IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 3, November 2026
Impact Factor: 8.2

Ergo Analyzer: Ergonomic Risk Assessment System Using Computer Vision & REBA

Khushi Kaduskar, Riddhi Gole, Sakshi Khade, Vaishnavi Mane
Department of Artificial Intelligence & Data Science
Jaihind Comprehensive Educational Institute, Pune, India
khushikaduskar@gmail.com, riddhigole194@gmail.com
khadesakshi204@gmail.com, manevaishnavi174@gmail.com

Abstract : *Work-related musculoskeletal disorders remain a leading cause of injury and lost workdays in manufacturing, yet ergonomic assessment on the shopfloor is still largely periodic, manual and subjective. This paper presents Ergo Analyzer, a contactless, camera-based system that converts ordinary video into real-time, actionable ergonomic insight. The framework detects workers in the video stream, estimates seventeen body keypoints using a real-time pose estimation model, computes joint angles and body alignment, and compares them against ergonomic thresholds. A machine learning classifier then assigns a color-coded risk level and a numerical score from zero to one hundred, together with body region sub-scores for the back, shoulders, wrists and legs. A recommendation engine converts the result into task-specific corrective feedback such as adjusting posture, raising the work surface or taking micro-breaks. A prototype built with Python and OpenCV and exercised in simulated shopfloor conditions assigned distinct risk levels to good posture, improper lifting and excessive bending, and, in a live qualitative demonstration, the risk score fell from 72 to 32 after the recommended posture was adopted. We describe the architecture, scoring method, implementation and limitations, and outline a validation plan and roadmap toward alerts, mobile applications and wearable integration.*

Keywords: ergonomics, human pose estimation, musculoskeletal disorders, computer vision, risk assessment, occupational safety.

I. INTRODUCTION

Manual handling, repetitive motion and awkward postures are everyday features of manufacturing work. Over time they lead to musculoskeletal injuries, lost workdays and higher costs for both workers and employers. Musculoskeletal conditions are a leading contributor to disability worldwide [10], and work-related musculoskeletal disorders are among the most common occupational health problems in Europe, with substantial costs for workers and employers [11].

Despite this burden, ergonomic monitoring on the shopfloor is rarely continuous or objective. Established observational methods such as RULA, REBA and OWAS [1]–[3] and the NIOSH lifting equation [4] depend on a trained assessor who watches a task, scores postures by eye and records the result on paper. The process is slow, covers only a small sample of tasks and shifts, and is open to inter-rater variation [9]. Wearable motion sensors improve objectivity but add cost, calibration effort and discomfort, which limits large-scale use.

Advances in deep learning based human pose estimation [5], [6] make a different approach possible: ordinary cameras can recover body keypoints in real time without contact. This paper presents Ergo Analyzer, a system that builds on this capability to provide continuous, objective and scalable ergonomic risk analysis. The main contributions are:

(1) an end-to-end pipeline that goes from video input to a color-coded risk score and corrective advice in a single flow;

(2) a transparent scoring scheme that combines joint angle analysis with a machine learning classifier and reports region-level sub-scores for the back, shoulders, wrists and legs;

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 3, November 2026 Impact Factor: 8.2

(3) a task-specific recommendation engine that tells workers and safety managers what to change, not only that a risk exists; and
(4) a working prototype and a set of simulated shopfloor case studies, together with a clear account of limitations and a validation plan.

Four design goals guided the work. The system must be contactless, so that no sensor has to be worn or calibrated for each worker. It must run at interactive rates, so that feedback arrives while a posture is still being held. It must be interpretable, so that every score can be traced to a joint angle and an ergonomic threshold. And it must be actionable, so that the output is a concrete change rather than a number alone. These goals motivate a lightweight pose estimator, an angle-based scoring backbone that can be inspected by an ergonomist, and a recommendation layer driven by the region sub-scores.

The rest of the paper is organized as follows. Section II reviews related work. Section III describes the method. Section IV covers implementation, Section V presents prototype results, Section VI discusses limitations, Section VII outlines future work, and Section VIII concludes.

