

Article

Context-Aware Alerting in Elderly Care Facilities: A Hybrid Framework Integrating LLM Reasoning with Rule-Based Logic

Nazmun Nahid ^{1,*} , Md Atiqur Rahman Ahad ^{2,*}  and Sozo Inoue ¹ 

¹ Graduate School of Life Science and Systems Engineering, Kyushu Institute of Technology, 2-4 Hibikino, Wakamatsu Ward, Kitakyushu 808-0135, Japan; sozo@brain.kyutech.ac.jp

² Department of Engineering & Computing, School of Architecture Computing and Engineering, University of East London, London E16 2RD, UK

* Correspondence: nahid.nazmun389@mail.kyutech.jp (N.N.); mahad@uel.ac.uk (M.A.R.A.)

Abstract

The rising demand for elderly care amid ongoing nursing shortages has highlighted the limitations of conventional alert systems, which frequently generate excessive alerts and contribute to alarm fatigue. The objective of this study is to develop a hybrid, context-aware nurse alerting framework for long-term care (LTC) facilities that minimizes redundant alarms, reduces alarm fatigue, and enhances patient safety and caregiving balance during multi-person care scenarios such as mealtimes. To do so, we aimed to intelligently suppress, delay, and validate alerts by integrating rule-based logic with Large Language Model (LLM)-driven semantic reasoning. We conducted an experimental study in a real-world LTC environment involving 28 elderly residents (6 high, 8 medium, and 14 low care levels) and four nurses across three rooms over seven days. The proposed system utilizes video-derived skeletal motion, care-level annotations, and dynamic nurse–elderly proximity for decision making. Statistical analyses were performed using F1 score, accuracy, false positive rate (FPR), and false negative rate (FNR) to evaluate performance improvements. Compared to the baseline where all nurses were notified (100% alarm load), the proposed method reduced average alarm load to 27.5%, achieving a 72.5% reduction, with suppression rates reaching 100% in some rooms for some nurses. Performance metrics further validate the system’s effectiveness: the macro F1 score improved from 0.18 (baseline) to 0.97, while accuracy rose from 0.21 (baseline) to 0.98. Compared to the baseline error rates (FPR 0.20, FNR 0.79), the proposed method achieved drastically lower values (FPR 0.005, FNR 0.023). Across both spatial (room-level) and temporal (day-level) validations, the proposed approach consistently outperformed baseline and purely rule-based methods. These findings demonstrate that the proposed approach effectively minimizes false alarms while maintaining strong operational efficiency. By integrating rule-based mechanisms with LLM-based contextual reasoning, the framework significantly enhances alert accuracy, mitigates alarm fatigue, and promotes safer, more sustainable, and human-centered care practices, making it suitable for practical deployment within real-world long-term care environments.

Keywords: alarm fatigue; large language model; fall detection; context-aware systems; nurse alerting; long-term care; nurse care



Academic Editor: Toshiyo Tamura

Received: 11 September 2025

Revised: 9 October 2025

Accepted: 16 October 2025

Published: 24 October 2025

Citation: Nahid, N.; Ahad, M.A.R.; Inoue, S. Context-Aware Alerting in Elderly Care Facilities: A Hybrid Framework Integrating LLM Reasoning with Rule-Based Logic. *Sensors* **2025**, *25*, 6560. <https://doi.org/10.3390/s25216560>

Copyright: © 2025 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the Creative Commons Attribution (CC BY) license

(<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

With aging, the human body undergoes inevitable physiological and cognitive changes. Elderly individuals experience mobility impairment, chronic illnesses, and cognitive de-

cline. This makes them more prone to injuries and falls. Among the causes of injury, fall incidents are prominent as a major public health concern. Each year, nearly 36 million older adults experience a fall, with 3 million requiring emergency care [1]. Falls often result in severe outcomes, including fractures (64%), post-fall trauma or fear (44%), and hospitalizations (32%) [2,3]. These episodes can also lead to long-term consequences, such as reduced autonomy, immobility, dependence, depression, and an overall decline in functional capacity. Despite increasing awareness, the effectiveness of traditional fall prevention strategies—such as scheduled check-ins or individualized physical training—has been limited [4,5]. For healthcare personnel, the need to continuously monitor vulnerable elderly patients adds substantial physical and cognitive workload, especially in long-term care (LTC) environments where nurse-to-patient ratios are often stretched. Consequently, the development of proactive and intelligent fall detection and alert systems has garnered significant interest across both clinical and scientific domains.

With the advancement of information and communication technologies, a variety of smart healthcare solutions have emerged to support fall detection and early warning systems [6–9]. While such systems showed promising results [8,10,11], the existing alert systems often interrupt clinical workflow [12,13], and may compromise safety outcomes [14–16]. The reason for safety concern is that the alerts often overwhelm caregivers. This phenomenon, widely known as ‘alarm fatigue’, is now recognized as a critical threat to patient safety and caregiver well-being [17,18]. Alarm fatigue occurs when healthcare workers are exposed to a large volume of alerts—most of which are either irrelevant, false, or not actionable [18]. These interruptions fragment caregiver workflow, elevate stress levels, and introduce a high cognitive burden that diminishes the quality of attention paid to actual patient care [19,20]. As a result, caregivers may become desensitized, delay response, or entirely ignore alerts, even when they are valid [21,22].

The main reason behind this situation lies in the pattern that the traditional monitoring systems follows. Most of the systems to ensure patients’ safety tend to send alerts to all the nurses. This leads to an over- and unnecessary alert situation. Moreover, these alert systems lack consideration of the care context to convey the priority of the situation. As health professionals face increasing demands for multitasking and constant responsiveness, the cognitive load caused by alarm over-saturation has become a crisis. False alarms and poorly prioritized alerts fragment clinical workflows, reduce trust in alert systems, and contribute to staff burnout and techno-stress. Many nurses report difficulty discerning alert priority, locating the source of alarms, and responding effectively due to environmental noise or non-clinical disruptions [23]. The problems that arise due to poor clinical alarm management can be categorized mainly in three categories [24]:

- **False Alarms:** False alarms can be both positive and negative. A false positive alarm occurs when an alert is triggered even though there is no actual clinical need or emergency. A false negative alarm refers to a situation where a clinically important event occurs, but the system fails to detect or alert it.
- **Unreachable Alarms:** An unreachable alarm refers to an alert that is sent out but fails to reach the appropriate caregiver in time. This can be due to routing inefficiencies where an alert is sent to a nurse at such a time when the time to reach the elderly patient is not enough.
- **Unattainable Alarms:** An unattainable alarm is triggered and sent to a caregiver who is occupied or too far away to respond at the moment the event occurs.

These factors severely limit the clinical utility of existing alert systems and highlight a critical need for intelligent, context-aware frameworks that issue only the most relevant alerts to the right caregiver, at the right time, to reduce alarm fatigue but ensure care quality, as shown in Figure 1. In the figure, the baseline method is the traditional notify-all nurses

approach; the rule-based method includes both rule-based suppression and delay and rule-based suppression, delay, and heuristic validation; and the proposed hybrid model integrates rule-based logic with LLM-driven semantic reasoning.

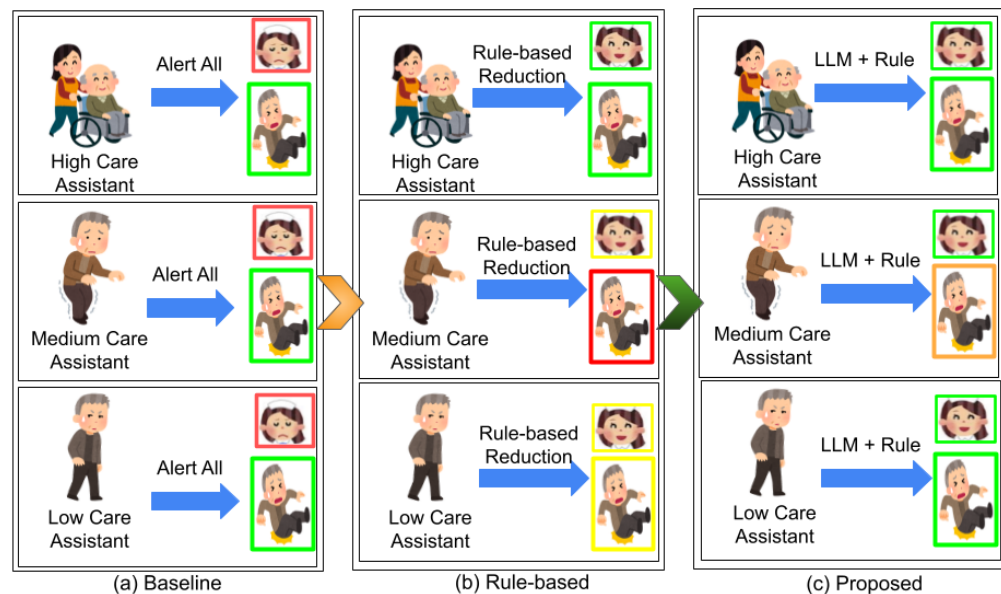


Figure 1. Illustration demonstrating the need for a context-aware alarm system. Panel (a) represents the baseline scenario where all nurses receive alerts, (b) depicts reduced alerts through rule-based logic, and (c) shows further refinement when LLM-based validation is integrated with rule-based methods. For nurses, ‘red’ indicates overload from frequent alarms, ‘yellow’ reflects a moderate reduction in alarm load, and ‘green’ signifies minimal alerts received only when necessary. For elderly residents, ‘red’ denotes a high fall risk, ‘yellow’ indicates a reduced fall risk, ‘orange’ suggests an almost negligible fall risk, and ‘green’ represents no fall risk. Care assistance levels are defined as ‘high’ for residents with care levels 4–5, ‘medium’ for care level 3, and ‘low’ for care levels 1–2.

To address these challenges, this paper presents a feasibility study toward a context-aware nurse alerting system for elderly monitoring in long-term care settings. The proposed system leverages video-derived skeletal data to detect elderly motion, estimate nurse proximity, and intelligently route alerts based on real-time activity analysis and contextual prioritization. Instead of relying solely on rule-based binary thresholds, the system incorporates motion features, estimated distance between the resident and available caregivers, and behavioral markers to dynamically assess risk with the help of an LLM. This approach enables a multi-layered alerting strategy: when a potentially risky activity—such as a resident attempting to stand up unassisted—is detected, an alert is issued to the most suitable nurse based on spatial–temporal metrics; then, a forecast of the resident’s future skeletal motion is generated to assess the likelihood of instability or fall and also analyze nurse responsiveness prompting a second alert if needed. In this study, we addressed three key questions to develop a context-aware nursing alert system:

- Research question 1 (R1): How well does a hybrid alert framework consisting of rule-based logic and LLM reasoning minimize alert fatigue without compromising care quality?

Contribution: This study introduces a hybrid alerting framework that combines rule-based logic with LLM reasoning to deliver context-aware alerts, aiming to reduce alarm fatigue without compromising care quality. Compared to the baseline scenario where all nurses were alerted (100% alarm load), the proposed framework lowered the average load to 27.5% but achieved an accuracy of 0.98 and a macro F1 score of 0.97, significantly surpassing the baseline (accuracy 0.21, macro F1 score 0.18).

- Research question 2 (R2): To what extent can the inclusion of clinically interpretable information help minimize false alarms in long-term care settings?
Contribution: Our method leverages clinically interpretable parameters like resident care condition, nurse–resident distance, and mobility to generate alerts, minimizing false alarms. Compared to the baseline (FPR 0.20, FNR 0.79), the method achieves markedly lower error rates (FPR 0.005, FNR 0.023).
- Research question 3 (R3): Is a flexible, context-driven alert validation strategy more effective than static rule-based thresholds in addressing the unique needs of elderly care?
Contribution: We introduce an LLM-based validation module that incorporates spatial, temporal, and clinical context, instead of rule-based thresholding. Our model achieved an accuracy of 0.98 and a macro F1 score of 0.97, outperforming the rule-based approach (accuracy 0.78, macro F1 score 0.79).

