

A Review of Human-Robot-Machine Collaboration in Industry 5.0

How a factory that listens to its workers could change the way we think about human-robot collaboration: the theory, the evidence from the code, and a proposed fix.

Ms. Pranali Nitnavare¹, Rutusha Somkuwar², Aryan Wadibhasme³, Deepak Sharma⁴, Yash Ingole⁵

^{1, 2, 3, 4}UG Scholars, Department of Industrial IoT, Priyadarshini College of Engineering, Nagpur, India Affiliated to Rashtrasant Tukdoji Maharaj Nagpur University, Nagpur - 440019, Maharashtra, India

Abstract - This paper presents the Industry 5.0 Human-Robot Machine Collaboration (HRMC) Cyber-Physical Digital Twin Platform, a Python 3 program with an embedded HTML/JavaScript dashboard. It simulates a miniature factory with three stations, three workers and three collaborative robots (cobots). For each worker, a Human Digital Twin generates synthetic physiological signals, and an edge classifier turns them into a Mental Pressure Index (MPI) and a Physical Fatigue Index (PFI). Depending on these indices, each cobot slows down, offers help or stops, while a live web dashboard shows a 3D factory view and four charts. Over 8-hour virtual shifts, the settings as supplied kept workers in the nominal state for only 0.1% of the time and in emergency rest for 75.2%, and output fell by about 89% compared with a no-intervention run. A proposed redesign using personal baselines, relative thresholds and a recovery credit raised nominal time to 76.4%, cut the defect rate from 22.1% to 16.7% and cost 12.4% of output over 30 simulated shifts. All worker data are synthetic, so the results describe how the program behaves, not how real people behave. This document also supports an Indian Copyright Office application for the computer programme.

Key Words: Industry 5.0, human-robot collaboration, digital twin, collaborative robot (cobot), worker fatigue monitoring, edge classifier

1. INTRODUCTION

1.1 Overview

Consider a collaborative robot (a cobot) working beside a person on an assembly line. In most factories today, the cobot moves at the speed programmed some time earlier and has no knowledge of whether the adjacent worker is calm, tired, or overwhelmed.

This work investigates a simple question: **what if the robot could sense the state of the worker beside it?** It builds a miniature factory with three stations, three workers and three cobots. For each worker it simulates heart rate, heart-rate variability, skin conductance, blood pressure, breathing, posture and further signals. An "edge" classifier turns these signals into two indices, a Mental Pressure Index (MPI) and a Physical Fatigue Index (PFI). The cobot then slows down, offers help, or stops, depending on their values. A live web page shows the 3D factory and the graphs.

1.2 Summary of Findings

- **The idea works as a concept.** Within about 17-28 seconds of simulated work, every cobot starts adapting to its worker, and the first full halts follow within 63-297 seconds.
- **As supplied, the settings are too jumpy.** Over an 8-hour shift the workers spend only 0.1% of the time in "nominal" and 75.2% in emergency rest. Output falls by about 89% compared with a no-intervention run.
- **A small redesign fixes the balance.** Personal baselines, relative thresholds and a recovery credit bring nominal time to 76.4%, cut the defect rate from 22.1% to 16.7%, and cost 12.4% of output (simulated, 30 shifts).

2. PROBLEM STATEMENT

Human-robot collaboration in industrial environments traditionally treats the worker as a relatively fixed resource operating at a predetermined pace, with safety primarily focused on preventing physical contact between humans and robots. However, a worker's physiological and physical state can change continuously due to factors such as fatigue, strain, posture, and workload. Existing approaches that evaluate these conditions only after a work period or rely on fixed

thresholds may therefore fail to support timely, individualized interventions.

To address this limitation, this research proposes a **Human-State-Aware Collaboration (HSAC)** approach, in which the worker is treated as a dynamic signal whose state can be continuously monitored and incorporated into robot decisionmaking. The proposed system uses a **Human Digital Twin** to generate multiple human-state signals and an edge-based classifier operating over a short sliding window to estimate worker strain within seconds. Based on the estimated state, the collaborative robot can adapt its operating behaviour, such as adjusting its feed rate or initiating interventions including throttling, assistance, or rest.

The approach further considers **individual working baselines** rather than applying identical fixed thresholds to all workers, recognizing that physiological responses vary between individuals. It also emphasizes that intervention should be accompanied by an appropriate recovery mechanism, so that reducing robot activity contributes to the worker's recovery rather than creating a continuous alarm condition. The central research problem is therefore to develop a human-state-aware collaborative system capable of **monitoring worker condition in real time and adapting robot operation according to the worker's current state**, while addressing the combined requirements of productivity, safety, fatigue management, and worker well-being.

