Journal of Computer Science and Technology Studies
ISSN: 2709-104X
DOI: 10.32996/jcsts
Journal Homepage: www.al-kindipublisher.com/index.php/jcsts

JCSTS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

From Periodic Inspection to Continuous Monitoring: A Case Study of Camera-Based Joint-Angle Alerting for Ergonomic Risk on a Printed Circuit Board Assembly Line

Mariane Aguiar de Melo[1], Tiago Carvalho de Almeida[2], Tony Douglas Martins Pena[3], Adriano Martins de Melo[4], Renato Moreira Teixeira Junior[5]✉, Alessandra Castro Rocha[6], Daniel Angelo Oliveira de Abreu[7] e Arthur Drumond Parente[8]

[1]Technologist in Systems Analysis and Development from. Conecthus Institute of Technology and Biotechnology of Amazonas, Manaus, Amazonas, Brazil. E-mail: mariane.melo@conecthus.org.br.
[2]Technologist in Systems Analysis and Development. Conecthus Institute of Technology and Biotechnology of Amazonas, Manaus, Amazonas, Brazil. E-mail: tiago.almeida@conecthus.org.br.
[3]Bachelor in Computer Science. Conecthus Institute of Technology and Biotechnology of Amazonas, Manaus, Amazonas, Brazil. E-mail: tony.pena@conecthus.org.br.
[4]Postgraduate Specialist in Teaching for Professional and Technological Education. Conecthus Institute of Technology and Biotechnology of Amazonas, Manaus, Amazonas, Brazil. E-mail: adriano.melo@conecthus.org.br.
[5]Postgraduate Specialist in Advanced Manufacturing Engineering for Industry 4.0. Conecthus Institute of Technology and Biotechnology of Amazonas, Manaus, Amazonas, Brazil. E-mail: renato.junior@conecthus.org.br
[6]Postgraduate Specialist in Innovation Management with an Emphasis on Research, Development, and Innovation (P&DI). Conecthus Institute of Technology and Biotechnology of Amazonas, Manaus, Amazonas, Brazil. E-mail: alessandra.rocha@conecthus.org.br.
[7]MBA in Information Management, Technology, and Innovation. Conecthus Institute of Technology and Biotechnology of Amazonas, Manaus, Amazonas, Brazil. E-mail: daniel.abreu@conecthus.org.br.
[8]Project Management Specialist. Conecthus Institute of Technology and Biotechnology of Amazonas, Manaus, Amazonas, Brazil. E-mail: arthur.parente@conecthus.org.br.

Corresponding Author: Renato Moreira Teixeira Junior, **E-mail**: renato_moreirat@hotmail.com

| ABSTRACT

Ergonomic risk management in electronics manufacturing still relies largely on manual, periodic inspections that offer little continuous visibility of operators' postural exposure. This case study reports the design, implementation and pilot-line validation of a Web-based ergonomic monitoring platform at an electronics manufacturing plant in Manaus, Brazil, developed over a seven-month applied research and development cycle (January–July 2026) by a technology institute in partnership with the manufacturer. The platform detects operators in industrial camera streams with a YOLOv8l detector, extracts body and hand landmarks with MediaPipe Pose and Hand Landmarker, computes joint angles, classifies each monitored joint into safe, light, moderate or high-risk angle bands with ordinal scores, and raises an alert only when a high-risk angle persists beyond a configurable holding time. Angle bands and persistence rules were defined with an NR-17-oriented ergonomics consultancy and are editable through the interface. Events flow through a containerised micro-service pipeline (RTSP ingestion, RabbitMQ messaging, REST API, SQL Server) to a Web interface with real-time monitoring, dashboards, PDF reports and a 3-D skeleton view. On a 500-frame sample from the surface-mount technology line, compared with human inspection, operator detection achieved 94.5% precision, 98.2% recall and a 96.3% F1-score; the reported hit rate on critical joints was 91.8% and the continuous hand-tracking rate 89.4%. Between 22 May and 5 August 2026 the platform recorded 409,527 alert events, with a mean uninterrupted tracking persistence of 72.0 s at 15 frames per second and a mean joint angle of 129.14° for high-risk events. Under approximately 250 concurrent requests the pipeline sustained 1.64 frames per second with about 650 ms end-to-end latency and no critical failures during a 152 s stress run, and the system passed formal acceptance testing with zero open issues (self-assessed Technology Readiness Level 6). The case shows that pose-based joint-angle alerting can be integrated into a live industrial IT environment, but also that the evidence produced does not yet establish ergonomic validity: no comparison

Copyright: © 2026 the Author(s). This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC-BY) 4.0 license (https://creativecommons.org/licenses/by/4.0/). Published by Al-Kindi Centre for Research and Development, London, United Kingdom.

with RULA/OWAS expert scoring, no joint-angle error against a reference measurement and no health-outcome data were obtained. These gaps define the agenda for moving from prototype to deployment.

| KEYWORDS

Human pose estimation; occupational ergonomics; industrial case study; computer vision; joint-angle thresholds; real-time alerting; musculoskeletal disorders; NR-17; Industry 4.0; printed circuit board assembly.

| ARTICLE INFORMATION

ACCEPTED: 20 October 2026 **PUBLISHED:** 20 November 2026 **DOI:** 10.32996/jcsts.2026.8.18

1. Introduction

Work-related musculoskeletal disorders (WMSDs) remain one of the most prevalent causes of occupational ill-health and lost working time in manufacturing, and awkward or sustained postures, repetitive movements and prolonged static loading of the upper limbs are among their best-established physical risk factors (Punnett & Wegman, 2004; David, 2005). Electronics assembly and inspection of printed circuit boards (PCBs) concentrate exactly these exposures: operators perform highly repetitive, visually demanding tasks at fixed workstations, with sustained elbow and wrist flexion, neck flexion toward small components and prolonged sitting.

In Brazil, employers are legally required to assess and manage such exposures under Regulatory Standard NR-17 (Ergonomics), which mandates a preliminary ergonomic assessment and, where indicated, a full Ergonomic Work Analysis covering posture, repetitiveness, effort, furniture, environment and work organisation (Brasil, 2022). In practice, however, as the project documents analysed in this study describe for the plant investigated here, ergonomic management in electronics factories is still predominantly manual, reactive and dependent on periodic inspections, which offers low real-time visibility of postural conditions and little capacity to correlate physical effort, exposure time and productivity.

