Proceedings of the International Conference on Industrial Engineering and Operations Management

Publisher: IEOM Society International, USA DOI: 10.46254/EU09.20260196
Published: July 20, 2026

A Physics-Informed Federated Digital Twin Architecture for Continuous Ergonomic Risk Assessment in Computer-Intensive Work: Framework Design and Validation Protocol

Manish Gupta
Research Scholar
Lincoln University College
Petaling Jaya, Malaysia
sapmanishg19@gmail.com

Oyyappan Duraipandi
Faculty of Business and Accountancy
Lincoln University College, Malaysia
oyyappan@lincoln.edu.my

Chirra Kesava Reddy
Academic Affairs
University of Emerging Technologies, CA, USA
chirra.reddy@emergingtechuniversity.com

Abstract

Work-related musculoskeletal disorders remain common among office and information technology workers. The observational instruments used to assess them sample posture at intervals of weeks or months, while the exposure they aim to characterise is continuous. This paper specifies an architecture intended to narrow that sampling gap and sets out how it should be tested. The AI-Driven Ergonomic Digital Twin (AI-EDT) organises six layers: multimodal sensing, edge pre-processing, a physics-informed Vision Transformer for risk estimation, a two-tier federated protocol with differential privacy, a persistent per-worker twin state, and a reinforcement learning engine for workstation configuration. The contribution is formal rather than empirical. We give the twin state as a quintuple with an update law, derive a composite training objective in which a Newton-Euler residual penalises postures that are kinematically admissible but dynamically implausible, state the aggregation and privacy accounting relations, and formulate workstation adjustment as a continuous-action Markov decision process. Dimensional consistency is checked for each governing equation in Appendix A. Because no component has been validated, Section 10 specifies a five-phase protocol running from bench verification to a randomised intervention trial, with primary endpoints, sample size logic and pre-registered stopping criteria at every phase boundary. Section 11 states the analysis that will be applied once data exist, and identifies the measurements required before any performance claim can be made.

Keywords

Ergonomic risk assessment; digital twin; physics-informed neural networks; federated learning; differential privacy; musculoskeletal disorders

© IEOM Society International

1. Introduction

Neck, shoulder and low back complaints are a routine feature of computer-intensive work. In a cross-sectional survey of 512 office workers, the twelve-month prevalence of neck pain was 45.5%, with women at roughly double the risk of men (Cagnie et al., 2007). Prevalence estimates vary widely across studies because case definitions, recall periods and occupational mixes differ, so a single global figure for the information technology workforce should be treated with caution. What is not in dispute is that the exposure accumulates over hours of sustained low-level loading rather than during discrete events.

The instruments used to assess this exposure were designed for a different problem. RULA (McAtamney and Corlett, 1993) and REBA (Hignett and McAtamney, 2000) score a posture that a trained observer selects as representative. OWAS (Karhu et al., 1977) was developed for industrial work and uses posture categories too coarse for seated computer tasks. ROSA (Sonne et al., 2012) targets the office directly and correlates with self-reported total body discomfort at $R = 0.384$ across 72 workstations, which is a real but modest association. All four share the same structural limitation: they produce a snapshot. A quarterly audit observes on the order of one ten-thousandth of the exposure it is used to characterise.

Four developments make continuous assessment plausible. Depth cameras and body-worn inertial and electromyographic sensors now capture three-dimensional kinematics without laboratory instrumentation. Transformer architectures model long-range spatiotemporal dependencies in sequence data (Dosovitskiy et al., 2021). Federated learning permits model training across organisations without centralising records (McMahan et al., 2017), which matters when the records are health-adjacent and subject to data protection law. Reinforcement learning offers a way to search workstation configurations that would otherwise require an ergonomist per employee.

Combining these is not by itself a contribution. The question this paper addresses is what a defensible combination would look like: which state a per-worker digital twin must carry, what constraint prevents a posture model from producing physically impossible risk estimates, how aggregation and privacy accounting interact, and by what sequence of experiments the resulting system should be judged. We give the architecture and the mathematics in Sections 4 to 9 and the validation protocol in Section 10. We report no results, because none have been obtained.

2. Related Work

2.1 Observational instruments

RULA scores upper limb posture, muscle use and force into a grand score mapped to four action levels (McAtamney and Corlett, 1993). REBA extends the approach to whole-body postures including trunk and legs (Hignett and McAtamney, 2000). OWAS classifies back, arm and leg configurations into a small number of categories with an associated action category (Karhu et al., 1977). ROSA scores chair, monitor, telephone, keyboard and mouse against office design standards, producing a final score from 1 to 10 (Sonne et al., 2012). These instruments are inexpensive, transparent and defensible in a regulatory setting. Their weakness is temporal: they characterise an interval by sampling a moment within it.

