Real-time IoT-Driven Physical Fatigue Monitoring for Human-Centric Business Processes in Industry 5.0

Milo Duffeleer[1], Emile Goris[1], Mathis Wyffels[1], Estefanía Serral[1], Irene Vanderfeesten[1]

[1] *Research Centre for Information Systems Engineering, KU Leuven, Warmoesberg 26, Brussel, 1000, Belgium*

Abstract

Industry 5.0 reframes industrial innovation around human-centricity, positioning worker well-being as a core design objective. Within this paradigm, physical fatigue stands out as a key well-being dimension that increases the risk of workplace accidents and long-term health risks in Industry 5.0 business processes. Integrating real-time biometric fatigue data from wearable IoT sensors into Business Process Management (BPM) systems would enable fatigue-aware resource allocation decisions, allowing timely interventions before well-being consequences materialize on the work floor. However, while BPM research underscores the need to incorporate worker well-being into process execution, existing work on wearable-based fatigue monitoring has concentrated on offline predictive modeling, and few systems are capable of processing streaming physiological data in real-time. Consequently, fatigue insights are not fed into BPM systems promptly enough to inform timely interventions. This research addresses that gap by developing an end-to-end data pipeline that ingests, processes, and analyzes streaming physiological data from the Movesense HR2 chest sensor to produce real-time physical fatigue classifications. The pipeline integrates data ingestion, signal processing, feature engineering, and a rule-based classification model. An end-to-end demonstration under simulated working conditions confirmed stable real-time operation. The resulting pipeline provides the real-time information required to support fatigue-aware resource allocation within a BPM system.

Keywords

Real-Time Data Pipeline, Physical Fatigue, Wearable, Resource Allocation, Industry 5.0

1. Introduction

In 2021, the term 'Industry 5.0' was formally introduced by the European Commission (EC), which highlighted the importance of a human-centric and sustainable approach, defining Industry 5.0 as a resilient provider of future prosperity [1], [2]. As first introduced by [3] and later built upon in the EC's description of Industry 5.0, the 'Healthy Operator', who uses IoT-based wearable devices to track their well-being, is one tangible way in which technology can be used to put human health at the center of the production process [1]. By continuously capturing physiological indicators, wearables enable real-time worker health data to inform the dynamic allocation of human resources to business process tasks, supporting optimal working conditions and timely interventions that mitigate operational disruptions [4], [5].

While worker well-being encompasses many dimensions, physical fatigue stands out as a particularly critical concern in manufacturing environments. It is widely recognized as an occupational risk in industrial settings [6], with well-documented short-term consequences including discomfort, diminished motor control, and reduced strength capacity. As defined by [7], physical fatigue is a sub-optimal psychophysiological condition caused by exertion, a state of progressive physiological decline in which the body's capacity to sustain physical output deteriorates over time. Left unaddressed, it

Proceedings Name, Month XX–XX, YYYY, City, Country
EMAIL: milo.duffeleer@student.kuleuven.be (A. 1); emile.goris@student.kuleuven.be (A. 2); mathis.wyffels@kuleuven.be (A. 3)
ORCID: 0009-0002-3021-3544 (A. 1); 0009-0002-3516-6900 (A. 2); 0009-0005-4279-128X (A. 3); 0000-0001-7579-910X (A. 4); 0000-0002-6670-5519 (A. 5)

increases the risk of workplace injuries and accidents and contributes to longer-term health problems [9], making timely detection of physical fatigue essential in manufacturing contexts.

This research therefore contributes to the adoption of Industry 5.0's principle of human-centricity in BPM by developing an end-to-end data pipeline that continuously monitors streaming physiological data on physical fatigue using wearable IoT devices and generates fatigue classifications within short, operationally relevant time windows. A data pipeline is understood here as a series of interconnected processing stages, ranging from data acquisition and preparation to analysis and interpretation, with the ultimate goal of generating actionable insights from raw data [9]. The aim is to enable a BPM system to use these real-time fatigue classifications as a basis for smarter, more sustainable resource allocation decisions, intervening before the consequences of physical fatigue materialize on the work floor.

To achieve these goals, this research aims to answer the following overarching research question: How can an end-to-end data pipeline be developed to collect, process, and analyze streaming physiological data for real-time worker physical fatigue monitoring that can enable better resource allocation in Industry 5.0 processes?

The pipeline's development is considered successful if it functions reliably in a simulated working environment, processing streaming physiological data and generating fatigue insights within operationally relevant time windows. To validate this criterion, an experimental demonstration assesses the pipeline's real-time functioning under simulated working conditions, in which participants wear the selected device and perform physically demanding tasks while the system ingests, processes, and analyzes the resulting streaming data. Performance is evaluated against the concrete solution objectives developed in Section 4.

The remainder of this paper is organized as follows. Section 2 positions the work against related studies. Section 3 presents the research methodology. Section 4 analyses worker physical fatigue, identifies its physiological indicators, and derives the concrete solution objectives. Section 5 details the design and development of the end-to-end pipeline. Section 6 describes the demonstration of the resulting artefact, and Section 7 discusses operationalization within Industry 5.0 and limitations. Section 8 concludes with a summary of the study.

2. Related Work

As described in Section 1, IoT-based sensors enable the tracking of worker well-being, from which changes in process execution or resource allocation can be suggested to maintain optimal working conditions [10]. This need has been explicitly recognized in the IoT-meets-BPM manifesto, and is further elaborated in Challenge 7 of the Digital Twins of Business Processes manifesto, which identifies the human factor as a key open challenge at the intersection of IoT and BPM [5]. Integrating real-time well-being data, such as physical fatigue, into dynamic resource allocation mechanisms would directly address this challenge by enabling timely, well-being-driven interventions in process execution [4]. However, as the following review demonstrates, a real-time pipeline capable of enabling this integration remains absent from literature.

This section therefore reviews directly comparable studies along the two dimensions central to this gap: 1) the development of a complete end-to-end data pipeline and 2) the ability to process streaming physiological data from wearables in real-time to enable timely predictions. As defined by [9], a complete data pipeline comprises the following stages: data acquisition, data preparation, data storage, feature engineering, modelling, training, evaluation and prediction. Table 1 summarizes the comparison of these studies and additionally reports the monitored well-being dimension and sensor setup. The studies are divided into two sensor groups: chest-worn heart rate sensors (ECG) combined with motion sensors (IMU) and wrist-worn multimodal devices.