Observational ergonomic methods such as RULA (McAtamney & Corlett, 1993), REBA (Hignett & McAtamney, 2000) and OWAS (Karhu et al., 1977) score postures from joint angles and load categories, but they depend on trained analysts observing or filming work for short windows, which makes them costly to apply continuously and sensitive to observer variability (Lowe et al., 2019). Over the last decade, markerless human pose estimation from RGB or RGB-D cameras has matured to the point where 2-D and 3-D landmarks of the human body can be extracted in real time on commodity hardware (Cao et al., 2021; Bazarevsky et al., 2020), and several studies have shown that joint angles computed from such landmarks can feed RULA, REBA or OWAS scoring automatically (Diego-Mas & Alcaide-Marzal, 2014; Manghisi et al., 2017; Kim et al., 2021; Lin et al., 2022). Most of this work, however, has been validated in laboratory settings or on short recordings; comparatively little has been published on systems that run continuously on a real production line, integrated with the plant's IT infrastructure, with alerts consumed by safety staff rather than by researchers.

This article presents a case study of an applied research, development and innovation (RD&I) project carried out at a large electronics manufacturing plant in Manaus, in the Brazilian Amazon, by a technology institute in partnership with the manufacturer, within a publicly coordinated Industry 4.0 programme. The project set out to develop and validate, in a relevant operational environment (Technology Readiness Level 6), a functional prototype of a Web and mobile system that captures video from industrial cameras, estimates operators' body landmarks, calculates joint angles, compares them with predefined ergonomic thresholds and persistence rules, and issues graded alerts and reports aligned with NR-17. The case is of interest because it documents, with actual field figures, what it takes to move ergonomic monitoring from periodic manual inspection to a continuous, camera-based process inside a running factory, and where the evidence produced by such a project stops short of ergonomic validation.

The contributions of the study are threefold. First, it documents a complete end-to-end architecture, from RTSP camera ingestion to dashboards, in which pose-based joint-angle monitoring is operationalised through a configurable threshold model (angle bands, ordinal scores, persistence and re-detection intervals per joint) rather than a fixed scoring table. Second, it reports the quantitative indicators actually measured on the pilot line: detection performance against human inspection on a 500-frame sample, database-level volumetrics of more than 400,000 alert events over roughly ten weeks, and computational load and latency figures under stress. Third, it offers a critical appraisal of what this evidence does and does not demonstrate, identifying the methodological gaps that must be closed before the approach can be considered ergonomically validated. Throughout, the authors distinguish between what was demonstrated by the project itself, as recorded in its technical reports, kick-off documentation and interface evidence; what has been reported by previous research; and interpretations or recommendations proposed by the authors.

The remainder of the paper is organised as follows. Section 2 reviews observational ergonomic methods, markerless pose estimation and prior vision-based ergonomic assessment. Section 3 describes the study design, the system architecture, the pose and angle pipeline, the threshold model and the validation procedures. Section 4 presents the results and the limitations identified by the authors, and Section 5 concludes with recommendations for further validation.

2. Literature Review

2.1 Ergonomic risk factors and observational assessment methods

Epidemiological evidence links awkward postures, repetition, force and insufficient recovery time to disorders of the neck, shoulder, elbow and wrist (Punnett & Wegman, 2004). Observational assessment tools were developed to make such exposures auditable with minimal instrumentation. OWAS classifies back, arm and leg postures together with load into four action categories (Karhu et al., 1977). RULA scores upper-arm, lower-arm, wrist, neck, trunk and leg positions using angular ranges (for example, lower-arm flexion between 60° and 100° scores 1, whereas flexion below 60° or above 100° scores 2) and combines them with muscle-use and force factors into a grand score from 1 to 7 (McAtamney & Corlett, 1993). REBA extends this logic to whole-body tasks (Hignett & McAtamney, 2000). ISO 11226 provides complementary guidance on static working postures, including holding-time limits as a function of trunk and neck inclination (ISO, 2000). David (2005) and Lowe et al. (2019) review the practical limitations of these methods: they sample short observation windows, require trained raters, exhibit inter-observer variability and are rarely applied continuously.

2.2 Markerless human pose estimation

Deep-learning pose estimators infer body keypoints directly from images. OpenPose introduced bottom-up multi-person 2-D estimation with part affinity fields (Cao et al., 2021); HRNet improved keypoint precision through high-resolution representations (Sun et al., 2019); ViTPose showed that plain vision transformers achieve state-of-the-art accuracy on COCO and related benchmarks (Xu et al., 2022). For edge and real-time use, Google's MediaPipe framework (Lugaresi et al., 2019) provides BlazePose, a lightweight detector–tracker that outputs 33 body landmarks with 2-D image coordinates and a 3-D metric estimate (Bazarevsky et al., 2020), and a hand-landmark model that returns 21 keypoints per hand (Zhang et al., 2020). Object detectors of the YOLO family (Redmon et al., 2016), including the Ultralytics YOLOv8 series (Jocher et al., 2023), are widely used to localise persons before landmark extraction. Standard accuracy metrics are the Percentage of Correct Keypoints (PCK) at a normalised distance threshold (Yang & Ramanan, 2013) for 2-D estimation and the Mean Per-Joint Position Error (MPJPE) in millimetres for 3-D estimation, usually reported on Human3.6M (Ionescu et al., 2014). Reference values from these benchmarks are obtained on posed or laboratory data and do not transfer automatically to factory conditions with occlusions, unusual viewpoints and variable illumination.

2.3 Vision-based ergonomic assessment

Research on automating observational methods with pose data began with depth sensors. Diego-Mas and Alcaide-Marzal (2014) used Kinect skeletons to compute OWAS categories and reported good agreement for trunk and arm postures but difficulties with wrists and self-occlusion. Plantard et al. (2017) validated Kinect-based RULA scoring in real workplaces, showing that occlusion-related errors could be partially corrected by trajectory-based reconstruction. Manghisi et al. (2017) implemented real-time RULA with the Kinect v2 and found substantial agreement with expert raters for most body regions. Abobakr et al. (2019) trained a deep network on RGB-D data to predict RULA-relevant angles directly. With RGB-only estimators, Kim et al. (2021) evaluated OpenPose for RULA and REBA scoring and reported acceptable agreement for large joints but larger errors for the wrist and for oblique views; Lin et al. (2022) presented a real-time system that selects and applies RULA, REBA or OWAS from pose data. Wearable inertial approaches have also been used to trigger real-time posture warnings (Yan et al., 2017), but they require instrumenting each worker, which limits acceptance on high-turnover lines.

Three recurring findings from this literature are directly relevant to the present case: wrist and hand angles are the least reliable outputs of body-pose models, which motivates dedicated hand-landmark models; single-camera oblique views degrade angle accuracy, so 3-D or multi-view estimation is preferable; and validation is usually performed against expert scoring of the same frames, using agreement statistics, rather than against independent health outcomes.