II. RELATED WORK

A. Observational Ergonomic Assessment

OWAS [3] classifies work postures into coded combinations of back, arm and leg positions. RULA [1] focuses on the upper limbs, neck and trunk, and REBA [2] extends the idea to the whole body and to load handling. The NIOSH equation [4] estimates a recommended weight limit for lifting tasks. A review comparing OWAS, RULA and REBA [9] notes that the methods differ in scope and sensitivity and that all rely on human observation. These methods supply the thresholds and risk logic that Ergo Analyzer automates, but they do not provide continuous monitoring.

Table I summarizes the scope of the four methods. Because each was designed for a human observer, their inputs are posture categories, load and coupling ratings and activity descriptors that are scored on a worksheet. Two consequences follow for automation. The angle bands used by RULA and REBA can be reproduced from keypoint geometry, which is why they supply the thresholds used in Section III-E. In contrast, factors that cannot be seen in an image, such as the mass of a handled load or the quality of a grip, cannot be recovered from posture alone and must be entered manually or measured with additional sensors.

TABLE I: OBSERVATIONAL ERGONOMIC ASSESSMENT METHODS

Method	Body scope	Typical inputs	Output
OWAS [3]	Back, arms, legs	Posture codes, load	Action category (4 levels)
RULA [1]	Upper limbs, neck, trunk	Posture, muscle use, force	Grand score (1 to 7)
REBA [2]	Whole body	Posture, load, coupling, activity	Final score (1 to 15)
NIOSH [4]	Manual lifting	Load, distances, asymmetry, frequency, coupling	Recommended weight limit

B. Vision-Based Pose Estimation

Bottom-up approaches such as OpenPose [5] detect keypoints for all people in an image using part affinity fields. RTMPose [6], part of the MMPose toolbox, is designed for real-time multi-person estimation on modest hardware, which suits deployment on the shopfloor. Models are commonly trained on datasets such as COCO [7], which defines the seventeen-keypoint skeleton used in this work.

Top-down methods instead detect each person first and then estimate keypoints inside the cropped box. This generally gives higher keypoint accuracy per person, at a cost that grows with the number of people in view. Because a shopfloor camera usually sees a limited number of workers, and because joint angles are sensitive to small keypoint errors, Ergo Analyzer adopts a top-down model.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 3, November 2026
Impact Factor: 8.2

C. Gap Addressed

Existing observational tools are manual and intermittent, while generic pose estimators do not interpret posture in ergonomic terms. Ergo Analyzer bridges the two by coupling real-time keypoint detection with ergonomic scoring and by closing the loop with corrective recommendations delivered through a monitoring dashboard.

III. PROPOSED METHOD

A. System Overview

Fig. 1 shows the six-stage pipeline. Video from safety cameras or mobile devices is passed to a detector that locates each worker. A pose model extracts seventeen keypoints per person, from which joint angles and body alignment are computed and compared with ergonomic standards. The resulting features are classified into a risk level, and a recommendation engine generates corrective feedback.

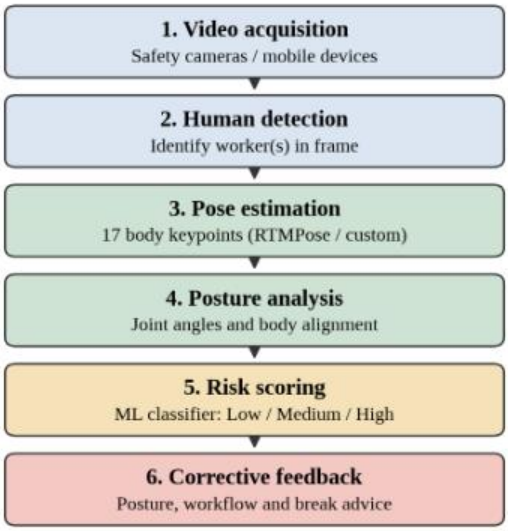


Fig. 1. Six-stage Ergo Analyzer pipeline from video input to corrective feedback.

Formally, for each worker the input at frame t is a set of seventeen keypoints with image coordinates and confidences. The system maps this set, together with a short history of earlier frames and the task context, to three outputs: a continuous risk score S between 0 and 100, a risk class (Low, Medium or High) and a short list of corrective messages. The mapping is split into a geometric part, which turns keypoints into smoothed joint angles, and a decision part, which turns angles and context into scores, classes and advice.