This study advances current elderly monitoring research by introducing a hybrid, context-aware nurse alerting framework that combines rule-based suppression and delay logic with LLM-driven contextual reasoning. Unlike existing systems that rely solely on fixed thresholds or opaque learning models, our approach balances interpretability and adaptability, enabling reliable alert decisions in real-world LTC settings. By integrating spatial (nurse–resident distance), temporal (motion trajectory), and clinical (care level and mobility) information, the framework ensures alerts are both clinically relevant and context-sensitive. Quantitative results demonstrate clear improvements over the baseline and rule-based methods. In alignment with the study’s objectives, the system (1) reduces alarm fatigue through intelligent suppression, (2) enhances alert precision via LLM-based validation, (3) maintains interpretability, and (4) supports balanced, patient-centered caregiving. Overall, this work contributes a scalable, human-centered AI approach that improves safety and efficiency in LTC environments.

2. Related Works

Alarm management in healthcare has been a persistent challenge in long-term care (LTC) settings. According to Imhoff et al. [25], the quality of medical device alarms is unsatisfactory, affecting quality of care and patient safety. Numerous strategies have been explored to improve the clinical validity, relevance, and responsiveness of alarm systems. The proposed approaches can be broadly categorized into three categories. In this section, we will elaborately discuss these categories.

2.1. Alarm Suppression Algorithms

Alarm suppression algorithms aim to reduce false or clinically non-actionable alerts by enhancing the quality of alarm-generating logic. Several machine learning and statistical approaches have been proposed to address this. Several machine learning methods have been applied for alarm classification and suppression. Random forests and support vector machines have been used to distinguish between true and false alarms, with sensitivity control through voting thresholds in complex physiological time series data [26]. Neural networks, though promising, often face challenges such as long training times, instability in learning for dynamic patient states, and interpretability issues [27]. Fuzzy logic systems are applied in cases where alarm thresholds are not precisely defined, enabling nuanced alert generation [20]. Statistical techniques like autoregressive models, Kalman filters, and spectral regression have been used to model physiological time series for anomaly detection [28]. Time-series-based segmentation and median filtering are used to smooth signals and reduce transient spikes that may trigger false alarms [29]. Autoregressive modeling and Kalman filtering have also been effective in early studies for trend detection

and change-point analysis [30]. Feature engineering and dimensionality reduction, such as principal component analysis (PCA) and graphical models, have been adopted to identify the most informative features from high-dimensional physiological data [31]. A lack of clinical interpretability and high false negative risk due to aggressive suppression make these approaches inefficient to deploy in LTC facilities. In summary, while suppression algorithms reduce non-actionable alarms, their limited interpretability and risk of missing critical events make us look for a better context-aware approach [32].

2.2. Middleware Solutions and Notification Delay-Based Strategies

To reduce alarm overload, middleware solutions introduce a delay between signal detection and alert transmission. These systems act as intermediaries, filtering alerts before sending them to clinicians' devices. Notification delays using middleware or internal device logic filter out transient anomalies, allowing only persistent or critical events to reach staff. Studies have shown that introducing such delays can reduce false alarms [33]. Research Article [34] implemented reasoning-based middleware to defer alerts until conditions warranting escalation are met. Research Article [35] reported alarm-reduction delay strategies and alarm parameter reconfiguration. While effective in reducing alert sending frequency, this approach risks delaying true alarms, especially in rapidly deteriorating conditions. Moreover, middleware configurations require complex rule sets and maintenance, increasing system overhead. Hence, while delay-based strategies offer a practical way to reduce alarm volume, they must be complemented with real-time contextual analysis to avoid critical oversights. Overall, middleware and delay-based strategies can ease alarm overload but risk unsafe delays and require complex maintenance [36]. These highlight the demand for more adaptive, context-aware alternatives.

2.3. Alarm Validation Approaches

Alarm validation refers to determining the clinical or technical validity of an alarm before notifying caregivers. This is especially crucial in elderly care settings where patient variability is high, and false positives are common. Statistical process control (SPC) methods, such as CUSUM and Shewhart charts, monitor deviations from individualized baselines but often ignore autocorrelation or patient-specific thresholds [37]. Knowledge-based systems attempt to integrate clinical rules and expert knowledge for alarm interpretation [38], yet these systems are often brittle and require constant rule updates. Bayesian networks and fuzzy logic systems have shown potential for probabilistic reasoning and soft classification of alarm relevance, particularly in uncertain or imprecise clinical environments [39]. Bayesian networks allow for probabilistic validation of alarms using dependencies between physiological parameters [40]. Multivariate logic-based systems use combinations of signals and thresholds to identify alarm-relevant scenarios with greater specificity of alerts [41]. However, these methods often overlook behavioral and environmental context, such as nurse–resident distance and mobility status [20,42–44], patient-specific care dependency and response needs, and forecasted behavior (e.g., a slow attempt to stand up that may lead to a fall). In short, alarm validation approaches add clinical meaning to alarms. However, their lack of behavioral and environmental context restricts their effectiveness in LTC, motivating richer, context-driven methods.

2.4. Challenges and Motivation

Alarm management remains a persistent challenge in long-term care (LTC) facilities, where elderly individuals require nuanced, personalized attention and caregivers operate under high-workload conditions. Despite considerable efforts to address alarm fatigue in clinical settings, especially in long-term care (LTC) environments, most existing alarm

management strategies remain insufficient. We identify three primary limitations across current alarm systems.

- Clinically non-interpretable suppression often leads to false negatives in LTC, where subtle transitions carry significant risk. Our proposed approach addresses this by discarding unattainable alerts by rules based on patient condition, nurse-to-patient distance, and nurse and patient speed.
- Rigid or middleware-based delays can endanger timely responses and impose technical overhead. We replace these with dynamic, context-driven delay mechanisms that balance responsiveness and workload.
- Context-unaware validation fails to capture care priorities, ignoring workload and resident-specific needs. Hence, we integrate spatial, temporal, and clinical context to ensure alerts are both technically valid and clinically actionable.

Together, these improvements form our hybrid framework integrating rule-based and LLM reasoning, designed to reduce false alarms, enhance clinical relevance, and ensure more technically valid and clinically actionable responsive care in LTC facilities.

3. Problem Statement

Alarm systems in LTC often generate excessive, non-specific alerts, leading to false positives, missed events, and caregiver fatigue. We propose a hybrid Context-Aware Alert System that integrates rule-based suppression, context-sensitive delays, and LLM-driven validation using clinical, spatial, and temporal features. By assigning alerts to the most suitable caregiver based on proximity, response time, and activity context, the system reduces alarm overload and ensures timely, efficient, and relevant interventions. Let the set of alarms generated over a monitoring period be denoted by $\mathcal{A} = \{a_1, a_2, \dots, a_N\}$. Each alarm a_i is represented as a tuple:

$$a_i = (t_i, x_i, u_i, c_i, r_i) \quad (1)$$

where t_i is the temporal context, x_i is the spatial metric, u_i is the urgency measure, c_i is the clinical context, and r_i is the resource context. The system comprises three essential decision functions used to optimize alert handling in long-term care environments:

Suppression Function ($S(a_i)$):

The suppression function $S(a_i) \in \{0, 1\}$ determines whether an alarm a_i is unreachable or non-actionable and filters it out before further processing. Suppression is based on interpretable rules derived from three dimensions:

- Distance Constraint: alarms are suppressed if the nurse is beyond a maximum reachable distance θ_x .
- Urgency Constraint: alarms triggered by very slow or clinically insignificant motion, i.e., trying to stand up slowly or moving to a side ($u_i < \theta_u$) are suppressed.
- Care Level Constraint: if the care level c_i of the elderly resident indicates low dependency or low fall risk, alerts are suppressed unless overridden by urgent context.

This eliminates unreachable or non-actionable alerts early in the pipeline and improves both efficiency and clinical interpretability.

Delay Function ($D(a_i)$):

The delay function $D(a_i) \in \mathbb{R}_{\geq 0}$ computes a personalized delay before the alarm is sent to caregivers. It balances system responsiveness against caregiver burden by dynamically calculating delay using:

- Urgency: faster patient movement (e.g., abrupt standing up) implies shorter delays.

- Nurse Workload: overloaded nurses should receive alerts slightly later to reduce interruptions if the event is non-urgent.
- Resident Stability: higher risk of fall or instability reduces permissible delay.

This formulation optimizes the trade-off between timely intervention and alert volume. Low-urgency alerts are delayed to avoid unnecessary interruption, while high-urgency ones are dispatched promptly.

Validation Function ($V(a_i)$):

The validation function $V(a_i) \in \{0, 1\}$ determines whether an alarm is clinically meaningful and should be escalated after suppression and delay. This validation is performed by computing a dynamic priority score using an LLM-based reasoning module. Due to the logical reasoning ability we opted for LLM-based validation. The target of this function is to validate only those alerts that are both technically triggered and clinically warranted. This step ensures alert escalation aligns with patient care plans and the operational context, reducing alert fatigue and enhancing caregiver focus.

Optimization Goal:

Definition of the set of effective alerts:

$$\mathcal{A}^* = \{a_i \in \mathcal{A} \mid S(a_i) = 1 \wedge t_i + D(a_i) \leq T_{\max} \wedge V(a_i) = 1\} \quad (2)$$

Our system aims to minimize false negatives ($\sum_{a_i \in \mathcal{A}} \mathbb{1}[y_i = 1 \wedge V(a_i) = 0]$), minimize false positives ($\sum_{a_i \in \mathcal{A}} \mathbb{1}[y_i = 0 \wedge V(a_i) = 1]$), minimize alert volume ($|\mathcal{A}^*|$), and maximize timely coverage (coverage(\mathcal{A}^*) within T_{\max}). This hybrid, context-aware nurse alerting framework leverages structured suppression logic, personalized delay, and intelligent LLM-driven validation to issue fewer but more meaningful alerts. By integrating spatial, temporal, and clinical information into the decision-making pipeline, the system improves both the precision of alarm delivery and the timeliness of caregiver intervention—addressing the pressing challenge of alarm fatigue while supporting safer elderly care.

4. Context-Aware Alert System

In this section, we provide a detailed explanation of the anomaly detection, suppression, and validation process, along with the notification procedure. We have used YOLOv7 custom object detection-based person identification [45] and tracking [46] to identify and track the elderly and nurses. We used YOLOv7-pose [45] to extract 2D keypoints from the video. With the help of Facebook prophet [47], anomalous activities are detected for each frame. The alert set is initialized for all the abnormal activities detected in a frame. Each alarm tuple contains information such as time, motion intensity, urgency level, care level, and nurse–resident distance. An empty alert set is initialized to store only the effective alerts. Each alarm is first checked to remove false or unnecessary ones. Alerts with low urgency, high stability, or low care priority are discarded to reduce false alarms and nurse fatigue. For the remaining alerts, a delay time is calculated using factors such as urgency, care level, and distance between nurse and resident. Alerts that would take too long to process (beyond the acceptable acknowledgment time) are discarded. Each valid alert is then evaluated using an LLM, which determines its clinical priority and recommends the most suitable nurse based on the situation. Alerts with low priority scores are ignored, ensuring that only meaningful alerts are sent. The selected nurse receives the alerts, and a timer starts to monitor acknowledgment. If no response is received within the set time limit, the alert is automatically escalated to the next available nurse. After acknowledgment, the system updates nurse workload and adds the handled alert to the effective alert set. The final output is a list of all processed and escalated alerts. The complete algorithm

is presented in Algorithm 1, with key component notations summarized in Table 1. The overall system workflow is illustrated in Figure 2.

Algorithm 1 Context-Aware Alert Decision Framework

```

1: Input: Set of alarms  $\mathcal{A} = \{a_1, a_2, \dots, a_N\}$ , where  $a_i = (t_i, x_i, u_i, c_i, r_i)$ 
2: Initialize: Effective alert set  $\mathcal{A}^* \leftarrow \emptyset$ 
3: for all  $a_i \in \mathcal{A}$  do
4:   // Suppression Check
5:   if  $x_i > \theta_x$  or  $u_i < \theta_u$  or  $care(c_i) \in \{\text{low-priority}\}$  then
6:     Discard  $a_i$  and log suppression
7:     continue
8:   end if
9:   // Context-Aware Delay Assignment
10:  Compute delay:  $D(a_i) = \alpha_1(1 - u_i) + \alpha_2 r_i + \alpha_3 c_i$ 
11:  if  $t_i + D(a_i) > T_{ack}$  then
12:    Discard  $a_i$  due to unacceptable delay
13:    continue
14:  end if
15:  // LLM-Based Validation and Nurse Assignment
16:  Generate prompt:  $Prompt(a_i) = T(t_i, x_i, u_i, c_i, r_i)$ 
17:  Query LLM to get:  $priority(a_i)$  and assigned nurse  $n^*$ 
18:  Compute dynamic threshold  $\tau(c_i, x_i, r_i, t_i)$ 
19:  if  $priority(a_i) < \tau(c_i, x_i, r_i, t_i)$  then
20:    Discard  $a_i$  (Not clinically significant)
21:    continue
22:  end if
23:  // Notification Dispatch
24:  Send alert  $a_i$  to nurse  $n^*$ 
25:  Start acknowledgment timer  $T_{ack}$ 
26:  if No acknowledgment within  $T_{ack}$  then
27:    Escalate to next-best nurse  $n'$ 
28:  end if
29:  Update workload and availability for assigned nurse
30:  Add  $a_i$  to effective alert set:  $\mathcal{A}^* \leftarrow \mathcal{A}^* \cup \{a_i\}$ 
31: end for
32: Output: Escalated and handled alert set  $\mathcal{A}^*$ 

```

Table 1. Key components of the context-aware alert system.