2.4 Explainability, privacy and maturity frameworks

The adoption of AI in occupational safety raises requirements of interpretability and data protection. Post-hoc attribution methods such as Grad-CAM (Selvaraju et al., 2017) and SHAP (Lundberg & Lee, 2017) are commonly cited for explaining vision and tabular models; in ergonomic monitoring, however, the most direct form of explanation is often the visual overlay of the estimated skeleton together with the explicit rule (joint, angle, threshold, duration) that triggered an alert. Brazil's General Data

Protection Law (LGPD) treats biometric data as sensitive, so camera-based worker monitoring must justify its legal basis and minimise identification (Brasil, 2018). Applied RD&I programmes commonly characterise progress with the Technology Readiness Level (TRL) scale, where TRL 6 denotes a prototype demonstrated in a relevant environment (Mankins, 1995), and with the acatech Industrie 4.0 Maturity Index, whose stages progress from computerisation and connectivity through visibility and transparency to predictive capacity and adaptability (Schuh et al., 2020). Human-centred manufacturing, in which digital technologies are oriented toward worker well-being, is also the guiding principle of the European Industry 5.0 agenda (Breque et al., 2021).

3. Methodology

3.1 Study design and context

The study is a single-site case study of an applied research and experimental development project carried out between 5 January and 5 August 2026 (technical execution period; financial execution extended to 30 October 2026). Work was organised in four macro-activities, namely project management, requirements survey and analysis (months 1–2), design and development (months 2–7) and testing, homologation and delivery (months 3–7), following a hybrid PMBOK/Scrum approach with sprints, a Jira backlog, GitLab repositories with CI/CD pipelines and monthly status reports. The industrial context is the surface-mount technology (SMT) PCB assembly and inspection line of the case-study plant, an electronics manufacturer located in Manaus. The kick-off meeting of 15 January 2026 recorded that the scope was directed at the generation of technical knowledge, controlled experimentation, technological validation on a pilot line and strategic recommendations, and explicitly not at definitive deployment in production.

A specialised consultancy in NR-12 (machine safety) and NR-17 (ergonomics) was engaged from month 3 to translate the regulatory requirements into computational parameters, to define angular limits and acceptable isometric holding times, to structure the severity matrix for alert triggering, to advise on LGPD compliance and, in month 6, to audit the algorithms' behaviour on the factory floor, including false-positive and false-negative rates, and to guide the fine calibration of angle and time thresholds. The consultancy endorsed the use of anonymous 3-D skeletons without facial biometric identification as the basis for monitoring.

