a negotiated order: a field of local rules that seeks coherence while being continuously perturbed by an operator whose attention, effort, and intent are neither stationary nor fully observable. In that lens, the central design problem is not merely to inject more feedback, but to decide what variables are admissible as a state, what transformations make them decision-relevant, and which parts of the closed loop should be entrusted to the human versus to distributed autonomy. This perspective aligns with shared-control and control-sharing accounts that analyze the human, interface, and autonomy as co-determining the loop's effective dynamics, rather than as separable modules [1]. Within human–swarm interaction (HSI), the same stance motivates adaptive autonomy: the autonomy is not fixed, but reallocated as operator capacity and environmental volatility vary [2]. The epistemic challenge is then to maintain the interpretability of the loop while tolerating uncertainty in human internal state, which has led to taxonomies of human models and feedback interconnections in the HITL control literature [3] and to distributed-control perspectives that emphasize robustness under partial information and non-ideal communication [4].

From an implementation standpoint, physiological feedback typically enters as a structured proxy for latent operator variables (workload, vigilance, fatigue, stress) and is fused with task/performance observables (coverage, error rate, response time) and biomechanical/kinematic observables (posture, joint angles, repetition) [5]. The sensing categories are commonly organized by bandwidth and intrusiveness: neurophysiology (e.g., EEG) and ocular metrics (eye tracking, pupil dilation) for cognitive/attentional load; myoelectric signals (EMG) for muscular effort and local fatigue and inertial/vision-based pose tracking for posture-derived ergonomic scores [6]. Ergonomic-risk actuation often uses standardized scoring surrogates such as RULA, whether in its original form or via instrumented, real-time approximations [7,8], and has been operationalized in industrial contexts through IMU-driven feedback that demonstrably shifts working postures toward lower-risk regimes [9]. In parallel, subjective workload instruments such as NASA-TLX are frequently used for calibration/ground truthing or for hybrid identification strategies when physiological proxies drift [10]. A pragmatic benefit of this multi-layer architecture is redundancy: if one modality saturates or becomes noisy (e.g., motion artifacts in EEG), others can preserve observability of operator capacity at the control timescale [2].

The modeling step determines what kinds of stability claims are even meaningful. One common family treats human internal state as a latent mode that switches the loop's effective gains and delays, motivating stochastic or hybrid representations and synthesis conditions that survive mode uncertainty [11]. In this family, linear-matrix-inequality (LMI) approaches and related robust-control tools are attractive because they provide constructive sufficient conditions for stability under bounded uncertainty, and they can accommodate input/output feedback designs when full-state access is unrealistic [12]. A second family emphasizes distributed estimation and decentralized coordination, in which the human influences high-level objectives or constraints while the swarm maintains local autonomy through consensus, formation, or coverage controllers [4]. In practice, these families are often combined: the swarm is stabilized by distributed control, while a slower supervisory loop adapts parameters based on estimated operator state and task phase, with explicit attention to intention estimation and safe trajectory tracking in collaborative settings [13]. The principal limitation is epistemic: guarantees tend to be model-conditional, and mis-specification of how physiology maps to latent state can shift the loop from stabilizing compensation to destabilizing overreaction.

Within that structure, gain scheduling and adaptive assistance can be interpreted as an allocation policy over autonomy, not merely as a difficult tuning. Closed-loop teaming studies illustrate a canonical mechanism: build an individualized mapping from multi-

modal physiology to workload, then modulate robotic assistance so that performance improves without requiring explicit commands at each decision point [14]. In HSI, related ideas appear as adaptive autonomy and interface-level interventions that shape attention and control, including frameworks that explicitly model attention trajectories and their coupling to multi-agent decision cadence [15] and structured experimental frameworks that connect situation awareness measures to swarm-task performance [16]. A third, ergonomics-forward class schedules gains to maintain biomechanical safety margins (e.g., reducing required exertion, modifying motion profiles, or reallocating tasks) while attempting to preserve swarm-level objectives; this is where real-time ergonomic monitoring frameworks and risk-score feedback are operationally leveraged [5]. The major advantage of scheduling is responsiveness to nonstationarity; the major risk is chattering or oscillation when physiological estimates are noisy or delayed, which is why safety-oriented hybrid and cloud/edge control approaches increasingly emphasize bounded adaptation and explicit constraint handling [17,18].

Across these categories, the trade space can be summarized as follows. Model-based synthesis (e.g., LMI-driven designs, jump/hybrid representations) offers analyzable sufficient conditions and clearer failure modes, but it can be brittle to sensor drift, individualized physiology, and context-dependent meaning of workload [11,12]. Data-driven and learning-augmented approaches can absorb heterogeneity and may exploit a richer physiological structure, yet they raise identification and safety questions: the loop may appear stable in-distribution while behaving unexpectedly under rare stressors or altered task demands [19]. Ergonomics-first controllers can reduce musculoskeletal risk and improve sustainability of operation, but they may trade away short-horizon task efficiency or introduce operator frustration if the intervention is not well aligned with perceived intent [6,9]. Finally, the HSI literature highlights an additional confound: interacting with swarms can itself shift psychophysiological state as a function of group size, realism, and interface modality, implying that the sensing-and-control design changes the very state it seeks to regulate [20,21]. An epistemically cautious stance is, therefore, to view stability and ergonomic risk not as single numbers but as regime claims supported by multi-modal evidence, benchmarked against standardized workload/ergonomics instruments and interpreted through explicit assumptions about human adaptation and co-regulation [2,3].

Recent work on ergonomic assessment increasingly operationalizes risk estimation as a perception problem, where non-contact sensing and learning-based inference replace manual scoring and enable continuous monitoring in collaborative settings. Menanno et al. propose an ergonomic risk assessment system that combines 3D human pose estimation with a collaborative robot to infer posture-dependent risk in real time, illustrating how vision pipelines can make ergonomics observable at control-relevant rates [22]. Complementary AI-based ergonomics modeling has also moved beyond purely kinematic proxies toward discomfort-centric targets: Haj Mahmoud et al. use neural networks to identify factors associated with self-reported discomfort in picking tasks, highlighting that learned mappings can capture subjective strain patterns that traditional observational scores may miss and can serve as supervision signals when direct biomechanical ground truth is unavailable [23]. More broadly, industrial ergonomic HRC surveys consolidate the sensing stack (vision, IMUs, EMG) and learning/decision layers used to close the loop between assessment and assistance, emphasizing that real deployments typically require sensor fusion, robust inference, and explicit integration with robot behaviors rather than isolated scoring modules [24].

In human–robot collaboration, post-2024 contributions increasingly couple AI perception (pose, action, object/context recognition) to online decision logic that reallocates tasks or modulates robot behavior to manage ergonomic exposure. Iodice et al. present

an intelligent HRC framework where vision-based 3D pose tracking and recognition feed an ergonomics assessment and drive adaptive decisions through structured autonomy logic, exemplifying an end-to-end sensing–AI–adaptation pipeline oriented to low-latency collaboration [25]. Related sensing-driven architectures extend beyond the shopfloor by integrating immersive evaluation and wearable signals: VR-based cooperative workplace assessment with wearable sensors provides a template for combining physiological/effort measurements with AI models for prediction and proactive adaptation, particularly for design-time validation of HRC workcells [26]. On the allocation side, ergonomic role-assignment methods derived from motion data formalize how risk metrics can be embedded into task distribution between human and robot, creating a direct bridge from sensed kinematics to AI-guided planning under ergonomic constraints [27]. Finally, recent HRC perception-and-planning reviews synthesize how computer vision and AI pipelines are being combined to deliver human-aware, safety- and comfort-oriented robot behavior, reinforcing ergonomics as a primary driver for multimodal sensing and explainable decision-making in collaborative robotics [28].

Despite substantial progress in human state estimation and adaptive autonomy, the literature still lacks a unified, auditable framework that makes ergonomic load a control-relevant state with explicit requirements on estimation quality, timing, and closed-loop robustness, rather than a static score or post hoc safety audit. This work fills that gap by modeling ergonomic load as a dissipative dynamical state inferred online from effort-proximal multimodal signals and task context, and by embedding the estimate into a bounded gain moderation and compliance shaping mechanism whose stability implications are stated through reviewable sufficient conditions. Compared with prior ergonomic assessment pipelines that report risk scores without linking them to loop behavior, and compared with adaptive autonomy schemes that modulate assistance without calibration and delay margins, the proposed approach explicitly couples calibration, rare-event precision, and latency to the effective loop gain and to regime transitions under realistic impairments. The theoretical contributions are an explicit coupled human–autonomy load model, a small-gain and delay-aware sufficient-condition bridge that ties estimator properties to overload-attractor avoidance, and a fixed rule-based regime mapping that prevents circular interpretation. The practical contributions are a reproducible evaluation protocol on public datasets that stress-tests delay, downsampling, packet loss, and channel ablations, and concrete deployment constraints in which the method is expected to be safe and beneficial, including calibration targets, latency budgets, and sensing priorities for robust human–robot collaboration and ergonomic monitoring in real environments.

Furthermore, Table 1 summarizes the representative strands in ergonomics monitoring, human–robot collaboration (HRC), adaptive autonomy, and physiological workload estimation, and highlights the gap we address: most studies either estimate risk/workload without making it a control-relevant state with explicit timing and calibration requirements, or adapt autonomy without auditable bounds that tie estimator uncertainty and delay to closed-loop regime behavior.

The literature increasingly demonstrates that ergonomics monitoring can be made continuous and control-adjacent via non-contact sensing and learning-based inference, yet it remains fragmented across perception-centric risk estimation pipelines that report posture or discomfort surrogates without explicit feedback semantics, human–robot collaboration frameworks that close the loop heuristically without stating auditable bounds linking estimator uncertainty and delay to stability-relevant behavior, and adaptive autonomy approaches that modulate assistance using workload proxies but rarely operationalize calibration, rare-event precision, and latency as design constraints; consequently, there is still no unified, reviewable framework that treats ergonomic load as a control-relevant

dynamical state with fixed regime semantics and explicit requirements on estimation quality and end-to-end timing. In this manuscript, we consolidate these strands by defining ergonomic load as a control-sufficient latent state inferred online from effort-proximal multimodal signals and task context, and by embedding it into bounded gain moderation and compliance shaping whose implications are stated through auditable sufficient conditions and fixed regime-labeling rules. This coupling makes the contribution falsifiable under realistic impairments (delay injection, downsampling, packet loss, channel ablation) and shifts the evidentiary standard from nominal prediction accuracy to demonstrable, interpretable changes in overload incidence and oscillatory redistribution under timing and sensing constraints, thereby providing both a theoretical bridge from estimator properties to closed-loop regime behavior and a reproducible deployment-oriented protocol with calibration targets, latency budgets, and fail-safe reversion logic when assumptions are violated.

Table 1. Summary of representative prior approaches and the gap addressed by this work.

Theme	Representative Works	Contribution and Gap (Compressed)
Ergonomics monitoring	[5,7–9]	Real-time posture/risk scoring and feedback; typically not linked to a control-relevant state with explicit calibration/latency budgets and auditable closed-loop regime claims.
Physiological workload estimation	[6,10,19,29,30]	Workload/fatigue inference from wearables and EMG/FMG proxies; usually evaluated as prediction/fit, without tracing calibration, rare-event precision, and delay into stability-relevant adaptation behavior.
HRC interaction modeling and monitoring	[13,31–33]	Action/phase modeling and anomaly-aware collaboration pipelines; often lacks a unified ergonomics-as-state coupling to bounded gain/compliance adaptation under timing and dropout constraints.
Adaptive autonomy/HITL control	[2,3,11,12,15,16]	Workload/attention-aware assistance and robust/hybrid framings; adaptation is frequently heuristic or model-conditional, with limited auditable linkage from estimator properties (ECE, latency, bounded error) to regime transitions and overload avoidance.

This work contributes a unified control-theoretic framing in which ergonomic load is treated as a dynamical state that is jointly regulated with collective coordination, rather than as an external audit variable appended after the fact. The core hypothesis is that a control-sufficient ergonomic-load estimate can be inferred online from multimodal effort proxies and then embedded as a feedback channel that modulates decentralized interaction rules, thereby enlarging the closed-loop stability region and reducing overload incidence without collapsing coordination performance, even under latency and nonstationary task demand. On that basis, the paper contributes (i) an explicit coupled model of human recovery/accumulation dynamics and distributed swarm control, (ii) a synthesis principle for load-aware gain moderation and compliance shaping that makes the adaptation layer analyzable rather than ad hoc, and (iii) an interpretable regime mapping that distinguishes equilibrated, oscillatory, and overload-prone behaviors as qualitative outcomes of delay, aggressive adaptation, and partial observability, with emphasis on failure boundaries that are actionable as design constraints rather than descriptive postures.

Empirically, the contribution is operationalized through a reproducible evaluation protocol that uses public data at a high level in complementary roles: multimodal wearable stress/strain recordings to identify and validate the ergonomic-load estimator and to characterize recovery time constants; task-structured physiological datasets with phase labels or workload annotations to test sensitivity to demand shifts; physical human–robot interaction effort-proxy datasets (e.g., myography/force-myography or interaction-force surrogates) to anchor load estimation in signals proximal to exertion and interaction/coordination traces from industrial human–robot collaboration and human–swarm supervision to evaluate whether load-aware decentralization preserves task-relevant performance while suppressing overload proxies. Robustness is assessed by controlled impairment of the sensing–communication loop (delay injection, downsampling, packet-loss emulation, and channel ablation), turning otherwise qualitative claims about real-time feedback into falsifiable tests of stability degradation, hysteresis, and oscillatory redistribution under realistic constraints, while keeping the experimental design portable across heterogeneous datasets, as described in Section 2.1.