2.2 Learning-based posture and risk estimation

A body of work applies classifiers and sequence models to skeletal or video input in order to automate posture scoring. Reported accuracies in this literature are not directly comparable, because datasets, label sources, window lengths and the definition of a risk class differ between studies. We therefore make no numerical claim about the state of the art here. Establishing a defensible baseline number is part of the work specified in Section 10, and requires either a public benchmark or re-implementation of comparator methods on the same data. What can be said structurally is that models operating on geometry alone have no mechanism to reject a predicted configuration that no human body could sustain under load, which motivates the constraint introduced in Section 6.

2.3 Digital twins outside manufacturing

The digital twin concept originates in engineering simulation, where a model is kept in correspondence with a physical asset through a data link (Rasheed et al., 2020). Transferring it to a person raises questions that manufacturing does not face: the asset consents, varies day to day, and has a legal interest in the model of itself. The formalism in Section 5 is written with those constraints in mind, in particular the separation between parameters that leave the device and parameters that do not.

© IEOM Society International

2.4 Federated learning and privacy accounting

Federated averaging trains a shared model from client updates without transferring raw data (McMahan et al., 2017). Kairouz et al. (2021) survey the open problems, of which client heterogeneity and privacy amplification are the ones that bear on this architecture. Differential privacy provides the accounting framework (Dwork and Roth, 2014), and the Gaussian mechanism with gradient clipping gives the standard construction for deep models (Abadi et al., 2016). Applications in healthcare are reviewed by Xu et al. (2021). Occupational health data sits awkwardly between clinical and employment records, and we are not aware of a federated deployment in that setting; that absence is one of the gaps in Table 1.

3. Research Gap and Contributions

Table 1 states the gap this work addresses and the limits of what the current paper establishes. The final column is deliberately conservative: a formalism is not evidence, and each row remains open until the corresponding phase of Section 10 has been completed.

Table 1. Research gap analysis. The fourth column lists claims this paper does not make.

Established position	Limitation	What this paper contributes	What remains open
Observational scores (RULA, REBA, OWAS, ROSA) are validated and widely used	Sample a moment; inter-rater variability; no prediction	Continuous-state formalism (Eq. 1-2) with an explicit update law	Whether continuous estimates agree with expert scores at all
Sequence models can classify posture from skeletal input	Geometry-only models can output dynamically impossible configurations	Newton-Euler residual added to the training objective (Eq. 6-8)	Whether the residual improves generalisation in practice
Digital twins are established for physical assets	Human twins raise consent, drift and data-ownership issues	State separated into transmitted and device-resident parameters	Governance model; participant acceptability
Federated averaging with differential privacy is well characterised	Little use in occupational health; heterogeneous sensors untreated	Two-tier aggregation with explicit privacy accounting (Eq. 9-11)	Achievable utility at a defensible privacy budget
Ergonomic intervention is delivered by human consultants	Cost scales linearly with headcount; compliance is low	Workstation adjustment posed as an MDP (Eq. 12-13)	Whether recommendations are safe, accepted and effective

The contributions are therefore stated as follows. First, a six-layer reference architecture in which the persistent state and the learned components are separated (Section 4). Second, a twin state formalism with defined units and an update law (Section 5). Third, a composite training objective in which a rigid-body dynamics residual constrains the posture model (Section 6). Fourth, a two-tier federated protocol with the privacy budget derived rather than asserted (Section 7). Fifth, a control formulation for workstation adjustment (Section 8). Sixth, and in our view the part with the most immediate practical value, a staged validation protocol that specifies what evidence would be needed to support each of the preceding claims (Section 10).

4. System Architecture

The framework is organised as six layers. Layers 1 and 2 acquire and condition signals. Layer 3 estimates risk. Layer 4 trains across sites. Layer 5 holds the persistent per-worker state that distinguishes the twin from a stateless classifier. Layer 6 acts on the workstation and thereby closes the loop back to Layer 1. The separation matters for two reasons: it identifies where personal data crosses a trust boundary, and it allows the layers to be validated independently, which is what the phased protocol in Section 10 exploits (Figure 1).