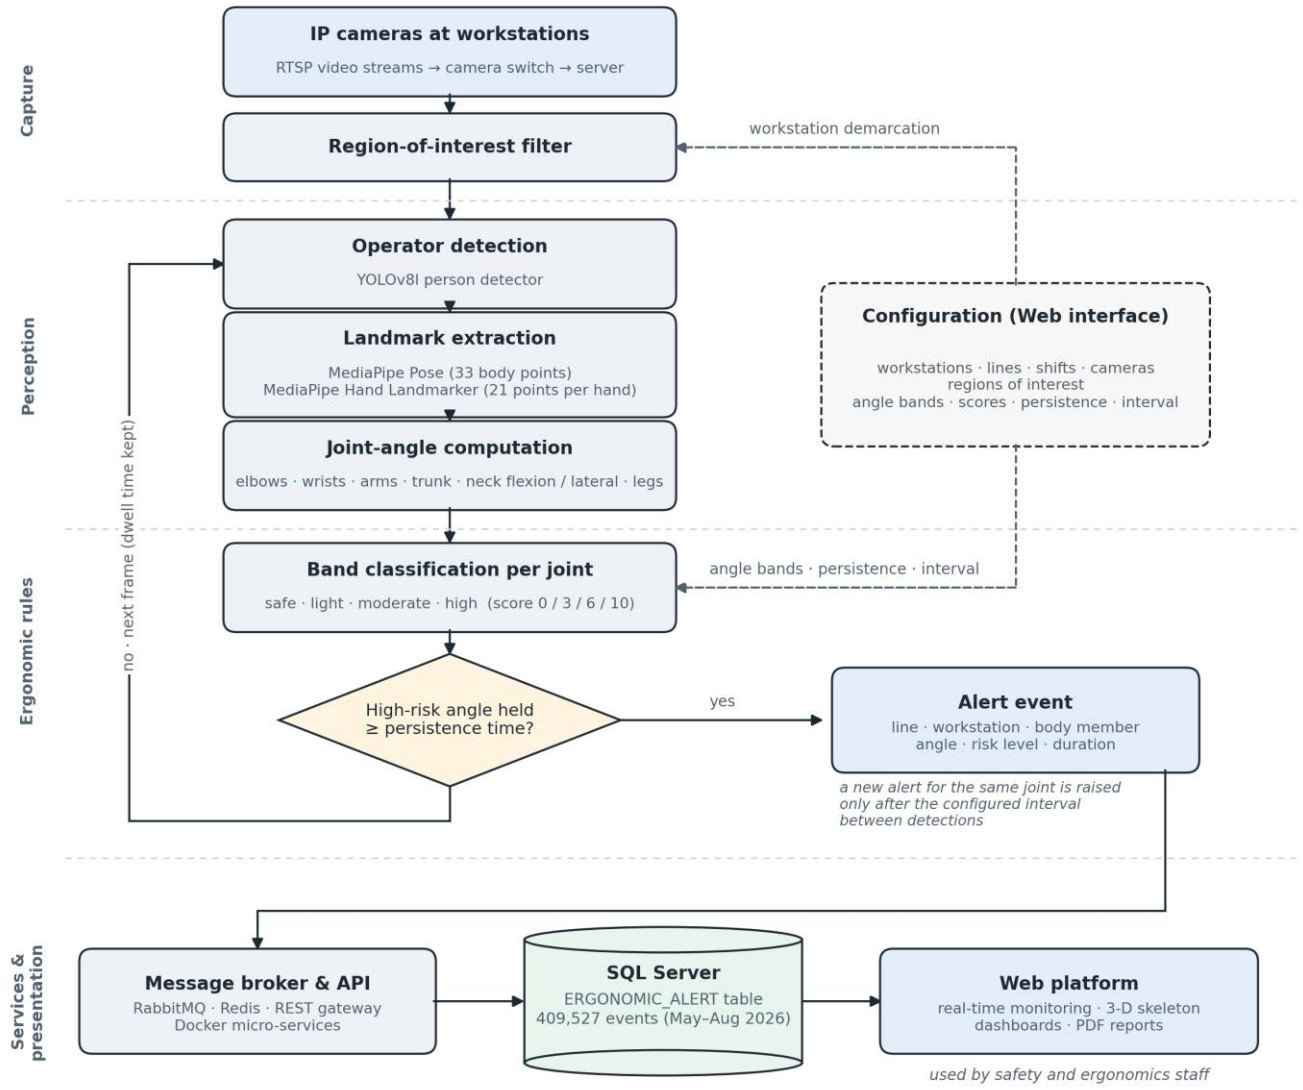
3.2 Requirements elicitation

Month 1 combined executive planning with the manufacturer (priority functionalities, pilot line, technical constraints and TRL 6 validation criteria), a review of scientific and technical literature on AI and computer vision for ergonomics, and analysis of the plant's operational processes, work instructions and workstation data to identify biomechanical variables observable by camera. Three technical work streams were defined: ergonomics (posture and repetitive-movement analysis based on NR-17, RULA and OWAS), layout and continuous improvement (compatibility with the existing line layout) and data security and privacy (LGPD, anonymisation, access control). The requirements phase concluded that the feasibility of the solution would depend on the correct mathematical modelling of joint angles, postural exposure time and frequency of repetitive movements. Month 2 produced the conceptual hardware/software architecture, user stories registered in Jira and low-fidelity wireframes.

3.3 System architecture and processing pipeline

Figure 1 summarises the processing pipeline as reconstructed from the project's architecture drawing, container evidence and interface screens. The architecture comprises four layers.

Figure 1. Processing pipeline of the ergonomic monitoring platform, from camera capture to alerts and dashboards. Dashed elements indicate configuration inputs.



Source : Conecthus (2025).

Capture and edge layer. IP cameras positioned at the monitored workstations deliver video over RTSP through a dedicated camera switch to the processing server. The approved equipment list included industrial and 4K IP cameras, industrial PoE switches, a core switch, Cat-6 and optical cabling and a rack server (Lenovo ThinkSystem SR650 V3) as the processing core. Monitored areas are delimited in the Web interface by drawing regions of interest ("workstation demarcation") over each camera's field of view, so that people outside productive areas are not analysed.

Processing and intelligence layer. Inference runs in a containerised (Docker) set of Python micro-services. The container evidence shows an API gateway (Gunicorn), a line aggregator, an inference service (Uvicorn), a visualiser service and an ergonomics processor, supported by Redis and RabbitMQ 3.11 for messaging. Operators are localised with a YOLOv8l detector; MediaPipe Pose and MediaPipe Hand Landmarker then extract body and hand landmarks for each detected operator. The ergonomics processor computes joint angles from the landmarks, evaluates them against the configured thresholds and persistence rules and emits alert events.

Data and services layer. Alerts and monitoring data are persisted in a Microsoft SQL Server database (table ERGONOMIC_ALERT) and exposed through a RESTful API. The documented endpoints cover monitoring (real-time status per workstation and

member, alert history, camera and workstation streams), angle parameterisation (CRUD operations, activation toggle, AI rules and offset degrees), workstations and their demarcations, production lines, shifts, dashboards, reports and user profiles. Authentication is token-based with refresh tokens and role-based access control; images are anonymised and no facial biometric identification is performed.

Presentation layer. A React/Node.js Web application provides the modules Dashboard, Real-Time Ergonomic Monitoring, Reports (with PDF export), Workstations, Production Shifts, Production Lines, Devices, Angle Parameterisation, Workstation Demarcation, Users and Profiles. A mobile front end was part of the approved scope; the reports describe integration between Web, mobile and API modules but provide interface evidence only for the Web application.

One deviation from the original plan should be recorded. The proposal named YOLOv11-Pose and ViTPose as the intended pose estimators and Grad-CAM/SHAP as explainability layers, and anticipated an annotated dataset of more than 20,000 images for training. The month-2 design report states that, after evaluating candidate models for factory performance, latency and hardware compatibility, the team selected YOLO for person detection and MediaPipe for landmark extraction. No evidence of dataset construction, annotation or model fine-tuning appears in the technical reports; the detection and pose models therefore appear to be pre-trained models applied without domain-specific retraining, which the authors identify as a methodological gap (Section 4.7).

3.4 Pose estimation, joint-angle computation and 3-D visualisation

The proof of concept implemented in month 2 demonstrated the extraction of joint angles (elbows are given as the example) from MediaPipe landmarks and their display in a three-dimensional plot with per-joint angle and score annotations. The production interface provides, per monitored workstation, a live camera view with the skeleton overlaid on the operator, a table of analysed ergonomic factors with their current alert level and dwell time, a 3-D skeletal view and an alert history. The interface legend lists eight monitored factors: right elbow, left elbow, trunk, neck flexion, neck lateral inclination, right arm, left arm and crossed legs; the parameterisation module additionally lists right and left wrist among the registered parameters.

The technical reports describe the computation only at the level of calculating joint angles from anatomical points and translating pose inference into mathematical calculations tied to the NR-17 risk matrix. The exact formulation is not specified in the available documentation: whether angles are computed in the 2-D image plane or from MediaPipe's 3-D world landmarks, which landmark triplets define each joint, how left/right assignment and occlusion are handled, and whether temporal smoothing is applied before thresholding are all left undescribed. The parameterisation API exposes an "offset degrees" resource, which suggests that per-camera angular offsets can be configured, but its semantics are not documented. The authors flag this as a reporting gap that limits reproducibility (Section 4.7).

3.5 Ergonomic thresholds and alert logic

Threshold definition is data-driven and configurable rather than hard-coded. For each anatomical point, the parameterisation form requires a description and, for each of three risk levels (light, moderate, high), an angle condition and an ordinal score on a 1–12 scale; for the high-risk level it additionally requires a persistence time (seconds) and an interval between detections (seconds). A fourth "safe" band covers the remaining range. Table 1 reproduces the parameters that are legible in the interface evidence (created on 27 May 2026 and last edited on 4 August 2026; see Figure 2 in Section 4.4).

Table 1. Registered angle parameterisations for wrists and elbows (created 27 May 2026; last edited 4 August 2026).

Anatomical point	Safe (score 0)	Light (score 3)	Moderate (score 6)	High (score 10)	High-risk persistence	Interval between detections
Right wrist	0°–15°	16°–29°	30°–44°	45°–180°	10 s	5 s
Left wrist	0°–15°	16°–29°	30°–44°	45°–180°	10 s	5 s
Right elbow	0°–30°	30°–59°	60°–89°	90°–180°	15 s	5 s
Left elbow	0°–30°	30°–59°	60°–89°	90°–180°	15 s	5 s

Source: angle-parameterisation module of the platform (project technical report and Figure 2). The module listed eight registered items in total; the parameters for trunk, neck and arm factors were not legible in the available evidence and are therefore not reported.