In this context, the research moves beyond the conventional question of "How fast can the robot operate safely?" toward the more adaptive question of "How fast should the robot operate right now for this particular worker?" This perspective forms the basis of the proposed HSAC framework and its testable propositions regarding rapid strain detection, statedependent robot adaptation, personalized thresholds, and recovery-aware intervention.

3. OBJECTIVES

The objectives below follow directly from the problem statement and from the four testable propositions (P1 to P4) of the HSAC framework.

1. To build a simulated factory of three stations, three workers and three cobots, in which a Human Digital Twin generates synthetic physiological signals for each worker.
2. To estimate worker strain within seconds, through an edge classifier that converts the signals into a Mental Pressure Index (MPI) and a Physical Fatigue Index (PFI) (P1, rapid strain detection).
3. To let each cobot adapt its behaviour to the worker's current state by throttling its feed rate, offering task assistance or halting for rest (P2, state-dependent robot adaptation).
4. To test personal baselines and relative thresholds against one-size-fits-all thresholds (P3, personalized thresholds).
5. To add a recovery mechanism so that reduced robot activity helps the worker recover instead of creating a continuous alarm condition (P4, recovery-aware intervention).
6. To provide a live dashboard, with a 3D factory view and four charts, that makes the workers' states and the cobots' responses visible, and to evaluate the settings as supplied against a proposed redesign (v2) over repeated 8-hour virtual shifts.

4. LITERATURE REVIEW

Before developing the theory, it is useful to see where the field stands. The following points from the literature matter most for this work.

4.1 Previous Studies

Industry 5.0 puts the worker back at the centre

Industry 5.0 asks cyber-physical systems to adapt to the worker rather than the other way round. A 2026 survey finds that worker-state monitoring (cameras, wearables, fatigue and cognitive-load estimation) is already technically mature, yet it rarely feeds real-time operational decisions, and workers almost never take part in the adaptive processes that shape their tasks [1]. This is precisely the gap this work tries to close.

A systematic review of 46 studies on well-being and cognitive workload in human-cobot work confirms how much attention this area now gets [3]. A neuroergonomics study found higher mean heart rate when workers lacked support

during a task, a sign of sympathetic activation linked to cognitive load [4].

Why heart rate variability and skin conductance

A systematic review of wearables for stress found that studies consistently use heart-rate variability (HRV) for detecting stress and anxiety. It also found that average heart rate alone, as used in many consumer devices, is less accurate than HRV, electrodermal activity (EDA) and possibly respiratory rate. EDA often showed high accuracy but questionable reliability [7]. In a separate study with 17 volunteers, consumer wearables tracked laboratory sensors well in telling mental states apart [8]. This matches the choice made here to weight HRV (40%) and EDA (30%) most heavily in the MPI. **4.2 Existing Approaches**

Fatigue-aware cobots already exist, in pieces

A 2026 open-source ROS 2 system on a UR3 cobot combines gesture recognition, intention prediction and fatigue monitoring from wearable eye-tracking glasses (blink rate, eye closure and pupil size). Its authors argue that fatigueadaptive behaviour extends collaborative-safety thinking into the operator's psychophysiological state, something current standards do not address [2]. A broader review notes that some systems already reassign tasks as physical fatigue grows, using heart rate, muscle activity or posture [5], and a review of personalisation in human-robot collaboration points the same way [9].

Safety standards deal with the body, not the mind

ISO/TS 15066 describes four collaborative methods: safety-rated monitored stop, hand guiding, speed and separation monitoring, and power and force limiting [6]. The last two limit robot speed or contact force. They react to distance and contact, not to how the person is feeling. This work borrows the idea of slowing the robot, but triggers it from the worker's state. It is **not** a safety-rated implementation of the standard, and the dashboard's mention of ISO/TS 15066 should be read as inspiration only.

5. RESEARCH GAP

Taken together, the studies above show progress on separate fronts but leave several limitations in place:

- Worker-state monitoring is technically mature, yet it rarely feeds real-time operational decisions, and workers seldom take part in the adaptive processes that shape their tasks [1].
- Fatigue-aware cobots exist only in pieces: individual systems use eye tracking, heart rate, muscle activity or posture, but not as one integrated loop [2], [5].
- Current safety standards react to distance and contact, not to the person's psychophysiological state [2], [6].
- Approaches that evaluate worker condition only after a work period, or that rely on one fixed threshold for everybody, cannot support timely, individualized interventions.
- Interventions are rarely paired with a recovery mechanism, so slowing the robot may raise an alarm continuously instead of helping the worker recover.