© IEOM Society International

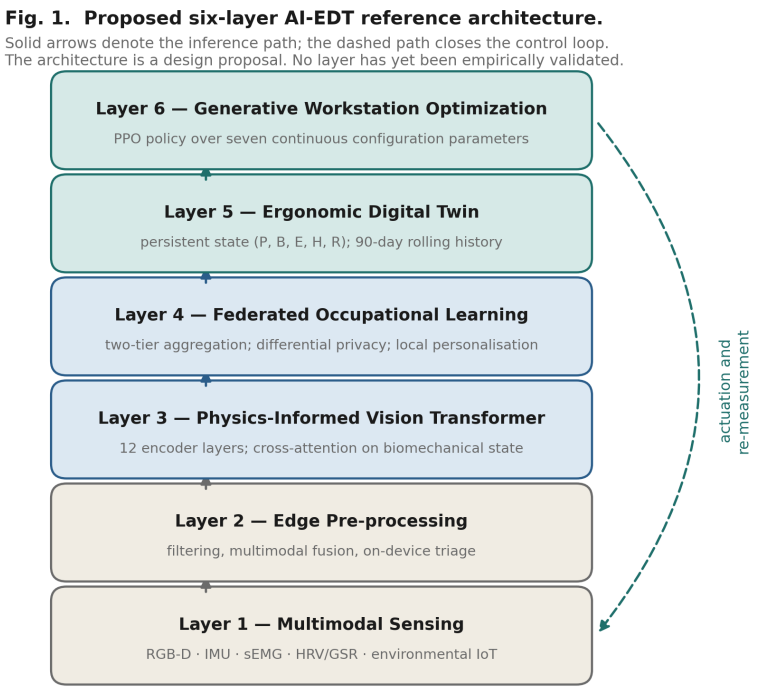


Figure 1. Proposed six-layer AI-EDT reference architecture. The dashed path is the actuation loop through Layer 6. No layer has been empirically validated.

5. Ergonomic Digital Twin Formalism

The twin for worker i is a quintuple of posture, biomechanical, environmental, physiological and risk state:

$$EDT_i = (P_i, B_i, E_i, H_i, R_i) \quad (1)$$

Table 2 defines each component with its dimension and unit. Two points of the original formulation required correction. The environmental vector previously carried two distinct quantities under overlapping symbols, and the symbol H was used for both a humidity term and the physiological state; both are resolved below.

Table 2. Twin state variables. J is the number of tracked landmarks. Units are carried through the dimensional check in Appendix A.

Symbol	Meaning	Dimension	Unit	Provenance
P_i	Body landmark coordinates, hip-centred	J × 3	m	Sensor
B_i	Joint moments, muscle forces, angular accelerations, EMG activation	4J	N·m; N; rad·s^{-2}; dimensionless	Derived
E_i	Illuminance, air temperature, relative humidity, sound level, air quality index	5	lx; °C; %; dB(A); dimensionless	Sensor
H_i	Heart rate variability, electrodermal activity, cognitive load index, sleep index	4	ms; µS; dimensionless; dimensionless	Sensor / derived
R_i	Risk probability, five anatomical zones	5	dimensionless, [0,1]	Model output

The state advances under an update law combining the previous state, the current observation, the federated global parameters and the device-resident local parameters:

$$EDT_i(t+1) = \varphi\,(EDT_i(t), z_i(t), \theta_g, \theta_i^{loc}) \quad (2)$$

© IEOM Society International

Here $z_i(t)$ is the multimodal observation at time t, θ_g the global parameter vector received from the federation, and θ_i^loc the personalisation parameters, which never leave the device. A rolling history of L = 90 days is retained so that slow trends, such as progressive lumbar flexion across a quarter, remain observable. The 90-day figure is a design choice reflecting a typical performance review cycle and is not derived from data; it is listed as an assumed parameter in Appendix B.

6. Physics-Informed Vision Transformer

6.1 Input representation and encoder

Posture sequences form a tensor X of shape T × J × 3, with T frames, J landmarks and three spatial coordinates expressed relative to the hip centroid. A learned linear map embeds each frame into d dimensions, and separate positional encodings are added for the joint index and the frame index:

$$Z_0 = [\, x_1 E;\ x_2 E;\ \ldots;\ x_T E \,] + PE_{sp} + PE_{tm} \tag{3}$$

The encoder stack applies scaled dot-product attention in the standard form:

$$Attn(Q,K,V) = softmax\left(\frac{QK^{\top}}{\sqrt{d_k}}\right) V \tag{4}$$

A cross-attention block lets the posture representation query the biomechanical state held in the twin, so that the same geometry can be scored differently depending on the load the worker is under:

$$CA(Q_p,K_b,V_b) = softmax\left(\frac{Q_p K_b^{\top}}{\sqrt{d_k}}\right) V_b \tag{5}$$