Chest-worn ECG and IMU systems. Three studies [8], [11], [12] pair a chest ECG monitor with multiple body-worn IMUs to model physical fatigue. They all concentrate their effort on offline predictive performance rather than pipeline implementation. The earliest [11] applies penalized regression to predict binary fatigue occurrence, with prediction folded into a one-off evaluation on two held-out participants and acquisition limited to offline onboard recording. Its direct follow-up [12] retains the same sensor setup but replaces the penalized regression with ensemble methods and adds three downstream analytical phases. These additional phases are downstream analytical layers rather than implementations of the missing pipeline stages, so pipeline coverage is unchanged. The most recent of the three [8] predicts continuous fatigue scores with XGBoost and a custom loss function across one chest sensor and five upper-body IMUs, but its "closed-loop real-time" dashboard is presented only as a mock-up. In every case, data was recorded during controlled lab sessions and analyzed retrospectively.

Table 1
Overview of related work on wearable-based worker well-being monitoring

Study	Monitoring Well-Being Dimension	Sensors Used	Missing Pipeline Stages	Real-Time
[11]	Physical fatigue	Chest-worn ECG + 4 IMUs	No Storage, Prediction	No
[12]	Physical fatigue	Chest-worn ECG + 4 IMUs	No Storage, Prediction	No
[8]	Physical fatigue	Chest-worn ECG + 5 IMUs	No Storage, Prediction	No
[13]	Physical fatigue	Wrist-worn multimodal	No Storage, Prediction	No

Wrist-worn multimodal systems. The remaining study [13] instead uses a commercial wrist-worn device that captures a broader set of signals. The work predicts physical fatigue from an Empatica EmbracePlus smartwatch (heart rate, skin conductance, temperature, motion) and contextual features. Its central contribution is a fuzzy-logic labeling method that eliminates mid-shift self-reports, while real-time scoring and the alert policy remain illustrated only in figures.

Two patterns hold across all four studies. First, the full pipeline is never implemented: research effort concentrates on offline predictive performance, leaving stages such as data storage and prediction absent. Second, none performs real-time processing of streaming physiological data, even though several explicitly identify it as the desired endpoint. By focusing on a real-time end-to-end data pipeline, this research addresses both gaps, aiming to deliver fatigue insights to the work floor in time to inform interventions.

3. Research Methodology

The research question formulated in Section 1 is design oriented in nature. It asks how a novel artefact (end-to-end data pipeline) can be developed to address a practical problem. Such questions are well-served by the Design Science Research Methodology (DSRM), which is concerned with constructing innovative artefacts for specific application contexts [14], [15], [16]. Figure 1 shows the DSRM, adapted for this study.

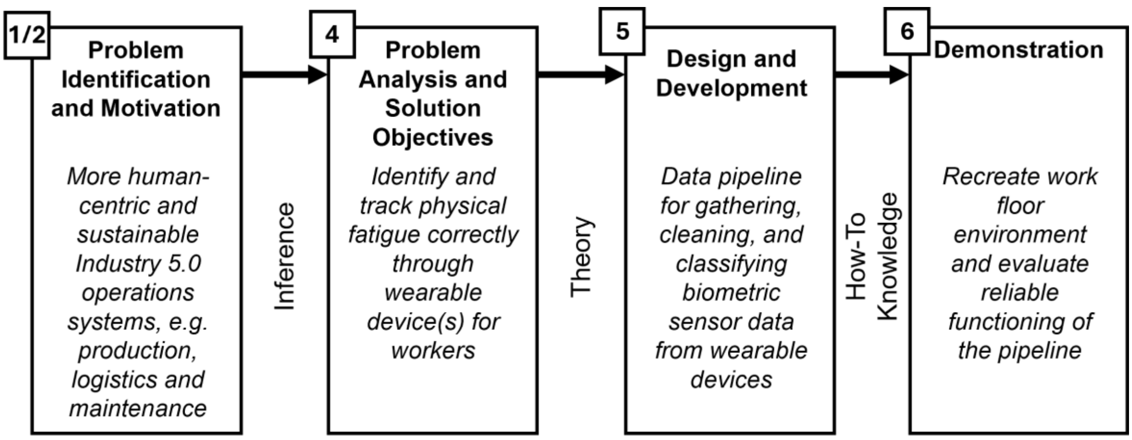

Figure 1: Design Science Research Methodology adapted for this research with corresponding sections in white squares

Activity 1: Problem Identification and Motivation. This activity is covered in Sections 1 and 2. Section 1 establishes the relevance of the research within the contexts of Industry 5.0 and BPM and formulates the research question concerning the development of an end-to-end data pipeline for real-time worker physical fatigue monitoring. Section 2 identifies the research gap by reviewing and evaluating the existing literature.

Activity 2: Problem Analysis and Solution Objectives Definition. This activity analyzes the problem and specifies the concrete objectives of the solution. Through a literature review, it identifies the physiological indicators that reliably reflect worker physical fatigue in industrial environments and selects a wearable device capable of measuring them. These, together with the problem framing, ground the concrete solution objectives that guide the design and development in Activity 3.

Activity 3: Design and Development. Building on the results of Activity 2, this activity develops an end-to-end data pipeline that gathers, cleans, aggregates, and interprets wearable sensor data into a meaningful real-time evaluation of worker physical fatigue.

Activity 4: Demonstration. This activity demonstrates the reliable functioning of the developed pipeline under simulated working conditions. A lab environment is set up to emulate an industrial work floor, where participants wear the selected device and perform physically demanding tasks while the pipeline ingests, processes, and analyzes the resulting streaming data.

As part of the DSRM communication activity, this report communicates the problem, its relevance, the artefact, and its demonstration and evaluation, while discussing its potential for decision-making in an Industry 5.0 setting and acknowledging its limitations in addressing the research question.

4. Problem Analysis and Solution Objectives

This section analyzes the problem and specifies the concrete objectives of the solution, according to the DSRM. Building on a literature review, it first identifies the physiological indicators that reliably reflect worker physical fatigue and can be measured non-intrusively, and then selects a wearable device capable of capturing them. Together, these ground the solution objectives. A full description of the literature reviews underpinning the selection of physiological indicators and the wearable device are available in the online project repository.[2]

4.1. Problem Analysis

Physiological Indicators of Physical Fatigue. The physiological indicators examined in this research, as identified by [7], are brain activity, muscular activity, cardiovascular activity, and movement dynamics.

These indicators are evaluated as potential physical fatigue metrics based on two criteria, which must both be satisfied. First, any monitoring pipeline is only as trustworthy as its underlying measurements. Therefore, the scientific validity of each indicator is evaluated. Second, since these indicators must be measurable in a non-intrusive way within an industrial setting, they are also assessed based on the intrusiveness of their required measurement techniques. This is important as established measurement techniques are often of limited practical use in industrial environments due to their intrusiveness [17].