The descriptions attached to the wrist parameters ("flexed or supported for a long time") and elbow parameters ("supported and flexed for a long time (bent/closed)") indicate that the bands model sustained flexion rather than instantaneous deviation. The

alert logic, as far as it can be reconstructed from the interface and reports (Figure 1, decision block), operates as follows: for each frame in which an operator is tracked, the angle of each configured joint is classified into a band; a high-risk alert is raised only if the angle remains in the high band for at least the persistence time; and a new alert for the same joint is not raised until the configured interval has elapsed. The Real-Time Monitoring view displays, for each factor, the current level (safe, light, moderate, high) and the time already spent in that state; the sidebar aggregates the number of collaborators monitored, active workstations, the distribution of collaborators across risk levels, a count of critical alerts and an overall risk index expressed as a percentage. The reports state that the alert severity matrix and the final calibration of angle and time thresholds were defined with the NR-12/NR-17 consultancy; the documents do not give the derivation of each numeric band from the normative texts or from RULA/OWAS tables.

3.6 Data storage, reporting and explainability

Each alert record contains the date and time, the type of risk (for example, left elbow flexion, neck flexion, trunk inclination or left wrist deviation), its duration and its risk level, together with the identifiers of the production line, workstation and affected body member. The Reports module consolidates, per workstation and period, the total number of alerts, the cumulative time in risk exposure, that time as a percentage of monitored time, the number of monitored hours, a ranking of risk factors with occurrence counts and cumulative exposure, and a ranking of production lines. The Dashboard shows alert trends per hour and per weekday, workstations with most alerts and a ranking of postural risk types, with filters by period, shift, production line and workstation.

The project characterises its explainability approach as "visual XAI": rather than post-hoc attribution maps, the system projects the estimated 3-D skeleton over the operator's image, annotates each joint with the measured angle and its score, and isolates the root cause of every alert by workstation identifier, line identifier and member name; the stored alert frames in Figure 3 (Section 4.4) illustrate this overlay, with the offending segment highlighted and faces blurred. The authors note that this constitutes rule transparency (the triggering joint, angle, threshold and duration are exposed) rather than a model-level explanation of the pose estimator itself; Grad-CAM and SHAP, foreseen in the proposal, were not implemented according to the reports.

3.7 Validation procedures

Validation proceeded in stages that the project mapped onto TRL levels: a laboratory proof of concept (TRL 3, month 2); integrated validation of Web, API, SQL Server, AI and workstation-management components in a controlled environment (TRL 4); a quantitative field assessment on a 500-frame sample from the production line compared with human inspection (TRL 5); and demonstration of the integrated prototype on the SMT line with continuous operation, a stress test and a formal QA approval (TRL 6, month 7). Functional, regression and user-acceptance tests (UAT) were run by the internal QA team and with end users on the pilot line during months 3–7, and the consultancy audited algorithm behaviour on the factory floor in month 6.

Several limitations of the validation protocol are evident from the documentation and are discussed in Section 4.7: the human-inspection ground truth for the 500-frame sample is not described in terms of number of annotators, annotation instructions or inter-rater agreement; the "base indicators" of the models (Section 4.1) are not attributed to a source and were not produced by the study; the mean absolute error (MAE) formula for joint angles is stated but no MAE value or reference measurement is reported; and no comparison of the system's risk classification with RULA/OWAS scores assigned by ergonomists was performed, although such a comparison was among the specific objectives.

4. Results and Findings

4.1 Detection and landmark-extraction performance

Table 2 summarises the indicators reported in the final technical report. The "base" values are presented there as reference indicators of the models; their provenance (benchmark dataset, evaluation protocol) is not stated, and they should be read as reported reference figures rather than as measurements of the present study. The "operational" values were obtained on the 500-frame field sample from the SMT line against human inspection.

Table 2. Reported model indicators.

Component	Metric	Reported value	Origin
YOLOv8l (operator detection)	Precision	88.9%	Base indicator (source not stated)
YOLOv8l (operator detection)	Recall	84.5%	Base indicator (source not stated)

Component	Metric	Reported value	Origin
YOLOv8l (operator detection)	F1-score	86.6%	Base indicator (source not stated)
YOLOv8l (operator detection)	mAP50	90.2%	Base indicator (source not stated)
MediaPipe Pose & Hand Landmarker	MPJPE	25 mm	Base indicator (source not stated)
MediaPipe Pose & Hand Landmarker	PCK@0.2	> 92%	Base indicator (source not stated)
MediaPipe Pose & Hand Landmarker	Hand-landmark mean precision	95.7%	Base indicator (source not stated)
YOLOv8l (operator detection)	Operational precision	94.5%	Field sample, 500 frames, SMT line
YOLOv8l (operator detection)	Operational recall	98.2%	Field sample, 500 frames, SMT line
YOLOv8l (operator detection)	Operational F1-score	96.3%	Field sample, 500 frames, SMT line
MediaPipe Pose	Hit rate on critical joints	91.8%	Field sample, 500 frames, SMT line
MediaPipe Hands	Continuous tracking rate	89.4%	Field sample, 500 frames, SMT line

Source: Conecthus (2025).

The operational F1-score was computed as F1 = 2 × (0.945 × 0.982)/(0.945 + 0.982) ≈ 96.3%, exceeding the ≥ 90% AI-precision target set at kick-off. Operational precision and recall on the field sample are higher than the reported base values, which is plausible for a constrained scene with few operators at fixed workstations, but which also underlines that the sample is small and that its representativeness (number of operators, workstations, times of day, occlusion conditions) is not documented. The "hit rate on critical joints" and "continuous tracking rate" are project-specific metrics whose operational definitions (what counts as a hit, which joints are critical, how continuity is measured) are not given in the reports.

4.2 Alert volumetrics and temporal stability

Analysis of the ERGONOMIC_ALERT table for the period 22 May–5 August 2026 recorded 409,527 processed ergonomic events. The mean persistence of alerts, defined as uninterrupted tracking of the same deviation, was 72.0 s; at the operational rate of 15 frames per second this corresponds to 72.0s × 15 fps = 1,080 consecutively processed frames per detected deviation. For deviations classified as high risk, the mean angle of the monitored joints was 129.14°. The report states that the geometric coherence of these angles was audited with the mean absolute error MAE = (1/n) Σ |calculated angle - reference angle|, but no MAE value, sample size or reference-angle source is reported.