B. Video Acquisition and Human Detection

Frames are read from fixed safety cameras or handheld mobile devices using OpenCV [8]. A detector identifies the worker or workers in each frame and crops the region of interest, so that pose estimation runs only on people. This keeps the approach contactless and allows several workers to be monitored from one viewpoint.

C. Pose Estimation

A top-down model, RTMPose or a custom-trained alternative, predicts seventeen keypoints for each detected worker: nose, eyes, ears, shoulders, elbows, wrists, hips, knees and ankles. Each keypoint has image coordinates and a confidence value. Keypoints below a confidence threshold are discarded for that frame, and the affected angles are held at their last reliable value.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 3, November 2026
Impact Factor: 8.2

The seventeen-keypoint skeleton has no spine or neck joint. The trunk is therefore defined by the mid-shoulder point and the mid-hip point, computed as the means of the left and right shoulder and hip keypoints, and the neck angle is measured between this trunk segment and the line from the mid-shoulder point to the ears. The skeleton also has no hand or finger keypoints, so wrist-region scores can only be inferred from forearm orientation and are coarser than those of the other regions.

D. Joint Angle and Alignment Analysis

For three keypoints p_i, p_j and p_k, with p_j at the joint, the joint angle is

$$\theta_j = \arccos((a \cdot b) / (|a|\,|b|)) \qquad (1)$$

where $a = p_i - p_j$ and $b = p_k - p_j$. Trunk flexion is measured between the hip-to-shoulder segment and the vertical, and similar angles are computed at the neck, shoulders, elbows, wrists and knees. To suppress jitter from frame-to-frame keypoint noise, angles are smoothed by exponential filtering to give the smoothed angle φ:

$$\varphi_t = \alpha\theta_t + (1 - \alpha)\,\varphi_{t-1} \qquad (2)$$

where α in (0, 1] controls responsiveness.

A small α suppresses jitter but delays the response to genuine posture changes, whereas a value close to 1 follows the raw angle. The choice is therefore a trade-off between stability and responsiveness and should be set according to the frame rate. Fig. 2 illustrates the geometry: the trunk angle is measured from the vertical at the hip, and the other joint angles are measured at the middle keypoint of each triplet in Eq. (1).

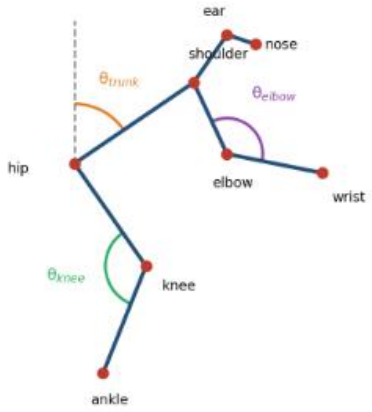


Fig. 2. Keypoint geometry for angle computation. Trunk flexion is measured between the mid-hip to mid-shoulder segment and the vertical; the knee and elbow angles are measured at the middle keypoint of each triplet (Eq. 1). The posture shown is illustrative.

E. Ergonomic Risk Scoring

Each body region r (back, shoulders, wrists, legs) receives a sub-score that grows as its angle moves from a safe value θ_r^{safe} toward a limit value θ_r^{max} taken from the ergonomic standards:

$$s_r = 100 \cdot \min(1, \max(0, (\varphi_r - \theta_r^{safe}) / (\theta_r^{max} - \theta_r^{safe}))) \qquad (3)$$

The overall score is a weighted combination of the region scores:

$$S = \Sigma_r\, w_r\, s_r\,, \quad \Sigma_r\, w_r = 1 \qquad (4)$$

so that S lies between 0 and 100. The region sub-scores, the overall score and task context such as posture duration form the feature vector given to a machine learning classifier, which outputs the risk class. Table II lists the color-coded bands used in the prototype. The band limits shown are illustrative defaults that can be configured per site and should be calibrated against expert assessment.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 3, November 2026
Impact Factor: 8.2

As an illustration, take trunk flexion with θsafe = 20° and θmax = 60°. A smoothed angle of 45° gives s = 100 × (45 − 20) / (60 − 20) = 62.5, which lies in the Medium band, and any angle of 60° or more gives 100 (Fig. 3). If the back, shoulder, wrist and leg sub-scores are 62.5, 40, 10 and 30 with weights 0.4, 0.2, 0.2 and 0.2, then S = 25 + 8 + 2 + 6 = 41, a Medium score. The weights should reflect the relative contribution of each region and can be set from ergonomic guidance or tuned on expert-labelled data, with equal weights as a neutral default. The limits and weights in this example are illustrative and are not the calibrated values of the prototype.