Within the manufacturing context of this study, a measurement technique is specifically regarded as intrusive if it interferes with normal worker movements, task execution, or the use of personal protective equipment, or if it introduces discomfort to the operator.

The literature review concludes that cardiovascular (such as measured through ECG) and movement-based (measured through IMU) metrics are the only indicators that are both scientifically valid and measurable in a non-intrusive way. They are therefore retained as candidate indicators for the end-to-end data pipeline.

Wearable Device Selection. Even though some studies employ combinations of multiple wearable devices to capture a broader set of physiological signals, this research considered single-device solutions only. Minimizing the number of devices worn per operator reduces operator burden, lowers

[2] https://doi.org/10.5281/zenodo.20669891

costs, and enhances scalability, all of which are essential for deployment in Industry 5.0 environments. More concretely, devices that can perform ECG, PPG, or motion tracking, were considered.

This assessment rests on four criteria, each grounded in earlier parts of the study. First, data accuracy: the pipeline can only produce valid predictions from accurate sensor data. Second, real-time data and programmatic access (e.g., via SDKs), which follow from the research question and enable real-time monitoring and pipeline integration. Third, industrial suitability: whether the device operates reliably on the work floor, reflecting the study's Industry 5.0 framing. Fourth, cost-effectiveness, since scalable deployment is unrealistic if per-operator costs are excessive.

From all devices identified (see online repository), the device that best satisfied these criteria and is thus selected as the wearable device for this research, is the Movesense HR2 [18]. It provides high-quality single-lead ECG data, includes a 9-axis IMU for complementary motion data [19], and supports real-time access to raw signals. Importantly, the availability of a proprietary Python client facilitates the integration into the data pipeline enormously [20]. Additionally, worn as a chest strap, it is comfortable and non-intrusive for operators in industrial workflows, and cost-effective enough for potential large-scale adoption [19].

4.2. Solution Objectives

The solution objectives presented in this section are derived from three sources. The first is the overall problem identification and research gap, which frames the artefact as a real-time end-to-end pipeline for monitoring worker physical fatigue in industrial environments. The second is the set of selected physiological indicators. The third is the selected wearable device (Movesense HR2) together with its associated assessment criteria.

- O1: The artefact must support the real-time acquisition of physiological data, integrating with the Movesense Python client.
- O2: The artefact must reliably ingest and process streaming raw sensor data and be capable of handling multiple devices simultaneously.
- O3: The artefact must filter out erroneous sensor readings so that only physiologically plausible data informs downstream feature extraction and fatigue estimation.
- O4: The artefact must transform raw physiological signals into meaningful fatigue-related features. Thus, the system must support individual calibration through a baseline measurement, which serves both as reference points for feature computation and as the basis for person-specific classification thresholds.
- O5: The artefact must convert extracted features into fatigue classifications, completing the pipeline's end-to-end path from raw physiological signal to actionable monitoring output. The emphasis here lies on the design and implementation of a functional, real-time classification pipeline rather than on developing the most accurate predictive model.
- O6: The artefact must operate within an operationally acceptable time window. Millisecond-level responsiveness is not required, but fatigue-related insights must be produced quickly enough to support timely industrial decision-making.
- O7: The artefact must rely on a single, cost-effective wearable to minimize operator burden and support scalable industrial deployment.
- O8: The artefact must ensure robustness and fault tolerance in the data streaming process. Since the reliance on Bluetooth Low Energy (BLE) makes temporary connection losses plausible, the pipeline must handle short interruptions in connectivity.

5. Design and Development

The design of the end-to-end data pipeline outlined in this section is based on [9]. Figure 2 illustrates how this framework is translated into the concrete architecture of the artefact. It organizes the implementation into the active runtime pipeline gathering and processing the sensor data, the offline model development phase for physical fatigue evaluation, and the underlying data flow.

During live monitoring, the active runtime pipeline operates through four sequential code modules (Module 1-4). Module 1 handles the continuous acquisition of raw ECG and IMU data via BLE. This data is subsequently grouped into 5-minute chunks and passed to Module 2 for data cleaning and filtering. Next, Module 3 computes the participant's baseline and extracts all the HR, HRV and temporal features. Finally, Module 4 feeds the extracted features for each incoming observation into the deployed model to generate real-time fatigue classifications. The deployed model itself is constructed earlier, during the offline Model Development phase.

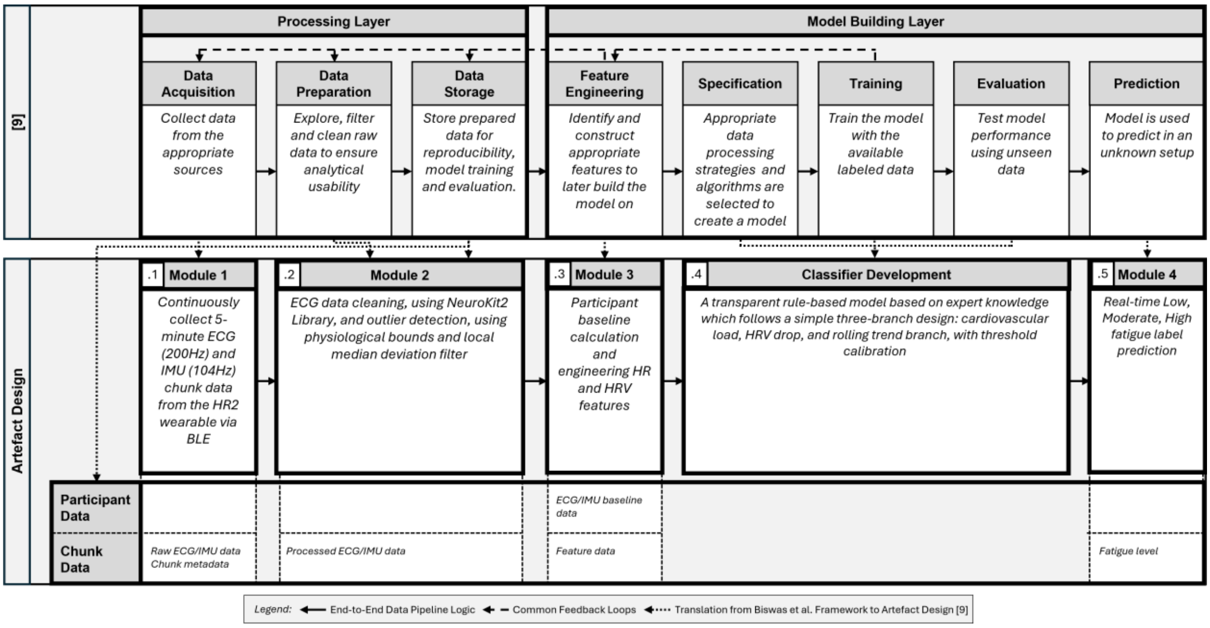

Figure 2: Data pipeline design from [9], and translation to artefact design for real-time fatigue prediction with corresponding subsections in white squares