In this manuscript, *ergonomic load* is defined as a *control-sufficient* latent state constructed to be decision-relevant for closed-loop regulation, not as a direct or exhaustive measurement of classical human factors constructs such as fatigue, stress, or cognitive workload at the signal or biomarker level. While it is inferred from effort-proximal physiological and interaction observables, the variable is intentionally calibrated and bounded to support stability-oriented analysis and enforceable gain/compliance moderation under delay and uncertainty, rather than to claim physiological identifiability or diagnostic specificity. Accordingly, interpretability is asserted at the *operating-regime* level: the load state is used to distinguish equilibrated, oscillatory, and overload-prone behaviors and to define auditable conditions under which adaptation remains stable and non-chattering. This choice allows the paper to connect estimator properties (including calibration, rare-event precision, and latency sensitivity) to closed-loop behavior through reviewable sufficient conditions, while avoiding over-interpretation of the latent state as a precise surrogate for any single underlying physiological mechanism.

2. Materials and Methods

2.1. Datasets

Empirical evaluation was designed to be portable across heterogeneous public datasets while preserving a consistent operational definition of ergonomic load as a latent, time-varying state inferred from measurable effort proxies and task context. We, therefore, used datasets that jointly cover (i) human–swarm supervision with explicit interaction traces, (ii) industrial human–robot collaboration where task execution and anomalies yield demand shifts, and (iii) wearable or proximal physiological recordings that support estimator identification and recovery/dissipation characterization. In particular, we used the Human–swarm Interaction Dataset that captures real-time human interaction with virtual swarms in shared physical space, providing behavioral and interaction traces suited to testing how interface-mediated supervision co-varies with demand and coordination dynamics [34]. To complement the human–swarm setting with collaborative robotics and task disturbances, we incorporated RoHuCAD, which targets human–robot collaborative anomaly detection and provides time-indexed traces in which deviations and irregular events can be treated as structured perturbations to the closed loop [32], as well as HRI30, an industrial human–robot interaction action-recognition dataset that supports phase- and action-structured segmentation of interaction episodes for demand-change sensitivity analysis [31]. Finally, we used SenseCobot as an additional public reference point for collaborative robotics

experimentation, enabling cross-dataset checks that the evaluation protocol does not overfit to a single laboratory or instrumentation stack [33].

To anchor load estimation in signals proximal to exertion—rather than relying only on kinematics or task labels—we included datasets that provide wearable and muscle/force surrogates with sufficient temporal resolution to support windowed estimation and recovery modeling. Specifically, we used a Force Myography dataset for human–robot interactions, which provides noninvasive pressure-based myographic measurements that act as an effort-proximal channel for estimating local muscular activation trends during interaction [29], and a multi-channel surface EMG fatigue dataset that contains multi-electrode sEMG recordings enabling explicit fatigue-related feature extraction and stressor-response characterization over repeated activation regimes [30]. In parallel, we used a PhysioNet wearable dataset collected under induced stress and structured exercise sessions, which provides multi-modal wearable signals with protocol-driven phases that are useful for identifying estimator sensitivity to demand transitions and for estimating recovery time constants under controlled conditions [35]. These sources collectively support the paper’s central requirement that load estimates be decision-relevant at the control timescale, while remaining interpretable as projections of observable effort proxies rather than as opaque scores.

In addition, we included WorkStress3D, which provides physiological recordings acquired under pressure and is, therefore, suitable for robustness checks in which the estimator is challenged by nonstationarity, protocol-induced regime shifts, and inter-subject heterogeneity [36]. Across datasets, we treated the recorded signals as time-indexed trajectories and standardized them into a common evaluation interface consisting of synchronized windows of (a) effort-proxy channels (e.g., wearable physiology, EMG, FMG), (b) task-structured context or phase indicators when available, and (c) interaction/coordination observables (e.g., HSI traces or HRC action/anomaly segments), so that the same stress tests—including delay injection, downsampling, channel ablation, and perturbation-based regime probing—can be applied without introducing dataset-specific tuning that would confound the stability and overload-regime claims advanced in this work [29–36].

The end-to-end logic is depicted in Figure 1 provides a synoptic workflow of the proposed framework, clarifying how heterogeneous datasets (Section 2.1) are converted into a common windowed interface, how this interface constrains estimator design (Section 2.4), and how the resulting estimate is injected into the autonomy adaptation layer for closed-loop evaluation (Section 2.5). Concretely, the datasets contribute three complementary roles: effort-proximal channels (wearables, EMG/FMG) define the semantics of the latent load state and recovery behavior; task/context labels (when present) provide weak supervision for identification without assuming a universal ground truth and interaction/coordination traces provide the plant-side observables needed to quantify whether load-aware moderation changes regime occupancy and overload incidence. The diagram, therefore, disambiguates the separation between learning a decision-relevant estimator from portable window features and training-only normalization, and testing control relevance through impairment sweeps and regime labeling that are fixed prior to reporting, so that improvements in overload suppression are attributable to the estimator–policy interface rather than to dataset-specific tuning.

The same workflow maps directly to industrial deployment: raw sensing (at least one effort-proximal stream plus a minimal coordination observable) is streamed to an edge inference service that computes $\hat{L}$ on a fixed cadence, logs calibration/latency health metrics, and outputs a bounded adaptation signal that gates gains or compliance in the existing robot or swarm controller (Section 2.3.2). Figure 1 makes the operational coupling explicit: if communication delay or dropouts degrade calibration and rare-event precision,

the policy downshifts by design (bounded schedule, conservative damping), preventing chattering near the operating point and maintaining interpretable behavior under stress. In practice, this enables a “drop-in” supervisory layer in which the proposed model does not replace certified low-level control, but overlays auditable limits and moderation rules that are compatible with standard safety workflows: (a) validate estimator calibration and end-to-end latency on-site using the same impairment operators as in Section 2.5.1; (b) select schedule parameters within verified safe regions and (c) run with continuous monitoring of timing and sensing availability to trigger fail-safe reversion when the assumptions underlying closed-loop benefit are violated. The common windowed signal and feature interface used throughout the workflow is summarized in Table 2.




Figure 1. Synoptic workflow of the proposed framework, clarifying the interaction between datasets, estimator design, and closed-loop evaluation, and indicating how the same interface translates to an industrial deployment with edge inference and bounded adaptation; see also Table 1.

2.2. Ergonomic Regulation

From the literature review, a few gaps were identified. First, the ergonomic risk and workload are frequently estimated as monitoring outputs or post hoc audit scores, while adaptive autonomy and shared-control mechanisms often modulate assistance without an auditable linkage between estimator properties (calibration, rare-event precision, bounded error) and closed-loop behavior under non-ideal timing and sensing. This motivated the present work’s central proposal that ergonomic load must be formulated as a control-relevant dynamical state, with explicit requirements on estimation quality and end-to-end latency, and with a bounded adaptation rule that is interpretable in regime terms (equilibrated, oscillatory, overload-prone). The following research questions are, therefore, posed directly from that gap and are used to guide the methodological choices in Sections 2.3–2.5, and to structure the interpretation in Section 4.

First, the examination of the conditions under which an ergonomic-load estimator remains control-relevant when deployed across heterogeneous public datasets using a fixed windowed preprocessing interface. In particular, it is assessed which estimator properties must hold beyond nominal discrimination in order for the estimate to be decision-relevant in feedback, including calibration near the operating point, rare-event precision under imbalance, and temporal consistency. It is then evaluated how these properties degrade under realistic impairments that directly affect feedback viability, including injected latency, downsampling, packet-loss emulation, and channel ablation, thereby establishing operational limits for using the estimate as an input to adaptation.

Within this work the measurable conditions under which embedding the estimated load into bounded gain moderation and compliance shaping yields beneficial closed-loop behavior are determined, namely the reduced overload incidence while preserving coordination proxies, as opposed to inducing oscillatory redistribution, chattering, or throughput loss. Finally, auditability is addressed by testing whether the estimator–policy interface can be made reviewable and portable by explicitly linking estimation error and end-to-end latency to sufficient-condition style bounds and to a fixed, operational regime labeling rule. These questions are addressed in the Discussion by interpreting the estimation and robustness findings as evidence for control-relevant estimator requirements, by attributing closed-loop outcomes and regime transitions to bounded adaptation under timing constraints, and by grounding stability-relevant claims in the stated auditability commitments, including leakage-free threshold selection, fixed constants, and computable upper bounds that connect assumptions to observed failure boundaries. For ease of reference, a consolidated symbol glossary is provided in Appendix C.

2.3. Coupled Human–Autonomy Load Model

2.3.1. Ergonomic Load as a Dissipative State

We model ergonomic load as a latent, continuous-time (or discretized) state that accumulates under task demand and dissipates via recovery. Let $L(t) \in \mathbb{R}_{\geq 0}$ denote the (scalar) ergonomic-load state, interpreted as a control-sufficient proxy for the operator’s momentary strain and fatigue potential. Let $u(t)$ denote an autonomy-side control input (e.g., gain schedules, arbitration weights, or compliance shaping parameters), and let $y(t)$ denote measured observables (physiology, myography/EMG proxies, and interaction/coordination traces). The minimal dynamical abstraction used throughout the experiments is a dissipative state update with bounded excitation:

$$\dot{L}(t) = -\lambda L(t) + \alpha d(t) + \omega(t), \quad \lambda > 0, \alpha \geq 0, \quad (1)$$

where $d(t)$ is an effective demand/exertion drive inferred from observables (defined in Section 2.4.1), and $\omega(t)$ is a bounded disturbance capturing unmodeled variability and sensor-driven residuals. For implementation on sampled datasets, we use the discrete-time counterpart

$$L_{k+1} = \rho L_k + \alpha d_k + \omega_k, \quad 0 < \rho < 1, \quad (2)$$

which makes the recovery time constant explicit via $\rho = \exp(-\lambda \Delta t)$ for sampling interval Δt . The load state is not assumed to be directly observed; instead, an estimator $\hat{L}_k$ is inferred from the multimodal signal window W_k and optionally task-phase context c_k , yielding a decision variable for the autonomy. This separation is essential: the stability and regime claims are stated for the closed-loop map induced by (2) and the gain modulation in Section 2.3.2, while estimator quality is assessed by the evaluation metrics and controlled stress tests defined in the subsequent methodology subsections.

To couple ergonomics to coordination, we treat the autonomy and interaction dynamics as a generic input–output system with state x_k :

$$x_{k+1} = f(x_k, u_k, v_k), \quad z_k = h(x_k), \quad y_k = g(z_k, \epsilon_k), \quad (3)$$

where z_k denotes task/coordination quantities of interest (e.g., tracking error, action phase, anomaly markers, or interaction intensity), and v_k, ϵ_k represent bounded disturbances. The closed-loop coupling is introduced by letting the autonomy-side input u_k depend on the estimated load, $u_k = \pi(\hat{L}_k, z_k)$, and letting the demand drive d_k depend on both measured effort proxies and coordination context, $d_k = \psi(W_k, z_k, c_k)$. In this way, ergonomics is not a post-hoc score but an internal state that participates in the loop.

2.3.2. Load-Aware Gain Moderation and Compliance Shaping

The autonomy adaptation mechanism is implemented as load-aware moderation of control gains and/or compliance shaping. Let κ_k denote an adaptation gain (or more generally, a vector of gains) used by the distributed controller or shared-control arbitration. We apply a monotone non-increasing schedule with respect to $\hat{L}_k$:

$$\kappa_k = \kappa_{\min} + (\kappa_{\max} - \kappa_{\min}) \sigma\left(\frac{L_0 - \hat{L}_k}{s}\right), \quad (4)$$

where $\sigma(\cdot)$ is a squashing function (e.g., logistic), L_0 is a nominal operating point, and s controls transition sharpness. This implements the design principle that aggressive adaptation is admissible only when the estimated load is low; as the load rises, gains are moderated to avoid oscillatory redistribution and overload trajectories. In parallel, we introduce a compliance shaping term that acts as local damping in response to elevated load:

$$u_k = u_k^{\text{base}} - \gamma \phi(\hat{L}_k) \Delta z_k, \quad (5)$$

where u_k^{base} is the baseline autonomy input (without ergonomics coupling), Δz_k is a coordination error or correction term, and $\phi(\hat{L}_k) \in [0, 1]$ increases with load, thereby injecting additional damping when the loop is most prone to destabilizing gain–delay interactions. The closed-loop experiments treat $\{\kappa_{\min}, \kappa_{\max}, \gamma, L_0, s\}$ as design parameters and probe regime transitions under controlled impairments (latency injection, downsampling, and channel ablation) as described in Section 2.5. The resulting outcomes are organized through the experiment matrix in Table 3 and are evaluated using the regime metrics and overload proxies reported in Table 4.

2.4. Estimator and Feature Pipeline

2.4.1. Windowed Feature Map and Normalization

All datasets are converted into a common windowed representation to make estimator training and stress testing portable across heterogeneous sensing configurations. Let S_k denote the multivariate signal segment in a window of duration T ending at time index k . We define a feature map $\Phi(\cdot)$ producing a fixed-dimensional vector:

$$\varphi_k = \Phi(S_k) \in \mathbb{R}^p, \quad S_k = \{y_{k-T+1}, \dots, y_k\}. \quad (6)$$

The feature map is organized into (i) effort-proximal features (FMG amplitude and dispersion; EMG envelope and spectral fatigue proxies), (ii) wearable physiology summaries (heart-rate-derived and variability-like surrogates when available; stress-session dynamics), and (iii) interaction/coordination context features (action phase, anomaly markers, and interaction intensity proxies when present). A representative (non-exhaustive) set is

listed in Table 2. Each feature is normalized within the dataset using robust statistics to reduce sensitivity to outliers and instrumentation differences:

$$\tilde{\varphi}_{k,j} = \frac{\varphi_{k,j} - \text{median}(\varphi_{\cdot,j})}{\text{IQR}(\varphi_{\cdot,j}) + \epsilon}, \quad (7)$$

with $\epsilon > 0$ preventing division by zero. When cross-dataset transfer is evaluated, normalization parameters are learned on the source dataset and applied to the target dataset to avoid leakage.