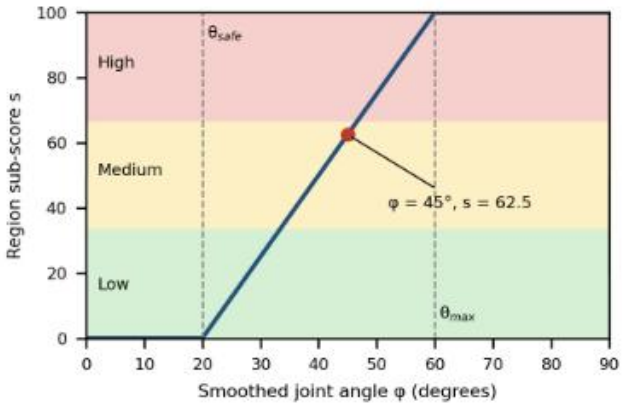


Fig. 3. Region sub-score as a function of the smoothed joint angle (Eq. 3) for illustrative limits θsafe = 20° and θmax = 60°, with the risk bands of Table II.

TABLE II: RISK BANDS AND DISPLAY

Band	Score	Meaning and action
Low	0 to < 34	Acceptable posture; maintain
Medium	34 to < 67	Improve posture or workflow soon
High	67 to 100	Immediate corrective action

F. Corrective Recommendation Engine

The dominant risk region and the detected task type select feedback from three categories: adjust posture (for example, keep the back straight and avoid deep bending), adjust workflow (raise the work surface or use assistive equipment) and take breaks (micro-breaks to reduce fatigue and strain). For a lifting task with a high back score, the system advises keeping the load close, bending at the knees rather than the back and using leg strength to lift.

Table III gives examples of how the dominant region maps to feedback. The sub-scores also show where the strain is concentrated, so a worker who stands well but twists the wrists receives different advice from one who bends the back. When a region stays in the Medium or High band for a prolonged period, a micro-break suggestion can be added to the posture or workflow advice. Feedback is worded as short instructions that a worker can act on immediately, and each message is traceable to the region and angle that triggered it.

TABLE III: EXAMPLES OF REGION-DRIVEN FEEDBACK

Dominant region	Typical cue	Example feedback
Back	Deep trunk flexion while lifting	Keep the load close; bend at the knees, not the back
Shoulders	Arms raised or held away from the body	Lower the work surface or bring the task closer
Wrists	Sustained wrist deviation	Reposition the part or use a tool that keeps the wrist neutral

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 3, November 2026
Impact Factor: 8.2

Dominant region	Typical cue	Example feedback
Legs	Prolonged deep knee flexion or static stance	Use a support or stool, change stance, take a micro-break

G. Dashboard and Architecture

A backend service runs the video, pose and scoring stages and streams results to a frontend dashboard. The dashboard displays the live feed with the skeleton overlay, the ergonomic risk score and gauge, region-level risk and the current corrective suggestions, giving safety managers continuous visibility across the workforce.

H. Per-Frame Processing Summary

For each frame, the system (1) reads the frame and detects the workers present; (2) predicts seventeen keypoints and confidences for each worker; (3) discards keypoints below the confidence threshold and holds the affected angles at their last reliable values; (4) computes the joint angles of Eq. (1) and the trunk flexion angle; (5) smooths each angle with Eq. (2); (6) maps the smoothed angles to region sub-scores with Eq. (3) and combines them with Eq. (4); (7) forms the feature vector from the sub-scores, the overall score and the posture duration and classifies the risk level; and (8) selects corrective feedback from the dominant region and task type and sends the result to the dashboard. Steps 4 to 6 are deterministic functions of the keypoints, which keeps the angle and score computation auditable.