The foundation for this entire process is the data flow, which is depicted at the bottom of Figure 2, which illustrates the progressive transformation of data from raw sensor output to predicted fatigue labels. It is organized into two parallel streams. The first is participant data, consisting of static, person-specific information such as the personal baseline, which is used in Module 3 to compute individualized features. The second is chunk data, the dynamic stream of raw and processed physiological measurements flowing continuously through Modules 1 to 4. Together, these two streams ensure that each fatigue prediction is both grounded in the individual's personal baseline and continuously updated as new sensor data arrives. The complete Python implementation of the pipeline is openly available in the online project repository.[3]

5.1. Module 1

Module 1, which corresponds to the data acquisition phase, consists of gathering the ECG and IMU data produced by the wearable device. Its realization was facilitated by the Python client provided by Movesense [21]. The system connects to the wearable device and streams sensor data directly into CSV format. The connection with the wearable is established via BLE using the Bleak library. Several extensions are added to the original Python client to meet the requirements of this research.

First, the storage of the raw sensor data in CSV format was enabled. The incoming raw data stream is grouped into 5-minute data chunks over which, in a later stage, the fatigue prediction is made. The chunk length must be carefully chosen. It needs to be long enough to ensure that the calculated features are stable, yet short enough to assess worker fatigue frequently. Five minutes satisfies both: it is the standard short-term heart-rate variability (HRV) recording window [22], providing enough time for

[3] https://doi.org/10.5281/zenodo.20669891

these variable features to achieve sufficient stability, while still producing predictions at a practically useful time interval to apply resource allocation interventions.

Second, the code was extended to allow multiple wearables to connect simultaneously, which is a key feature for deployment in industry.

Third, a reconnection logic was implemented so that the system automatically attempts to reconnect to the sensor in case of a disconnection.

5.2. Module 2

Module 2 focuses on data cleaning, initial processing, and storage, encompassing both the data preparation and data storage steps.

Data Preparation. The raw ECG data was cleaned using NeuroKit2. NeuroKit2 is an open-source Python package for processing physiological signals, such as ECG [23]. R-peaks were subsequently detected from the cleaned signal using NeuroKit2's peak detection algorithm. The detected R-peaks serve as the basis for all subsequent ECG feature derivation.

RR intervals (temporal distances between consecutive R-peaks) are derived from these detected R-peaks, and an instantaneous heart rate (HR) value is computed from each interval. Subsequently, they are subjected to an outlier detection procedure to ensure data quality. First, physiological bounds are applied to the RR intervals. Only values that correspond to realistic heart rates between 30 and 220 beats per minute are retained. RR intervals falling outside this range are considered physiologically implausible. In addition, a local median deviation filter is applied to detect irregular beats. This approach is adapted from the procedure described by [24], who demonstrated that removing RR intervals based on a deviation from the local mean is an effective approach. While their model utilized a three-beat window and the local mean, this study employs a larger 11-beat window and the local median. The goal of these modifications is to increase robustness against the clumping of outliers. If the relative deviation from this local median exceeds 20%, the interval is classified as an outlier. This threshold is consistent with limits evaluated by [24]. All identified outliers are subsequently removed before further analysis.

The processing script continuously monitors the data directory for newly generated raw ECG data chunks and automatically processes them in real-time when they appear.

Data Storage. To maintain a clear structure and facilitate downstream processing, the data is automatically organized into a hierarchical directory system. This structure reflects the pipeline's four sequential modules (1–4), which communicate via the filesystem. Each module monitors the output directory of the preceding stage and processes new files as they appear. At the highest level, all data from a single recording session of a single person is grouped within a session folder. This folder is further subdivided to categorize the raw data, processed data, and chunk metadata. Within these specific subdirectories, the incoming data stream is continuously saved into 5-minute CSV files. Additionally, in the session folder, a file with the participant's personal baseline is saved. Finally, a participant's data folder is automatically copied to a timestamped backup directory whenever a session is stopped.

The chunk metadata subdirectory contains a JSON file for each chunk. This file records its index, actual duration, closure reason (5-minute timer or disconnection), corrected start and end timestamps, and row counts. This metadata is used downstream to verify chunk completeness and to filter out incomplete chunks that are too short for analysis.

5.3. Module 3

Module 3 corresponds to the feature engineering step and involves processing each 5-minute data segment, representing a single observation, to extract features used to predict the operator's current fatigue level.

The first chunk of each session is used as an ECG baseline. It is recorded while the participant stands still. It establishes personal resting values for HR and HRV metrics. Because individuals differ substantially in cardiovascular characteristics, computing features relative to a personal baseline is essential for cross-participant generalization [25]. Three types of features are engineered:

(1) HR Features. Elevated cardiovascular load is a physiological indicator of physical fatigue [7], [26]. The mean HR is therefore computed for each chunk and used to derive the percentage of HR reserve (%HRR), the central cardiovascular feature, defined as:

$$\%HRR = \left(\frac{HR_{mean} - HR_{rest}}{HR_{max} - HR_{rest}}\right) \times 100 \tag{1}$$

where HR_{mean} is the mean HR of the chunk, HR_{rest} is the personal resting HR measured during the baseline chunk, and HR_{max} is the theoretical maximum HR based on [27]. This metric expresses the current cardiac workload as a fraction of the individual's available cardiovascular range. It takes inter-individual differences in resting HR and maximum capacity into account.

(2) HRV Features. Physical fatigue also leads to significant reductions in HRV [28]. Two time-domain HRV metrics, RMSSD and SDNN, are computed per chunk [22]. They are not used directly but feed the temporal features below.

(3) Temporal Features. Per-chunk metrics are placed in temporal context through two feature groups, since physical fatigue develops progressively over a session. Baseline-drop features express how far RMSSD and SDNN have fallen relative to the personal resting baseline (as a percentage). Rolling-trend features fit a linear slope for HR, ln(RMSSD) (a log transform that stabilizes the skewness of raw RMSSD [22], [29]), and SDNN over the most recent four chunks. The goal of these features is to signal within-session deterioration.

The wearable's IMU stream is acquired but not converted into features, for reasons set out in Section 5.4.

5.4. Classifier Development

The classifier development module contains the specification, configuration and evaluation of the rule-based physical fatigue classifier.