Components	Mathematical Expressions
Time in an instant	t_i
Nurse-to-patient distance	x_i
Urgency context	u_i
Clinical context	c_i
Nurse availability	r_i
Alarm instance	a_i
Maximum reachable distance threshold	θ_x
Minimum urgency threshold	θ_u
Suppression component	$S(a_i)$
Delay component	$D(a_i)$
Dynamic threshold computed by the LLM	τ
Nurse alert acknowledgment time window	T_{ack}

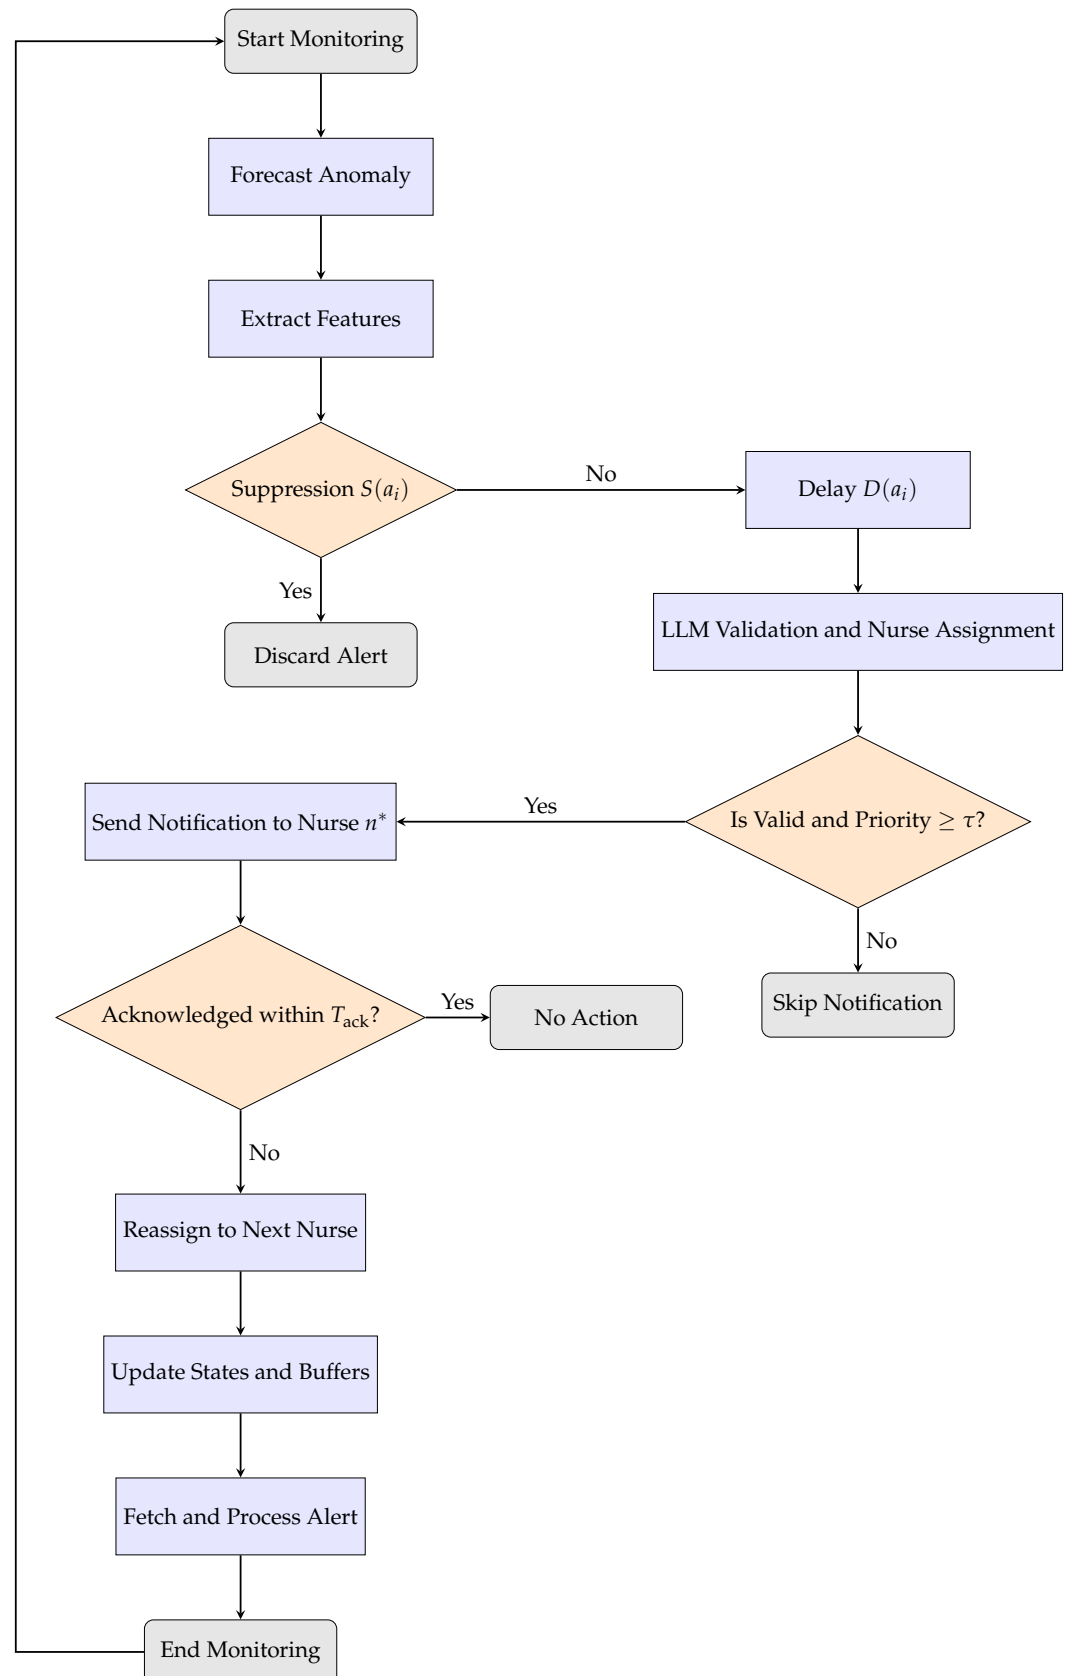


Figure 2. Flowchart of context-aware alerting system.

4.1. Anomalous Event Detection

Because of asynchrony in real-time monitoring, there is often a noticeable delay between observing an event and responding to it. Instead of depending only on the current

motion path, it becomes crucial to predict the person's next movement and the activity they are likely to perform. Since abnormal behaviors can result from physical or environmental factors, this study addresses the issue by forecasting elderly motion through joint trajectory prediction and activity state estimation, allowing caregivers to act proactively. In our context, abnormal behavior mainly refers to pre-fall conditions such as sudden attempts to stand up or unstable standing posture.

Let $q(t)$ be the 3D position vector of a key joint (e.g., hip, knee, shoulder) at time t . The input sequence to the forecasting model is:

$$X(t) = [q(t - N), q(t - N + 1), \dots, q(t - 1)] \in \mathbb{R}^{N \times 3} \quad (3)$$

We use an LSTM model \mathcal{F} to predict the joint's future position:

$$\hat{q}(t + \Delta t) = \mathcal{F}(X(t)) \quad (4)$$

After the future trajectory is forecasted, this data is fed to Facebook prophet to identify abnormalities before it even occurs. Abnormal behavior is identified by examining the magnitude or pattern of these residuals. Large residuals or patterns that deviate significantly from the typical behavior of residuals are considered abnormal behavior in the data. Let $y(t)$ be the observed value at time t , and $\hat{y}(t)$ be the predicted value by Prophet. The residual at time t is given by:

$$\text{Residual}(t) = y(t) - \hat{y}(t) \quad (5)$$

Mathematically, the abnormality score is $A(t)$ for each time point t based on the residuals. It is calculated by:

$$A(t) = |\text{Residual}(t)| \quad (6)$$

A threshold τ_A is used to classify behavior as abnormal:

$$\text{Abnormal if } A(t) > \tau_A \quad (7)$$

4.2. Feature Extraction and Representation Module

Following the detection or forecasting of an anomalous event, the system enters the Feature Extraction and Representation Module, which plays a critical role in constructing a high-dimensional, context-rich feature representation for each alarm event. This representation is essential for the subsequent suppression, delay assignment, and validation modules that rely on clinical, spatial, and temporal information.

Given a detected or forecasted anomaly at time t , the system extracts multi-contextual features from real-time monitoring data, motion forecasts, care records, and nurse schedules. Each alarm a_i is encoded as a structured tuple, $a_i = (t_i, x_i, u_i, c_i, r_i)$, as shown in Equation (1), where t_i is the timestamp of the detected or predicted anomalous event, x_i represents the spatial context, u_i encodes the urgency context, c_i captures the clinical context of the elderly, and r_i includes resource-specific attributes. Each component of the tuple is derived using specialized extraction procedures:

(1) Temporal Context (t_i):

The temporal feature t_i contains the timestamp of the detected or predicted anomalous event and nurse speed. The YOLOv7 pose model detects human figures in each video frame and estimates the 2D positions of keypoints (such as head, shoulders, elbows, hips, knees, and ankles). These positions are usually expressed in pixel coordinates within the frame. To compute the nurse speed of motion, we analyze how the positions of these keypoints

change over time (i.e., across frames). Since video is a time-sequential medium, changes in keypoint positions can reflect movement. By measuring how much a person's keypoints have shifted from one frame to the next and knowing the time difference between frames (based on the video frame rate), we can estimate how fast the body is moving.

Let N be the total number of keypoints per person, $P_i^{(t)} = (x_i^{(t)}, y_i^{(t)})$ be 2D coordinates of keypoint i at frame t , f be the frame rate in frames per second (fps), $\Delta t = \frac{1}{f}$ be the time between two frames in seconds, and s be the scale factor in meters per pixel. For each keypoint i , compute the Euclidean distance between frames t and $t + 1$:

$$\Delta P_i^{(t)} = \|P_i^{(t+1)} - P_i^{(t)}\| = \sqrt{(x_i^{(t+1)} - x_i^{(t)})^2 + (y_i^{(t+1)} - y_i^{(t)})^2} \quad (8)$$

Convert pixel displacement to meters using the scale factor:

$$P_{i,\text{meters}}^{(t)} = s \cdot P_i^{(t)} \quad (9)$$

Calculate speed in meters per second:

$$v_i^{(t)} = \frac{P_{i,\text{meters}}^{(t)}}{\Delta t} = s \cdot f \cdot P_i^{(t)} \quad (10)$$

YOLOv7-pose provides 2D keypoint coordinates for each detected person per video frame. To calculate the full-body motion speed, we analyze how these keypoints move over time, convert pixel displacement into meters, and compute speed using frame rate.