I. Deployment Considerations

Angle accuracy depends on how the camera views the worker. Trunk and knee flexion are best observed from a side view, whereas shoulder elevation and lateral trunk bending are better observed from the front or rear, so a single fixed camera favors some tasks over others. Where two views are available, the view with the higher mean keypoint confidence can be used for each region. The camera should keep the whole body in frame, because the knee and ankle keypoints are needed for the leg sub-score. The frame rate should be high enough to resolve the posture changes of the task: a few frames per second suffices for slow, sustained postures, whereas lifting motions need more. Posture duration, which enters the feature vector, is accumulated over the frames in which a region stays in the same band.

IV. IMPLEMENTATION

The prototype is written in Python. OpenCV handles capture and frame processing, the pose model produces seventeen keypoints per worker, and a machine learning classifier assigns the risk level. Table IV summarizes the components.

TABLE IV: TECHNOLOGY STACK

Component	Technology
Video processing	OpenCV and Python, real-time frames
Pose detection	RTMPose or custom model, 17 keypoints
Risk classification	Machine learning classifier
System	Integrated backend and frontend dashboard

The design is contactless and automated: no worker-worn hardware is required, and additional cameras can be added without changing the analysis pipeline.

The scoring stages operate on keypoints rather than pixels, so they are independent of the particular detector or pose model, and a heavier model on a GPU or a lighter model on an edge device can be exchanged without changing the angle and scoring code. The per-frame outputs (keypoints, angles and scores) are compact compared with video, which makes skeleton-only storage practical.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 3, November 2026
Impact Factor: 8.2

The per-frame cost is dominated by detection and pose estimation. Angle computation, filtering and scoring involve a handful of vector operations per worker and add negligible time by comparison, so the end-to-end frame rate is set mainly by the choice of detector and pose model and by the number of workers in view.

V. PROTOTYPE RESULTS

A. Simulated Shopfloor Scenarios

The prototype was exercised on simulated shopfloor tasks. Fig. 4 shows the risk scores produced. Good posture scored 18 out of 100 (Low), improper lifting scored 68 (High) and excessive bending scored 84 (High). In the live monitoring demonstration on a manufacturing line, the baseline posture scored 72 with a high lower back strain warning, and the system suggested bending at the knees, keeping the back neutral, bringing the work closer and adjusting table height. When the recommended posture was adopted, the score fell to 32 (Low), a reduction of 40 points.

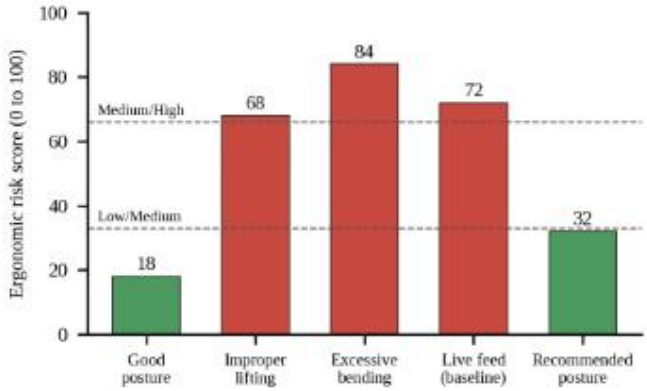


Fig. 4. Risk scores from the prototype in simulated shopfloor scenarios, with the Low/Medium and Medium/High band limits.

The ordering good posture, improper lifting, excessive bending matches the expectation that larger and more sustained trunk flexion yields a higher risk. However, three scenarios and one live demonstration cannot show how the system behaves across different workers, tasks, clothing, lighting or camera views, and the scores should not be read as a validated measure of accuracy.

B. Region-Level Scorecard

For the improper lifting case the scorecard in Table V reports an overall score of 68 in the High risk zone, with the back identified as the main contributor. This breakdown lets the recommendation engine target the back and legs rather than issuing generic advice.

TABLE V: SCORECARD FOR THE IMPROPER LIFTING CASE

Body region	Risk level
Back	High
Shoulders	Medium
Wrists	Low
Legs	Medium
Overall (68 / 100)	High

These results indicate that the pipeline separates distinct postures and responds to posture changes in real time. They are qualitative prototype outcomes; a formal quantitative evaluation is planned, as described next.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 3, November 2026
Impact Factor: 8.2

C. Planned Validation

To establish accuracy, we will (i) measure keypoint and joint angle error against a motion capture or manually annotated reference, (ii) compare Ergo Analyzer risk classes with REBA or RULA scores from certified ergonomists on a common set of recorded tasks, reporting agreement with Cohen's kappa and confusion matrices, and (iii) measure end-to-end latency and frame rate on target hardware. Results of this evaluation have not yet been collected and will be reported in future work.