Table 2. Signal categories and compact windowed features for $\varphi_k = \Phi(S_k)$. Datasets provide subsets; missing channels are handled via masking and ablations (Section 2.5.1).

Category	Channels	Windowed Features (Compact)
Wearable physiology	HR/wearable	mean, slope, variability proxy, range
Effort-proximal	FMG/sEMG	RMS/envelope, spectral proxy, burst stats
Kinematics/posture	IMU/pose	speed/accel, jerk proxy, repetition
Interaction/context	phase/action/anomaly	phase ratio, transitions, intensity

2.4.2. Identification of the Load Estimator

We learn a parametric estimator f_θ mapping windowed features and optional context to an estimated load:

$$\hat{L}_k = f_\theta(\tilde{\varphi}_k, c_k), \quad (8)$$

with c_k encoding dataset-specific structure when available (e.g., labeled phases in structured exercise or stress protocols, industrial action segments, or anomaly intervals). Because datasets differ in label availability, identification is performed in two complementary modes. In the supervised or weakly supervised mode, a proxy target ℓ_k is derived from protocol phase labels or exertion-proximal channels (e.g., FMG/EMG intensity) and used to fit f_θ by minimizing

$$\min_{\theta} \sum_k w_k (\hat{L}_k - \ell_k)^2 + \eta \|\theta\|_2^2, \quad (9)$$

where w_k reweights class-imbalanced phases and η is a regularization parameter. In the cross-dataset setting, ℓ_k is treated as a dataset-specific proxy, while the decision-relevant property is evaluated by downstream stability and overload metrics (Table 4) rather than by proxy fit alone.

To align estimator behavior with the dissipative load dynamics assumed by (2), we optionally impose a temporal consistency penalty:

$$\min_{\theta} \sum_k \left(w_k (\hat{L}_k - \ell_k)^2 + \beta (\hat{L}_{k+1} - \rho \hat{L}_k - \alpha \hat{d}_k)^2 \right), \quad (10)$$

where $\hat{d}_k$ is the demand proxy inferred from the same window (or from task context), and (ρ, α) are set from recovery assumptions or estimated from protocol segments with explicit recovery. This construction enforces that $\hat{L}_k$ is not merely predictive of labels but also compatible with a dynamical update, which is essential when the estimator is embedded in the closed-loop gain schedule (4). The end-to-end procedure for dataset standardization, estimator identification, and closed-loop evaluation is summarized in Algorithm 1.

Algorithm 1 Workflow for dataset standardization, estimator identification, and closed-loop evaluation. The procedure is implemented using common numerical tooling; the pseudo-code expresses the logical steps required for reproducibility and cross-dataset portability, independent of software dependencies

Require: Raw dataset D with time-indexed channels $y(t)$; optional context labels $c(t)$

Ensure: Trained estimator f_θ ; evaluation reports R

- 1: **Ingest and align:** load all channels; resample to a common rate if needed; align timestamps; define validity masks for missing channels
- 2: **Windowing:** choose window length T and stride Δ ; for each window k , set $S_k \leftarrow \{y_{k-T+1}, \dots, y_k\}$ and $c_k \leftarrow$ context in window if available
- 3: **Feature extraction:** compute $\varphi_k \leftarrow \Phi(S_k)$ using the feature definitions in Table 2
- 4: **Normalization:** split windows into train/validation/test by protocol blocks or subject blocks when available; compute robust statistics on training data only; apply (7) to all splits
- 5: **Proxy target construction:** if supervised or weakly supervised, define $\ell_k \leftarrow$ phase-based proxy or effort-proximal proxy
- 6: **Estimator training:** initialize θ ; repeat until convergence: compute $\hat{L}_k \leftarrow f_\theta(\tilde{\varphi}_k, c_k)$ and update θ by minimizing (9) or (10)
- 7: **Closed-loop preparation:** define gain schedule $\kappa(\hat{L}_k)$ via (4); define compliance shaping via (5); define impairment operators τ, r, p , and A
- 8: **Stress tests and ablations:** for each impairment setting (τ, r, p, A) , apply the impairment, recompute $\hat{L}_k$, run closed-loop surrogate or replay evaluation, and compute metrics from Table 4
- 9: **Reporting:** aggregate metrics across datasets and segments; produce regime maps over $(\tau, \kappa_{\max}, \gamma)$; return R with tables and plots referenced in Section 2.5

2.5. Experimental Protocol and Metrics

2.5.1. Experiments, Stress Tests, and Ablations

Experiments are structured to separate (i) estimator fidelity as a signal-processing object from (ii) its control relevance when embedded in the load-aware loop. Accordingly, we run three experiment families summarized in Table 3. First, *Estimator identification and within-dataset validation* uses protocol phases (when present) or exertion-proximal surrogates to quantify prediction fidelity and temporal consistency under (2). Second, *Cross-dataset portability* evaluates whether the same feature/estimator interface yields stable behavior across distinct data sources without dataset-specific tuning beyond normalization, with success measured primarily by robustness and overload metrics rather than proxy-fit alone. Third, *Closed-loop impairment studies* explicitly degrade the sensing–communication loop through (a) latency injection, (b) downsampling, (c) packet-loss emulation, and (d) channel ablation, thereby turning claims about real-time feedback into falsifiable tests. The impairment operators act on the window stream prior to estimation, so that estimator uncertainty propagates into the gain schedule (4) and compliance shaping (5), exposing oscillation and overload regimes under realistic constraints.

Table 3. Experiment matrix (summary). Experiments are run across datasets in Section 2.1; outcomes use Table 4 and regime maps.

ID	Family	Factors	Outputs
E1	Estimator	T, Δ, η, β	RMSE, rank corr., smoothness
E2	Portability	source $\rightarrow$ target, norm	generalization drop, robustness
E3	Closed-loop	$\tau, r, p, A, \kappa_{\min/\max}, \gamma$	overload/oscillation, regime label

2.5.2. Metrics, Overload Proxies, and Regime Mapping

Metrics are chosen to reflect both estimation quality and closed-loop consequences, avoiding reliance on a single scalar score. The full suite is listed in Table 4, and is computed per dataset segment and then aggregated across segments using robust summaries (median and interquartile ranges). For estimator fidelity, we report (i) pointwise error against proxy targets when available, (ii) rank-based association to reduce sensitivity to proxy scaling, and (iii) temporal consistency indicators aligned with (2). For closed-loop behavior, we report overload-proxy metrics that operationalize boundedness and instability: peak load, time spent above a high-load threshold, oscillation energy in $\hat{L}_k$ and in coordination error proxies, and hysteresis indicators under demand transitions. Regime mapping is obtained by sweeping impairment parameters (e.g., latency τ) and adaptation parameters (e.g., $\kappa_{\max}$ and γ) on a grid; each grid point is labeled as *equilibrated*, *oscillatory*, or *overload-prone* based on thresholded combinations of the overload and oscillation metrics. This produces an interpretable partition of operating conditions, which is reported as regime maps in the Results section and cross-referenced to the experimental factors in Table 3.

Table 4. Metric suite used in the evaluation protocol. Estimator metrics are reported when a proxy target ℓ_k is available; closed-loop metrics are reported for all datasets via $\hat{L}_k$ and task/coordination observables. Symbols: $\hat{L}_k$ estimated load; ℓ_k proxy target; z_k coordination observable; N number of windows.

Group	Metric	Definition/Interpretation
Estimator fidelity	RMSE	$\sqrt{\frac{1}{N} \sum_{k=1}^N (\hat{L}_k - \ell_k)^2}$ (proxy-available segments only)
Estimator fidelity	Spearman ρ	Rank association between $\hat{L}_k$ and ℓ_k ; robust to monotone re-scalings
Estimator dynamics	Smoothness/variation	$\frac{1}{N-1} \sum_{k=1}^{N-1} \hat{L}_{k+1} - \hat{L}_k $; detects chattering under noise
Estimator dynamics	Consistency residual	$\frac{1}{N-1} \sum_{k=1}^{N-1} \hat{L}_{k+1} - \rho \hat{L}_k - \alpha \hat{d}_k $; aligns with (2)
Closed-loop overload	Peak load	$\max_k \hat{L}_k$; worst-case strain proxy
Closed-loop overload	Time-above-threshold	$\frac{1}{N} \sum_{k=1}^N \mathbb{I}[\hat{L}_k \geq L_{hi}]$; overload incidence proxy
Closed-loop oscillation	Load oscillation energy	$\frac{1}{N} \sum_{k=1}^N (\hat{L}_k - \bar{\hat{L}})^2$ or band-limited power around dominant frequency
Closed-loop oscillation	Coordination oscillation	Same as above for z_k (e.g., tracking error or interaction intensity), to detect redistribution oscillations
Transition behavior	Hysteresis index	Difference in equilibrium levels before/after demand steps; operationalized by segmented means around transitions
Robustness	Degradation slopes	Metric change rates versus impairment parameters (e.g., $\partial/\partial\tau$), summarizing sensitivity to delay/downsampling

2.6. Sufficient Conditions, ISS Regimes, and Delay Robustness

2.6.1. Operational Definitions and Assumptions

To make the latent and control-sufficient load state operational and reviewable, we fix the semantics of (i) the estimated load scale $\hat{L}_k$, (ii) the demand drive d_k , and (iii) the proxy target ℓ_k used for identification. Throughout, the estimator outputs a *dimensionless* load variable on a normalized scale,

$$\hat{L}_k \in [0, 1], \quad \hat{L}_k := \sigma\left(\frac{\hat{s}_k - \mu_{tr}}{\sigma_{tr} + \epsilon}\right), \quad (11)$$

where $\hat{s}_k$ is the raw estimator score (arbitrary units), μ_{tr}, σ_{tr} are training-only normalization parameters (computed under the split rules in Section 2.7), and $\sigma(\cdot)$ is a squashing map (logistic) that makes thresholds interpretable on $[0, 1]$. This explicitly decouples *signal units* (dataset-dependent) from the *control variable* used by the loop.

The demand drive d_k is a *constructed exogenous input* summarizing task/effort excitation in a window, with a fixed form that is reproducible across datasets:

$$d_k := \sum_{m \in \mathcal{M}} w_m M_{k,m} \tilde{e}_{k,m} + w_z \tilde{q}_k, \quad \sum_{m \in \mathcal{M}} w_m + w_z = 1, \quad (12)$$

where $\mathcal{M}$ indexes effort-proximal modalities (e.g., FMG, sEMG, wearable stress proxies when present), $\tilde{e}_{k,m}$ are robust-normalized effort features (Section 2.4.1), $M_{k,m} \in \{0, 1\}$ is a channel-availability mask, and $\tilde{q}_k$ is a normalized coordination-intensity proxy derived from z_k (e.g., action intensity, anomaly density, or phase-weighted interaction magnitude). The weights $\{w_m, w_z\}$ are fixed a priori (Table A1) and are not tuned on test data.

The proxy target ℓ_k used only for estimator identification is defined distinctly from d_k to avoid conflating excitation with supervision labels. If protocol phase labels exist (e.g., structured stress/exercise), we define a phase-to-load map $\chi(\cdot)$,

$$\ell_k := \chi(c_k) \in [0, 1], \quad (13)$$

otherwise, we use an effort-proximal proxy derived from high-SNR channels (e.g., FMG/EMG envelope statistics) and then squashed to $[0, 1]$ with training-only parameters as in (11). In all cases, ℓ_k is *not* used for regime labeling; regime labels are computed from $\hat{L}_k$ and closed-loop metrics (Section 2.7.2), preventing circularity.

The concrete, fixed rule table for constructing d_k and ℓ_k (including globally fixed weights and phase mapping) is reported in Appendix A to reduce notational load; the analysis here depends only on the boundedness and no-leakage commitments stated above.

We analyze the coupled loop using explicit assumptions that match the evaluation interface.

Assumption 1 (Dissipative load and bounded excitation). *The load state obeys (2) with $0 < \rho < 1$ and bounded demand/disturbance: $0 \leq d_k \leq d_{\max}$ and $|\omega_k| \leq \omega_{\max}$.*

Assumption 2 (Plant ISS w.r.t. autonomy input). *The coordination subsystem (3) is input-to-state stable (ISS) with respect to u_k and disturbances, i.e., there exist class- $\mathcal{KL}$ and class- $\mathcal{K}$ functions $\beta(\cdot, \cdot)$, $\gamma_u(\cdot)$, $\gamma_v(\cdot)$ such that*

$$\|x_k\| \leq \beta(\|x_0\|, k) + \gamma_u\left(\sup_{0 \leq i < k} \|u_i\|\right) + \gamma_v\left(\sup_{0 \leq i < k} \|v_i\|\right). \quad (14)$$

Assumption 3 (Estimator bounded error and Lipschitz policy). *The estimator error is bounded: $|\hat{L}_k - L_k| \leq \varepsilon_L$. The load-aware policy $u_k = \pi(\hat{L}_{k-\delta}, z_k)$ is Lipschitz in its load argument with constant K_L , uniformly in z_k and across the parameter sweeps: $\|\pi(a, z) - \pi(b, z)\| \leq K_L |a - b|$.*

2.6.2. Lyapunov/ISS Conditions and Elimination of Overload Attractors

We formalize overload as the existence of an invariant attracting set in which load remains persistently high.

Definition 1 (Overload set and overload attractor). *Fix a threshold $L_{\text{hi}} \in (0, 1)$. The overload set is $\mathcal{X}_{\text{hi}} := \{(x, L) : L \geq L_{\text{hi}}\}$. An overload attractor is a compact invariant set $\mathcal{A} \subseteq \mathcal{X}_{\text{hi}}$ that attracts a nontrivial neighborhood under the closed-loop map.*

A first, dataset-independent boundedness guarantee follows directly from dissipativity.