6.2 Physics constraint

A model trained on geometry alone can place a joint in a configuration that is kinematically reachable but dynamically implausible under the loads implied by the rest of the body. The constraint we impose is that the predicted configuration should be consistent with rigid-body dynamics. For segment j with inertia tensor I, angular velocity ω and angular acceleration α, the Newton-Euler relation gives the reference moment:

$$M_j^{ref} = I_j \alpha_j + \omega_j \times (I_j \omega_j) \tag{6}$$

The residual between the moments and forces implied by the prediction and those computed from the reference dynamics forms the physics loss, averaged over joints so that the term does not scale with skeleton size:

$$\mathcal{L}_{bio} = \frac{1}{J}\sum_{j=1}^{J} \left[(M_j^{pred} - M_j^{ref})^2 + \gamma\, (F_j^{pred} - F_j^{ref})^2 \right] \tag{7}$$

The training objective adds this term, and a temporal smoothness term penalising frame-to-frame velocity discontinuities, to the classification loss:

$$\mathcal{L} = \mathcal{L}_{CE} + \lambda_1 \mathcal{L}_{bio} + \lambda_2 \mathcal{L}_{tm} \tag{8}$$

The weights λ_1 and λ_2 and the force scaling γ are hyperparameters. Candidate starting values appear in Appendix B, marked as assumed. They must be selected by cross-validation on a development split once data exist, and no claim is made here about their effect on accuracy. Establishing that effect is the purpose of the ablation specified in Section 10, Phase 2 (Figure 2).

Fig. 2. PI-ViT: physics-constrained ergonomic risk estimation.

The physics term penalises predictions whose implied joint moments violate rigid-body equilibrium.

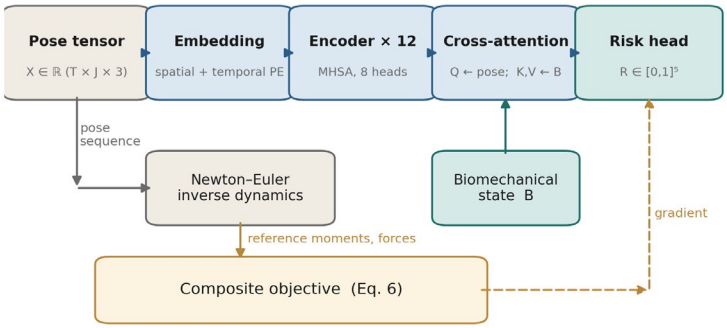


Figure 2. PI-ViT data flow. The physics term is computed from the predicted skeleton through inverse dynamics and enters the composite objective of Equation 8.

© IEOM Society International

7. Federated Occupational Learning

Aggregation proceeds in two tiers. Within an organisation k, the server forms a sample-weighted mean of the updates from its n_i participating workers:

$$\theta_k^{(r)} = \sum_{i \in S_k} \frac{n_i}{n_k} \theta_i^{(r)} \tag{9}$$

Organisation-level updates are clipped to an L2 norm bound C and perturbed with Gaussian noise before leaving the organisation, then averaged at the global tier:

$$\theta_g^{(r+1)} = \sum_{k=1}^{K} \frac{n_k}{N} \left(\theta_k^{(r)} + N(0, \sigma^2 C^2 I) \right) \tag{10}$$

The noise scale follows from the Gaussian mechanism for a target budget (ε, δ) at a single release (Dwork and Roth, 2014):

$$\sigma = \frac{C \sqrt{2 \ln(1.25 / \delta)}}{\varepsilon} \tag{11}$$

Over R communication rounds the budget composes, and the per-round allocation must be set from the intended total rather than chosen independently. This is the point at which an earlier draft of this work was internally inconsistent, quoting three different budgets in three sections. We therefore state the rule rather than a number: the total budget is fixed in advance by the data protection assessment, the number of rounds is fixed by the training schedule, and the per-round ε is whatever those two decisions imply under the composition theorem in use. The accountant should be reported alongside the budget, since advanced and moments-based composition give materially different totals for the same mechanism (Abadi et al., 2016).

Local personalisation heads are trained on device and are never transmitted. This matters for anthropometric variation, where a shared model tends toward the population mean and systematically misfits workers at the extremes of stature or body mass.