(2) Spatial Context (x_i):

The spatial feature x_i is extracted using the Euclidean distance between the elderly and nurses. The nurse to elderly distance is calculated based on the distance between the centroid of their respective bounding boxes. A bounding box in a two-dimensional image or video frame is typically represented by its top-left and bottom-right coordinates of a person or object. The centroid is the average of the bounding box's horizontal and vertical boundaries. This gives the center point of the rectangle, which is especially useful in applications such as object detection, tracking, and activity recognition. By computing this central point, we can monitor object movement across frames or analyze motion trajectories more effectively. Let a bounding box be defined by its top-left and bottom-right corner coordinates:

$$(x_{\min}, y_{\min}) \quad \text{and} \quad (x_{\max}, y_{\max}) \quad (11)$$

The centroid (x_c, y_c) of a bounding box is calculated as:

$$x_c = \frac{x_{\min} + x_{\max}}{2}, \quad y_c = \frac{y_{\min} + y_{\max}}{2} \quad (12)$$

Let Box A and Box B have centroids $C_A = (x_A, y_A)$ and $C_B = (x_B, y_B)$, respectively. Then, the Euclidean distance x_i between the two centroids is given by:

$$x_i = \sqrt{(x_B - x_A)^2 + (y_B - y_A)^2} \quad (13)$$

Substituting centroid expressions for Box A and Box B:

$$x_A = \frac{x_{\min}^A + x_{\max}^A}{2}, \quad y_A = \frac{y_{\min}^A + y_{\max}^A}{2} \quad (14)$$

$$x_B = \frac{x_{\min}^B + x_{\max}^B}{2}, \quad y_B = \frac{y_{\min}^B + y_{\max}^B}{2} \quad (15)$$

Therefore, the Euclidean distance between two bounding boxes is:

$$x_i = \sqrt{\left(\frac{x_{\min}^B + x_{\max}^B}{2} - \frac{x_{\min}^A + x_{\max}^A}{2}\right)^2 + \left(\frac{y_{\min}^B + y_{\max}^B}{2} - \frac{y_{\min}^A + y_{\max}^A}{2}\right)^2} \quad (16)$$

(3) Urgency Context (u_i):

The urgency score quantifies how quickly an event may require a nurse's intervention. It is derived from motion characteristics and historical incident risk. Let v_i be the velocity of the elderly, a_i be the acceleration (rate of change of movement), s_i be the semantic activity category, and f_i be the number of recent abnormal behavior flags. The urgency score for the elderly resident i is defined as:

$$\text{urgency}(u_i) = \gamma_1 \cdot \sigma(v_i) + \gamma_2 \cdot \sigma(a_i) + \gamma_3 \cdot I(s_i \in R) + \gamma_4 \cdot f_i \quad (17)$$

where $I(s_i \in R)$ is an indicator function that returns 1 if the activity s_i belongs to the set of risky activities R (e.g., trying to stand up), and 0 otherwise, and $\gamma_j \in [0, 1]$ are weighting parameters such that $\sum_{j=1}^4 \gamma_j = 1$. Currently, the weights are manually set based on domain knowledge. Frequent anomaly occurrence is considered the most important, $\gamma_4 = 0.35$, followed by the activity category, $\gamma_3 = 0.25$. The elderly resident's velocity and acceleration are given equal importance, $\gamma_1 = 0.2$ and $\gamma_2 = 0.2$.

(4) Clinical Context (c_i):

The care level (CL) is extracted from resident records and encoded on a scale of 1 to 5 where lower care levels indicate higher autonomy and higher care levels suggest increased dependency. We also added scores of assistance requirements for five key activities: meals, bath, excretion, movement, and dress-up. The scores of assistance levels (ALs) are given as follows: full assistance = 5, partial assistance = 4, reminder = 3, check = 2, and independence = 1. We define a clinical priority mapping:

$$\text{priority}(c_i) = \begin{cases} \text{high,} & \text{if } (\text{CL} \in \{3 > 2 > 1\}) \wedge (\text{AL} \in \{4 > 3 > 2 > 1\}) \wedge (\text{CL and AL} \neq 1) \\ \text{mid,} & \text{if } \text{CL} = 1 \wedge \text{AL} = 1 \\ \text{low,} & \text{if } \text{CL} \in \{5 > 4\} \wedge \text{AL} = 5 \end{cases} \quad (18)$$

(5) Resource Context (r_i):

The resource context is computed by tracking real-time or logged data on nurse availability. If $nurses_{active}$ is the set of on-duty nurses and $task(n)$ indicates whether a nurse is busy, we define nurse workload as:

$$r_i = \frac{1}{|nurses_{active}|} \sum_{n \in nurses_{active}} task(n) \quad (19)$$

where $task(n) = 1$ if the nurse is currently engaged, and 0 otherwise.

(6) Feature Vector Construction:

These extracted components are then combined into a joint feature vector:

$$\mathbf{f}_i = [t_i, x_i, u_i, c_i, r_i] \in \mathbb{R}^d \quad (20)$$

This unified vector \mathbf{f}_i is used for all downstream decisions, including suppression, delay adjustment, and alert validation. This module transforms raw physiological and contextual data into a structured, multi-dimensional representation that captures the clinical,

environmental, and operational significance of each alarm. It forms the backbone of the decision-making pipeline in the proposed context-aware alert system.

4.3. Rule-Based Suppression Module

The Rule-Based Suppression Module is the first decision gate in the hybrid nurse alerting framework. Its purpose is to eliminate clinically non-actionable or physically unattainable alarms before they propagate through the system—thereby preventing unnecessary caregiver distraction, reducing system load, and minimizing the risk of alarm fatigue. This component operates based on interpretable, clinician-informed decision rules, rather than black-box logic, to ensure transparency and trust in safety-critical environments like LTC facilities. This module suppresses alarms that fall into the following categories:

- **Unreachable Alarms:** This condition arises when a caregiver cannot reach the elderly within a clinically acceptable time window. In our study, we define this as the nurse failing to arrive within $\Delta t = 1$ s, a value chosen empirically based on the nurse movement speed in our case study. It is adjustable depending on the requirements of the deployment setup. Here, $\Delta t = t_{\text{elderly-stand-up-time}} - t_{\text{nurse-to-elderly-reaching-time}}$, and the acceptable reaching time is expressed as $t_a = t_{\text{elderly-stand-up-time}} - \Delta t$.
- **Unattainable Alarms:** Alarms that are technically triggered but represent conditions that do not require further intervention. For our considered scenario, this condition is applied for fully dependent elderly residents. As there is always one attendant nurse beside them, there is no necessity for the alarm to be generated.
- **Non-Urgent Context:** Alarms where resident movement or posture change is within the normal behavioral profile based on their care level and history.

Thus, the suppression function is:

$$S(a_i) = \begin{cases} 0, & \text{if } x_i > \theta_x \quad (\text{Too far}) \\ 0, & \text{if } u_i < \theta_u \quad (\text{Not urgent}) \\ 0, & \text{if } \text{care}(c_i) \in \{\text{low priority}\} \quad (\text{fully dependent}) \\ 1, & \text{otherwise} \end{cases} \quad (21)$$

where x_i is the nurse-to-resident distance, calculated by $t_{\text{nurse-to-elderly-reaching-time}} \times v_{\text{nurse-speed}}$, u_i is the urgency score derived from movement speed, posture risk, etc., θ_x is the maximum reachable distance threshold ($t_a \times v_{\text{nurse-speed}}$), θ_u is the minimum urgency threshold (e.g., motion entropy above baseline), and $\text{care}(c_i)$ is the mapping of care context priority. If $S(a_i) = 0$, the alarm is discarded and not passed to subsequent modules. If $S(a_i) = 1$, the alarm proceeds to the Delay Function and Validation Function. All suppressed alarms are stored in a shadow log for retrospective analysis and refinement.

4.4. Context-Aware Delay Assignment Module

The Context-Aware Delay Assignment Module is designed to mitigate alert fatigue in long-term care (LTC) environments by strategically deferring non-urgent alerts without compromising clinical safety. Unlike static delay mechanisms that treat all alerts equally, this module assigns dynamic, personalized delays based on real-time evaluation of the care environment. In LTC settings, not every alert requires an immediate response. For example, a stable patient slowly rising from a chair may not demand the same urgency as a resident with a fall risk attempting to walk unassisted. Therefore, timing becomes a key variable: this module ensures urgent alerts remain interruptive, while lower-priority alarms are queued for deferred attention based on workload, patient stability, and urgency context. The delay function $D(a_i)$ determines how long to defer its notification:

$$D(a_i) = \alpha_1 \cdot (1 - \text{urgency}(u_i)) + \alpha_2 \cdot \text{resource}(r_i) + \alpha_3 \cdot \text{care}(c_i) \quad (22)$$

where urgency (u_i) $\in [0, 1]$ is a normalized urgency score (higher = more critical), resource (r_i) $\in [0, 1]$ indicates current nurse workload, and care(c_i) $\in [0, 1]$ represents elderly physical and clinical condition-based priority (for mid priority, we choose a longer delay compared to high priority condition). Tunable coefficients $\alpha_1, \alpha_2, \alpha_3 \in [0, 1]$ are $\alpha_1 + \alpha_2 + \alpha_3 = 1$. For our case we chose these parameters from domain knowledge: $\alpha_1 = 0.5$, $\alpha_2 = 0.25$, and $\alpha_3 = 0.25$. Thus, the delay time before alert delivery is $D(a_i) \in \mathbb{R}_{\geq 0}$. For high urgency, lower delay is assigned; for high nurse workload, higher delay (if safe to defer) is assigned; and for mid priority, high delay (lower intervention need) is assigned. This module ensures that care teams are not overwhelmed by alarms while safeguarding that critical interventions are never delayed beyond clinically acceptable thresholds.

4.5. LLM-Driven Validation and Prioritization Module

The LLM-Driven Validation and Prioritization Module serves as the final and most adaptive stage of the alert decision process. Its main purpose is to ensure that only clinically relevant and contextually important alerts are sent to the most appropriate nurse, thereby reducing false positives, filtering out unnecessary information, and maintaining confidence in the alert system. Unlike conventional validation methods based on static thresholds, this module employs LLM reasoning to analyze complex, multi-dimensional contextual data—including spatial, temporal, and clinical information—to assess both the urgency and necessity of an alert in real time. When multiple abnormal events occur simultaneously, the LLM allocates nurses intelligently based on their current availability, preventing overload on any single caregiver. The module constructs structured and interpretable prompts using these real-time contextual features to support transparent and explainable decision making. The prompt generation function is defined as:

$$\text{Prompt}(a_i) = T(t_i, x_i, u_i, c_i, r_i) \quad (23)$$

where $T(\cdot)$ maps structured input into a fixed natural language template. Prompt structure and mapping is shown in Table 1. This prompt is designed to be concise yet context-rich, tuned for long-term care (LTC) environments. GPT-4 has been used in our study. The LLM gives outputs if this alarm should be handled as high priority and which nurse should be given priority to send the alert. To support numerical scheduling and nurse selection, the LLM output is mapped to a scalar priority score:

$$\text{priority}(a_i) = f_{\text{LLM}}(c_i, x_i, u_i, r_i) \quad (24)$$

where f_{LLM} is a prompted or learned mapping that produces: $\text{priority}(a_i) \in [0, 1]$. The alert is escalated only if:

$$\text{priority}(a_i) \geq \tau(c_i, x_i, r_i, t_i) \quad (25)$$

where $\tau(\cdot)$ is a dynamic threshold computed by the LLM based on clinical severity, nurse-to-patient distance, nurse availability or workload, time to stand up, and time to reach the elderly resident.

4.6. Notification Sending and Preparation for Next Alert

Once a validated alert a_i has been assigned to a specific nurse n^* , the system proceeds to dispatch the notification. The notification is sent to the selected nurse n^* using a multi-modal delivery mechanism optimized for responsiveness and clarity. Once the notification is dispatched, the system waits for acknowledgment within a time window T_{ack} . If no

acknowledgment is received, the system triggers escalation to reassign to the next-ranked nurse n' . Upon completion of alert handling (successful response or fallback), the system resets internal buffers and performs the following:

- Releases the nurse from temporary lock if acknowledgment was successful.
- Updates the workload index ω_j and availability status $a_{n,j}$.
- Clears the current alert context and fetches the next candidate alert a_{i+1} from the queue.
- Invokes the full validation–assignment–dispatch loop for a_{i+1} .

The alert analysis module processes each video frame to detect and handle multiple abnormal behaviors simultaneously. When several irregular events occur in the same frame, the system uses the LLM to assign the most suitable nurse for each situation instead of defaulting to the nearest or a single nurse. This ensures that alerts are distributed intelligently, preventing any individual nurse from being overloaded. As a result, each alert is prioritized and handled in a balanced and clinically appropriate order, maintaining both patient safety and an efficient caregiver workflow.

5. Experiment and Analysis

In this section, we outline the data collection process, detail the validation of our proposed method, and discuss its limitations, contributions, and future directions.