Interface evidence provides two illustrative per-workstation reports (individual alert records and frames are shown in Figure 3, Section 4.4). For a test workstation ("Posto teste 1", one production line, 1–20 March 2026), 19 alerts were recorded in 160.17 monitored hours, with 5 min 19 s of cumulative risk exposure (0.06% of monitored time); the leading factors were left elbow flexion (8 occurrences, 2 min 21 s), neck flexion (7, 1 min 56 s), right elbow flexion (2, 31 s) and trunk inclination (2, 30 s), and all listed alerts were of high severity, with durations from 1 s to 42 s. For a second test workstation ("Equipe DEV", 26–30 May 2026), 3,749 alerts were recorded in 44 monitored hours, with 76 h 54 min of cumulative exposure reported as 174.79% of monitored time; right elbow flexion dominated (2,898 occurrences, 69 h 33 min), followed by left elbow flexion (841, 7 h 19 min), left wrist deviation (6, 1 min 15 s) and trunk inclination (4, 24 s). The authors interpret the exposure percentage above 100% as the result of summing exposure across concurrently active factors (for example, both elbows) without de-duplication in time; this is an observation from the evidence, not a statement made in the reports. Both examples come from development or test workstations rather than from the production line and are shown only to illustrate the report structure and the behaviour of the alert logic; they should not be read as ergonomic findings about the SMT operators.

4.3 Computational performance, latency and stress testing

Final integration tests were run in the containerised environment under synthetic concurrent load. With approximately 250 simultaneous requests, the effective throughput at CPU saturation was 1.64 fps and end-to-end latency at peak was approximately 650 ms. A continuous stress run of 152 s with uninterrupted requests and an internal 3 s timeout produced zero out-of-memory events and zero critical failures in the API gateway or in RabbitMQ. Earlier, in month 6, database queries had been refactored from in-memory processing to native aggregations in SQL Server, which the reports describe as significantly reducing the response times of critical queries; no before/after figures are given.

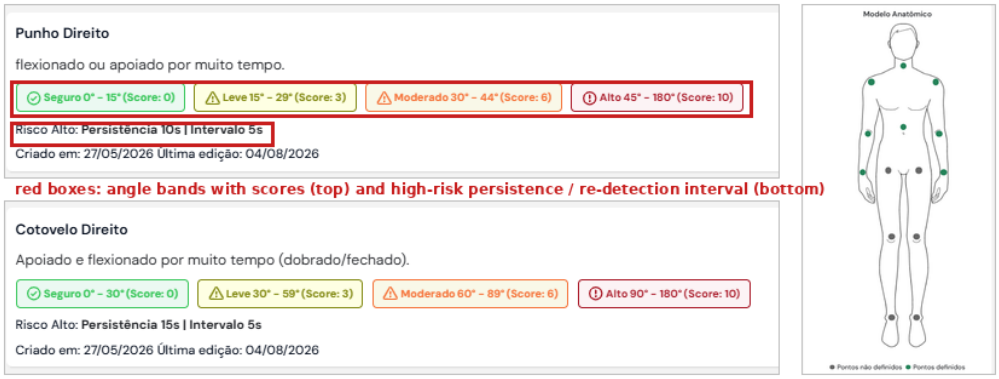
The relationship between these figures and the 15 fps operational rate used in Section 4.2 deserves attention. The 1.64 fps figure describes system throughput under a saturation load of concurrent API requests, whereas 15 fps describes the frame rate of the

video-processing path in normal operation; the reports do not specify how many camera streams were processed concurrently at 15 fps, the GPU/CPU resources used, or the latency from frame capture to alert display under normal (non-saturated) conditions. Without these data, the real-time claim is supported qualitatively, in that alerts were displayed live during UAT on the pilot line, rather than by a measured capture-to-alert latency.

4.4 Platform evidence: configuration, alert records and dashboards

Figures 2 and 3 reproduce screens captured from the platform in operation. They complement the numerical indicators with evidence of how thresholds and alerts are presented to users, and they are the basis for several observations made elsewhere in this paper.

Figure 2. Angle-parameterisation module showing the registered right-wrist and right-elbow parameters (the left-side parameters are identical). Each anatomical point carries a description, four colour-coded angle bands with ordinal scores and, for the high-risk band, a persistence time and an interval between detections; the anatomical model (right) marks the points for which parameters are defined.



Source : Conecthus (2025).

Figure 2 shows the angle-parameterisation module from which Table 1 was transcribed. Each anatomical point carries a textual description, four colour-coded angle bands with their ordinal scores and, for the high-risk band, the persistence time and the interval between detections; a toggle enables or disables each parameter, and the anatomical model on the right marks which points have parameters defined. The screen records that the wrist and elbow parameters were created on 27 May 2026 and last edited on 4 August 2026, which is consistent with the consultancy-guided fine calibration of angle and time thresholds reported for the final months of the project.

Figure 3. Alert records with the corresponding frames, as stored by the platform: (a) right-wrist deviation and (b) trunk inclination at a test workstation (July 2026); (c) right-elbow flexion and (d) left-elbow flexion at SMT-line workstations (September 2026). The estimated skeleton is overlaid on the operator, the segment responsible for the alert is highlighted and faces are blurred.



Source : Conecthus (2025).

Figure 3 shows four alert records with the corresponding frames. Each record carries the date and time, the type of risk, the duration and the risk level, and the stored frame shows the estimated skeleton overlaid on the operator with the body segment responsible for the alert highlighted and the face blurred. Panels (a) and (b) were captured at a development/test workstation in July 2026 (a right-wrist deviation held for 12 s while the operator's hand rested behind the head, and a trunk inclination of 11 s while leaning toward the desk), and panels (c) and (d) at SMT-line workstations in September 2026: a right-elbow flexion of 37 s and a left-elbow flexion of 28 s during seated board handling by operators in cleanroom garments. All four durations exceed the persistence times in Table 1 (10 s for wrists, 15 s for elbows), as the alert logic requires. Panels (c) and (d) are also the only evidence available to the authors of the system operating on the production line after the formal close of the project on 5 August 2026. A dashboard snapshot for 9 September 2026 (not reproduced) is consistent with this: it reported 32 alerts across 7 monitored workstations, with 2 of 13 registered cameras active and 9 h 2 min of monitored time, all attributed to right-elbow (20) and left-elbow (12) flexion.

4.5 Functional homologation and acceptance

Functional tests in months 3–5 validated navigation, dashboards, dynamic filters (period, shift, production, workstation), PDF report export and workstation management; month 6 added regression batteries, load simulations, MES/WMS interface tests, token-based authentication with session refresh, password-recovery fixes and inactive-user notifications; month 7 concluded UAT cycles on the pilot line covering authentication, real-time monitoring, environment parameterisation, device management, dashboards, reports and user administration, plus integrated tests of the full RTSP → Docker inference → SQL Server → Web path. The Test Report Summary v1.0.0 (Web CMS / API / AI), dated 1 July 2026, recorded a verdict of APPROVED with zero open items (0 impediments; 0 high-, normal- or low-severity failures). The kick-off had set targets of ≥ 95% test success and ≥ 80% user approval; the reports document the zero-open-issue QA verdict but do not report a quantified user-approval survey.