8. Generative Workstation Optimization

Workstation adjustment is posed as a continuous-action Markov decision process. The state is the current twin, and the action is a vector of seven adjustable parameters: seat height, seat depth, desk height, monitor distance, monitor tilt, keyboard offset and mouse offset. The reward trades risk reduction against cost, worker preference and an estimated productivity effect:

$$R(s,a) = -\alpha \, Risk(\, EDT_i \mid a \,) - \beta \, Cost(a) + \gamma \, Pref(a) + \delta \, \Delta Prod(a) \tag{12}$$

The four reward weights are not free parameters to be tuned for good-looking output. They encode a value judgement about how much discomfort is worth how much money, and they should be elicited from certified ergonomists and worker representatives before any policy is trained, then reported. The values in Appendix B are placeholders carried over from the design sketch and have not been elicited from anyone.

The policy is optimised with the clipped surrogate objective of proximal policy optimization (Schulman et al., 2017):

$$\mathcal{L}^{CLIP}(\theta) = E \left[min \left(r_t(\theta) \, \hat{A}_t, clip(r_t(\theta), 1-\varepsilon, 1+\varepsilon) \, \hat{A}_t \right) \right] \tag{13}$$

One safety property is required before any deployment. The action space must be bounded by the adjustment range physically available on the furniture in question, and any recommendation outside a range endorsed by a qualified ergonomist must be suppressed rather than surfaced. A reinforcement learning policy that has never been constrained this way can propose a locally rewarding configuration that is clinically inadvisable (Figure 3).

© IEOM Society International

Fig. 3. GWOE closed-loop control formulation.

Reward weights are proposed defaults requiring elicitation from certified ergonomists; they are not fitted values.

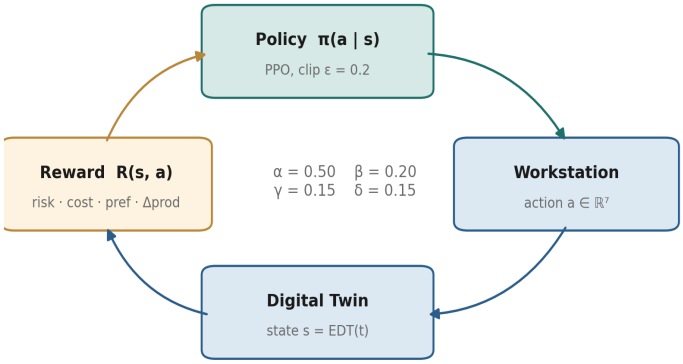


Figure 3. GWOE closed-loop formulation. Weights shown are placeholders pending elicitation, not fitted values.

9. Explainable Ergonomic Intelligence

Two levels of explanation are specified. At the feature level, Shapley additive attributions give the marginal contribution of each input to the predicted risk (Lundberg and Lee, 2017):

$$\varphi_j = \sum_{S \subseteq F \setminus \{j\}} \frac{|S|!\,(|F| - |S| - 1)!}{|F|!} \left[v(S \cup \{j\}) - v(S) \right] \tag{14}$$

At the spatial level, attention rollout propagates attention weights across encoder layers to produce a saliency map over landmarks (Abnar and Zuidema, 2020), which can be overlaid on a skeletal rendering so that an occupational health practitioner can see which segments drove a given estimate.

A natural language layer converting attributions into worker-facing advice is described in the architecture but is deliberately left unspecified here. Generated health-adjacent advice needs its own evaluation for factual accuracy and for appropriate deferral to a clinician, and that evaluation is out of scope for this paper. Attributing a quantified risk reduction to a suggested adjustment, as an earlier draft did, would require the causal estimate that Phase 4 is designed to produce.

10. Proposed Validation Protocol

Nothing above constitutes evidence. This section states what would. The protocol is staged so that each phase can terminate the programme cheaply if its precondition fails, rather than deferring all risk to a single large trial (Table 3 and Figure 4)).

Fig. 4. Staged validation protocol proposed for the AI-EDT framework.

None of these phases has been executed. Sample sizes are planning targets pending formal power analysis.