This work addresses these gaps by closing the loop from simulated worker state to cobot behaviour, using personal baselines, relative thresholds and a recovery credit, and by testing the result openly, including where the original settings fail.

6. METHODOLOGY

6.1 Methods and Algorithms: The Two Indices

The edge classifier combines the signals into two indices, as set out in Table 1.

Table -1: MPI and PFI as implemented

Index	Formula as implemented	What drives it
Mental Pressure Index (MPI)	$0.40 \times \text{vagal deficit} + 0.30 \times \text{EDA strain} + 0.15 \times \text{ratepressure strain} + 0.15 \times \text{breathing strain}$	HRV falling below 65% of the worker's baseline, skin conductance above 1.5 uS, high blood pressure x heart rate, fast breathing
Physical Fatigue Index (PFI)	$0.45 \times \text{heart-rate strain} + 0.25 \times \text{SpO2 deficit} + 0.30 \times \text{posture strain}$	Heart rate above 70 bpm, oxygen saturation below 99%, RULA posture score

6.2 From Index to Action

The intervention logic is a first-match-wins rule set, shown in Table 2.

Table -2: Intervention logic (first match wins)

Condition (first match wins)	Level	Cobot action	Feed rate
MPI > 0.85, or PFI > 0.85, or systolic BP > 165	CRITICAL	EMERGENCY_REST (safety halt)	0%
MPI > 0.68, or workstation occlusion > 75%	ELEVATED	TASK_ASSIST	45%
MPI > 0.48, or PFI > 0.52	ELEVATED	THROTTLE_COBOT	$100\% - 0.65 \times \text{MPI}$, min 35%
Otherwise	NOMINAL	Normal collaborative work	100%

The Automatic Stress Engine adds random disruptions (conveyor micro-stops, upstream delays, heavy-part misalignment, and so on) lasting 12 to 28 seconds. Their chance rises with task difficulty, with time into the shift, and when the worker's MPI is already above 0.55.

6.2 Signals Collected but Not Yet Used

The telemetry also includes gaze fixation stability, blink rate and core temperature. The current indices do not use them. Given the eye-tracking evidence in the literature [2], these are natural next inputs, and they connect this project to the author's smart-glasses work.

6.3 Process: Experimental Design

To see what the intervention logic really does, each configuration was run over many 8-hour virtual shifts with fixed random seeds, so that the comparisons are paired. The three configurations are listed in Table 3. **Table -3: Experimental configurations**

Setup	What it is
No intervention	The same code, but the cobot always runs at 100% (indices are still computed)
As supplied	The code exactly as it behaves, 30 shifts
Proposed v2	Three changes described below, 30 shifts

7. PROPOSED SYSTEM

7.1 Architecture

The program is organised in four layers, from sensing to dashboard, shown in Fig. 1.




Fig -1: Software architecture: four layers from sensing to dashboard (Mermaid flowchart).

7.2 Modules

The main modules, as described in this work, are:

- **Human Digital Twin:** generates the synthetic physiological signals for each worker, including heart rate, heart-rate variability, skin conductance, blood pressure, breathing and posture.
- **Edge classifier:** works over a short sliding window and turns the signals into the Mental Pressure Index (MPI) and the Physical Fatigue Index (PFI).
- **Intervention logic:** maps the indices to a cobot action (normal work, throttle, task assist or emergency rest) and a feed rate.
- **Automatic Stress Engine:** adds random disruptions of 12 to 28 seconds to the simulated workload.
- **Live dashboard:** a web page showing a 3D factory view, worker cards, intervention states and four live graphs.

7.3 Flowchart

The flowchart in Fig. 2 follows one second of simulated time. Signals are generated, assessed, and turned into a cobot action. The same loop repeats for each of the three stations.




Fig -2: Control loop of the HRMC digital twin (Mermaid flowchart).

The proposed closed loop (v2)

Three changes are proposed, none of which alter the original modules' structure. They were applied as wrappers around the original classes, which keeps the original code intact. The proposed closed loop is drawn in Fig. 3, and the changes are summarized in Table 4.




Fig -3: The proposed closed loop.