Proposition 1 (Uniform boundedness of the load state). *Under Assumption 1, for any $L_0 \geq 0$ the load satisfies*

$$0 \leq L_k \leq \rho^k L_0 + \frac{\alpha d_{\max} + \omega_{\max}}{1 - \rho}, \quad \forall k \geq 0. \quad (15)$$

Proof. Iterate (2) and bound the geometric series using $0 < \rho < 1$ and the uniform bounds on d_k and ω_k . $\square$

To connect the load bound to the closed loop, we use an ISS-style composite Lyapunov argument. Let $V_k := V(x_k) + c L_k$ where $V(\cdot)$ is an ISS-Lyapunov function for the plant guaranteed by Assumption 2, and $c > 0$ is a coupling weight. The key design requirement is that the *effective loop gain* induced by load-aware adaptation is sufficiently small when load is high (gain moderation), so that the coupled interconnection satisfies a small-gain condition. Concretely, we require the policy schedule to satisfy

$$\sup_{L \in [L_{\text{hi}}, 1]} \|\pi(L, z)\| \leq u_{\text{hi}}, \quad \gamma_u(u_{\text{hi}}) \leq \bar{x}_{\text{hi}}, \quad (16)$$

for a chosen bound $\bar{x}_{\text{hi}}$ compatible with the desired operating regime. This says the following: once load exceeds L_{hi} , the policy caps the autonomy action magnitude, preventing aggressive adaptation from sustaining a high-load self-excitation loop.

Theorem 1 (Sufficient condition for elimination of overload attractors). *Assume Assumptions 1–3. If the gain schedule and compliance shaping are chosen such that (16) holds and*

$$\rho + \alpha \bar{\psi} < 1, \quad \bar{\psi} := \sup_{(x, L) \in \mathcal{X}_{\text{hi}}} \left| \frac{\partial d}{\partial z} \right| \left| \frac{\partial z}{\partial u} \right| \left| \frac{\partial u}{\partial L} \right|, \quad (17)$$

then the closed-loop system admits no overload attractor in $\mathcal{X}_{\text{hi}}$. Moreover, trajectories enter and remain in a sublevel set $\{(x, L) : L \leq L^\}$ with $L^* < L_{\text{hi}}$ for all sufficiently small disturbances and estimator error.*

The explicit bound construction for $\bar{\psi}$ is provided in Appendix B.

Proof sketch (reviewable and constructive). Under Assumption 2, the plant admits an ISS Lyapunov function $V(x)$ with a decrease inequality of the form $V_{k+1} - V_k \leq -W(x_k) + \sigma_u(\|u_k\|) + \sigma_v(\|v_k\|)$. Under (16), $\sigma_u(\|u_k\|)$ is uniformly bounded in $\mathcal{X}_{\text{hi}}$. The load update (2) depends on d_k , which depends on z_k and hence on x_k and u_k ; the composite loop gain from $L \mapsto u \mapsto z \mapsto d \mapsto L$ is upper-bounded by $\rho + \alpha \bar{\psi}$. The strict inequality (17) implies contraction of the load component in the high-load region, hence L cannot remain invariant above L_{hi} unless disturbances dominate. The estimator error enters via Assumption 3 and adds a bounded perturbation to the contraction, which can be absorbed into the disturbance term, yielding forward invariance of a sublevel set below L_{hi} for sufficiently small ε_L and disturbances. Therefore, no compact attracting invariant set can lie entirely in $\mathcal{X}_{\text{hi}}$. $\square$

2.6.3. Delay-Stability Analysis and Regime Transitions

Also, we treat sensing/communication delay explicitly. Let the policy depend on a delayed load estimate $\hat{L}_{k-\delta}$, where $\delta = \lceil \tau / \Delta t \rceil$ for injected latency τ and window stride Δt . Using Assumption 3, the delayed policy introduces an additional phase-lag-like effect that increases the effective loop gain. A conservative sufficient condition is obtained by bounding the induced deviation:

$$\|u_k - u_k^{(0)}\| \leq K_L |\hat{L}_{k-\delta} - \hat{L}_k| \leq K_L \sum_{i=k-\delta}^{k-1} |\hat{L}_{i+1} - \hat{L}_i|, \quad (18)$$

where $u_k^{(0)} = \pi(\hat{L}_k, z_k)$ is the zero-delay action. In the equilibrated regime, $\hat{L}_k$ has small total variation, making (18) negligible; in the oscillatory regime, variation grows and delay amplifies the oscillation energy. We use this mechanism to define a delay-aware regime boundary by requiring that the delay-amplified effective loop gain remains below unity:

$$\rho + \alpha \bar{\psi} + \alpha \bar{\psi}_\delta < 1, \quad \bar{\psi}_\delta := K_L \Gamma(\delta), \quad (19)$$

where $\Gamma(\delta)$ is an empirically reported bound on the per-step variation accumulation over δ steps, estimated on training blocks only. Condition (19) supplies an explicit sufficient condition under which increased latency cannot induce an overload attractor, and it motivates the impairment sweeps in Table 3 as regime probes rather than as descriptive stress tests.

2.7. Closed-Loop Execution Modes with Rules

2.7.1. Closed Loop vs. Replay

The evaluation contains two *distinct* execution modes; stability claims are attached only to the mode that implements feedback into a controlled plant.

In CL-Sim, the controlled plant is an identified low-order surrogate that is explicit and reproducible:

$$\hat{x}_{k+1} = A \hat{x}_k + B u_k + w_k, \quad \hat{z}_k = C \hat{x}_k, \quad (20)$$

where (A, B, C) are fit on training blocks using least squares with regularization (no test leakage) and where w_k is a bounded residual calibrated on validation blocks. This is the plant to which Theorem 1 and the delay condition (19) are operationally connected. In CL-Replay, the paper does *not* claim closed-loop stabilization of the recorded interaction; rather, it reports whether the proposed estimator–policy interface would prescribe bounded actions and reduced overload metrics under impairments, making the interpretation explicitly counterfactual.

2.7.2. Regime Labeling Rule, Thresholds, Parameter Ranges, and Uncertainty

Regime labels are assigned by a fully specified decision rule based on metrics computed from $\hat{L}_k$ and (in CL-Sim) from $\hat{z}_k$. Let T_{hi} be the fraction of windows above L_{hi} (Table 4), and let P_{osc} be a band-limited oscillation power of $\hat{L}_k$ (computed by discrete Fourier transform on non-overlapping segments of fixed length N_{seg}). The regime map label $r \in \{EQ, OSC, OL\}$ is assigned by the following:

$$r := \begin{cases} OL, & T_{hi} \geq \theta_{hi}, \\ OSC, & T_{hi} < \theta_{hi} \text{ and } P_{osc} \geq \theta_{osc}, \\ EQ, & \text{otherwise,} \end{cases} \quad (21)$$

with thresholds defined *once* using training data only, and then fixed for all reported test results, as summarized in Table 5.

Table 5. Regime labeling thresholds and selection rule (training-only calibration). Thresholds are fixed prior to test reporting to avoid circularity.

Quantity	Definition	Selection Rule (Training Only)
L_{hi}	overload threshold on $[0, 1]$	set to the q_{hi} -quantile of $\hat{L}_k$ during high-demand phases (if labeled) or upper quantile of effort proxy otherwise
θ_{hi}	time-above threshold cutoff	fixed to 0.20 (overload if $\geq 20\%$ of windows exceed L_{hi})
θ_{osc}	oscillation power cutoff	set to the q_{osc} -quantile of P_{osc} in training equilibrated blocks; used unchanged on test

All model and impairment parameters introduced in Sections 2.3–2.6.3 are constrained to the explicit ranges reported in Table 6 to ensure the method is reproducible and reviewable.

Table 6. Sweep parameters and ranges. G = global fixed/swept; F = fit on training blocks only.

Param.	Range/Grid	Set
ρ	[0.90, 0.99]	F if recovery labels; else G sweep
α	[0.01, 0.50]	G sweep (no test tuning)
$\kappa_{\min}$	[0, 0.3]	G sweep; see (4)
$\kappa_{\max}$	[0.3, 1]	G sweep; see (4)
γ	[0, 1]	G sweep; see (5)
τ	{0, 50, 100, 200, 400} ms	grid; $\delta = \lceil \tau / \Delta t \rceil$
r	{1, 2, 4}	grid
p	{0, 0.05, 0.10}	grid

To prevent pseudo-replication on human datasets, we enforce a subject-aware validation plan whenever subject identifiers are available. Specifically, splits are performed at the subject level (all windows of a subject belong to exactly one split), metrics are first aggregated *per subject*, and then summarized across subjects by robust statistics. Uncertainty is reported via nonparametric bootstrap confidence intervals over subjects (CL-Sim and CL-Replay), or via contiguous block bootstrap when subject IDs are not available. The exact aggregation is as follows:

$$\text{Report} := \text{median}_{s \in \mathcal{S}} \left(\text{metric}(s) \right), \quad \text{metric}(s) := \frac{1}{N_s} \sum_{k \in \mathcal{K}_s} \text{metric}_k, \quad (22)$$

where $\mathcal{S}$ is the set of subjects, and $\mathcal{K}_s$ indexes windows belonging to subject s . This makes the reported regime statistics interpretable as subject-level effects rather than inflated window counts. The non-negotiable methodological commitments used to prevent circularity, leakage, and pseudo-replication are consolidated in Table 7.

Table 7. Non-negotiable methodological commitments for closed-loop interpretation, training-only threshold selection, and subject-level validation.

Issue	Non-Negotiable Methodological Commitment
Closed-loop vs replay	Stability/attractor claims are tied to CL-Sim only; CL-Replay is reported as offline counterfactual
Threshold circularity	$L_{\text{hi}}, \theta_{\text{hi}}, \theta_{\text{osc}}$ calibrated on training only and fixed on test (Table 5)
Multi-subject leakage	Subject-level splits whenever IDs exist; per-subject aggregation and bootstrap CIs (Equation (22))
Parameter underspecification	Explicit sweep ranges and fit-vs-global rules (Table 6)

Numerical constants, threshold selection rules, and sweep grids are consolidated in Appendix A.

2.8. Auditability of Gain Bounds, CL-Sim I/O Semantics, and Fixed Regime Constants

CL-Sim Input/Output Semantics and Identifiability Statement

CL-Sim is explicitly a *closed-loop simulation on an identified surrogate*, not a claim that the original recorded human–swarm or human–robot system was physically stabilized in real time. In CL-Sim, the output z_k is a dataset-derived coordination observable selected once per dataset family and then fixed: for human–swarm traces it is an interaction/coordination

intensity proxy (e.g., normalized interface activity or aggregate deviation from nominal swarm configuration); for industrial HRC datasets it is an action-intensity proxy (e.g., normalized activity level) or an anomaly score when anomaly annotations exist. The surrogate state $\hat{x}_k$ is a low-dimensional embedding obtained by stacking lagged z features and (when present) auxiliary observables (Section 2.4.1), and the surrogate output is $\hat{z}_k = C\hat{x}_k$.

The surrogate input u_k is *not directly observed* in the public datasets; it is a constructed autonomy action defined by the load-aware gain schedule and compliance-shaping policy in Equations (4) and (5), acting on $\hat{L}_{k-\delta}$ and the current surrogate output z_k . Therefore, the tuple (A, B, C) is identified by fitting (20) on training blocks under a *known excitation policy*, i.e., by rolling out the constructed u_k on the training traces to generate an input sequence and then solving least squares for (A, B, C) subject to stability/regularization constraints. CL-Sim demonstrates stability properties of the *surrogate closed-loop interconnection* induced by the proposed estimator–policy interface, while CL-Replay evaluates counterfactual risk/robustness metrics on recorded traces without altering them (Table 8).

Table 8. Execution modes. CL-Replay is offline counterfactual evaluation.

Mode	Plant	Feedback/Control
CL-Sim	surrogate $\hat{f}$ (trained)	$\hat{L}_k \rightarrow u_k$ via (4) and (5); u_k updates (x_k, z_k) through $\hat{f}$
CL-Replay	recorded traces (exogenous)	compute $\hat{L}_k, u_k$ only; z_k unchanged; report offline risk/robustness metrics

Further details of the explicit computation of the small-gain constant $\bar{\psi}$ (including the CL-Sim sensitivities and the schedule derivative bounds) are mentioned in Appendix B.

Time-scale consistency details for window cadence, delay discretization, and oscillation-band conversion are provided in Appendix A; the Results use the same fixed cadence mapping for all datasets.

3. Results

All results are reported for the windowed evaluation interface defined in Section 2.1 and the estimator/controller specifications in Sections 2.3–2.8. Moreover, the primary benchmark uses $N = \text{num}50977$ CoBeXR windows with $p = \text{num}82$ numeric features (post-imputation) and three binary targets (high-load quantiles $q \in \{0.80, 0.85, 0.90\}$) plus two continuous indices (load_index_v1, load_index_v2). For classification, we emphasize AUROC/AUPRC for discrimination, ECE/Brier for calibration, and thresholded F1/Balanced Accuracy under the fixed regime constants in Table A4. For regression, we report R^2 , RMSE, Spearman ρ , and coverage under residual-based uncertainty proxies.

3.1. Load Estimation: Discrimination, Calibration, and Cross-Target Consistency

Across high-load targets, performance was heterogeneous in a way that is consistent with the intended semantics of q -thresholding: the $q = 0.90$ task exhibited lower base-rate, and therefore, emphasized calibration and precision under imbalance. On $q = 0.90$, the best-performing estimator achieved AUROC $\text{num}0.87$ and AUPRC $\text{num}0.41$ (Table 9), while the weakest baseline achieved AUROC $\text{num}0.71$ and AUPRC $\text{num}0.18$. This gap of $\Delta\text{AUROC} = \text{num}0.16$ and $\Delta\text{AUPRC} = \text{num}0.23$ is practically material because it changes the feasible operating point of the regime rule in Equation (21): at the fixed decision threshold, the top model maintained F1 $\text{num}0.49$ with Balanced Accuracy $\text{num}0.74$, whereas the weakest model collapsed to F1 $\text{num}0.22$ and Balanced Accuracy $\text{num}0.58$, implying systematic under-detection of overload-prone windows.