Classifier Specification. The classifier development module assigns each completed 5-minute chunk to one of three operational fatigue levels: Low, Moderate, or High. These levels enable clear and traceable operationalization of physical fatigue within the proposed dynamic resource allocation mechanism. A simple, transparent, rule-based classifier was developed as a first step to complete the pipeline.

The classifier follows a three-branch design, each branch capturing a distinct, physiologically grounded aspect of fatigue. First, the cardiovascular-load branch is driven by %HRR and reflects the current cardiac workload relative to the individual's capacity. Second, the HRV-drop branch flags meaningful reductions in RMSSD and SDNN relative to the personal resting baseline. Third, the rolling-trend branch fits linear slopes for HR, ln(RMSSD), and SDNN over the four most recent chunks to detect active within-session deterioration. The three branch outputs are combined through a small set of priority-ordered rules into a single classification visible in Table 2, which keeps it interpretable.

Table 2
Decision table for rule-based model

Rule	Condition	Fatigue classification	Rationale
R1: Extreme exertion	%HRR = High	High	Cardiovascular ceiling alone is sufficient at this intensity
R2: Systemic escalation	%HRR = Moderate AND HRV = crash OR trend = deteriorating trend	High	Moderate exertion confirmed by an HRV decrease (crash) or deteriorating trend warrants escalation
R3: Moderate exertion	%HRR = Moderate	Moderate	HR slightly elevated but no decrease in HRV
R4: Healthy fallback	None of the above matched	Low	Resting or light activity with no signs of HRV deterioration

High cardiovascular-load alone is sufficient to classify a chunk as High (Rule 1). The HRV-drop and or rolling-trend branches, by contrast, serve a confirming role and cannot independently trigger a High classification. These can only escalate a chunk already flagged as Moderate by the cardiovascular-load branch (Rule 2). Two considerations motivate this asymmetry. First, short-term HRV-based metrics are sensitive to minor technical artifacts [22] that can produce false positives. Second, the rolling-trend branch is only meaningful relative to the underlying load: a deteriorating trend at near-resting load does not indicate fatigue, whereas the same trend on top of a moderate %HRR does. Rule 2 captures exactly this case, in which cardiovascular-load alone cannot distinguish a worker operating sustainably from one in whom fatigue is starting to develop, either a clear crash in the HRV-drop branch or a deteriorating trend in the rolling-trend branch resolves this uncertainty by providing a second confirmation of fatigue. If neither the HRV-drop nor rolling-trend branch indicates a crash or deteriorating trend, but cardiovascular load is Moderate, the chunk is classified as Moderate (Rule 3). If cardiovascular load is not Moderate, the chunk defaults to Low (Rule 4). The calibrated thresholds for each branch are reported in the following section (Classifier Configuration).

%HRR serves as the decision logic's primary driver. This is derived from the Continuous Performance Limit framework (CPL), which identifies exertion above the CPL as fatigue-inducing [25]. In this framework, HR serves as a suitable cardiac parameter for this purpose. Because the HR used to determine the CPL is highly individualized, it should always be evaluated relative to a person's specific HRrest and HRmax. The %HRR metric, normalized by exactly these values, captures this requirement directly.

All three branches draw exclusively on cardiovascular signals. This reflects the transparency motivation behind the rule-based approach: a rule is only meaningful when the signal it acts on can be tied to physiologically interpretable thresholds. Heart rate and HRV allow this, as elevated cardiac load and within-session HRV reduction are well-documented and interpretable fatigue responses [7], [26], [28]. Movement, although also a recognized indicator of fatigue [7], [11], cannot be grounded in interpretable thresholds as readily. A rule defined over IMU-derived metrics would therefore be an empirical pattern fitted to the data rather than a physiologically interpretable mechanism, so movement data is excluded from the model and, accordingly, from feature engineering.

Classifier Configuration. A controlled experiment was conducted to collect a labeled dataset for the calibration of the thresholds used in the rule-based system (Table 2): the %HRR level (cardiovascular-load), the HRV-drop relative to the personal baseline, and the rolling-trend slope.

The labels were Borg Rating of Perceived Exertion (RPE) [30] scores collected each chunk and mapped to three fatigue categories: Low (6–10), Moderate (11–14), and High (15–20). Since physical fatigue cannot be observed directly, this short-term, self-reported measure is used as a practical proxy. It is widely adopted in ergonomics field studies and is ideal in the setting of this research because of its simplicity and unobtrusiveness alongside wearable sensing [8], [11]. The experimental task mimicked a common industrial warehouse activity. The 60 participants repeatedly lifted, carried, and placed weighted boxes between two platforms located 10 meters apart over a 50-minute period. The weight of the boxes was pre-calibrated per-person load and varied across the session to elicit a range of fatigue states.

Table 3

Final thresholds for the rule-based model

Decision branch	Metric	Threshold	Branch outcome
Cardiovascular-load	%HRR	≥ 20%	Moderate
		≥ 50%	High
HRV-drop	RMSSD Drop	≥ 20%	HRV Crash
	SDNN Drop	≥ 20%	
Rolling-trend (slopes)	HR	≥ 1	Deteriorating Trend (dual confirmation required)
	Ln(RMSSD)	≤ - 0.2	
	SDNN	≤ - 1	

Threshold values were selected via an automated procedure combining greedy sequential optimization and pairwise grid search, optimizing Cohen's Kappa. This metric was chosen for its robustness to class imbalance: it corrects for chance agreement, penalizing models that simply over-predict the majority class [31], and naturally extends to multi-class problems, making it well-suited for the three-tier fatigue classification.

The rule-based model was calibrated and evaluated using a cross-validation approach in which each participant is assigned exclusively to either the training or test set. This design eliminates both interpersonal and temporal data leakage and tests generalization to entirely unseen workers, reflecting the model's intended use as a fixed, shared decision policy. The training set comprised 48 participants (434 chunks, 80%), on which GroupKFold(5) cross-validation was applied to calibrate the thresholds. The remaining 12 participants (108 chunks, 20%) formed the held-out test set. A Leave-One-Session-Out stability check confirmed consistent thresholds across folds. The final calibrated thresholds are reported in Table 3.

Classifier Evaluation. After configuration, the rule-based classification approach was evaluated using the previously outlined cross-validation procedure. Table 4 presents the confusion matrix while Table 5 presents an overview of the performance metrics of the rule-based model.