Table -4: Three proposed changes (v2)

Change	What it does	Setting used
1. Personal baseline and relative thresholds	Takes the median MPI of each worker's first 10 minutes as their normal working point. Interventions trigger relative to it.	Throttle at +0.10, assist at +0.20, halt at +0.30 (capped at 0.85, 0.90, 0.95)
2. Recovery credit	While any intervention is active, the fatigue pool drains faster, representing slower pace and rest.	0.0024 per second extra recovery
3. Fewer random disruptions	Only a fraction of the original stress-event chances are allowed to fire, so disruption is not continuous.	20% of original chance

8. IMPLEMENTATION

8.1 Technologies and Tools

- **Python 3** for the simulation, the Human Digital Twin, the edge classifier and the server.
- **HTML and JavaScript** for the embedded live dashboard, which shows a 3D factory view and four charts.
- **Mermaid** for the flowcharts and architecture diagrams.
- A **headless browser**, run against the real server, to capture the dashboard screenshots.
- **Fixed random seeds** (for example, seed 42) so that comparisons between configurations are paired.

8.2 Actual Implementation

The platform simulates a three-station factory with three workers (Sarah, Dev and Elena) and three cobots. The workers have different tasks: Dev has the heavy physical task, while Sarah and Elena have more cognitive tasks. The dashboard colour-codes each worker's state: green is nominal, amber is elevated and red is critical, and the cobot tool tips follow the

same colours. Figs. 4 to 6 are screenshots taken from the running system.




Fig -4: The factory about 20 seconds into the shift: two stations nominal (green), one already elevated (amber).




Fig -5: The factory about 15 minutes in: one station elevated (amber), two critical (red) with cobots halted.




Fig -6: Full dashboard at about 15 minutes: worker cards, intervention states and the four live graphs.

In the full view, the right-hand panel reports each worker's MPI, PFI and heart rate next to the cobot mode, such as "EMERGENCY_REST (0%)".

9. RESULTS AND DISCUSSION

9.1 Results

The first ten minutes

The workers start relaxed (MPI about 0.2). Within seconds, task load pushes the sympathetic tone up and the indices rise.

In the seed-42 run, the first throttle came at 17 s (Sarah), 26 s (Dev) and 28 s (Elena). Task assist followed at 27, 44 and 68 s, and the first emergency halt at 63, 64 and 297 s.




Fig -7: The same four graphs as the dashboard, over the first 600 seconds (seed 42). Dashed lines mark the MPI thresholds.

Three things stand out. Dev, who has the heavy physical task, shows PFI climbing steadily toward 0.9 and heart rate drifting above 130 bpm. Sarah and Elena, with more cognitive tasks, show MPI oscillating as stress events come and go. And the cobot feed-rate graph shows the cobots living almost entirely at 45% or 0%, never returning to 100%. **Behaviour across an 8-hour shift**




Fig -8: Five-minute averages across an 8-hour shift (seed 42).

There is no recovery phase. The averages plateau high and stay there. Over the shift the Automatic Stress Engine fired 2026 times, roughly 253 events per hour across the three workers, so disruptions are almost continuous.




Fig -9: Stress events per half hour in the 8-hour run.

Tier-3 analytics at the end of the shift

Table -5: Tier-3 cardiovascular analytics at the end of the shift

Worker	Mean HR (bpm)	Mean systolic BP (mmHg)	Mean HRV RMSSD (ms)	Rate-pressure product	Classification
Sarah	131.95	220.13	7.15	290.47	HIGH_MYOCARDIAL_RISK
Dev	134.66	212.77	8	286.52	HIGH_MYOCARDIAL_RISK
Elena	132.4	216.46	8.49	286.6	HIGH_MYOCARDIAL_RISK

A sustained systolic pressure above 210 mmHg is not a realistic working state for a healthy person. It tells us the simulated physiology is saturating: the fatigue pool fills within minutes and never drains, because recovery in the model is far slower than the build-up for any physical task. This is a calibration issue to fix before anyone treats the output as physiological evidence.

Predictions and findings

Table -6: Results per 8-hour shift (3 stations)

Result per 8-hour shift (3 stations)	No intervention	As supplied	Proposed v2
Units completed	6040	685	5290
Change in output vs no intervention	-	-88.7%	-12.4%
Defects	1336	133	883
Defect rate	22.1%	19.5%	16.7%
Time in NOMINAL	-	0.1%	76.4%
Time in THROTTLE / ASSIST / HALT	-	0.5% / 24.3% / 75.2%	16.3% / 7.1% / 0.2%
Average MPI	-	0.88	0.73




Fig -10: Same shift, same seed: MPI saturates as supplied (left) but stays in a workable range with v2 (right).




Fig -11: Share of the shift spent in each intervention state, as supplied versus v2.