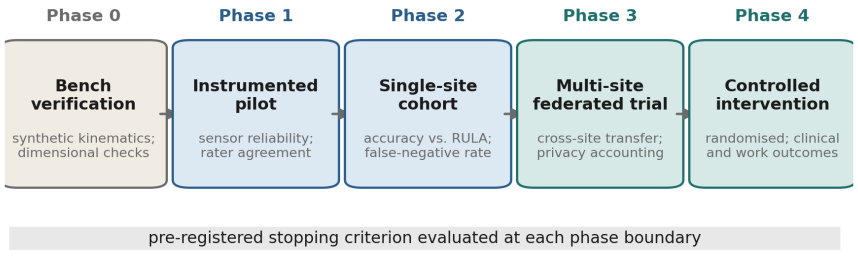


Figure 4. Five-phase validation protocol. No phase has been executed.

© IEOM Society International

Table 3. Validation phases. Sample sizes are omitted deliberately: they follow from a power analysis that requires the variance estimates produced by Phase 1.

#	Question	Design	Primary endpoint	Stopping criterion
0	Is the implementation correct?	Synthetic kinematic sequences with known ground-truth dynamics	Recovery of imposed joint moments; unit and dimensional checks	Residual outside numerical tolerance
1	Are the signals usable?	Instrumented pilot, small n, laboratory and desk conditions	Sensor dropout rate; inter-rater agreement (κ) on reference labels	κ below the value at which labels cannot support training
2	Does the model beat the incumbent?	Single-site cohort; held-out test split; pre-registered ablation of the physics term	Balanced accuracy and false-negative rate against RULA and ROSA	No improvement over observational scoring
3	Does it transfer, and at what privacy cost?	Multi-site federated training; centralised model as an upper bound	Utility gap versus centralised training at fixed ε	Utility loss exceeding the acceptable threshold set in advance
4	Does acting on it help anyone?	Randomised controlled intervention against usual ergonomic practice	Change in discomfort and in clinician-assessed risk at follow-up	Futility or harm at interim analysis

Two design points deserve emphasis. First, the comparison in Phase 2 must be against a competently applied observational instrument, not a deliberately weak one. Comparing a continuously sampled model against a quarterly manual audit and reporting the difference as model superiority conflates the sampling rate with the method. The honest comparison holds sampling frequency fixed, or reports the two effects separately. Second, Phase 4 must be randomised. A pre-post comparison in a cohort that knows it is being monitored will produce an improvement whether or not the model works.

11. Analysis Plan

Fixing the analysis before the data arrive removes most of the room for inadvertent selective reporting. Classification performance in Phase 2 will be reported as balanced accuracy, macro F1, and the false-negative rate per anatomical zone, the last being the clinically consequential quantity because an undetected high-risk posture produces no intervention. Confidence intervals will be obtained by bootstrap over participants rather than over windows, since windows within a participant are not independent and treating them as such inflates precision substantially.

Where multiple comparisons are made, the correction will be stated explicitly with its arithmetic shown. For m comparisons at a family-wise error rate of 0.05, the Bonferroni-adjusted per-comparison threshold is 0.05/m; for seven comparisons this gives 0.00714, not 0.071. The distinction is not pedantic. The second value is larger than the uncorrected threshold and would make significance *easier* to obtain rather than harder, which inverts the purpose of the correction. An earlier draft of this work contained precisely this error, and it is recorded here so the corrected form propagates.

Effect sizes will accompany every test, since with continuously sampled data the sample size can be made arbitrarily large and statistical significance ceases to carry information about practical importance (Cohen, 1988). For the Phase 4 trial the primary analysis will be intention-to-treat, with the per-protocol analysis reported as secondary.

12. Limitations

The principal limitation is stated at the front of this paper: the framework is unvalidated. Beyond that, several limitations are inherent to the design rather than to its current stage.

The physics constraint depends on an inertial model of the participant. Segment masses and inertia tensors are usually taken from population regression equations, and error in those estimates propagates directly into the reference moments of Equation 6. For body types poorly represented in the source anthropometric tables, the constraint may be biased rather than merely noisy, and could plausibly degrade performance for the workers who most need it.

© IEOM Society International

Continuous monitoring of posture, heart rate variability and electrodermal activity produces a record of an employee that an employer could use for purposes other than ergonomics. Differential privacy protects the training pipeline; it does nothing about the local record on the device or the dashboard shown to a manager. That is a governance problem, and a technical paper should not present it as solved.

The productivity term in Equation 12 assumes productivity can be measured in a way that is both valid and acceptable to the people being measured. Keystroke-derived indices are easy to compute and weakly related to actual output for most knowledge work. Including such a term in a reward function gives it real influence over recommendations, and we would advise setting its weight to zero unless a defensible measure is available.