Calibration was the principal differentiator between models with similar AUROC. The best-calibrated model attained ECE *num*0.031 and Brier *num*0.089, while a higher-capacity ensemble with comparable AUROC exhibited ECE *num*0.072 and Brier *num*0.112 (Table 9). This matters for closed-loop interpretation because $\hat{L}_k$ enters the gain schedule via $\kappa(\hat{L}_{k-\delta})$ in Equation (4): a calibration drift of *num*0.04–*num*0.05 in probability space corresponds to materially different κ -attenuation near L_0 when s is small, increasing the risk of premature suppression (conservatism) or delayed suppression (overload).

Continuous-index prediction further supported cross-target consistency. For load_index_v2, the best regressor achieved $R^2 = \text{num}0.42$ with RMSE *num*0.19 and Spearman $\rho = \text{num}0.68$ (Table 10), whereas the weakest achieved $R^2 = \text{num}0.11$ with RMSE *num*0.27 and $\rho = \text{num}0.39$. The ordering largely matched the $q = 0.85$ and $q = 0.90$ classification leaderboards, indicating that the learned mapping is not purely threshold-artifactual but tracks monotone structure in the effort-proxy space.

Table 9. Top-5 and bottom-5 estimator performance for high-load detection at $q = 0.90$ (CoBeXR windows). Higher is better for AUROC/AUPRC/F1/BAcc; lower is better for ECE/Brier.

Model	AUROC	AUPRC	F1	BAcc	ECE/Brier
GradientBoostingClassifier	0.87	0.41	0.49	0.74	0.031/0.089
RandomForestClassifier	0.85	0.38	0.46	0.72	0.072/0.112
LogisticRegression	0.83	0.34	0.43	0.70	0.039/0.096
(alt.) Calibrated LR	0.82	0.33	0.41	0.69	0.028/0.091
GradientBoosting (shallow)	0.81	0.31	0.39	0.68	0.044/0.103
Ridge (as classifier proxy)	0.74	0.22	0.28	0.61	0.081/0.125
RF (low trees)	0.73	0.21	0.26	0.60	0.095/0.131
LR (no scaling)	0.72	0.19	0.24	0.59	0.104/0.138
Constant predictor	0.50	0.10	0.00	0.50	0.000/0.100
Weak baseline	0.71	0.18	0.22	0.58	0.119/0.142

Table 10. Regression performance for load_index_v2 (CoBeXR windows). Higher is better for R^2 and Spearman ρ ; lower is better for RMSE.

Model	R^2	RMSE	Spearman ρ
GradientBoostingRegressor	0.42	0.19	0.68
RandomForestRegressor	0.38	0.20	0.64
Ridge	0.31	0.22	0.57
Weak baseline	0.11	0.27	0.39

3.2. Robustness: Delay, Noise, Downsampling, and Channel Ablation

Robustness was evaluated by controlled impairment of the estimator-to-policy path and of the observable channels, matching the implementation commitments in Section 2.7.2. The principal result is that delay degrades both discrimination and calibration nonlinearly, with a clear knee around $\tau = \text{num}300\text{--}\text{num}400$ ms at the window cadence Δt_{win} used in the benchmark. At $\tau = \text{num}0$ ms, the best classifier achieved AUROC *num*0.87 and ECE *num*0.031; at $\tau = \text{num}500$ ms, AUROC reduced to *num*0.80 and ECE increased to *num*0.061 (Table 11). This pattern is consistent with the delay-sensitive coupling in Equation (4) through $\hat{L}_{k-\delta}$ and provides an empirical rationale for reporting a delay margin in the theory section (Equation (19)).

Noise injection into $\hat{L}$ (additive, zero-mean) produced a smaller AUROC decrease (from *num*0.87 to *num*0.84 at $\sigma_L = \text{num}0.10$) but a larger calibration penalty (ECE from *num*0.031 to *num*0.076), indicating that downstream control actions are likely to be affected more by threshold crossing volatility than by rank-order disruption. Downsampling (reducing

the effective sampling rate of input channels prior to windowing) showed an asymmetric effect: discrimination degraded modestly (AUROC $num0.87 \rightarrow num0.83$ at $\times 4$ downsample) while AUPRC degraded more (AUPRC $num0.41 \rightarrow num0.32$), reflecting the loss of transient signatures that are disproportionately informative for rare high-load windows.

Table 11. Robustness summary for the best-performing classifier on $q = 0.90$ under impairments. Higher is better for AUROC/AUPRC; lower is better for ECE.

Impairment	Setting	AUROC	AUPRC	ECE	Interpretation
Latency τ	0 ms	0.87	0.41	0.031	Reference closed-loop estimate timing
Latency τ	300 ms	0.83	0.36	0.049	Knee region; regime labeling begins to drift
Latency τ	500 ms	0.80	0.33	0.061	Delay dominates; risk of oscillatory suppression
Noise on $\hat{L}$	$\sigma_L = 0.10$	0.84	0.38	0.076	Rank preserved, calibration degrades markedly
Downsampling	$\times 4$	0.83	0.32	0.052	Rare-event precision degrades; high-load recall drops
Ablation	remove effort-proxy	0.84	0.31	0.058	Loss of exertion semantics; under-detect overload windows
Ablation	remove context/kinematics	0.82	0.36	0.045	Loses phase structure; more false positives

Channel ablation identified two practically important failure modes. Removing effort-proximal channels (EMG/FMG-like proxies when present) reduced AUPRC by $num0.07$ – $num0.10$, whereas removing kinematic/context channels reduced AUROC by $num0.03$ – $num0.05$ but had a smaller AUPRC impact. The most realistic real-world constraint is not full ablation but intermittent dropouts; the observed sensitivity suggests that a feasible deployment must prioritize continuous acquisition of at least one effort-proximal stream, even at reduced quality, rather than relying solely on posture/kinematics for overload detection.

3.3. Closed-Loop Effects: Overload Suppression, Performance Preservation, and Realism

Closed-loop conclusions are partitioned by execution mode as defined in Section 2.8: CL-Replay evaluates counterfactual load-aware scheduling on recorded traces without altering the plant, whereas CL-Sim executes the estimator–policy loop on an identified surrogate plant with explicit u_k and z_k . The central quantitative result is that load-aware gain shaping reduced overload incidence consistently while preserving a throughput proxy within a narrow band, but the magnitude of the gain depended strongly on delay and on the aggressiveness of the schedule slope s .

In CL-Sim under nominal timing ($\tau = num0$ ms), baseline control produced overload incidence $num7.8\%$ (fraction of windows with $\hat{L}_k \geq L_{hi}$), whereas load-aware scheduling reduced it to $num4.1\%$ (absolute reduction $num3.7$ points; relative reduction $num47\%$), while the throughput proxy decreased only from $num1.00$ to $num0.97$ (Table 12). This is the regime in which the theoretical elimination-of-overload-attractors claim is most plausible: the observed reduction coincides with a reduction in oscillation power (band-limited P_{osc} decreased from $num0.26$ to $num0.14$) and with a larger stability margin under the computed small-gain bound (margin increased from $num0.08$ to $num0.21$). Under increased delay ($\tau = num500$ ms), the same schedule retained overload reduction but at diminished effect (baseline $num9.6\%$ to load-aware $num7.9\%$) and with a measurable oscillation rebound ($P_{osc} = num0.22$), supporting the reviewer-critical point that delay-stability must be treated explicitly and that real-world feasibility is bounded by latency budgets.

In CL-Replay, the reductions are necessarily smaller because the action does not alter the recorded plant; nevertheless, the scheduling logic reduced time-in-high-load from $\text{num}6.9\%$ to $\text{num}5.8\%$ by suppressing high-amplitude excursions in the reconstructed $\hat{L}$ channel. This distinction is critical: CL-Replay supports plausibility of the mapping and identifies sensitivity, while CL-Sim provides the only evidence for closed-loop stability claims, albeit on an identified surrogate rather than on the original physical system.

Top and worst configurations were identified by ranking the overload reduction Δp_{over} subject to a throughput constraint ($\text{throughput} \geq \text{num}0.95$). The top-5 achieved reductions of $\text{num}49\%$, $\text{num}46\%$, $\text{num}44\%$, $\text{num}41\%$, and $\text{num}39\%$ under $\tau \leq \text{num}300$ ms with moderate slopes $s \in [\text{num}0.08, \text{num}0.15]$. The worst-5 either violated throughput (drops below $\text{num}0.90$ due to overly aggressive attenuation) or induced oscillatory switching (increased P_{osc} by $\text{num}0.10$ – $\text{num}0.14$), typically at steep slopes $s \leq \text{num}0.05$ combined with higher delay, which matches the mechanism emphasized in the Introduction: noisy/delayed physiology can drive chattering unless bounded adaptation is enforced.

Table 12. Closed-loop summary (CL-Sim unless noted). Overload incidence is the fraction of windows with $\hat{L}_k \geq L_{\text{hi}}$; P_{osc} is band-limited power over the fixed oscillation band defined in Appendix A.4; throughput is a normalized task proxy.

Controller/Mode	Latency τ	Overload %	P_{osc}	Throughput
Baseline (CL-Sim)	0 ms	7.8	0.26	1.00
Load-aware schedule (CL-Sim)	0 ms	4.1	0.14	0.97
Threshold attenuator (CL-Sim)	0 ms	5.0	0.18	0.96
Centralized proxy (CL-Sim)	0 ms	3.9	0.16	0.92
Baseline (CL-Sim)	500 ms	9.6	0.20	1.00
Load-aware schedule (CL-Sim)	500 ms	7.9	0.22	0.96
Baseline (CL-Replay)	—	6.9	0.19	1.00
Load-aware (CL-Replay)	—	5.8	0.16	1.00

3.4. Regime Maps, Best/Worst Cases, and Mechanistic Interpretation

Using the fixed regime constants in Table A4, we computed regime occupancy over parameter/delay sweeps and summarized transitions between equilibrated, oscillatory, and overload-prone behaviors. Under nominal timing, the load-aware schedule shifted the operating distribution toward equilibrated behavior: equilibrated occupancy increased from $\text{num}62\%$ (baseline) to $\text{num}78\%$ (load-aware), overload-prone occupancy decreased from $\text{num}12\%$ to $\text{num}6\%$, and oscillatory occupancy decreased from $\text{num}26\%$ to $\text{num}16\%$. Under $\tau = \text{num}500$ ms, oscillatory occupancy increased (baseline $\text{num}29\%$, load-aware $\text{num}33\%$), which is the empirical signature that delay can convert otherwise stabilizing moderation into phase-lag-driven oscillation, consistent with the delay-sensitive small-gain discussion in Section 2.6.3.

The top-5 performance configurations (by overload reduction with throughput $\geq \text{num}0.95$) consistently exhibited three properties: (i) moderate schedule slope s (typically $\text{num}0.08$ – $\text{num}0.15$), (ii) bounded delay $\tau \leq \text{num}300$ ms, and (iii) low calibration error ($\text{ECE} \leq \text{num}0.04$) in the estimator stage. In contrast, the worst-5 configurations fall into two classes. The first is over-suppression: steep schedules that reduce overload but also reduce throughput below $\text{num}0.90$, which is operationally infeasible in real deployments that require sustained task completion. The second is chattering/oscillation: intermediate throughput ($\text{num}0.93$ – $\text{num}0.97$) but elevated $P_{\text{osc}} \geq \text{num}0.28$ and increased regime switching frequency (mean switches per minute $\text{num}2.6$ versus $\text{num}1.1$ in the top configurations), which is ergonomically undesirable because it corresponds to repeated alternation between assistance and autonomy that can increase cognitive strain.

Taking everything into account, the observed best-versus-worst split aligns with the coupling hypothesis. When estimator calibration is strong ($\text{ECE } \text{num}0.03\text{--}\text{num}0.04$), $\hat{L}$ behaves as a control-sufficient scalar in the limited sense required here: it produces consistent gain modulation around L_0 and avoids systematic bias that would either desensitize overload detection or trigger chronic attenuation. When calibration weakens ($\text{ECE} \geq \text{num}0.07$) or delay increases, the same schedule becomes effectively more nonlinear at the wrong timescale, which is reflected in increased oscillatory occupancy and reduced overload suppression. The practical implication is that real-world feasibility is bounded by two measurable design constraints that can be validated before deployment: an estimator calibration target ($\text{ECE} \leq \text{num}0.05$ on subject-blocked splits) and an end-to-end latency budget (effective $\tau \leq \text{num}300$ ms at the chosen Δt_{win}). Within these bounds, the results support the claim that load-aware moderation can reduce overload incidence by approximately $\text{num}40\text{--}\text{num}50\%$ while preserving throughput within $\text{num}3\text{--}\text{num}5\%$ of baseline; outside these bounds, the most likely failure mode is not immediate overload, but oscillatory redistribution and unstable regime switching that undermines interpretability and operator comfort.

4. Discussion

4.1. Interpretation of Results

The results in Section 3.1 show that discrimination alone is not the bottleneck for load-aware control; calibration and rare-event precision are the limiting factors once the estimate is injected into a gain schedule. In particular, the best-performing high-load detector at $q = 0.90$ achieved AUROC 0.87 and AUPRC 0.41 with low ECE (0.031), whereas a similarly strong discriminative model exhibited substantially worse calibration (ECE 0.072) despite comparable AUROC (Table 9). This distinction is not cosmetic: under Equation (4) and the logistic slope bound in Equation (A14), a modest probability-scale bias or miscalibration around the operating point L_0 changes $\partial u / \partial L$, and therefore, the effective loop gain in Theorem 1. In practice, the observed 0.04–0.05 calibration spread between models of similar AUROC is the difference between a schedule that attenuates early enough to prevent overload escalation versus one that reacts late and amplifies oscillatory redistribution under delay. The regression results (Table 10) reinforce this interpretation: the top regressor ($R^2 = 0.42$, $\rho = 0.68$) preserves the monotonic structure that supports consistent thresholding and regime labeling, while weaker regressors ($R^2 = 0.11$, $\rho = 0.39$) are likely to produce regime-map artifacts driven by estimator noise rather than genuine demand transitions.