9.2 Discussion: What the Findings Mean

- **The as-supplied loop is always alarmed.** It is in NOMINAL only 0.1% of the time. Even with no disruption, the base MPI for a working person sits above the 0.48 throttle line, so the thresholds are set below the normal working point.
- **The loop is open.** In the code, the cobot's slowdown never changes the worker's physiology. Slowing down therefore cannot help the person recover, it can only cost output. The lower defect rate in the as-supplied run comes from the model's defect formula, not from a healthier worker.
- **So P2 and P4 are not yet testable in the original code,** and P3 is untested because thresholds are one-size-fits-all. v2 was designed from the diagnosis above and run with settings chosen in advance, not tuned afterwards, but the three changes were tested together, so it is not possible to say how much each one contributes. And all of it is simulation: v2 shows that the loop can be balanced in this model, not that real workers would respond this way.

10. CONCLUSION

The HRMC digital twin works as a concept: within about 17-28 seconds of simulated work every cobot starts adapting to its worker, and the first full halts follow within 63-297 seconds. As supplied, the settings are too jumpy: over an 8-hour shift the workers spend only 0.1% of the time in nominal and 75.2% in emergency rest. The proposed v2 redesign (personal baselines, relative thresholds and a recovery credit) brings nominal time to 76.4% and cuts the defect rate from 22.1% to 16.7%, at a cost of 12.4% of output (simulated, 30 shifts). All of this is simulation: it shows that the loop can be balanced in this model, not that real workers would respond this way.

10.1 Limitations

- All physiology is generated by equations inside the program. Nothing here is measured from people.
- The indices use fixed weights; none were fitted to data or compared with a validated workload scale.
- Real collaborative-robot safety needs safety-rated functions and a formal risk assessment [6].
- Monitoring workers' bodies raises consent, privacy and fairness questions. A real deployment needs informed consent, clear data-use rules and a way for workers to override or opt out, which the literature also calls for [1].

11. FUTURE SCOPE

The following validation path is suggested:

- Fix the model calibration (Section 9.1) and add a robot-to-human feedback path so that P2 and P4 become testable.
- Compare MPI and PFI against a validated workload questionnaire such as NASA-TLX, using a small pilot with consenting volunteers and a consumer wearable [8].
- Add the unused signals (gaze stability, blink rate, temperature), ideally linking in the smart-glasses prototype.

- Test personal baselines (P3) against fixed thresholds on the pilot data.
- Put the effect on a real cobot cell only behind a certified safety controller.

REFERENCES

- 1) L. Pereira, L. M. D. F. Ferreira, and J. R. Paulo, "Human-centricity in Industry 5.0: A survey of worker sensing, adaptive operations, and human-in-the-loop systems," arXiv. [Online]. Available: <https://arxiv.org/pdf/2609.26245>
- 2) "Predictive fatigue-aware human-robot collaboration: A real-time, open-source ROS 2 instantiation of a three-module humancentric digital twin framework on a standard cobot," *Frontiers in Robotics and AI*, 2026, doi: 10.3389/frobt.2026.1897261.
- 3) "Understanding workers' well-being and cognitive load in human-cobot collaboration: Systematic review," *JMIR*, 2025, doi: 10.2196/75658.
- 4) "Neuroergonomic signatures of improved human-robot collaboration in LLM-supported industrial workflows," *J. Intell. Manuf.*, 2026, doi: 10.1007/s10845-026-02958-5.
- 5) "Exploring tasks and challenges in human-robot collaborative systems: A review," *Robot. Comput.-Integr. Manuf.* [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0736584525001565>
- 6) *Robots and Robotic Devices: Collaborative Robots*, ISO/TS 15066:2016. Its content has since been incorporated into ISO 102182:2025.
- 7) B. A. Hickey et al., "Smart devices and wearable technologies to detect and monitor mental health conditions and stress: A systematic review," *Sensors*, vol. 21, no. 10, 2021.
- 8) "Wearable technologies for mental workload, stress, and emotional state assessment during working-like tasks: A comparison with laboratory technologies," *Sensors*, vol. 21, no. 7, Art. no. 2332, 2021, doi: 10.3390/s21072332.
- 9) "A review of personalisation in human-robot collaboration and future perspectives towards Industry 5.0," arXiv. [Online]. Available: <https://arxiv.org/pdf/2506.20447>
- 10) The Copyright Rules, 2013, ch. XIII, r. 70 (Application for registration of copyright), as amended 2021. [Online]. Available: https://copyright.gov.in/Copyright_Rules_2013/chapter_xiii.html