5.1. Data Collection

The dataset was collected from Global Care, a long-term care (LTC) facility in Japan that specializes in dementia care and operates an in-house video monitoring system. In collaboration with the facility, we accessed recorded footage capturing natural lunchtime activities of residents over seven consecutive days across three dining areas (Rooms 1, 2, and 3) for the case study. Each session lasted approximately 30–60 min, corresponding to the regular lunch period. A total of 28 elderly participants aged between 64 and 95 years were included in the study, along with four nurses who provided care support during the sessions. The elderly participants were included if they were residents of the long-term care facility, capable of performing daily activities with or without assistance, and provided informed consent (or consent was obtained from their legal guardians). Participants were excluded if they did not wish to take part in the study. The residents were categorized into three care groups based on their care level: low (1–2), medium (3), and high (4–5) [48]. Low-level residents were generally independent, requiring minimal assistance or occasional reminders for meals, hygiene, or dressing, and were able to maintain an upright posture with minor age-related stooping. Medium-level residents required partial assistance for daily activities and movement, could stand or transfer with assistive devices, and often exhibited a forward-leaning or hunched posture due to reduced core strength. High-level residents were completely dependent on caregivers, mostly immobile, and showed severe postural instability without a consistent upright posture. Among the participants 6 were in the high-care group, 8 were in the medium-care group, and 14 were in the low-care group. Seating arrangements and interactions were entirely managed by facility staff, with no intervention from researchers, ensuring that all recorded behaviors reflected natural, real-world care routines.

Importantly, no instructions were provided to either the participants or staff, ensuring that all recorded behaviors occurred naturally. For example, in Room 2, Participants 7 and 8 were seated in varying positions across sessions, reflecting routine adjustments made by caregivers based on real-time care needs. These positional changes were carefully noted during the annotation process to maintain labeling accuracy. The spatial layout of each

room is depicted in Figure 3, with nurses assigned to individuals requiring higher care levels seated nearby. Their positions were also documented accordingly.




Figure 3. Room layout including fisheye camera placement position and sitting placement of residents. In each room, two nurses dedicatedly attended to the high care assistance required by patients by staying beside them. We also marked their positions to understand the validation later as these nurses are supposed to receive no alarm or are only alerted when there is no available option. Except for these nurses, other nurses roamed around the room to conduct various tasks. Additionally, in Room 2, residents 7 and 8 were seated in varying positions during the data collection duration so nurse 2 and 3 also changed their positions.

5.2. Validation Strategy

To rigorously assess the effectiveness of our proposed context-aware nurse alerting system, we employed a dual validation approach—spatially across rooms and temporally across days. This design allowed us to evaluate both the consistency and robustness of the system in a real-world long-term care setting. We performed a comprehensive ablation study using four alerting strategies to dissect the contribution of each component in the pipeline. The baseline method reflects the conventional ‘notify-all’ approach, where all nurses receive alerts whenever an abnormal event is detected, without regard for contextual relevance or nurse workload. This approach, while exhaustive, tends to overwhelm caregivers and generate unnecessary cognitive load. The R(S+D) strategy builds upon this by introducing rule-based suppression and delay—where alerts are blocked if contextually insignificant and delayed if premature—mimicking more practical judgment found in human care. Further extending this, the R(S+D+V) method adds a heuristic validation layer, checking the appropriateness of alerts before dispatch to enhance decision precision. Our Proposed Method combines the deterministic logic of suppression and delay with adaptive, semantic validation powered by a large language model (LLM). The rules guide initial alert filtering, while the LLM performs high-level contextual reasoning to confirm whether an alert should be sent. This hybrid approach retains interpretability while enabling flexibility in complex and ambiguous situations.

For quantitative comparison, we computed a set of performance metrics including accuracy, precision, recall, F1 score, false positives (FPs), and false negatives (FNs). These metrics were analyzed both globally and per nurse to capture variations in alert distribution, responsiveness, and error rates. Accuracy and F1 score were used to assess the overall correctness and balance between sensitivity and specificity, while FPs and FNs provided insight into the model's reliability in suppressing false alarms and minimizing missed critical events. Tests were also conducted to see temporal (day-based) and spatial (room-based) variation.

All data were tabulated and visualized using bar charts and confusion matrices to illustrate comparative performance among the four methods. Descriptive statistics (mean and standard deviation) were used to summarize results across days and rooms, ensuring comparability between temporal and spatial dimensions. Ground truth annotations were manually created from video recordings, identifying moments when facility staff intervened by alerting nurses in response to genuine abnormal events. This human-labeled data served as the benchmark for both forecasting accuracy and alert validation, ensuring our comparisons reflected actual clinical response behavior.

The analyses were conducted using Python (version 3.9) with the NumPy, SciPy, PyCM, and scikit-learn libraries for statistical computation and metric evaluation, and Matplotlib/Seaborn for data visualization. This ensured the reproducibility, transparency, and rigorous quantitative assessment of the alert system's performance.

5.3. Results

Figure 4 presents a comparative analysis of nurse-specific alarm loads across different alert strategies, highlighting the effectiveness of our proposed context-aware method. Compared to the baseline, which alerts all nurses indiscriminately, our model significantly reduces alert load by tailoring notifications based on situational context. In the overall case, the proposed method reduced alarm loads to 71% for Nurse 1, 71% for Nurse 2, 78% for Nurse 3, and 70% for Nurse 4. A room-wise breakdown shows consistent reductions. In Room 1, alarm loads decreased to 62% for Nurse 1, 68% for Nurse 2, 88% for Nurse 3, and, remarkably, 100% for Nurse 4. In Room 2, alert loads dropped to 59% for Nurse 1, 98% for Nurse 2, 96% for Nurse 3, and 71% for Nurse 4. In Room 3, alarm loads reduced to 100% for Nurse 1, 97% for Nurse 2, 67% for Nurse 3, and 59% for Nurse 4. These results demonstrate the model's adaptability across spatial contexts and care configurations. On average, the proposed system achieved a 72.5% reduction in overall alert volume. When analyzing this by individual nurse, the average reduction in alerts for Nurse 1 is 73.67%, for Nurse 2 is 87.67%, for Nurse 3 is 83.67%, and for Nurse 4 is 76.67%. Notably, the model achieved complete alert suppression for Nurse 1 in Room 3 and Nurse 4 in Room 1, showing its capacity to eliminate unnecessary alerts when nurses are not contextually suited to respond. These findings underscore the model's ability to reduce redundant notifications and distribute alerts intelligently, ultimately enhancing caregiver efficiency without compromising responsiveness.

As illustrated in Table 2, and Figures 5 and 6, the proposed context-aware alerting method consistently outperforms all other evaluated strategies across multiple evaluation metrics. The proposed model achieves a macro average F1 score of 0.97 and macro average accuracy of 0.98, underscoring its effectiveness in minimizing both false alarms and missed alerts. In terms of forecasting errors, the number of false positives (FPs) is zero for all nurses except Nurse 2, and false negatives (FNs) are zero for all nurses except Nurse 3, reflecting high precision and recall for most responders.

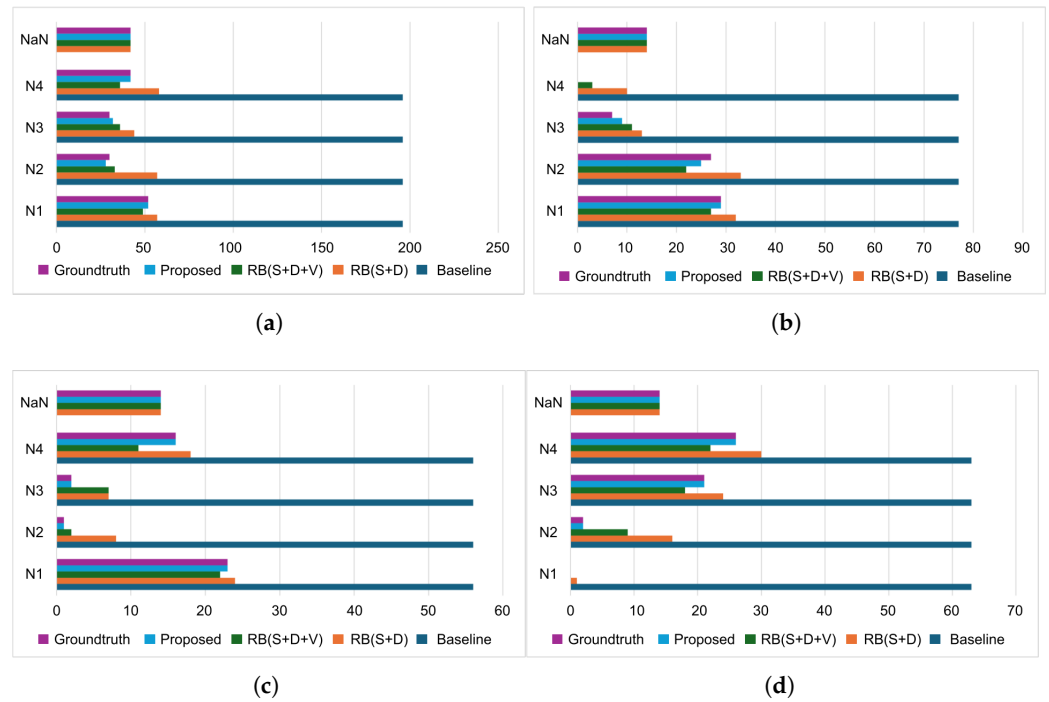


Figure 4. Nurse alarm load comparison for different alarm types. Here, (a) All = All over alarm load for each nurses for different methods, (b) R1 = Room 1 alarm load for each nurses for different methods, (c) R2 = Room 2 alarm load for each nurses for different methods and (d) R3 = Room 3 alarm load for each nurses for different methods. R(S+D) is rule-based suppression and delay; and R(S+D+V) adds validation layer with R(S+D). NaN = Occurrence when no alert is sent.

Table 2. Overall performance comparison. Here, the class represents alarms sent to each nurse. Nan represents the situation when no alarm is sent. R(S+D) is rule-based suppression and delay; and R(S+D+V) adds a validation layer with R(S+D).

Alarm Type	Class	Precision	Recall	F1 Score
Baseline	1	0.25	0.25	0.25
	2	0.18	0.25	0.21
	3	0.16	0.25	0.19
	4	0.24	0.25	0.25
	NaN	0.00	0.00	0.00
R(S+D)	1	0.91	0.80	0.85
	2	0.53	0.64	0.58
	3	0.68	0.73	0.71
	4	0.72	0.67	0.69
	NaN	1.00	1.00	1.00
R(S+D+V)	1	1.00	0.86	0.93
	2	0.60	0.70	0.65
	3	0.56	0.80	0.66
	4	0.83	0.60	0.70
	NaN	1.00	1.00	1.00
Proposed	1	1.00	1.00	1.00
	2	1.00	0.87	0.93
	3	0.87	1.00	0.93
	4	1.00	1.00	1.00
	NaN	1.00	1.00	1.00

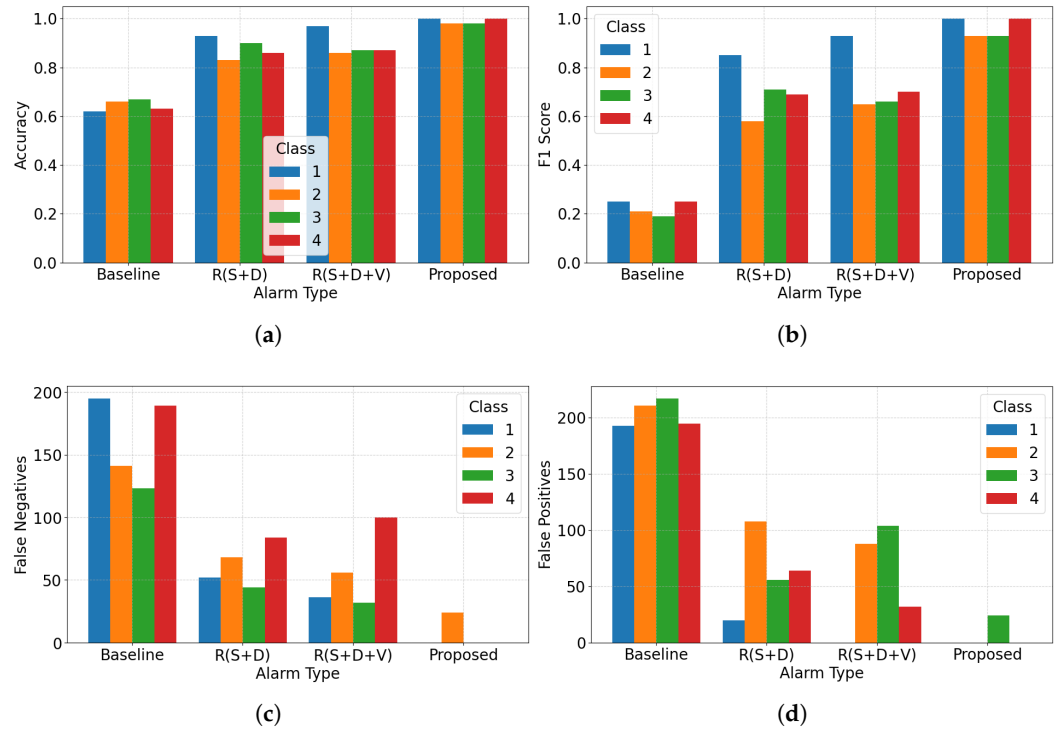


Figure 5. Overall performance comparison based on accuracy, F1 score, false negatives, and false positives for each nurse for different methods. R(S+D) is rule-based suppression and delay; and R(S+D+V) adds a validation layer with R(S+D). Here, (a) Accuracy; (b) F1 Score; (c) False Negatives and (d) False Positives.