Table 4

Confusion matrix rule-based model (Rows = true, Columns = predicted)

	Low	Moderate	High
Low	18	4	1
Moderate	25	19	8
High	6	12	15

The rule-based model achieved an overall bin accuracy of 48.1% and a Cohen's Kappa of 0.24, indicating fair agreement with the RPE ground truth. Performance was highest for the High fatigue class (F1 = 0.53, precision = 0.63), the most safety-critical category, while the Moderate class showed the weakest performance (F1 = 0.44), likely due to its intermediate and ambiguous nature. The main error pattern is the misclassification of Moderate chunks as Low (25/52), suggesting a tendency to underestimate moderate fatigue in the absence of HRV confirmation. Importantly, non-adjacent confusions between Low and High remain rare, indicating that severe misclassifications are largely avoided. The Kappa score suggests meaningful discriminative capability across fatigue states for unseen individuals, supporting the use of a transparent rule-based classifier when generalizing to an unseen population with limited training data. As such, this classifier can be considered as a first step to close the pipeline, that can be improved in future work.

Table 5

Overview of performance metrics for the rule-based model

Fatigue Bin	Precision	Recall	F1-Score	Support
Low	0.37	0.78	0.50	23
Moderate	0.54	0.37	0.44	52
High	0.63	0.45	0.53	33

5.5. Module 4

Once calibrated, Module 4 deploys the rule-based classifier to generate physical fatigue predictions for each incoming data chunk in real-time. The classification module operates in a live monitoring state: the moment the required features are calculated, the inference script reads them and instantaneously produces a Low, Moderate, or High fatigue classification. Each generated prediction is appended to a predictions log file within the session directory.

The real-time fatigue classifications generated by the pipeline can trigger resource allocation interventions when sustained high fatigue levels are detected. This fatigue information is subsequently incorporated into an optimization algorithm that determines the most appropriate intervention and passes it to the BPMS for execution and worker task list management. This is discussed in more detail in Section 7.1.

6. Demonstration

This section reports the demonstration of the proposed end-to-end data pipeline, covering DSRM Activity 4. By returning to the success criteria of Section 1 and the solution objectives of Section 4.1, it evaluates how well the developed artefact achieves what it set out to do.

To demonstrate the reliable functioning of the pipeline under simulated working conditions, a lecture room was adapted into a functional workspace in which six unseen participants performed the experimental task (as described in Section 5.4). Overall, the system performed as intended across all observed dimensions.

During these live monitoring runs, the pipeline successfully ingested and processed raw ECG data from up to four sensors simultaneously (O1 and O2). The connections remained stable throughout the session, and as a consequence the automated reconnection logic (O8) was not triggered, though it was extensively tested during development. While the fault tolerance behavior could therefore not be exercised live, the sustained connection does confirm the stability of the sensor-to-pipeline link in a controlled environment.

As each 5-minute data chunk was completed and saved to the hierarchical directory, the streaming preprocessing module was automatically triggered. This module effectively executed real-time signal quality assessments that filtered physiological outliers so that only plausible data informed downstream feature extraction (O3). The system then reliably generated all required features (O4). Various graphs were continuously plotted alongside the pipeline so that the researchers could visually inspect the extracted HR and HRV data at the end of every chunk.

Finally, the pipeline seamlessly transitioned to instantaneous inference. A fatigue classification appeared on the system's dashboard within seconds of the data chunk ingestion (O5 and O6). This confirms that the system operates within a practically acceptable time window. It generates insights fast enough to inform timely interventions on the work floor.

Operating entirely on the data stream of a single, cost-effective Movesense HR2 sensor, the artefact keeps hardware requirements minimal and operator burden low (O7).

Taken together, these observations confirm that the pipeline satisfies the functionality criterion within a controlled simulated work environment.

7. Discussion

This chapter reflects on the findings of this research. Section 7.1 discusses how the pipeline contributes to well-being-aware resource allocation, and Section 7.2 addresses its limitations and directions for future work.

7.1. Embedding Physical Fatigue in Dynamic Resource Allocation

As discussed in the introduction, current BPM Systems (BPMSs) lack the integration of real-time worker well-being information to support dynamic resource allocation decisions [4], [32]. To address this limitation, this research developed an end-to-end data pipeline capable of ingesting and processing streaming physiological data for real-time physical fatigue monitoring. As such, the pipeline constitutes a key building block for enabling human-centric resource allocation in Industry 5.0 environments.

More concretely, the developed pipeline constitutes the first of three components required to realize a fully well-being-aware resource allocation solution:

(1) A real-time well-being monitoring mechanism that generates intervention triggers when necessary. By translating raw physiological sensor data into categorical fatigue predictions (Low, Moderate, High), the pipeline provides the real-time signals required to initiate resource allocation interventions. A basic implementation of this trigger could rely on threshold-based rules, though because physiological data can fluctuate momentarily, the mechanism should act on sustained patterns

rather than isolated readings. For instance, an intervention could be triggered only when an operator's fatigue is classified as "High" for three consecutive 5-minute windows, rather than a single occurrence. In future work, additional well-being dimensions such as stress or concentration should be incorporated into these trigger decisions, enriching the inputs available to the broader resource allocation mechanism.

(2) A multi-objective optimization algorithm that determines the most suitable resource allocation intervention in response to a well-being trigger. This component is responsible for selecting the optimal intervention and correctly matching resources to tasks during reallocation, in accordance with the allocation objectives set by the mechanism [32].

(3) A BPMS integration layer that manages and executes the reallocation interventions while maintaining the worker task list [32].

While the end-to-end pipeline primarily serves as a monitoring and triggering mechanism, the real-time fatigue status it produces is also exposed as a resource attribute within the broader mechanism. This allows well-being information to directly inform resource-task matching, for example by preferentially selecting the least physically fatigued available worker during a reallocation decision.

As highlighted in the Digital Twins of Business Processes manifesto, ethical considerations should be an integral part of this line of research [5]. Future work should therefore uphold principles of human ethics and privacy by incorporating safeguards such as informed consent, (pseudo-)anonymization, and local and temporary data storage in practical implementations. Particular care should also be taken to prevent function creep, whereby fatigue data collected for safety purposes is repurposed for other ends, such as comparing workers against one another. This risk is mitigated in part by design, as the system exposes only the final reallocation intervention it proposes rather than the underlying physiological data, keeping individual worker data inaccessible for such comparisons. Beyond these technical safeguards, workers should retain autonomy over their participation, including the ability to pause or opt out of monitoring, so that the system is experienced as a means of support rather than surveillance.