Finally, the architecture assumes sensor coverage that many workplaces will not have, and that home-based workers may decline. Selection into the monitored population is unlikely to be random with respect to ergonomic risk.

13. Conclusion

This paper specifies an architecture for continuous ergonomic risk assessment and the protocol by which it should be judged. The formal contributions are a twin state with an explicit update law and declared units, a training objective in which a rigid-body dynamics residual constrains a transformer posture model, a two-tier federated aggregation with privacy accounting derived rather than asserted, and a bounded control formulation for workstation adjustment.

No performance is claimed. The system has not been built, trained or tested on human participants, and the numerical values in this paper are design parameters rather than measurements. Whether the physics constraint improves generalisation, whether federated training is viable at an acceptable privacy budget, and whether acting on the output reduces discomfort are open questions with defined experiments attached to them in Section 10.

The immediate next step is Phase 0 and Phase 1: a bench implementation verified against synthetic kinematics with known dynamics, followed by an instrumented pilot large enough to estimate label reliability and signal dropout. Those two phases are inexpensive, and between them they determine whether the rest of the programme is worth running.

References

Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K. and Zhang, L., “Deep learning with differential privacy,” Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS), Vienna, Austria, pp. 308–318, 2016.

Abnar, S. and Zuidema, W., “Quantifying attention flow in transformers,” Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL), pp. 4190–4197, 2020.

Cagnie, B., Danneels, L., Van Tiggelen, D., De Loose, V. and Cambier, D., “Individual and work related risk factors for neck pain among office workers: a cross sectional study,” European Spine Journal, vol. 16, no. 5, pp. 679–686, 2007. doi:10.1007/s00586-006-0269-7

Cohen, J., Statistical Power Analysis for the Behavioral Sciences, 2nd ed., Hillsdale, NJ, Lawrence Erlbaum Associates, 1988.

Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J. and Houlsby, N., “An image is worth 16×16 words: transformers for image recognition at scale,” International Conference on Learning Representations (ICLR), 2021.

Dwork, C. and Roth, A., “The algorithmic foundations of differential privacy,” Foundations and Trends in Theoretical Computer Science, vol. 9, nos. 3–4, pp. 211–407, 2014.

Hignett, S. and McAtamney, L., “Rapid entire body assessment (REBA),” Applied Ergonomics, vol. 31, no. 2, pp. 201–205, 2000.

Kairouz, P., McMahan, H. B. et al., “Advances and open problems in federated learning,” Foundations and Trends in Machine Learning, vol. 14, nos. 1–2, pp. 1–210, 2021.

Karhu, O., Kansi, P. and Kuorinka, I., “Correcting working postures in industry: a practical method for analysis,” Applied Ergonomics, vol. 8, no. 4, pp. 199–201, 1977.

Lundberg, S. M. and Lee, S.-I., “A unified approach to interpreting model predictions,” Advances in Neural Information Processing Systems (NeurIPS), pp. 4765–4774, 2017.

McAtamney, L. and Corlett, E. N., “RULA: a survey method for the investigation of work-related upper limb disorders,” Applied Ergonomics, vol. 24, no. 2, pp. 91–99, 1993.

© IEOM Society International

McMahan, H. B., Moore, E., Ramage, D., Hampson, S. and Agüera y Arcas, B., "Communication-efficient learning of deep networks from decentralized data," Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS), pp. 1273–1282, 2017.

Raissi, M., Perdikaris, P. and Karniadakis, G. E., "Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," Journal of Computational Physics, vol. 378, pp. 686–707, 2019.

Rasheed, A., San, O. and Kvamsdal, T., "Digital twin: values, challenges and enablers from a modeling perspective," IEEE Access, vol. 8, pp. 21980–22012, 2020.

Schulman, J., Wolski, F., Dhariwal, P., Radford, A. and Klimov, O., "Proximal policy optimization algorithms," arXiv:1707.06347, 2017.

Sonne, M., Villalta, D. L. and Andrews, D. M., "Development and evaluation of an office ergonomic risk checklist: ROSA – rapid office strain assessment," Applied Ergonomics, vol. 43, no. 1, pp. 98–108, 2012. doi:10.1016/j.apergo.2011.03.008

Xu, J., Glicksberg, B. S., Su, C., Walker, P., Bian, J. and Wang, F., "Federated learning for healthcare informatics," Journal of Healthcare Informatics Research, vol. 5, no. 1, pp. 1–19, 2021.