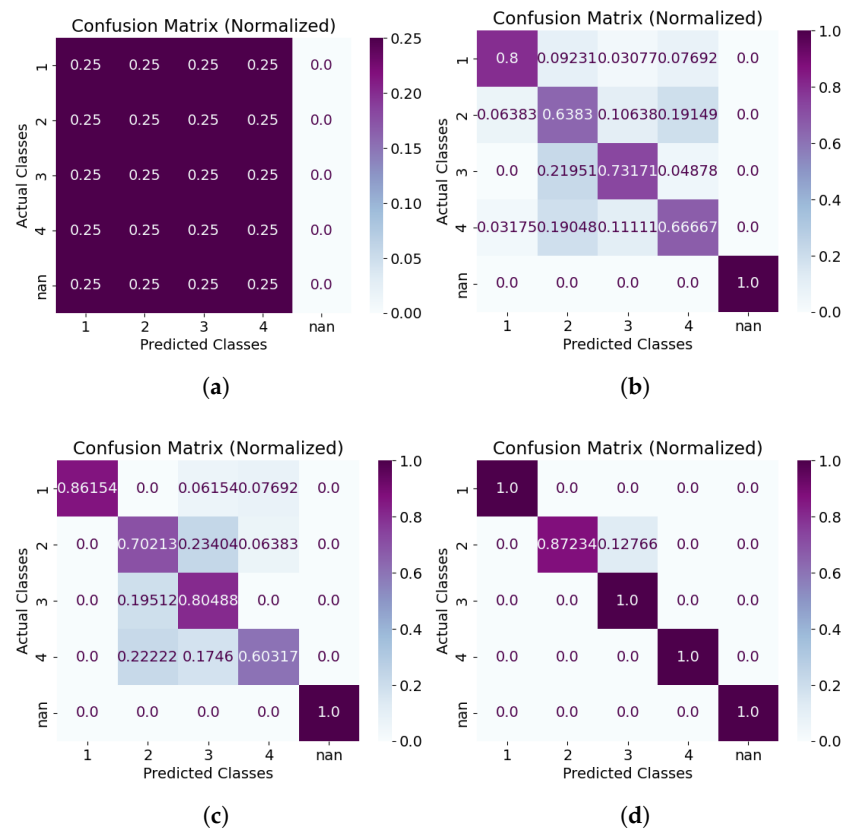


Figure 6. Confusion matrix for overall performance comparison for each nurse for different methods. R(S+D) is rule-based suppression and delay; and R(S+D+V) adds a validation layer with R(S+D). Here, (a) Baseline; (b) R(S+D); (c) R(S+D+V) and (d) proposed.

From the room-based validation in Figure 7, the proposed context-aware alerting system consistently outperforms all other strategies across multiple metrics. It achieves a macro average F1 score of 0.94 and accuracy of 0.98, effectively minimizing both false alarms and missed alerts. Accuracy across Rooms 1, 2, and 3 is 0.97, 1.00, and 0.97, while F1 scores are 0.91, 1.00, and 0.91, reflecting stable classification performance regardless of room context. Compared to day-based validation, accuracy variation is higher here, with a standard deviation of 1.73. Forecasting errors show zero false positives (FPs) except for Nurse 2 in Room 1, and zero false negatives (FNs) except for Nurse 3 in Room 1. Temporal validation over seven days (Figure 8) further confirms robustness, with the proposed method again outperforming all baselines. It achieves a macro average F1 score of 0.97 and accuracy of 0.99. Day-wise accuracy values are 1.00, 1.00, 0.97, 1.00, 0.97, 1.00, and 1.00, while F1 scores are 1.00, 1.00, 0.84, 1.00, 0.92, 1.00, and 1.00, demonstrating adaptability across operational conditions. F1 variation is higher here, with a standard deviation of 6.29. Forecasting errors show zero FNs except for Nurse 2 on Days 3 and 5, and zero FPs except for Nurse 3 on the same days. Overall, these results provide strong empirical evidence that the proposed hybrid system not only reduces unnecessary alarms but also ensures high reliability in alert recognition and nurse assignment, both spatially across rooms and temporally across daily sessions.

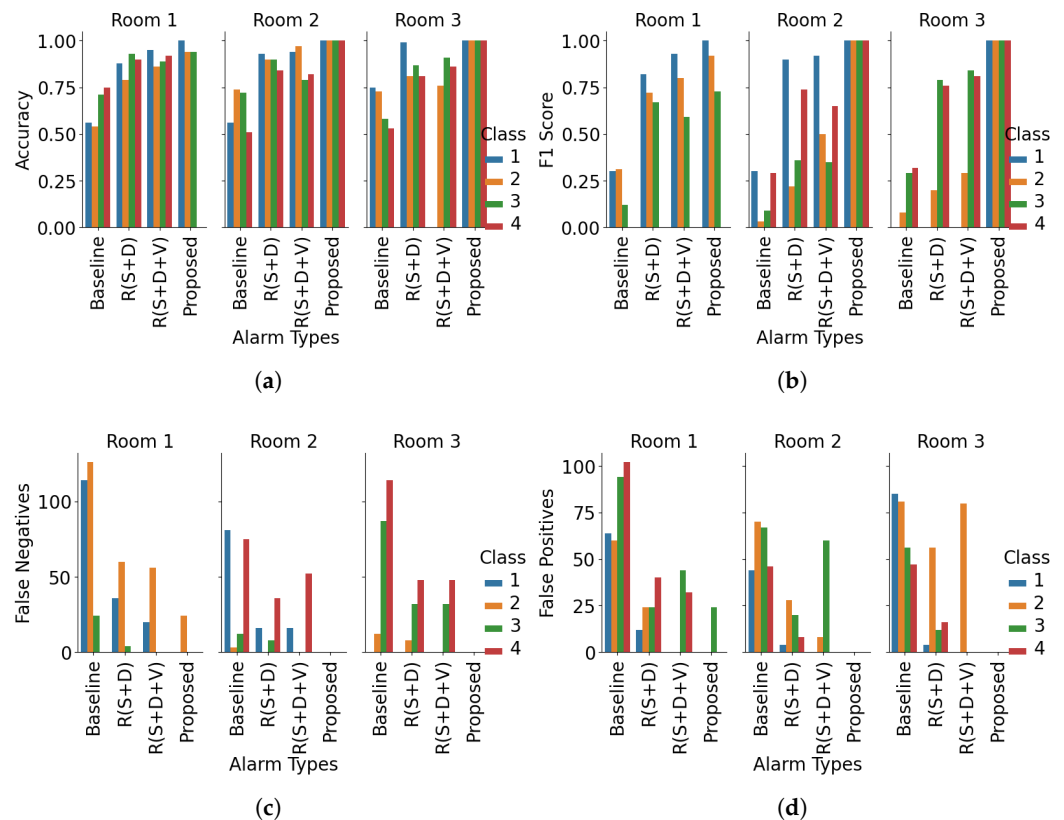


Figure 7. Room-wise performance comparison using accuracy (a), F1 score (b), false negatives (c), and false positives (d) as metrics. Each figure presents comparative results across three rooms and four nurses, highlighting performance variations influenced by spatial context. R(S+D) denotes rule-based suppression and delay, while R(S+D+V) represents the same with an added validation layer.