Table VI summarizes the protocol. Recordings will cover common shopfloor tasks such as lifting, bending and overhead work, captured from more than one camera position, with informed consent from the participating workers. Two certified ergonomists will rate each task segment independently, inter-rater agreement will be reported, and Ergo Analyzer will be compared with both raters and with their consensus. Band limits and region weights will be tuned on a development split and reported on a held-out split, so that the final figures are not fitted to the test recordings.

Angle accuracy is reported as the mean absolute error, $MAE = (1/N)\ \Sigma\ |\theta i - \theta i|$, between estimated and reference angles over N frames, per joint and per task. Agreement on the three risk classes is reported with quadratic-weighted Cohen's kappa, which penalizes a High versus Low disagreement more than a High versus Medium one, together with the confusion matrix. Latency is measured from frame capture to delivery of the score to the dashboard and is reported as a median and a 95th percentile, because occasional slow frames matter for real-time feedback.

TABLE VI: PLANNED VALIDATION PROTOCOL

Objective	Reference	Metric
Angle accuracy	Motion capture or manual annotation	Mean absolute angle error (degrees)
Risk classification	REBA or RULA by certified ergonomists	Accuracy, weighted Cohen's kappa, confusion matrix
Robustness	Same tasks under varied camera angle and occlusion	Change in angle error and in agreement
Temporal stability	Static postures held for a fixed time	Score variance with and without smoothing
Runtime	Target hardware	End-to-end latency (ms) and frame rate (fps)

VI. DISCUSSION AND LIMITATIONS

The main strengths of the approach are continuous visibility, objective and repeatable measurements, proactive prevention through real-time scoring, task-specific feedback that supports behavior change, and scalability across manufacturing and industrial environments.

Several limitations apply. First, monocular two-dimensional keypoints are sensitive to camera angle, and depth ambiguity can distort angles such as trunk twist. Second, occlusion by machinery or other workers lowers keypoint confidence. Third, vision alone does not measure the weight of a handled load or its force, so load-related risk is estimated from posture. Fourth, the current band limits and weights are illustrative and require calibration against expert ratings. Fifth, continuous video of workers raises privacy and consent concerns; deployments should follow local regulations, inform workers, and consider on-device processing, anonymized skeleton-only storage and role-based dashboard access. The system should support, not replace, qualified occupational health professionals.

Ergo Analyzer does not replace observational methods such as REBA and RULA. It automates their posture component and adds continuity, while load, coupling and many contextual factors still need human judgment or additional sensing.

Two further limitations concern the skeleton and the evidence. The seventeen-keypoint skeleton does not include hand keypoints, so wrist-region scores are less reliable than the others, and the standard thresholds were developed for observer-based assessment and may not transfer exactly to angles estimated from noisy keypoints. In addition, the scenarios were simulated and may not contain the postural variety, clothing, lighting and equipment of real shifts.

Scores describe a posture within a task and not the capability of a worker. They should be used to improve workstations, tools and task design and not to rank individuals or to measure productivity, because such use is outside the intended purpose and could discourage workers from accepting the system. Involving workers and safety

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 3, November 2026 Impact Factor: 8.2

representatives in the choice of thresholds and in the review of feedback messages is likely to improve both acceptance and the usefulness of the advice.

Three groups of parameters influence the score: the smoothing factor α, the limits θsafe and θmax, and the region weights. A planned sensitivity analysis will vary each group over a plausible range while holding the others fixed and will report the effect on the overall score and on band assignment, so that users can see which choices matter most and how stable the Low, Medium and High decisions are near the band boundaries.

Table VII places the approach among the alternatives. Wearable inertial sensors can measure motion and, with additional sensing, force, but they must be fitted and calibrated for each worker and can be uncomfortable over a full shift. Vision requires no contact and little per-worker setup, but it depends on camera placement, sees only what is visible and raises greater privacy concerns unless only skeletons are retained. The two are complementary, which is why wearable integration is part of the roadmap.