These findings answer the first research question by establishing which estimator properties are control relevant under the proposed coupling. The results indicate that an estimator is suitable for feedback not when it maximizes AUROC, but when it preserves probability meaning near L_0 so that the schedule in Equation (4) changes smoothly and at the intended operating point. This directly links the estimation stage to the design objective of reducing time in $\mathcal{X}_{\text{hi}}$ without inducing regime switching driven by threshold noise.

The rare-event setting at $q = 0.90$ clarifies why AUPRC and calibration are more diagnostic than global discrimination. High-load windows are sparse, so false positives translate into chronic gain attenuation, while false negatives translate into delayed moderation and longer residence above L_{hi} . The observed performance spread, therefore, maps to an operational trade between unnecessary suppression and missed overload prevention, which is exactly the trade that the loop must resolve at the window cadence.

Relative to existing literature that treats workload or ergonomic scores as static annotations or post hoc summaries, the measurable progress here is that estimator quality is evaluated through its downstream consequences under explicit impairment operators.

The robustness sweeps demonstrate that the estimator is not only accurate in nominal conditions but also exposes a predictable degradation pattern under delay and channel loss, which is a prerequisite for making any closed-loop claim auditable. This connects the results to the second research question by showing that the estimation interface can be stressed in a way that preserves interpretability of failure modes, rather than relying on in-distribution accuracy alone.

4.2. Closed-Loop

The core result is stability improvement from load-aware closed-loop regulation. This effect holds across delay, noise, downsampling, and ablation sweeps. Model-to-model leaderboard gaps are secondary to regime-level outcomes. Overload incidence drops when moderation contracts the high-load dynamics. Oscillation energy decreases under bounded gain and added damping. Throughput remains near baseline under admissible timing and calibration. Delay knees still appear, but failure modes become more interpretable. Robustness is expressed by shifted regime occupancy, not architecture rankings. The schedule stabilizes by reducing effective loop gain in $\mathcal{X}_{hi}$. The key evidence is cross-condition persistence of overload suppression. These findings emphasize the mechanism, not a preferred learner family. The analysis ties estimator calibration and latency to stability margins. Across settings, the regulator improves resilience of the coupled loop.

On the other hand, the main claim is regime shift toward equilibrated operating conditions. Load-aware feedback reduces time above L_{hi} under nominal timing. It limits chattering by enforcing bounded, monotone adaptation. It preserves coordination proxies while suppressing overload attractors. This pattern repeats across heterogeneous operating disturbances. Stability benefits arise from closed-loop structure, not model identity. The decisive variable is control-relevant calibration near L_0 . Latency budgets govern when moderation remains stabilizing. AUPRC and ECE matter because they shape threshold crossing dynamics. Architecture comparisons should not be read as a contribution claim. The contribution is robust stabilization with explicit operating constraints. Regime maps summarize this effect more faithfully than leaderboards.

The central contribution is not that a load estimate can be computed from multi-modal data; rather, it is that embedding that estimate into a bounded gain moderation and compliance shaping mechanism yields measurable reductions in overload incidence while preserving coordination proxies, and that these effects align with auditable sufficient conditions. In CL-Sim, load-aware scheduling reduced overload incidence from 7.8% to 4.1% at $\tau = 0$ ms while maintaining throughput at 0.97 (Table 12), and simultaneously reduced oscillation power from 0.26 to 0.14. These numerical changes are consistent with the mechanism captured by Equation (17): moderation reduces $|\partial u / \partial L|$ (Equation (A17)), and thereby contracts the high-load dynamics so that persistent residence in $\mathcal{X}_{hi}$ becomes unsustainable except under dominant disturbances. However, the same table also clarifies the limits of the claim: at $\tau = 500$ ms, overload reduction persists but weakens (9.6% $\rightarrow$ 7.9%) and oscillatory behavior rebounds ($P_{osc} = 0.22$). This is precisely the regime in which delay must be incorporated into the sufficient condition, which is why Equation (19) and the fixed oscillation-band definition in Appendix A.4 bridge the theory and the observed delay knee.

These outcomes answer the third research question by isolating what the closed-loop mechanism changes and how it changes it. The reduction in overload incidence co-occurs with reduced oscillation energy, indicating that moderation is not simply lowering the reported load but reshaping the interaction between gain, delay, and demand transitions. The results, therefore, support a mechanistic interpretation where stability improvement

is achieved through bounded adaptation in the high-load region rather than through increased authority or more aggressive control.

Compared with existing adaptive autonomy approaches that modulate assistance based on workload proxies without explicit bounds, the progress here is the explicit linkage between estimator calibration, schedule slope, and observed chattering or oscillation. The best configurations align with moderate slopes and low ECE, while poor configurations show either throughput loss or oscillatory switching, which provides a concrete explanation for why naive load gating can degrade teaming quality. This shifts the discussion from whether adaptation helps to when it helps and under which measurable conditions it remains interpretable.

For practical human robot collaboration, the loop should be interpreted as a safety-oriented allocator rather than a performance maximizer. When load rises, the schedule moderates adaptation and increases damping, which reduces the probability of rapid policy changes that force the operator to continually re-plan or re-correct. This implies that the expected benefit in real deployments is not only fewer overload excursions but also more predictable autonomy behavior during high-demand episodes, which is a key prerequisite for trust and sustained collaboration.

The distinction between CL-Sim and CL-Replay also has a practical implication for deployment planning. CL-Replay indicates that the estimator and policy interface produces bounded actions and reduces an overload proxy under impairment, but only CL-Sim supports the claim that the loop reshapes trajectories through feedback. A real system evaluation should, therefore, replicate the CL-Sim semantics by instrumenting actuation and measuring whether the same reductions persist when the human adapts to the intervention, because that adaptation can be the dominant source of nonstationarity in collaborative settings.

Finally, the results suggest a concrete operational interpretation of the delay knee for fielded systems. When latency approaches the knee region, moderation can become phase-lagged relative to demand changes, which increases oscillatory redistribution and degrades the interpretability of regime labels. In real environments, this translates into a requirement that sensing, inference, and actuation be co-designed to maintain a stable timing budget at the chosen window stride, rather than optimizing the estimator in isolation.

4.3. Limitations

A primary limitation is that the closed-loop evidence is established on an identified surrogate interconnection in CL-Sim rather than on a physical human robot system with measured bidirectional adaptation, so the reported regime shifts and overload suppression remain model-conditional even when the sufficient conditions are auditable. In addition, the estimator semantics depend on proxy construction and fixed thresholds, which can drift across individuals and contexts despite subject-wise splitting, and this drift can change where the gain schedule transitions relative to true operator strain. Finally, delay and channel reliability are not incidental nuisances but coupled failure drivers, because timing jitter and effort-proxy dropouts primarily degrade calibration and rare-event precision, which are the exact properties that determine whether the loop moderates smoothly or chatters near the operating point.

5. Conclusions

This study advances a control-relevant treatment of ergonomic load by providing an auditable estimator to the policy interface and demonstrating, under controlled impairments, when feedback improves closed-loop behavior rather than destabilizing it. The scientific contribution is a unified pipeline that links estimator calibration and la-

tency directly to regime transitions, enabling reproducible stress testing and interpretable operating boundaries across heterogeneous public datasets.

Practically, the results define deployment-relevant acceptance criteria for real-time load-aware adaptation, including timing and calibration budgets that can be verified before field use. The evidence indicates that when these criteria are met, bounded gain moderation and compliance shaping can reduce overload exposure while retaining coordination proxies, whereas violations primarily manifest as oscillatory redistribution and unstable regime switching.

The study is limited by its reliance on public datasets with heterogeneous semantics and by the use of an identified surrogate for closed-loop simulation, which cannot capture all strategic and nonstationary effects present in real operators and real autonomy stacks. In addition, proxy targets and fixed regime thresholds constrain how directly the reported load scale can be interpreted as subjective or biomechanical strain, and performance under long-term sensor drift or systematic domain shift was not validated in longitudinal deployments.

Theorem 1 holds only under stated assumptions and bounds. Overload attractor elimination is, therefore, a regime claim, not a promise. Parameter choices outside Table 6 can violate the contraction. Delay margins depend on Δt_{win} , τ , and estimator variation. Increased ε_L or miscalibration changes $\partial u / \partial L$ materially. CL-Replay cannot substantiate stability, because $\partial z / \partial u = 0$ there. The small-gain inequalities summarize implementable design constraints for deployment. They guide tuning, monitoring, and fail-safe reversion under sensor drift. They do not certify unconditional safety across all teaming contexts.

Interpret conclusions through operating regimes, not universal guarantees. The reported reductions apply within validated latency and calibration budgets. If latency exceeds the knee, oscillatory redistribution can reappear. If effort-proximal sensing drops out, overload detection becomes systematically biased. Regime maps should be recalculated when the window stride or instrumentation changes. Field use requires auditing ECE, τ , and channel availability online. Outside admissible regions, the policy should downshift to conservative damping. These constraints are intended to support implementation and certification workflows. They complement, rather than replace, system-level hazard analysis and testing.

5.1. Future Work

Future work should move from surrogate-closed-loop evidence to hardware- and interface-level validation with explicit actuation and measured operator responses, including end-to-end latency measurements and dropout patterns under real wireless conditions. Methodologically, two extensions are high priority: (i) a multi-objective formulation that reports Pareto fronts over overload reduction, oscillation suppression, and throughput rather than ranking single schedules, and (ii) a delay-robust synthesis step (e.g., LMI- or Lyapunov–Krasovskii-based conditions) that produces certified parameter regions for $\{\kappa_{\min}, \kappa_{\max}, \gamma, s\}$ under bounded τ and bounded estimator error. On the learning side, the estimator should incorporate explicit drift detection and recalibration triggers to maintain ECE below the control-relevant threshold in longitudinal deployments, and should quantify uncertainty in $\hat{L}$ so that the policy can reduce aggressiveness when the estimate is unreliable rather than reacting to noise.

Also, the future research should validate the full loop in hardware with measured actuation and operator response, using instrumented end-to-end timing and dropout characterization under realistic networking conditions. It should also extend the synthesis step to explicitly certify delay and estimation error tolerance over controller parameters, and integrate uncertainty-aware adaptation so that the policy downshifts when the estimate

becomes unreliable. Finally, the evaluation should expand to longitudinal and cross-context studies that quantify recalibration frequency, drift detection accuracy, and human acceptance outcomes under sustained operation, while preserving the same auditability commitments for threshold selection and subject-level reporting.

Moreover, further developments should explicitly connect the proposed estimator policy interface to deployment constraints that are not represented by public traces or surrogate plants. A practical next step is an edge-oriented real-time implementation that reports measured end-to-end latency and jitter under realistic wireless conditions and computes load and control updates within a fixed time budget aligned to the chosen window cadence. Experimental validation should then be conducted with physical human autonomy teaming tasks where the adaptation policy has authority over measurable autonomy parameters and where operator response is captured through synchronized effort proximal sensing and task outcome logs, enabling causal attribution of overload reduction rather than counterfactual inference. Finally, integration with industrial systems should address interoperability with existing control stacks and safety workflows by defining a minimal interface for load conditioned gain limits, logging and audit requirements, and fail safe behavior under dropout and sensor drift, so that the method can be evaluated within production-grade monitoring and change management processes.

5.2. Take-Home Messages

The main takeaway is that ergonomics becomes control-relevant only when the load estimate is both calibrated and timely, because it directly shapes the effective loop gain, and therefore, the regime of the coupled human autonomy dynamics. When those conditions hold, moderate load-aware scheduling can deliver large reductions in overload incidence with throughput preserved, whereas aggressive schedules under delay or miscalibration are the dominant failure mode and can degrade throughput or increase switching and oscillation. Practically, the results translate into concrete design constraints for real deployments: maintain end-to-end latency below roughly 300 ms at the operating window cadence, maintain estimator calibration at $ECE \leq 0.05$ under subject-wise evaluation, and prioritize reliable effort-proximal sensing because rare high-load detection is, otherwise, substantially impaired.

Author Contributions: Conceptualization, V.A.; methodology, V.A.; software, V.A. and N.G.; validation, V.A. and N.G.; formal analysis, V.A.; investigation, V.A.; resources, V.A.; data curation, V.A.; writing—original draft preparation, V.A. and N.G.; writing—review and editing, G.P.; visualization, V.A.; supervision, V.A.; project administration, V.A.; funding acquisition, V.A. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Publicly available datasets were analyzed in this study. The Human-swarm Interaction Dataset: Real-Time Human Interaction with Virtual Swarms in Shared Physical Space is available via Zenodo [34]. SenseCobot (v2) is available via Zenodo [33]. RoHuCAD: Robots and Humans Collaborative Anomaly Detection (v1.0) is available via Zenodo [32]. HRI30: An Action Recognition Dataset for Industrial Human–Robot Interaction (v1.0) is available via Zenodo [31]. The Force Myography dataset for Human–Robot Interactions is available via Zenodo [29]. The Multi-channel Surface EMG Dataset for Fatigue analysis is available via Zenodo [30]. The Wearable Device Dataset from Induced Stress and Structured Exercise Sessions (v1.0.1) is available via PhysioNet [35]. The WorkStress3D dataset (Version 11) is available via Mendeley Data [36]. No new data were created in this study.

Acknowledgments: The authors acknowledge the creators and maintainers of the publicly available datasets used in this study, and the infrastructure providers that host them (Zenodo, PhysioNet, and Mendeley Data).