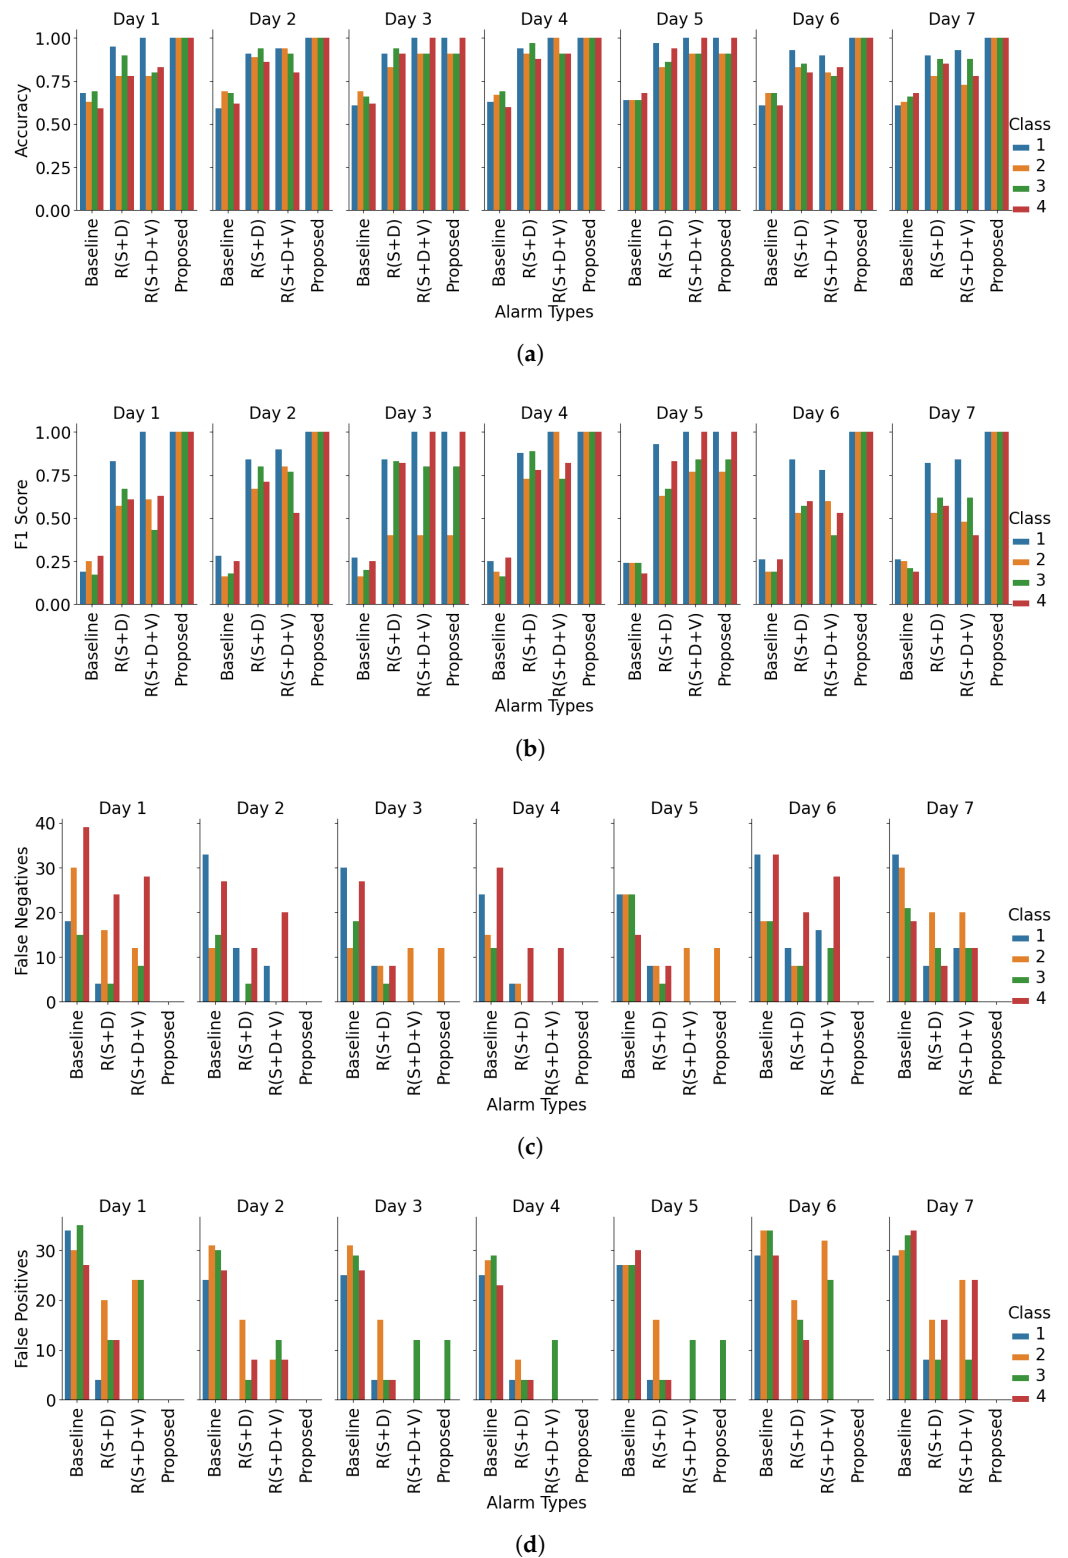


Figure 8. Day-wise performance comparison using accuracy (a), F1 score (b), false negatives (c), and false positives (d) as metrics. Each figure presents comparative results across three rooms and four nurses, highlighting performance variations influenced by temporal context. R(S+D) denotes rule-based suppression and delay, while R(S+D+V) represents the same with an added validation layer.

6. Discussion

This study investigated the design and evaluation of a hybrid, context-aware alerting framework for elderly care in long-term care (LTC) settings. The framework integrates deterministic rule-based suppression with flexible LLM-driven reasoning, enabling context-

sensitive, clinically meaningful alerts. The proposed context-aware alerting framework demonstrated significant improvements in both alert accuracy and reduction of unnecessary notifications. Evaluated using accuracy, precision, recall, F1 score, false positives (FPs), and false negatives (FNs), the model consistently outperformed baseline and rule-based methods. Ground truth was meticulously annotated from real staff interventions in video recordings, ensuring clinical relevance in validation.

6.1. Contributions

In the evaluation we could validate our initial research questions from the quantitative analysis. To answer research question ‘R1’ from Figure 4, we can see that our method reduced overall nurse alarm loads to 71% (Nurse 1), 71% (Nurse 2), 78% (Nurse 3), and 70% (Nurse 4) compared to the baseline. Room-wise reductions show even greater adaptability: in Room 1, Nurse 4 experienced 100% alert suppression; in Room 3, Nurse 1’s alerts were fully suppressed. These targeted reductions reduce the 100% alarm load (baseline) to 27.5% (proposed), achieving a 72.5% reduction on average. From Table 2, and Figures 5 and 6, the model achieved a macro average F1 score of 0.97 and accuracy of 0.98. Room-wise validation (Figure 7) further showed accuracies of 0.97, 1.00, and 0.97, and F1 scores of 0.91, 1.00, and 0.91 for Rooms 1, 2, and 3, respectively. Day-wise performance (Figure 8) was similarly robust: macro accuracy was ≥ 0.97 for all seven days, and F1 scores were ≥ 0.92 (except Day 3: 0.84). The overall accuracy of 0.98 and a macro F1 score of 0.97 far surpass the baseline accuracy of 0.21 and F1 score of 0.18. These results confirm that hybrid contextual alerting can effectively address alarm fatigue without sacrificing safety.

To answer research question ‘R2’ from Figure 4, we can see that compared to the baseline error rates (FPR 0.20, FNR 0.79), the method achieved drastically lower values (FPR 0.005, FNR 0.023). From Figure 7 we can see that FNs are zero except for Nurse 2 in Room 1, and FPs are zero except for Nurse 3 in Room 1, and from Figure 8 we can see that FNs are zero except for Nurse 2 on Days 3 and 5, and FPs are zero except for Nurse 3 on the same days. This confirms the utility of embedding interpretable, clinically grounded features for improving alert reliability.

To answer research question ‘R3’ from Figure 4, we can see that our proposed method achieved an accuracy of 0.98 and macro F1 score of 0.97, compared to the rule-based method (accuracy 0.78, F1 score 0.79). This validates the effectiveness of context-driven validation strategies in meeting the dynamic requirements of elderly care.

These results illustrate the clear benefits of contextual alerting in long-term care facilities. Despite the success, several areas remain for enhancement. The current model does not yet incorporate fatigue, emotional rapport, task load history, and in-depth temporal information, which significantly influence caregiver performance and trust. The lack of temporal information also reflects in the results as the day-based validation standard deviation of the F1 score is 6.29, which is higher compared to the room-based validation standard deviation of the F1 score (5.20). This indicates poorer performance in temporal variation. Caregiver specialization and dynamic task prioritization also warrant inclusion for broader applicability. Real-world deployment would require integration with feedback mechanisms, override options, and facility-specific routines. To that end, future work should focus on: embedding continuous learning from caregiver feedback, modeling longitudinal performance trends, and enabling interactive, bidirectional AI–human communication.

By combining spatial metrics with LLM-based reasoning, the system surpasses rigid rule-based filtering, enabling personalized nurse–patient matching, reducing alarm fatigue, and improving responsiveness. The hybrid approach balances interpretability and adaptability, making it well-suited for real-world elderly care, where it enhances alert efficiency, precision, and scalability.

6.2. Limitations and Future Challenges

While the proposed system demonstrates promising results for context-aware elderly monitoring and alerting in long-term care (LTC) environments, several limitations remain that open avenues for future improvement and practical enhancement. First, the current framework depends on prior patient information, such as care level assessment and average movement speed, to ensure accurate context modeling and alert generation. These parameters are essential for personalized monitoring, yet they pose challenges when applied to new residents without existing care history. To address this, future iterations of the system can integrate adaptive initialization modules that automatically estimate baseline parameters through short-term observation or unsupervised clustering of motion patterns during the first few days of monitoring. Such self-calibrating models would allow the system to dynamically personalize its thresholds and alerts for new patients without manual input. Second, the current study is confined to daytime group activity sessions, particularly during mealtimes, where multiple residents and nurses interact. Consequently, nighttime monitoring scenarios—such as sleeping posture analysis, nocturnal falls, or nighttime wandering—were beyond the present scope. However, because the proposed model operates on skeleton-based data, which are inherently privacy preserving, it can be extended for 24 h monitoring without compromising resident privacy. Future research could integrate low-light skeleton extraction and lightweight temporal models to extend system functionality to nighttime environments while maintaining operational continuity. Third, although broader challenges such as data privacy, budget constraints, and regulatory compliance remain important considerations, the proposed method already addresses some of these concerns by utilizing the existing video monitoring infrastructure of the facility without hardware modification. The only additional requirement is a GPU-enabled PC for processing, making deployment cost-effective and minimally intrusive. Additionally, the data utilized is skeleton data and the elderly and nurses are given a unique id for identification instead of their name, which to some extent mitigates the privacy issue. For future scalability, edge-computing integration or model optimization (e.g., pruning or quantization) could further reduce hardware dependency and energy consumption, enhancing accessibility for smaller care facilities. In summary, future extensions of this work will focus on adaptive parameter estimation, round-the-clock deployment, and resource-efficient scaling, ensuring that the system evolves into a fully autonomous, privacy-preserving, and sustainable solution for real-world elderly care environments.

7. Conclusions

This study presents a hybrid, context-aware nurse alerting framework for long-term care settings, aimed at mitigating alarm fatigue, caregiver overload, and delayed response in elderly monitoring. By combining rule-based filtering and delay mechanisms with semantic reasoning, the system achieved significant improvements in both efficiency and reliability. Evaluations using real-world video data demonstrated a 72.5% average reduction in nurse alert load and nearly complete elimination of false alarms, with macro F1 and accuracy scores of 0.97 and 0.98, respectively. These results validate the effectiveness of the interpretable suppression, adaptive delay, and contextual validation mechanisms proposed in this study. The framework offers a scalable path toward personalized and real-time alerting that enhances both resident safety and caregiver well-being. Future work will extend this approach to incorporate caregiver fatigue, specialization, and interactive AI-human collaboration to further advance human-centered elderly care.

Author Contributions: N.N.: methodology, data collection, visualization, writing—review and editing. M.A.R.A.: conceptualization, writing—review and editing, and supervision. S.I.: conceptualization, writing—review and editing, and supervision. All authors have read and agreed to the published version of the manuscript.

Funding: The data collection and experiment for this research were supported by the JST Challenging Research Program for Next Generation Researchers, Grant Number JPMJSP2154.

Institutional Review Board Statement: This work involved human subjects or animals in its research. Approval of all ethical and experimental procedures and protocols was granted by the Research Ethics Board at Kyushu Institute of Technology, under Application No. 24-15.

Informed Consent Statement: Informed consent was obtained from all subjects involved in the study.

Data Availability Statement: The data and source code will be made available upon request.

Acknowledgments: We would like to express our gratitude to Global Care, a dementia care facility in Kitakyushu, Japan, for supporting us in our experiment and providing us the opportunity to collect data on site.

Conflicts of Interest: The authors declare that they have no known competing financial interests or personal relationships that could have influenced the work reported in this study.

References

1. Dykes, P.C.; Bogaisky, M.; Carter, E.J.; Duckworth, M.; Hurley, A.C.; Jackson, E.M.; Khasnabish, S.; Lindros, M.E.; Lipsitz, S.R.; Scanlan, M.; et al. Development and validation of a fall prevention knowledge test. *J. Am. Geriatr. Soc.* **2019**, *67*, 133–138. [[CrossRef](#)]
2. Giordano, A.; Bonometti, G.P.; Vanoglio, F.; Paneroni, M.; Bernocchi, P.; Comini, L.; Giordano, A. Feasibility and cost-effectiveness of a multidisciplinary home-telehealth intervention programme to reduce falls among elderly discharged from hospital: Study protocol for a randomized controlled trial. *BMC Geriatr.* **2016**, *16*, 209. [[CrossRef](#)] [[PubMed](#)]
3. Yoshida, S. *A Global Report on Falls Prevention. Epidemiology of Falls*; World Health Organization: Geneva, Switzerland, 2007.
4. Røyset, B.; Talseth-Palmer, B.A.; Lydersen, S.; Farup, P.G. Effects of a fall prevention program in elderly: A pragmatic observational study in two orthopedic departments. *Clin. Interv. Aging* **2019**, *14*, 145–154. [[CrossRef](#)] [[PubMed](#)]
5. Florence, C.S.; Bergen, G.; Atherly, A.; Burns, E.; Stevens, J.; Drake, C. Medical costs of fatal and nonfatal falls in older adults. *J. Am. Geriatr. Soc.* **2018**, *66*, 693–698. [[CrossRef](#)] [[PubMed](#)]
6. Haddad, L.M.; Annamaraju, P.; Toney-Butler, T.J. Nursing Shortage. In *StatPearls [Internet]*; StatPearls Publishing: Treasure Island, FL, USA, 2024.
7. De Raeve, N.; Shahid, A.; De Schepper, M.; De Poorter, E.; Moerman, I.; Verhaevert, J.; Van Torre, P.; Rogier, H. Bluetooth-Low-Energy-Based Fall Detection and Warning System for Elderly People in Nursing Homes. *J. Sens.* **2022**, *2022*, 9930681. [[CrossRef](#)]
8. Harari, Y.; Shawen, N.; Mummidisetty, C.K.; Albert, M.V.; Kording, K.P.; Jayaraman, A. A smartphone-based online system for fall detection with alert notifications and contextual information of real-life falls. *J. Neuroeng. Rehabil.* **2021**, *18*, 124. [[CrossRef](#)]
9. Bugarin, C.A.Q.; Lopez, J.M.M.; Pineda, S.G.M.; Sambrano, M.F.C.; Loresco, P.J.M. Machine vision-based fall detection system using mediapipe pose with IoT monitoring and alarm. In Proceedings of the 2022 IEEE 10th Region 10 Humanitarian Technology Conference (R10-HTC), Hyderabad, India, 16–18 September 2022; pp. 269–274.
10. Saod, A.H.M.; Mustafa, A.A.; Soh, Z.H.C.; Ramlan, S.A.; Harron, N.A. Fall detection system using wearable sensors with automated notification. In Proceedings of the 2021 11th IEEE International Conference on Control System, Computing and Engineering (ICCSCE), Penang, Malaysia, 27–28 August 2021; pp. 182–187.
11. Madona, E.; Ulfa, S.; Nasution, A. Implementation of IMU sensor for Fall Detection in Dementia Syndrome Patients with Location Notification System. *J. Phys. Conf. Ser.* **2022**, *2406*, 012010.
12. Genco, E.K.; Forster, J.E.; Flaten, H.; Goss, F.; Heard, K.J.; Hoppe, J.; Monte, A.A. Clinically inconsequential alerts: The characteristics of opioid drug alerts and their utility in preventing adverse drug events in the emergency department. *Ann. Emerg. Med.* **2016**, *67*, 240–248. [[CrossRef](#)]
13. McFarlane, D.C.; Doig, A.K.; Agutter, J.A.; Brewer, L.M.; Syroid, N.D.; Mittu, R. Faster clinical response to the onset of adverse events: A wearable metacognitive attention aid for nurse triage of clinical alarms. *PLoS ONE* **2018**, *13*, e0197157. [[CrossRef](#)]
14. The Joint Commission. Sentinel Event Alert 50: Medical Device Alarm Safety in Hospitals. Available online: <https://www.jointcommission.org/en-us/knowledge-library/newsletters/sentinel-event-alert/issue-50> (accessed on 1 July 2025).