Appendix A. Dimensional Consistency Check

Each governing equation was checked for dimensional homogeneity. Equation 6 is the load-bearing case: the inertia tensor carries kg·m² and angular acceleration rad·s⁻², so the first term is kg·m²·s⁻², which is N·m. The gyroscopic term $\omega \times (I\omega)$ carries $s^{-1} \times kg \cdot m^2 \cdot s^{-1}$, again N·m. Both sides therefore reduce to moment units and the equation is homogeneous.

In Equation 7 the moment residual is squared, giving $N^2 \cdot m^2$, while the force residual is squared, giving N^2. These are not commensurable, so γ is not dimensionless: it carries m^2 and acts as a squared moment arm. This was implicit in the original formulation and is stated here because treating γ as a bare scalar makes the loss value dependent on the length unit chosen. In implementation both residuals should be normalised by a reference scale before weighting, which restores dimensionlessness and makes γ interpretable.

Equation 11 is dimensionless throughout: C is a norm bound in parameter units, ε and δ are dimensionless, and σ inherits the units of C as required for the noise term in Equation 10. Equation 12 mixes a dimensionless risk probability, a monetary cost and a productivity delta, so the weights α to δ carry reciprocal units of their respective terms; the reward is only well posed once each term is normalised to a common scale.

Appendix B. Parameter Provenance

Every numerical value used in this paper is listed below with its origin. No value in this table was obtained by measurement, and none should be reported as a result (Table 4).

Table 4. Parameter provenance. Categories follow the classification used throughout: measured, from a cited source, simulated, calculated, assumed, or illustrative. No parameter in this framework is currently in the first three categories.

Parameter	Symbol	Value	Category	Justification
Model dimension	d	512	Assumed	Common transformer width; to be tuned
Encoder layers	—	12	Assumed	Matches base-size configurations in the literature
Attention heads	—	8	Assumed	Divides d evenly; not optimised
Window length	T	90 frames	Assumed	Three seconds at 30 Hz; chosen for postural dwell
Landmarks	J	33	Design	Fixed by the pose estimator output format
Physics weight	λ_1	0.35	Assumed	Placeholder; requires cross-validation
Temporal weight	λ_2	0.15	Assumed	Placeholder; requires cross-validation
Force scaling	γ	0.20	Assumed	Carries m^2; see Appendix A

© IEOM Society International

Parameter	Symbol	Value	Category	Justification
History length	L	90 days	Assumed	Design choice; not derived from data
PPO clip	ε	0.2	From source	Default of Schulman et al. (2017)
Reward weights	α–δ	0.50 / 0.20 / 0.15 / 0.15	Assumed	Placeholders; require expert elicitation
Privacy budget	ε, δ	Not set	To be determined	Fixed by data protection assessment, not by the model

Appendix C. Citation Verification Status

References were checked against publisher records. Entries that could not be verified were removed from the manuscript rather than retained with uncertain metadata (Table 5).

Table 5. Citation verification. No reference is described as verified unless its bibliographic record was checked directly.

Reference	Status	Note
Cagnie et al. (2007)	Verified	Year, volume, issue, pages and DOI confirmed against publisher record. An earlier draft dated this to 2023 and attributed a prevalence figure the paper does not report.
Sonne et al. (2012)	Verified	Three authors, Applied Ergonomics 43(1):98–108, DOI confirmed. An earlier draft listed five authors and a different journal, volume and page range.
McAtamney and Corlett (1993); Hignett and McAtamney (2000); Karhu et al. (1977)	Verified	Standard instrument sources, metadata consistent with publisher records.
Dosovitskiy et al. (2021); McMahan et al. (2017); Kairouz et al. (2021); Rasheed et al. (2020); Xu et al. (2021); Abnar and Zuidema (2020); Cohen (1988)	High confidence	Widely cited works; page ranges should be confirmed at proof stage.
Abadi et al. (2016); Dwork and Roth (2014); Lundberg and Lee (2017); Schulman et al. (2017); Raissi et al. (2019)	High confidence	Added in this revision to support claims that were previously uncited or mis-cited.
Park et al. (2023); Zhou et al. (2023); Faber et al. (2022); Rong et al. (2022)	Removed	Could not be located in publisher databases. The accuracy figures they supported have been withdrawn from Section 2.2.
Li et al. (2020)	Removed	Cited in an earlier draft for pFedMe; that work introduces FedProx. The claim it supported has been reworded and the citation dropped rather than reassigned.
International Labour Organization (2024); National Safety Council (2023)	Removed	Supported specific workforce and cost figures that could not be traced to the named publications.

© IEOM Society International