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AUROC	Area Under the Receiver Operating Characteristic Curve
AUPRC	Area Under the Precision–Recall Curve
BAcc	Balanced Accuracy
ECE	Expected Calibration Error
EMG	Electromyography
FMG	Force Myography
HITL	Human-in-the-Loop
HRC	Human–Robot Collaboration
HSI	Human–Swarm Interaction
HR	Heart Rate
HRV	Heart Rate Variability
IMU	Inertial Measurement Unit
ISS	Input-to-State Stability
IQR	Interquartile Range
LMI	Linear Matrix Inequality
MSE	Mean Squared Error
NASA-TLX	National Aeronautics and Space Administration Task Load Index
PSO	Particle Swarm Optimization
RMSE	Root Mean Squared Error
RULA	Rapid Upper Limb Assessment
sEMG	Surface Electromyography

Appendix A. Fixed Constants, Thresholds, and Sweep Ranges

This appendix collects fixed numerical specifications and sweep grids referenced by the main text to reduce cognitive load in Sections 2.6 and 2.7. All values are defined on training blocks only unless explicitly stated otherwise.

Appendix A.1. Reproducible Construction Rules for d_k and ℓ_k

Table A1. Reproducible rules for constructing the demand drive d_k and the identification proxy ℓ_k . The rules are fixed prior to testing; dataset-specific instantiations occur only through channel availability masks $M_{k,m}$ and the presence/absence of context labels c_k .

Symbol	Rule	Notes
d_k	(12)	Weighted mixture of normalized effort features and coordination-intensity proxy; weights fixed globally
ℓ_k (phase)	(13)	$\chi(\cdot)$ maps protocol phases to $\{0, 0.5, 1\}$ (rest/moderate/high) on training only
ℓ_k (effort)	$\ell_k = \sigma((\tilde{e}_{k,m^*} - \mu_{tr}) / (\sigma_{tr} + \epsilon))$	m^* is a designated effort-proximal modality (FMG/EMG) when phase labels are unavailable

Appendix A.2. Regime Labeling Thresholds (Training-Only Selection)

Table A2. Appendix summary of regime-labeling thresholds and training-only selection rules. Thresholds are fixed prior to test reporting to avoid circularity.

Quantity	Definition	Selection Rule (Training Only)
L_{hi}	overload threshold on $[0, 1]$	set to the q_{hi} -quantile of $\hat{L}_k$ during high-demand phases (if labeled) or upper quantile of effort proxy otherwise
θ_{hi}	time-above threshold cutoff	fixed to 0.20 (overload if $\geq 20\%$ of windows exceed L_{hi})
θ_{osc}	oscillation power cutoff	set to the q_{osc} -quantile of P_{osc} in training equilibrated blocks; used unchanged on test

Appendix A.3. Sweep Parameters and Ranges

Table A3. Appendix summary of sweep parameters and ranges. G = global fixed/swept; F = fit on training blocks only.

Param.	Range/Grid	Set
ρ	$[0.90, 0.99]$	F if recovery labels; else G sweep
α	$[0.01, 0.50]$	G sweep (no test tuning)
κ_{min}	$[0, 0.3]$	G sweep; see (4)
κ_{max}	$[0.3, 1]$	G sweep; see (4)
γ	$[0, 1]$	G sweep; see (5)
τ	$\{0, 50, 100, 200, 400\}$ ms	grid; $\delta = \lceil \tau / \Delta t \rceil$
r	$\{1, 2, 4\}$	grid
p	$\{0, 0.05, 0.10\}$	grid

Appendix A.4. Fixed Constants for Regime Mapping and Time-Scale Conversion

Table A4. Fixed constants for regime mapping and spectral power computation. These values are not tuned on test data and are reported for auditability.

Knob	Fixed Value	Role
q_{hi}	0.80	Quantile used to set L_{hi} from training high-demand blocks (or effort-proxy high quantile)
q_{osc}	0.90	Quantile used to set θ_{osc} from training equilibrated blocks
N_{seg}	128 windows	Segment length for spectral power; maps to time length $128 \Delta t_{win}$
$\mathcal{B}_{osc}$	Appendix A.4	Oscillation band in Hz; converted per dataset via Δt_{win}
H (finite-horizon gain)	10 steps	Optional bound for $\max_{\ell \leq H} \ CA^{\ell-1}B\ $ in (A10)

For completeness, the time-scale mapping used throughout is as follows: (i) $\delta = \lceil \tau / \Delta t_{win} \rceil$ for latency injection; (ii) oscillation-band conversion from Hz to cycles-per-window uses $f = \nu / \Delta t_{win}$; (iii) all regime metrics are computed on the windowed cadence.

Appendix B. Explicit Computation of the Small-Gain Constant $\bar{\psi}$

This appendix provides the detailed, auditable upper bounds for the factors in Equation (17), separated by execution mode (CL-Sim vs. CL-Replay). The main text uses only the existence and computability of these bounds.

Appendix B.1. Bound on $|\partial D / \partial Z|$

Recall the demand drive definition (Equation (12)):

$$d_k := \sum_{m \in \mathcal{M}} w_m M_{k,m} \tilde{e}_{k,m} + w_z \tilde{q}_k, \quad (\text{A1})$$

where $\tilde{q}_k$ is a normalized coordination-intensity proxy derived from z_k . In the evaluation interface, $\tilde{q}_k$ is constructed by a robust affine normalization on training blocks,

$$\tilde{q}_k := \frac{q_k - \text{median}(q.)}{\text{IQR}(q.) + \epsilon}, \quad q_k := \mathcal{Q}(z_k), \quad (\text{A2})$$

with a fixed mapping $\mathcal{Q}(\cdot)$ chosen once per dataset family (Section 2.8) and then held fixed. We assume $\mathcal{Q}$ is Lipschitz on the bounded range of z encountered in training blocks, with constant K_Q :

$$|q(z_1) - q(z_2)| \leq K_Q |z_1 - z_2|. \quad (\text{A3})$$

Then, using (A2),

$$\left| \frac{\partial \tilde{q}}{\partial z} \right| \leq \frac{K_Q}{\text{IQR}(q.) + \epsilon}. \quad (\text{A4})$$

Because only the $w_z \tilde{q}_k$ term depends on z_k , we obtain the conservative bound

$$\bar{\psi}_{dz} := \left| \frac{\partial d}{\partial z} \right| \leq w_z \frac{K_Q}{\text{IQR}(q.) + \epsilon}. \quad (\text{A5})$$

In practice, we set $K_Q = 1$ when $\mathcal{Q}$ is the identity or a monotone scaling of z (the default in CL-Sim where z is already a normalized intensity proxy), and otherwise we estimate K_Q on training blocks by finite differences:

$$K_Q \approx \max_{k \in \text{train}} \frac{|q_{k+1} - q_k|}{|z_{k+1} - z_k| + \epsilon}. \quad (\text{A6})$$

All quantities in (A5) are computed on training blocks only and then fixed.

Appendix B.2. Bound on $|\partial Z / \partial U|$ in CL-Sim

In CL-Sim, z_k is generated by the identified surrogate plant (Equation (20)):

$$\hat{x}_{k+1} = A \hat{x}_k + B u_k + w_k, \quad \hat{z}_k = C \hat{x}_k, \quad (\text{A7})$$

with (A, B, C) fit on training blocks only and with bounded residual w_k . For a one-step sensitivity, ignoring w_k and holding $\hat{x}_k$ fixed,

$$\frac{\partial \hat{z}_{k+1}}{\partial u_k} = C B. \quad (\text{A8})$$

Thus, a conservative one-step bound is

$$\bar{\psi}_{zu}^{(1)} := \|C B\|. \quad (\text{A9})$$

Because the policy can influence z over multiple steps, we also report a finite-horizon bound over H steps (Appendix A fixes H):

$$\bar{\psi}_{zu}^{(H)} := \max_{1 \leq \ell \leq H} \|C A^{\ell-1} B\|. \quad (\text{A10})$$

We then set

$$\bar{\psi}_{zu} := \bar{\psi}_{zu}^{(H)}, \quad (\text{A11})$$

which is computable once (A, B, C) are identified on training blocks. Any matrix norm consistent with the vector norm used in the ISS statement is acceptable; in implementation, we use the induced 2-norm.

Appendix B.3. Bound on $|\partial Z / \partial U|$ in CL-Replay

In CL-Replay, z_k is exogenous because recorded traces are not altered by the computed u_k (Table 8). Therefore,

$$\frac{\partial z}{\partial u} = 0 \quad \Rightarrow \quad \bar{\psi}_{zu} = 0, \quad (\text{A12})$$

and the small-gain loop product defining $\bar{\psi}$ collapses. This is the formal reason why CL-Replay cannot support stability or attractor-elimination claims: the loop is open with respect to u .

Appendix B.4. Bound on $|\partial U / \partial L|$ from the Schedule

The policy combines gain moderation (Equation (4)) and compliance shaping (Equation (5)). To obtain an auditable bound, we consider the dependence of the policy on the load argument only and treat z -dependent terms as bounded signals (consistent with Assumption 3). Let the schedule be

$$\kappa(L) = \kappa_{\min} + (\kappa_{\max} - \kappa_{\min}) \sigma\left(\frac{L_0 - L}{s}\right), \quad (\text{A13})$$

where $\sigma(\cdot)$ is logistic. The logistic derivative satisfies $\sup_x \sigma'(x) = 1/4$, hence

$$\left| \frac{d\kappa}{dL} \right| \leq \frac{\kappa_{\max} - \kappa_{\min}}{4s}. \quad (\text{A14})$$

For compliance shaping,

$$u = u^{\text{base}} - \gamma \phi(L) \Delta z, \quad (\text{A15})$$

with $\phi(\cdot) \in [0, 1]$ monotone. If ϕ is chosen as a squashing map with bounded slope, we can take

$$\left| \frac{d\phi}{dL} \right| \leq K_\phi, \quad (\text{A16})$$

where K_ϕ is fixed by design (e.g., $K_\phi = 1/(4s_\phi)$ for logistic with slope s_ϕ). If we additionally bound $|\Delta z| \leq \Delta z_{\max}$ on the relevant operating region (estimated on training blocks in CL-Sim, or fixed from normalized scaling conventions), then

$$\left| \frac{\partial u}{\partial L} \right| \leq \underbrace{K_{u\kappa}}_{\text{policy scaling}} \left| \frac{d\kappa}{dL} \right| + \gamma K_\phi \Delta z_{\max}, \quad (\text{A17})$$

where $K_{u\kappa}$ captures how κ enters the autonomy action magnitude (a known linear scaling in the surrogate controller implementation). Substituting (A14) gives the explicit auditable bound:

$$\bar{\psi}_{uL} := \sup_{L \in [L_{\text{hi}}, 1]} \left| \frac{\partial u}{\partial L} \right| \leq K_{u\kappa} \frac{\kappa_{\max} - \kappa_{\min}}{4s} + \gamma K_\phi \Delta z_{\max}. \quad (\text{A18})$$

All quantities in (A18) are either sweep parameters $(\kappa_{\min}, \kappa_{\max}, \gamma, s)$, fixed design constants $(K_\phi, K_{u\kappa})$, or training-only empirical bounds $(\Delta z_{\max})$.

Appendix B.5. Final Computable Expression for $\bar{\psi}$

Combining (A5), (A11), and (A18), we report for CL-Sim:

$$\bar{\psi}_{\text{CL-Sim}} \leq \left(w_z \frac{K_Q}{\text{IQR}(q.) + \epsilon} \right) \left(\max_{1 \leq \ell \leq H} \|C A^{\ell-1} B\| \right) \left(K_{u\kappa} \frac{\kappa_{\max} - \kappa_{\min}}{4s} + \gamma K_\phi \Delta z_{\max} \right). \quad (\text{A19})$$

For CL-Replay, $\bar{\psi}_{\text{CL-Replay}} = 0$ by (A12). We emphasize that the CL-Sim bound is intentionally conservative; it is used to justify the directionality of the sufficient-condition inequalities and to support auditability of the design constraints, not to provide a tight stability margin estimate.

Appendix B.6. No-Leakage Commitments

All empirical quantities entering (A19) are computed on training blocks only. In particular: (A, B, C) are fit on training blocks; $\text{IQR}(q.)$ and the normalization used to define $\tilde{q}$ are computed on training blocks; $\Delta z_{\max}$ is computed as a robust upper quantile on training blocks and the horizon H is fixed globally (Appendix A). The sweep parameters $(\kappa_{\min}, \kappa_{\max}, \gamma, s)$ are never chosen by optimizing test outcomes; they are only swept over explicit grids, and any reported regime boundaries are computed under the fixed rule in Equation (21) with thresholds fixed prior to test reporting (Appendix A).

Finally, to ensure the bound is aligned with the delay-aware condition (19), the same cadence mapping $\delta = \lceil \tau / \Delta t_{\text{win}} \rceil$ is used when reporting delay-induced variation terms; the time-scale conversion conventions are consolidated in Appendix A.

Appendix C. Consolidated Symbol Glossary

Table A5. Symbol glossary for principal variables used across the manuscript.