7.2. Pipeline and Study Limitations

First, the classifier development stage was deliberately kept lightweight: while related work concentrates its effort on predictive modeling, this research prioritized closing the full pipeline end-to-end using a minimal, transparent rule-based classifier. The pipeline was explicitly architected to accommodate future modeling improvements: an organization can plug in a more sophisticated model, trained on a larger, representative dataset spanning its target workforce and tasks, without redesigning the surrounding acquisition, preparation, or feature-engineering stages. Future work will therefore focus on integrating a machine learning model to improve overall classification performance, with particular emphasis on maximizing recall for the High fatigue class, where misclassification carries the greatest safety consequences.

Furthermore, the pipeline was only demonstrated in a controlled lecture room rather than an operational industrial setting. The reliance on BLE is particularly relevant here, as its limited transmission range may not prove viable in all industrial environments. Wireless Body Area Networks (WBAN) could be an interesting alternative [33]. The stability and practical performance of the pipeline under such conditions therefore remains to be validated in future field studies.

Additionally, the system's reliance on a short, individualized baseline introduces a methodological constraint. The rule-based model depends on a brief resting measurement taken at the start of each session. If a worker calibrates while already fatigued, stressed, or heavily caffeinated, this baseline will not capture their true resting state and will skew all downstream computations. Future work should try to establish a more reliable baseline over a longer period of time and across multiple sessions.

Finally, the system assumes that deviations in cardiovascular activity are driven by physical exertion. Because it relies on cardiovascular indicators alone, it cannot natively distinguish an HRV reduction caused by physical load from one caused by acute cognitive or emotional stress. In future work, the pipeline could be extended to capture additional well-being dimensions beyond physical fatigue, such as psychological stress, and fuse them with the physical-fatigue signal into a single holistic worker well-being metric.

8. Conclusion

The transition toward Industry 5.0 demands a paradigm shift. It moves away from prioritizing pure manufacturing efficiency and toward a sustainable, human-centric approach that fosters worker well-being. To support this objective, this research aimed to address a notable gap in existing literature by answering how an end-to-end data pipeline can be developed for real-time physical fatigue monitoring. The resulting pipeline continuously ingests streaming data from the Movesense HR2 sensor, processes it, and produces fatigue classifications.

Returning to the success criteria of Section 1, the pipeline demonstrated stable real-time operation during evaluation: ingesting, processing, and analyzing concurrent streaming sensor data without failure. This was further validated through the successful evaluation of the concrete solution objectives during the pipeline demonstration.

Ultimately, the developed pipeline provides the real-time fatigue triggers needed to operationalize worker well-being monitoring in an industrial setting. By feeding these triggers into dynamic resource allocation mechanisms, BPMSs can intervene before fatigue-related disruptions occur, protecting both worker health and operational continuity. When supported by strict privacy safeguards, this system allows organizations to bring the "Healthy Operator" [3] concept to life, ensuring that Industry 5.0 innovations truly benefit the workers on the floor.

9. Acknowledgements

We sincerely thank the participants of our experiment, who volunteered their time and effort generously. Their willingness to contribute to the empirical part of this research is greatly appreciated. This study was funded by the Flemish Fund for Scientific Research (FWO) with grant number G042326N. The study complied with GDPR and was approved by the Social and Societal Ethics Committee (SMEC) of KU Leuven (approval no. G-2026-11179-R2).

10. References

[1] Maija. Breque, Lars. De Nul, and Athanasios. Petridis, *Industry 5.0 : towards a sustainable, human-centric and resilient European industry*. Publications Office of the European Union, 2021. doi: 10.2777/308407.

[2] B. Gladysz, T. Tran, D. Romero, T. van Erp, J. Abonyi, and T. Ruppert, “Current development on the Operator 4.0 and transition towards the Operator 5.0: A systematic literature review in light of Industry 5.0,” *J. Manuf. Syst.*, vol. 70, pp. 160–185, Oct. 2023, doi: 10.1016/j.jmsy.2023.07.008.

[3] D. Romero *et al.*, “Towards an Operator 4.0 Typology: A Human-Centric Perspective on the Fourth Industrial Revolution Technologies,” in *Proceedings of the 46th International Conference on Computers and Industrial Engineering (CIE46)*, Oct. 2016, pp. 1–11. [Online]. Available: https://www.researchgate.net/publication/309609488

[4] M. Wyffels, “Integrating Real-Time Worker Wellbeing into IoT-Enhanced Business Process Orchestration for Industry 5.0,” in *Intelligent Information Systems*, L. Pufahl, K. Rosenthal, S. España, and S. Nurcan, Eds., in Lecture Notes in Business Information Processing, vol. 557. Cham: Springer Nature Switzerland, 2025. doi: 10.1007/978-3-031-94590-8.

[5] F. Fornari *et al.*, “Digital Twins of Business Processes: A Research Manifesto,” 2024.

[6] Z. Sedighi Maman, M. A. Alamdar Yazdi, L. A. Cavuoto, and F. M. Megahed, “A data-driven approach to modeling physical fatigue in the workplace using wearable sensors,” *Appl. Ergon.*, vol. 65, pp. 515–529, Nov. 2017, doi: 10.1016/j.apergo.2017.02.001.

[7] E. Loizaga, A. T. Eyam, L. Bastida, and J. l. M. Lastra, “A Comprehensive Study of Human Factors, Sensory Principles, and Commercial Solutions for Future Human-Centered Working Operations in Industry 5.0,” *IEEE Access*, vol. 11, pp. 53806–53829, Jun. 2023, doi: 10.1109/ACCESS.2023.3280071.

[8] P. Mohapatra *et al.*, “Wearable network for multilevel physical fatigue prediction in manufacturing workers,” *PNAS Nexus*, vol. 3, no. 10, Oct. 2024, doi: 10.1093/pnasnexus/pgae421.

[9] S. Biswas, M. Wardat, and H. Rajan, “The art and practice of data science pipelines,” in *Proceedings of the 44th International Conference on Software Engineering*, New York, NY, USA: ACM, May 2022, pp. 2091–2103. doi: 10.1145/3510003.3510057.

[10] C. Janiesch *et al.*, “The Internet of Things Meets Business Process Management: A Manifesto,” *IEEE Syst. Man Cybern. Mag.*, vol. 6, no. 4, pp. 34–44, Oct. 2020, doi: 10.1109/msmc.2020.3003135.

[11] Z. Sedighi Maman, M. A. Alamdar Yazdi, L. A. Cavuoto, and F. M. Megahed, “A data-driven approach to modeling physical fatigue in the workplace using wearable sensors,” *Appl. Ergon.*, vol. 65, pp. 515–529, Nov. 2017, doi: 10.1016/j.apergo.2017.02.001.

[12] Z. Sedighi Maman, Y.-J. Chen, A. Baghdadi, S. Lombardo, L. A. Cavuoto, and F. M. Megahed, “A data analytic framework for physical fatigue management using wearable sensors,” *Expert Syst. Appl.*, vol. 155, p. 113405, Oct. 2020, doi: 10.1016/j.eswa.2020.113405.