15. Sowan, A.K.; Stagers, N.; Reed, C.C.; Austin, T.; Chen, Q.; Xu, S.; Lopez, E. State of Science in Alarm System Safety: Implications for researchers, vendors, and clinical leaders. *Biomed. Instrum. Technol.* **2022**, *56*, 19–28. [[CrossRef](#)]
16. Dee, S.A.; Tucciarone, J.; Plotkin, G.; Mallilo, C. Determining the impact of an alarm management program on alarm fatigue among ICU and telemetry RNs: An evidence based research project. *SAGE Open Nurs.* **2022**, *8*, 23779608221098713. [[CrossRef](#)]
17. Nyarko, B.A.; Yin, Z.; Chai, X.; Yue, L. Nurses' alarm fatigue, influencing factors, and its relationship with burnout in the critical care units: A cross-sectional study. *Aust. Crit. Care* **2024**, *37*, 273–280. [[CrossRef](#)]
18. Lewandowska, K.; Weisbrot, M.; Cieloszyk, A.; Mędrzycka-Dąbrowska, W.; Krupa, S.; Ozga, D. Impact of alarm fatigue on the work of nurses in an intensive care environment—a systematic review. *Int. J. Environ. Res. Public Health* **2020**, *17*, 8409. [[CrossRef](#)]
19. Shanmugham, M.; Strawderman, L.; Babski-Reeves, K.; Bian, L. Alarm-related workload in default and modified alarm settings and the relationship between alarm workload, alarm response rate, and care provider experience: Quantification and comparison study. *JMIR Hum. Factors* **2018**, *5*, e11704. [[CrossRef](#)] [[PubMed](#)]
20. Winters, B.D.; Cvach, M.M.; Bonafide, C.P.; Hu, X.; Konkani, A.; O'Connor, M.F.; Rothschild, J.M.; Selby, N.M.; Pelter, M.M.; McLean, N.; et al. Society for Critical Care Medicine Alarm and Alert Fatigue Task Force. Technological distractions (part 2): A summary of approaches to manage clinical alarms with intent to reduce alarm fatigue. *Crit. Care Med.* **2018**, *46*, 130–137. [[CrossRef](#)] [[PubMed](#)]
21. Cvach, M. Monitor alarm fatigue: An integrative review. *Biomed. Instrum. Technol.* **2012**, *46*, 268–277. [[CrossRef](#)]
22. Drew, B.J.; Harris, P.; Zègre-Hemsey, J.K.; Mammone, T.; Schindler, D.; Salas-Boni, R.; Bai, Y.; Tinoco, A.; Ding, Q.; Hu, X. Insights into the problem of alarm fatigue with physiologic monitor devices: A comprehensive observational study of consecutive intensive care unit patients. *PLoS ONE* **2014**, *9*, e110274. [[CrossRef](#)]
23. Sowan, A.K.; Gomez, T.M.; Tariela, A.F.; Reed, C.C. Changes in default alarm settings and standard in-service are insufficient to improve alarm fatigue in an intensive care unit: A pilot project. *JMIR Hum. Factors* **2016**, *3*, e5098. [[CrossRef](#)]
24. Criscitelli, T. Alarm management: Promoting safety and establishing guidelines. *AORN J.* **2016**, *103*, 518–521. [[CrossRef](#)]
25. Imhoff, M.; Kuhls, S. Alarm algorithms in critical care monitoring. *Anesth. Analg.* **2006**, *102*, 1525–1537. [[CrossRef](#)]
26. Sieben, W.; Gather, U. Classifying alarms in intensive care—Analogy to hypothesis testing. In Proceedings of the Conference on Artificial Intelligence in Medicine in Europe, Amsterdam, The Netherlands, 7–11 July 2007; Springer: Berlin/Heidelberg, Germany; pp. 130–138.
27. Imhoff, M.; Morik, K.; Gather, U. Development of Decision Support Algorithms for Intensive Care Medicine: A Hybrid System of Time Series Analysis and a Knowledge Base System. In Proceedings of the The AMIA 2001 Annual Symposium, Washington, DC, USA, 3–7 November 2001; p. 931.
28. Mira, J.J.; Lorenzo, S.; Carrillo, I.; Ferrús, L.; Silvestre, C.; Astier, P.; Iglesias-Alonso, F.; Maderuelo, J.A.; Pérez-Pérez, P.; Torijano, M.L.; et al. Lessons learned for reducing the negative impact of adverse events on patients, health professionals and healthcare organizations. *Int. J. Qual. Health Care* **2017**, *29*, 450–460. [[CrossRef](#)]
29. Funk, M.; Ruppel, H.; Blake, N.; Phillips, J. Research: Use of monitor watchers in hospitals: Characteristics, training, and practices. *Biomed. Instrum. Technol.* **2016**, *50*, 428–438. [[CrossRef](#)]
30. Allan, S.H.; Doyle, P.A.; Sapirstein, A.; Cvach, M. Data-driven implementation of alarm reduction interventions in a cardiovascular surgical ICU. *Jt. Comm. J. Qual. Patient Saf.* **2017**, *43*, 62–70. [[CrossRef](#)]
31. Cvach, M.M.; Frank, R.J.; Doyle, P.; Stevens, Z.K. Use of pagers with an alarm escalation system to reduce cardiac monitor alarm signals. *J. Nurs. Care Qual.* **2014**, *29*, 9–18. [[CrossRef](#)] [[PubMed](#)]
32. Schmid, F.; Goepfert, M.S.; Reuter, D.A. Patient monitoring alarms in the ICU and in the operating room. *Crit. Care* **2013**, *17*, 216. [[CrossRef](#)]
33. Chromik, J.; Klopfenstein, S.A.I.; Pfitzner, B.; Sinno, Z.C.; Arnrich, B.; Balzer, F.; Poncette, A.S. Computational approaches to alleviate alarm fatigue in intensive care medicine: A systematic literature review. *Front. Digit. Health* **2022**, *4*, 843747. [[CrossRef](#)] [[PubMed](#)]
34. López-Espuela, F.; Rodríguez-Martin, B.; Garcia, J.L.; Toribio-Felipe, R.; Amarilla-Donoso, F.J.; Rodríguez Almagro, J.J.; Ribeiro, A.S.F.; Fernandes, V.S.; Moran-García, J.M. Experiences and mediating factors in nurses' responses to electronic device alarms: A phenomenological study. *J. Nurs. Manag.* **2022**, *30*, 1303–1316. [[CrossRef](#)] [[PubMed](#)]
35. Piri, H.; Huh, W.T.; Shechter, S.M.; Hudson, D. Individualized dynamic patient monitoring under alarm fatigue. *Oper. Res.* **2022**, *70*, 2749–2766. [[CrossRef](#)]
36. Guglielmello, G.; Stawicki, S.P. Combating alarm fatigue: The quest for more accurate and safer clinical monitoring equipment. *Vignettes Patient Saf.* **2019**, *4*, 93.
37. Patel, N.Z.; Patel, D.V.; Phatak, A.G.; Patel, V.G.; Nimbalkar, S.M. Reducing false alarms and alarm fatigue from pulse oximeters in a neonatal care unit: A quality improvement study. *J. Neonatol.* **2022**, *36*, 135–142. [[CrossRef](#)]
38. Vieira, J.C.F. Human-Centred Design of Clinical Auditory Alarms. Ph.D. Thesis, Universidade de Lisboa, Lisboa, Portugal, 2022.

39. van Rossum, M.C.; Vlaskamp, L.B.; Posthuma, L.M.; Visscher, M.J.; Breteler, M.J.; Hermens, H.J.; Kalkman, C.J.; Preckel, B. Adaptive threshold-based alarm strategies for continuous vital signs monitoring. *J. Clin. Monit. Comput.* **2022**, *36*, 407–417. [[CrossRef](#)] [[PubMed](#)]
40. Isma'ila, U.A.; Danyaro, K.U.; Hassan, M.F.; Liew, M.S.; Maiwada, U.D.; Muazu, A.A. Evaluation on bot-IoT dataset enabled reducing false alarm rate for IoT threats. *Kepes* **2023**, *21*, 490–504.
41. Silveira, S.Q.; Nersessian, R.S.F.; Abib, A.D.C.V.; Santos, L.B.; Bellicieri, F.N.; Botelho, K.K.; de Oliveira Lima, H.; de Queiroz, R.M.; dos Anjos, G.S.; dos Santos Fernandes, H.; et al. Decreasing inconsistent alarms notifications: A pragmatic clinical trial in a post-anesthesia care unit. *Braz. J. Anesthesiol.* **2024**, *74*, 744456. [[CrossRef](#)]
42. Fernandes, C.O. Patient Monitoring and Caregivers Notification Based on a Reasoning Model to avoid Alarm Fatigue. Ph.D. Thesis, PUC-Rio, Gávea, Brazil, 2019.
43. Fernandes, C.; Miles, S.; Lucena, C.J.P. Detecting false alarms by analyzing alarm-context information: Algorithm development and validation. *JMIR Med. Inform.* **2020**, *8*, e15407. [[CrossRef](#)]
44. Mileski, M.; Brooks, M.; Topinka, J.B.; Hamilton, G.; Land, C.; Mitchell, T.; Mosley, B.; McClay, R. Alarming and/or alerting device effectiveness in reducing falls in long-term care (LTC) facilities? A systematic review. *Healthcare* **2019**, *7*, 51. [[CrossRef](#)]
45. WongKinYiu. (n.d.). GitHub—WongKinYiu/yolov7: Implementation of Paper—YOLOv7: Trainable Bag-of-Freebies Sets New State-of-the-Art for Real-Time Object Detectors. GitHub. Available online: <https://github.com/WongKinYiu/yolov7> (accessed on 1 February 2025).
46. Noorkhokhar. (n.d.). People-Tracking-yolov7/README.md at main · noorkhokhar99/people-tracking-yolov7. GitHub. Available online: <https://github.com/noorkhokhar99/people-tracking-yolov7/blob/main/README.md> (accessed on 1 February 2025).
47. Facebook. (n.d.). GitHub—Facebook/Prophet: Tool for Producing High Quality Forecasts for Time Series Data That Has Multiple Seasonality with Linear or Non-Linear Growth. GitHub. Available online: <https://github.com/facebook/prophet> (accessed on 1 February 2025).
48. Matsuda, S.; Muramatsu, K.; Hayashida, K. Eligibility Classification Logic of the Japanese Long Term Care Insurance. *Asian Pac. J. Dis. Manag.* **2011**, *5*, 65–74. [[CrossRef](#)]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

Reproduced with permission of copyright owner. Further reproduction prohibited without permission.