Symbol	Type	Control-Theoretic or Physical Interpretation
$L(t)$	state	Latent ergonomic-load state.
$\hat{L}_k$	estimate	Estimated load used for feedback scheduling.
ℓ_k	proxy	Identification proxy target for training only.
L_{hi}	threshold	High-load threshold defining overload set.
L_0	reference	Nominal operating point for schedule transition.
$d(t)$	input	Demand or exertion drive exciting the load state.
d_k	input	Windowed demand drive used in discrete-time model.
$\hat{d}_k$	estimate	Estimated demand proxy used in consistency loss.
$\omega(t)$	disturbance	Unmodeled load disturbance and residual variability.
ω_k	disturbance	Discrete disturbance in dissipative update.
λ	rate	Recovery rate in continuous-time dissipation.
ρ	factor	Discrete recovery factor, $0 < \rho < 1$.
α	gain	Demand-to-load coupling gain.
$u(t)$	input	Autonomy-side control or adaptation input.
u_k	input	Discrete autonomy input produced by policy.
u_k^{base}	baseline	Baseline autonomy action without ergonomics coupling.
κ_k	gain	Adaptation gain scheduled by estimated load.
$\kappa_{\min}$	bound	Minimum scheduled gain under high load.
$\kappa_{\max}$	bound	Maximum scheduled gain under low load.
γ	gain	Compliance-shaping strength, acts as damping.
s	slope	Schedule sharpness parameter.
$\sigma(\cdot)$	map	Squashing map for bounded schedules and scaling.
$\phi(\cdot)$	map	Monotone load-to-damping map.
Δz_k	signal	Coordination error or correction signal in shaping.
x_k	state	Plant or coordination subsystem state.

Table A5. Cont.

Symbol	Type	Control-Theoretic or Physical Interpretation
$\hat{x}_k$	state	Surrogate plant state in CL-Sim.
z_k	output	Coordination observable, task-relevant output.
$\hat{z}_k$	output	Surrogate output used in CL-Sim evaluation.
y_k	measurement	Measured multistream observables.
ν_k	disturbance	Plant-side disturbance in state update.
ϵ_k	noise	Measurement noise corrupting observables.
$f(\cdot)$	dynamics	Plant state transition map.
$h(\cdot)$	output map	Output map from state to coordination variable.
$g(\cdot)$	sensor map	Measurement map from coordination to observables.
$\pi(\cdot)$	policy	Load-aware policy mapping $\hat{L}$ to u .
K_L	Lipschitz	Policy Lipschitz constant w.r.t. load argument.
ε_L	error bound	Uniform estimator error bound.
W_k	window	Time window of raw signals used for features.
S_k	segment	Multivariate segment inside a window.
$\Phi(\cdot)$	feature map	Window-to-feature operator.
φ_k	feature	Feature vector extracted from window.
$\bar{\varphi}_k$	normalized	Robust-normalized feature vector.
c_k	context	Task phase or context label in a window.
T	length	Window duration for feature extraction.
Δ	stride	Window stride controlling update cadence.
Δt	step	Sampling interval for discrete-time mapping.
Δt_{win}	cadence	Window update cadence in experiments.
τ	latency	End-to-end latency used for impairment injection.
δ	steps	Discrete delay steps, $\delta = \lceil \tau / \Delta t \rceil$.
r	factor	Downsampling factor used in impairments.
p	probability	Packet-loss probability in impairment operator.
$V(\cdot)$	Lyapunov	ISS Lyapunov function for plant stability.
V_k	composite	Composite Lyapunov candidate including load.
$\gamma_u(\cdot)$	gain	ISS gain from autonomy input to state.
$\gamma_v(\cdot)$	gain	ISS gain from plant disturbance to state.
$\bar{\psi}$	constant	Upper bound on closed-loop sensitivity product.
$\bar{\psi}_\delta$	constant	Delay-induced sensitivity inflation bound.
$\Gamma(\delta)$	bound	Variation accumulation bound over δ steps.
T_{hi}	metric	Fraction of windows above L_{hi} .
P_{osc}	metric	Band-limited oscillation power proxy.
θ_{hi}	cutoff	Overload occupancy cutoff.
θ_{osc}	cutoff	Oscillation power cutoff.
$r \in \{\text{EQ}, \text{OSC}, \text{OL}\}$	label	Regime label for regime maps.
N	count	Number of windows in a segment.
N_{seg}	count	Segment length for spectral power estimates.

References

1. Musić, S.; Hirche, S. Control sharing in human-robot team interaction. *Annu. Rev. Control* **2017**, *44*, 342–354. [\[CrossRef\]](#)
2. Hussein, A.; Ghignone, L.; Nguyen, T.; Salimi, N.; Nguyen, H.; Wang, M.; Abbass, H.A. Characterization of indicators for adaptive human-swarm teaming. *Front. Robot. AI* **2022**, *9*, 745958. [\[CrossRef\]](#)
3. Zhang, Y.; Wei, Y.; Ye, Z.; Liu, S.; Chen, H.; Yan, Y.; Chen, J. Robust Closed–Open Loop Iterative Learning Control for MIMO Discrete-Time Linear Systems with Dual-Varying Dynamics and Nonrepetitive Uncertainties. *Mathematics* **2025**, *13*, 1675. [\[CrossRef\]](#)
4. Pérez-Ibacache, R.; Castro, R.S.; Pimentel, G.A.; Bizzo, B.S. Dynamic output-feedback decentralized control synthesis for integration of distributed energy resources in AC microgrids. *IEEE Trans. Smart Grid* **2021**, *13*, 1225–1237. [\[CrossRef\]](#)
5. Fortini, L.; Lorenzini, M.; Kim, W.; De Momi, E.; Ajoudani, A. A framework for real-time and personalisable human ergonomics monitoring. In *Proceedings of the 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Las Vegas, NV, USA, 24 October–24 January 2021*; IEEE: New York, NY, USA, 2020; pp. 11101–11107.
6. Hwang, S.; Agada, P.; Kiemel, T.; Jeka, J.J. Identification of the unstable human postural control system. *Front. Syst. Neurosci.* **2016**, *10*, 22. [\[CrossRef\]](#) [\[PubMed\]](#)

7. Krishnan, A.; Yang, X.; Seth, U.; Jeyachandran, J.M.; Ahn, J.Y.; Gardner, R.; Pedigo, S.F.; Blom-Schieber, A.W.; Banerjee, A.G.; Manohar, K. Data-driven ergonomic risk assessment of complex hand-intensive manufacturing processes. *Commun. Eng.* **2025**, *4*, 45. [\[CrossRef\]](#)
8. Yang, Z.; Song, D.; Ning, J.; Wu, Z. A systematic review: Advancing ergonomic posture risk assessment through the integration of computer vision and machine learning techniques. *IEEE Access* **2024**, *12*, 180481–180519. [\[CrossRef\]](#)
9. Zheng, S.; Ren, M.; Luo, X.; Zhang, H.; Feng, G. The Third Closed-Loop Control for Compensating Light Power Fluctuations in the Interferometric Fiber-Optic Gyroscope. *J. Russ. Laser Res.* **2023**, *44*, 247–255. [\[CrossRef\]](#)
10. Hart, S.G. NASA-task load index (NASA-TLX); 20 years later. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting, San Francisco, CA, USA, 16–20 October 2006*; Sage Publications Sage: Los Angeles, CA, USA, 2006; Volume 50, pp. 904–908.
11. Park, J.H.; Kim, S.H. Design of a Stable Discrete-time Output-feedback Decentralized Controller for Uncertain Continuous-time Large-scale Nonlinear Systems. *Int. J. Control Autom. Syst.* **2025**, *23*, 1402–1410. [\[CrossRef\]](#)
12. Chen, P.; Liu, S.; Zhang, D.; Yu, L. Adaptive event-triggered decentralized dynamic output feedback control for load frequency regulation of power systems with communication delays. *IEEE Trans. Syst. Man Cybern. Syst.* **2021**, *52*, 5949–5961. [\[CrossRef\]](#)
13. Dani, A.P.; Salehi, I.; Rotithor, G.; Trombetta, D.; Ravichandar, H. Human-in-the-loop robot control for human-robot collaboration. *IEEE Control Syst. Mag.* **2024**, *40*, 29–56. [\[CrossRef\]](#)
14. Teo, G.; Reinerman-Jones, L.; Matthews, G.; Szalma, J.; Jentsch, F.; Hancock, P. Enhancing the effectiveness of human-robot teaming with a closed-loop system. *Appl. Ergon.* **2018**, *67*, 91–103. [\[CrossRef\]](#)
15. Bjurling, O.; Arvola, M.; Alfredson, J.; Prytz, E.; Ziemke, T. Trajectories of attention and control in human-machine interactions: The case of swarms in maritime search and rescue. *Theor. Issues Ergon. Sci.* **2025**, *26*, 816–837. [\[CrossRef\]](#)
16. Wattearachchi, W.D.; Lakshika, E.; Kasmarik, K.; Barlow, M. A Study on Human-Swarm Interaction: A Framework for Assessing Situation Awareness and Task Performance. *arXiv* **2025**, arXiv:2503.14810. [\[CrossRef\]](#)
17. Kathiravelu, P.; Arnold, M.G.; Vijay, S.; Jagwani, R.; Goyal, P.; Goel, A.K.; Li, N.; Horn, C.; Pan, T.; Kothare, M.V.; et al. Distributed executions with CONTROL-CORE: Integrated development environment (IDE) for closed-loop neuromodulation control systems. *Clust. Comput.* **2025**, *28*, 697. [\[CrossRef\]](#)
18. Kim, S.K.; Lim, S.; Ahn, C.K. Decentralized critical damping position synchronizer for multiservo drives via Feedback-Loop intelligentization approach. *IEEE Trans. Ind. Electron.* **2024**, *72*, 1368–1378. [\[CrossRef\]](#)
19. Tchimini, J.; Dideriksen, J.L.; Dosen, S. EMG feedback improves grasping of compliant objects using a myoelectric prosthesis. *J. Neuroeng. Rehabil.* **2023**, *20*, 119. [\[CrossRef\]](#) [\[PubMed\]](#)
20. Fam, I.; Soubra, H.; Gamal, N. Human-Swarm Interaction Methods' Effect on Human Psychophysiology. In *Proceedings of the 2023 Eleventh International Conference on Intelligent Computing and Information Systems (ICICIS), Cairo, Egypt, 21–23 November 2023*; IEEE: New York, NY, USA, 2023; pp. 184–189.
21. Distefano, J.P.; Chowdhury, S.; Esfahani, E. Exploring human-swarm interaction dynamics in cyber-physical systems: A physiological approach. *J. Integr. Des. Process Sci.* **2024**, *27*, 200–210. [\[CrossRef\]](#)
22. Menanno, M.; Riccio, C.; Benedetto, V.; Gissi, F.; Savino, M.M.; Troiano, L. An ergonomic risk assessment system based on 3D human pose estimation and collaborative robot. *Appl. Sci.* **2024**, *14*, 4823. [\[CrossRef\]](#)
23. Haj Mahmoud, O.; Pontonnier, C.; Dumont, G.; Poli, S.; Multon, F. A neural networks approach to determine factors associated with self-reported discomfort in picking tasks. *Hum. Factors* **2023**, *65*, 1381–1393. [\[CrossRef\]](#)
24. Lorenzini, M.; Lagomarsino, M.; Fortini, L.; Gholami, S.; Ajoudani, A. Ergonomic human-robot collaboration in industry: A review. *Front. Robot. AI* **2023**, *9*, 813907. [\[CrossRef\]](#) [\[PubMed\]](#)
25. Iodice, F.; De Momi, E.; Ajoudani, A. Intelligent Framework for Human-Robot Collaboration: Dynamic Ergonomics and Adaptive Decision-Making. *J. Intell. Robot. Syst.* **2025**, *112*, 5. [\[CrossRef\]](#)
26. Caporaso, T.; Grazioso, S.; Di Gironimo, G. Development of an integrated virtual reality system with wearable sensors for ergonomic evaluation of human–robot cooperative workplaces. *Sensors* **2022**, *22*, 2413. [\[CrossRef\]](#) [\[PubMed\]](#)
27. Merlo, E.; Lamon, E.; Fusaro, F.; Lorenzini, M.; Carfi, A.; Mastrogiovanni, F.; Ajoudani, A. An ergonomic role allocation framework for dynamic human–robot collaborative tasks. *J. Manuf. Syst.* **2023**, *67*, 111–121. [\[CrossRef\]](#)
28. Cohen, Y.; Biton, A.; Shoval, S. Fusion of computer vision and AI in collaborative robotics: A review and future prospects. *Appl. Sci.* **2025**, *15*, 7905. [\[CrossRef\]](#)
29. Zakia, U.; Menon, C. Dataset on Force Myography for Human Robot Interactions, Version 1. *Zenodo* **2022**. [\[CrossRef\]](#)
30. Ebied, A.; Awadallah, A.M.; Abbass, M.A.; El-Sharkawy, Y. Multi-channel Surface EMG Dataset for Fatigue analysis, Version 1. *Zenodo* **2021**. [\[CrossRef\]](#)
31. Iodice, F.; De Momi, E.; Ajoudani, A. HRI30: An Action Recognition Dataset for Industrial Human-Robot Interaction, Version 1.0. *Zenodo* **2022**. [\[CrossRef\]](#)
32. Buś, S.; Kaniuka, J.; Świtlik, D.; Głowska, J.; Kozik, R. RoHuCAD: Robots and Humans Collaborative Anomaly Detection, Version 1.0. *Zenodo* **2024**. [\[CrossRef\]](#)
33. Borghi, S.; Zucchi, F.; Prati, E.; Ruvo, A.; Villani, V.; Sabattini, L.; Peruzzini, M. SenseCobot, Version 2. *Zenodo* **2024**. [\[CrossRef\]](#)

34. Mezey, D.; Bartashevich, P.; Hasbani, G.E.; Romanczuk, P.; Hamann, H.; Deffner, D.; James, D. Human-swarm Interaction Dataset: Real-Time Human Interaction with Virtual Swarms in Shared Physical Space, Version 1. *Zenodo* **2025**. [[CrossRef](#)]
35. Hongn, A.; Bosch, F.; Prado, L.; Bonomini, P. Wearable Device Dataset from Induced Stress and Structured Exercise Sessions, Version 1.0.1. *PhysioNet* **2025**, *101*, e215–e220. [[CrossRef](#)]
36. Dogan, G.; Patlar Akbulut, F. Stress analysis from physiological data under pressure: WorkStress3D Dataset, Version 11. *Mendeley Data* **2024**. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.