[13] C. A. Morillo, H. Shi, J. Suárez-Pérez, and M. Demichela, “Early detection of physical fatigue in industry using wearable sensors and contextual modeling,” *Saf. Sci.*, vol. 194, p. 107041, Feb. 2026, doi: 10.1016/j.ssci.2025.107041.

[14] G. Goldschmidt and B. Matthews, “Formulating design research questions: A framework,” *Des. Stud.*, vol. 78, p. 101062, Jan. 2022, doi: 10.1016/j.destud.2021.101062.

[15] N. Hoang Thuan, A. Drechsler, and P. Antunes, “Construction of Design Science Research Questions,” *Communications of the Association for Information Systems*, vol. 44, no. 1, pp. 332–363, Mar. 2019, doi: 10.17705/1CAIS.04420.

[16] A. R. Hevner, S. T. March, J. Park, and S. Ram, “Design Science in Information Systems Research,” *MIS Quarterly*, vol. 28, no. 1, pp. 75–106, Mar. 2004, doi: 10.2307/25148625.

[17] L. Cavuoto and F. Megahed, “Understanding Fatigue: Implications for Worker Safety,” *Prof. Saf.*, vol. 62, no. 12, pp. 16–19, Dec. 2017, Accessed: Nov. 29, 2025. [Online]. Available: https://foundation.assp.org/docs/BPCav_1217z.pdf

[18] R. Martín Gómez, E. Allevard, H. Kamstra, J. Cotter, and P. Lamb, “Validity and Reliability of Movesense HR+ ECG Measurements for High-Intensity Running and Cycling,” *Sensors*, vol. 24, no. 17, p. 5713, Sep. 2024, doi: 10.3390/s24175713.

[19] Movesense Ltd., “Shop Movesense.” Accessed: May 17, 2026. [Online]. Available: https://www.movesense.com/shop/

[20] Movesense Ltd., “Recording Data with Movesense Sensors: Five Free Tools to Get You Started,” Movesense. Accessed: Feb. 21, 2026. [Online]. Available: https://www.movesense.com/news/2025/11/recording-data-with-movesense-sensors-five-free-tools-to-get-you-started/

[21] Movesense Ltd., “Movesense sensor data Python client (movesense_sensor_data.py),” Bitbucket. Accessed: Mar. 18, 2026. [Online]. Available: https://bitbucket.org/movesense/movesense-device-lib/src/master/samples/gatt_sensordata_app/python_client/movesense_sensor_data.py

[22] F. Shaffer and J. P. Ginsberg, “An Overview of Heart Rate Variability Metrics and Norms,” Sep. 28, 2017, *Frontiers Media S.A.* doi: 10.3389/fpubh.2017.00258.

[23] D. Makowski *et al.*, “NeuroKit2: A Python toolbox for neurophysiological signal processing,” *Behav. Res. Methods*, vol. 53, no. 4, pp. 1689–1696, Aug. 2021, doi: 10.3758/s13428-020-01516-y.

[24] M. Karlsson, R. Hörnsten, A. Rydberg, and U. Wiklund, “Automatic filtering of outliers in RR intervals before analysis of heart rate variability in Holter recordings: a comparison with carefully edited data,” *Biomed. Eng. Online*, vol. 11, no. 1, p. 2, Dec. 2012, doi: 10.1186/1475-925X-11-2.

[25] S. Sammito, B. Thielmann, A. Klussmann, A. Deußen, K. M. Braumann, and I. Böckelmann, “Guideline for the application of heart rate and heart rate variability in occupational medicine and occupational health science,” Dec. 01, 2024, *BioMed Central Ltd*. doi: 10.1186/s12995-024-00414-9.

[26] D. N. F. P. Damit, S. M. N. A. Senanayake, O. A. Malik, and N. J. T. Tuah, “Neuromuscular Fatigue Analysis of Soldiers Using DWT Based EMG and EEG Data Fusion During Load Carriage,” in *Intelligent Information and Database Systems Part II*, vol. 10192, N. T. Nguyen, S. Tojo, L. M. Nguyen, and B. Trawiński, Eds., in Lecture Notes in Computer Science, vol. 10192. , Cham: Springer International Publishing, 2017, pp. 602–612. doi: 10.1007/978-3-319-54430-4.

[27] H. Tanaka, K. D. Monahan, and D. R. Seals, “Age-predicted maximal heart rate revisited,” *J. Am. Coll. Cardiol.*, vol. 37, no. 1, pp. 153–156, Jan. 2001, doi: 10.1016/S0735-1097(00)01054-8.

[28] P.-C. Kao and D. J. Cornell, “Effects of Induced Physical Fatigue on Heart Rate Variability in Healthy Young Adults,” *Sensors*, vol. 25, no. 17, p. 5572, Sep. 2025, doi: 10.3390/s25175572.

[29] K. E. Speer, S. Semple, N. Naumovski, and A. J. McKune, “Measuring Heart Rate Variability Using Commercially Available Devices in Healthy Children: A Validity and Reliability Study,” *Eur. J. Investig. Health Psychol. Educ.*, vol. 10, no. 1, pp. 390–404, Jan. 2020, doi: 10.3390/ejihpe10010029.

[30] G. Borg, “The Borg RPE Scale,” in *Borg’s perceived exertion and pain scales*, Human Kinetics, 1998, ch. 5, pp. 29–38. Accessed: Apr. 26, 2026. [Online]. Available: ISBN 0-88011-623-4

[31] P. Czodrowski, “Count on kappa,” *J. Comput. Aided. Mol. Des.*, vol. 28, no. 11, pp. 1049–1055, 2014, doi: 10.1007/s10822-014-9759-6.

[32] M. Wyffels, I. Vanderfeesten, and E. Serral, *Functional Requirements for Wellbeing-Aware Resource Allocation in Industry 5.0 BPM Systems*, vol. 594. in Lecture Notes in Business Information Processing, vol. 594. Cham: Springer Nature Switzerland, 2026. doi: 10.1007/978-3-032-28274-3.

[33] M. Yadav, P. Shoran, E. Saxena, A. Bijalwan, and J. G. Bijalwan, “IoTs-Based Wearable Health Monitoring Through Wireless Body Area Networks,” in *Healthcare Industry Assessment: Analyzing Risks, Security, and Reliability*, vol. 11, Kumar P., Singh P., Dikawar M., and Garg D., Eds., Cham: Springer, 2024, pp. 231–254. doi: 10.1007/978-3-031-65434-3_10.