4.6 Technology and Industry 4.0 maturity self-assessment

The project self-assessed TRL 1 through 6 as achieved, citing respectively the mapping of NR-12/NR-17 and the study of YOLO and MediaPipe; the definition of the architecture and pose-rule processing flow; the laboratory proof of concept for joint-angle calculation; integrated component validation in a controlled environment; the 500-frame quantitative field validation (F1 = 96.3%); and the demonstration of the integrated prototype on the plant's SMT line with 409,527 processed events, the 250-request stress test and the approved test summary. Against the acatech index, stages 1–4 were claimed: digitalisation of the ergonomic assessment process and normative rules (1), continuous integration of RTSP cameras, Docker AI services, back end and Web front end via APIs and CI/CD (2), real-time display of alerts and of the risk index in operational dashboards (3), and visual explainability with root-cause isolation by workstation, line and member (4). These levels are self-assessed by the executing team; no external TRL audit is reported, although the consultancy issued a final technical opinion confirming compliance with NR-12 and NR-17.

4.7 Limitations and methodological gaps identified by the authors

Based on a critical reading of the available evidence, the authors identify several aspects that could be further explored to strengthen the validation, generalisation, and scientific robustness of the proposed approach. This validation points to different directions of future research and development that would strengthen the system and support wider adoption.

Ergonomic validity. The field assessment on 500 frames covered operator detection. The next step is to check the risk-level classification against expert RULA, OWAS and REBA scoring of the same frames, using agreement statistics such as Cohen's kappa and the intraclass correlation coefficient. This would complete the comparison with traditional ergonomic analysis set out in the project objectives.

Joint-angle accuracy. The literature reports that wrist angles are the least reliable outputs of single-camera pose models, while the wrist carries the tightest bands in Table 1 (high risk from 45°). Measuring angular error against a reference system (goniometer, motion capture or inertial sensors) is therefore a priority, especially at the wrist.

Evaluation protocol and baselines. Future evaluations should document the annotation protocol in full: the number and qualification of inspectors, the guidelines, inter-rater reliability, and how the sample is split across operators, workstations, shifts and occlusion conditions. They should also report baseline metrics for YOLOv8l (precision, recall, mAP) and MediaPipe (MPJPE, PCK) measured in the target environment, not only the reference values.

Domain adaptation. The system ran on pre-trained models and still reached operational results. That leaves room to improve it through a domain-specific annotated dataset, fine-tuning, active learning and continuous retraining, all planned for later phases.

Threshold traceability. The angle bands, scores, persistence and alert intervals were set together with the ergonomics consultancy. Showing explicitly how each parameter maps to NR-17, ISO 11226 and the RULA/OWAS tables, for all monitored joints, would make the method easier to audit and transfer to other settings.

Real-time performance. Stress tests found saturation at 250 concurrent requests (1.64 fps, ≈650 ms). Future work should measure how many streams the system sustains at 15 fps and the capture-to-alert latency under normal load, to guide deployment sizing.

5. Conclusion

This case study documented the development and pilot-line validation of an ergonomic monitoring platform that combines a YOLOv8l person detector, MediaPipe body and hand landmark extraction, configurable per-joint angle bands with ordinal scores and persistence rules, and an integrated micro-service and Web stack, deployed on the SMT PCB assembly line of an electronics manufacturer in Manaus, Brazil. Within seven months the project moved from concept to a prototype demonstrated in a relevant environment: operator detection on a 500-frame field sample reached an F1-score of 96.3% against human inspection; the alert pipeline processed 409,527 events over roughly ten weeks with a mean tracking persistence of 72 s at 15 fps; the containerised back end tolerated 250 concurrent requests with about 650 ms latency and no critical failures; and the platform passed formal QA with zero open issues. The project's own assessment places the result at TRL 6 and acatech maturity stage 4.

What the study demonstrates, therefore, is engineering feasibility: a pose-estimation pipeline with configurable joint-angle thresholds can be integrated with industrial cameras, messaging, a relational database and role-based Web dashboards, and can run continuously on a production line while producing rule-transparent alerts (joint, angle, level, duration, workstation) that safety staff can audit. What it does not yet demonstrate is ergonomic validity or health impact. In the absence of a comparison with expert RULA/OWAS scoring, of a joint-angle accuracy study against a reference measurement, of a documented ground-truth protocol and of any outcome data, the alerts must currently be regarded as screening signals whose sensitivity and specificity for actual ergonomic risk are unknown.

The authors therefore recommend the following research steps before scaling: a criterion-validity study in which the system's per-joint bands and overall risk levels are compared with blinded RULA, REBA or OWAS ratings by at least two ergonomists on a stratified sample of frames from several workstations and shifts, reporting agreement statistics; a concurrent-validity study of joint-angle estimates against an inertial or optical reference, with particular attention to wrist angles and oblique camera views, and documentation of whether angles are computed in 2-D or 3-D; publication of the full threshold table with its normative justification and of the formula for the overall risk index; characterisation of real-time performance under normal load (streams per server at 15 fps and capture-to-alert latency); correction of exposure-time aggregation so that concurrent factors are not double-counted; and a prospective evaluation of user acceptance and of intermediate outcomes (alert rates over time, workstation modifications triggered by alerts and self-reported discomfort) before any claim about reductions in WMSD incidence is made. With these studies in place, the configurable-threshold, rule-transparent architecture described here offers a practical route from periodic manual ergonomic inspection to continuous, auditable, data-driven ergonomic management on electronics assembly lines.

Funding: This research was funded within the Priority Programme for Industry 4.0 and Industrial Modernisation (PPI 4.0), coordinated by CITS.AMAZONAS – Centro Internacional de Tecnologia de Software do Amazonas, under Plan of Use of Resources PUR 522/2025, with Foxconn MOEBG Indústria de Eletrônicos Ltda. as the beneficiary company and Instituto Conecthus as the executing institution.

Conflicts of Interest: The research was conducted in the context of an industry-funded RD&I project involving the executing institute, the coordinating entity, and the beneficiary company. On behalf of all parties involved, the authors declare that they have no conflicts of interest related to this research.

Ethics and Data Availability: The project reports state that image data were processed with anonymisation, that no facial biometric identification was performed (faces are blurred in the stored alert frames, Figure 3), that access was controlled through role-based permissions and that a specialised consultancy issued guidance on LGPD compliance in the context of occupational health and safety. The documents available to the authors do not record a research-ethics committee review or a formal informed-consent procedure for the operators observed on the pilot line; this should be clarified before publication. The alert database and video data belong to the beneficiary company and contain workplace footage; they are not publicly available. Aggregated indicators reported in this article were extracted from the project's technical reports.

Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers.

References

[1] Abobakr, A., Nahavandi, D., Hossny, M., Iskander, J., Attia, M., Nahavandi, S., & Smets, M. (2019). RGB-D ergonomic assessment system of adopted working postures. Applied Ergonomics, 80, 75–88.

[2] Bazarevsky, V., Grishchenko, I., Raveendran, K., Zhu, T., Zhang, F., & Grundmann, M. (2020). BlazePose: On-device real-time body pose tracking. arXiv. https://arxiv.org/abs/2006.10204

[3] Brasil. (2018). Lei nº 13.709, de 14 de agosto de 2018 – Lei Geral de Proteção de Dados Pessoais (LGPD). Diário Oficial da União.

[4] Brasil. Ministério do Trabalho e Previdência. (2022). Norma Regulamentadora NR-17 – Ergonomia (Portaria MTP nº 423, de 7 de outubro de 2021). Diário Oficial da União.

[5] Breque, M., De Nul, L., & Petridis, A. (2021). Industry 5.0: Towards a sustainable, human-centric and resilient European industry. European Commission, Directorate-General for Research and Innovation.

[6] Cao, Z., Hidalgo, G., Simon, T., Wei, S.-E., & Sheikh, Y. (2021). OpenPose: Realtime multi-person 2D pose estimation using part affinity fields. IEEE Transactions on Pattern Analysis and Machine Intelligence, 43(1), 172–186.

[7] David, G. C. (2005). Ergonomic methods for assessing exposure to risk factors for work-related musculoskeletal disorders. Occupational Medicine, 55(3), 190–199.

[8] Diego-Mas, J. A., & Alcaide-Marzal, J. (2014). Using Kinect sensor in observational methods for assessing postures at work. Applied Ergonomics, 45(4), 976–985.

[9] Hignett, S., & McAtamney, L. (2000). Rapid Entire Body Assessment (REBA). Applied Ergonomics, 31(2), 201–205.

[10] International Organization for Standardization. (2000). ISO 11226:2000 – Ergonomics: Evaluation of static working postures. ISO.

[11] Ionescu, C., Papava, D., Olaru, V., & Sminchisescu, C. (2014). Human3.6M: Large scale datasets and predictive methods for 3D human sensing in natural environments. IEEE Transactions on Pattern Analysis and Machine Intelligence, 36(7), 1325–1339.

[12] Jocher, G., Chaurasia, A., & Qiu, J. (2023). Ultralytics YOLOv8 (Version 8.0) [Computer software]. https://github.com/ultralytics/ultralytics

[13] Karhu, O., Kansi, P., & Kuorinka, I. (1977). Correcting working postures in industry: A practical method for analysis. Applied Ergonomics, 8(4), 199–201.

[14] Kim, W., Sung, J., Saakes, D., Huang, C., & Xiong, S. (2021). Ergonomic postural assessment using a new open-source human pose estimation technology (OpenPose). International Journal of Industrial Ergonomics, 84, 103164.

[15] Lin, P.-C., Chen, Y.-J., Chen, W.-S., & Lee, Y.-J. (2022). Automatic real-time occupational posture evaluation and select corresponding ergonomic assessments. Scientific Reports, 12, 2139.

[16] Lowe, B. D., Dempsey, P. G., & Jones, E. M. (2019). Ergonomics assessment methods used by ergonomics professionals. Applied Ergonomics, 81, 102882.

[17] Lugaresi, C., Tang, J., Nash, H., McClanahan, C., Uboweja, E., Hays, M., Zhang, F., Chang, C.-L., Yong, M. G., Lee, J., Chang, W.-T., Hua, W., Georg, M., & Grundmann, M. (2019). MediaPipe: A framework for building perception pipelines. arXiv. https://arxiv.org/abs/1906.08172

[18] Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems, 30, 4765–4774.

[19] Manghisi, V. M., Uva, A. E., Fiorentino, M., Bevilacqua, V., Trotta, G. F., & Monno, G. (2017). Real time RULA assessment using Kinect v2 sensor. Applied Ergonomics, 65, 481–491.

[20] Mankins, J. C. (1995). Technology Readiness Levels: A white paper. NASA Office of Space Access and Technology.

[21] McAtamney, L., & Corlett, E. N. (1993). RULA: A survey method for the investigation of work-related upper limb disorders. Applied Ergonomics, 24(2), 91–99.

[22] Plantard, P., Shum, H. P. H., Le Pierres, A.-S., & Multon, F. (2017). Validation of an ergonomic assessment method using Kinect data in real workplace conditions. Applied Ergonomics, 65, 562–569.

[23] Punnett, L., & Wegman, D. H. (2004). Work-related musculoskeletal disorders: The epidemiologic evidence and the debate. Journal of Electromyography and Kinesiology, 14(1), 13–23.

[24] Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016). You Only Look Once: Unified, real-time object detection. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 779–788.

[25] Schuh, G., Anderl, R., Dumitrescu, R., Krüger, A., & ten Hompel, M. (Eds.). (2020). Industrie 4.0 Maturity Index: Managing the digital transformation of companies – Update 2020 (acatech STUDY). acatech – National Academy of Science and Engineering.

[26] Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2017). Grad-CAM: Visual explanations from deep networks via gradient-based localization. Proceedings of the IEEE International Conference on Computer Vision (ICCV), 618–626.

[27] Sun, K., Xiao, B., Liu, D., & Wang, J. (2019). Deep high-resolution representation learning for human pose estimation. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 5693–5703.

[28] Xu, Y., Zhang, J., Zhang, Q., & Tao, D. (2022). ViTPose: Simple vision transformer baselines for human pose estimation. Advances in Neural Information Processing Systems, 35, 38571–38584.

[29] Yan, X., Li, H., Li, A. R., & Zhang, H. (2017). Wearable IMU-based real-time motion warning system for construction workers' musculoskeletal disorders prevention. Automation in Construction, 74, 2–11.

[30] Yang, Y., & Ramanan, D. (2013). Articulated human detection and pose estimation with flexible mixtures of parts. IEEE Transactions on Pattern Analysis and Machine Intelligence, 35(12), 2878–2890.

[31] Zhang, F., Bazarevsky, V., Vakunov, A., Tkachenka, A., Sung, G., Chang, C.-L., & Grundmann, M. (2020). MediaPipe Hands: On-device real-time hand tracking. arXiv. https://arxiv.org/abs/2006.10214