TABLE VII: QUALITATIVE COMPARISON OF MONITORING APPROACHES

Aspect	**Observation**	**Wearable sensors**	**Vision (this work)**
Worker contact	None	Sensors worn on the body	None
Setup per worker	Trained observer	Fitting and calibration	Camera placement
Continuity	Intermittent	Continuous	Continuous
Load and force	Observer judgment	Possible with added sensors	Not measured directly
Privacy exposure	Low	Body data collected	Video, unless only skeletons are kept

VII. FUTURE WORK

The roadmap comprises real-time alerts, mobile and web applications, integration with wearables to add force and motion data, and analytics reports that track risk trends across tasks, teams and time. Completing the validation in Section V-C is the first priority, followed by the evaluation of detection and tracking under occlusion and varied camera angles.

Technical extensions include monocular three-dimensional pose estimation or multi-view fusion to reduce depth ambiguity, hand keypoints for finer wrist assessment, temporal models that recognize the task and score risk over sequences instead of single frames, and load estimation from task context or wearable sensors. Calibration of the band limits and region weights against expert ratings, and a study of how workers respond to the feedback over time, are also planned.

VIII. CONCLUSION

This paper presented Ergo Analyzer, an AI-driven, contactless system that turns shopfloor video into real-time ergonomic risk scores and corrective recommendations. By combining real-time pose estimation, joint angle analysis, machine learning classification and a recommendation engine, it addresses the lack of continuous, objective ergonomic monitoring in manufacturing. Simulated shopfloor scenarios produced distinct scores for low and high risk postures, and a live demonstration showed a lower score when the recommended posture was followed. Future work will quantify accuracy against expert assessment and extend the platform with alerts, mobile applications and wearable integration to support safer and healthier workplaces.

ACKNOWLEDGMENT

The authors thank mentors and peers for their guidance and support. Early prototypes of this work were demonstrated at the Start Up Pitch Competition 2026 (E-Cell IIT Bombay and Team Navora), the Smart India Hackathon (SIH 2026) and the Anveshana Science & Engineering Fair. These were demonstrations only and are not archival publications.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 3, November 2026
Impact Factor: 8.2

REFERENCES

[1] L. McAtamney and E. N. Corlett, "RULA: A survey method for the investigation of work-related upper limb disorders," *Appl. Ergon.*, vol. 24, no. 2, pp. 91–99, 1993.

[2] S. Hignett and L. McAtamney, "Rapid entire body assessment (REBA)," *Appl. Ergon.*, vol. 31, no. 2, pp. 201–205, 2000.

[3] O. Karhu, P. Kansi, and I. Kuorinka, "Correcting working postures in industry: A practical method for analysis," *Appl. Ergon.*, vol. 8, no. 4, pp. 199–201, 1977.

[4] T. R. Waters, V. Putz-Anderson, A. Garg, and L. J. Fine, "Revised NIOSH equation for the design and evaluation of manual lifting tasks," *Ergonomics*, vol. 36, no. 7, pp. 749–776, 1993.

[5] Z. Cao, G. Hidalgo, T. Simon, S.-E. Wei, and Y. Sheikh, "OpenPose: Realtime multi-person 2D pose estimation using part affinity fields," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 43, no. 1, pp. 172–186, Jan. 2021.

[6] T. Jiang *et al.*, "RTMPose: Real-time multi-person pose estimation based on MMPose," 2023, arXiv:2303.07399.

[7] T.-Y. Lin *et al.*, "Microsoft COCO: Common objects in context," in *Proc. Eur. Conf. Comput. Vis.*, 2014, pp. 740–755.

[8] G. Bradski, "The OpenCV library," *Dr. Dobb's J. Softw. Tools*, 2000.

[9] D. Kee, "Systematic comparison of OWAS, RULA, and REBA based on a literature review," *Int. J. Environ. Res. Public Health*, vol. 19, no. 1, Art. no. 595, 2022.

[10] World Health Organization, "Musculoskeletal health," WHO Fact Sheets, Jul. 2022.

[11] European Agency for Safety and Health at Work (EU-OSHA), "Work-related musculoskeletal disorders: Prevalence, costs and demographics in the EU," Luxembourg: Publications Office of the European Union, 2019