Social Robots for Bed-Fall Detection in Hospitals

Luigi D'Arco*, Vincenzo Marotta, Silvia Rossi and Alessandra Rossi

Abstract— **Patients falling from their beds is still one of the major complications of the hospital's treatments. To address this issue, this research investigates the use of social robots for the identification of potential bed-related falls in hospitals. Using a robot camera and human pose estimation techniques, the patient's position in the bed is extracted, and a threshold-based algorithm is used to identify any anomalies that could indicate a fall. Due to the absence of publicly available datasets, a synthetic dataset was created using a simulation environment to develop and tune the detection algorithm. A user study was conducted to validate the proposed approach and evaluate people's perception of the robot. The system achieved an accuracy of 90.9% in the controlled setting using real data. Participants rated the robot as significantly more trustworthy and behaviorally aware when it detects a possible fall, suggesting that timely and meaningful assistance improved the perceived social competence of the robot. These findings highlight the feasibility of deploying social robots as monitoring systems in sensitive clinical settings, offering a cost-effective and socially acceptable solution.**

I. Introduction

The lowering of the birth rate and the increase in life expectancy are progressively leading to the aging of the world population [1]. Such aging leads to an increase in the need for healthcare and social support, with a significant impact on health and social systems [2], and an increase in the causes of hospitalization, such as chronic diseases, degenerative diseases, and domestic accidents [3]. This phenomenon translates into an increasingly exhausting workload for healthcare professionals, who must cope with a growing number of patients with complex and diverse needs. Among the various tasks that are assigned to medical staff, are those of monitoring and supervising patients in their beds. This task is particularly critical for older patients, who are more vulnerable to falls from bed [4]. Falls have been classified by the World Health Organization as the second leading cause of accidental death in the world, and represent 40% of deaths due to injury [5]. Consequently, to limit the burden on hospital professionals different technologies have been introduced and researched, including wearable sensors [6], [7], floor sensors [8], bed sensors [9], and camera-based solutions [10], [11]. However, such systems present several limitations, such as high costs, limited to a single environment, and privacy concerns [12].

L. D'Arco, V. Marotta, S. Rossi, and A. Rossi are with the Department of Electrical Engineering and Information Technologies of the University of Naples Federico II, Italy. (e-mail: luigi.darco@unina.it; vincenzo.marotta6@studenti.unina.it; silvia.rossi@unina.it; alessandra.rossi@unina.it).

This work has been supported by the European Union - Next Generation EU, Mission 4 Component 1, CUP E53D23016260001 PRIN 2022 PNRR ADVISOR (L D'Arco), and PRIN 2022 TrustPACTX CUP E53D23007850001 (A. Rossi and S. Rossi).

To address these limitations, this research investigates the use of Socially Assistive Robots (SARs) for the prevention of falls in hospitals. SARs are already used for monitoring patients [13], and their importance is constantly growing due to the lack of personnel in hospitals [14]. One of the advantages of SARs in hospitals is the ability to collect information about the patient's health status through sensors connected to the device; making them adaptable to different environments without the need for additional installations or modifications. SARs can also contribute to the mitigation of emotional and psychological problems, such as loneliness and depression, which can negatively affect the physical health of patients [15].

This study presents a robotic-based solution for bed-fall detection via the robot's camera to monitor the patient's bed position and detect a possible risk of falling from the patient's position. The system is designed to be real-time and multi-person, to ensure greater accuracy and reduce costs compared to sensor-based solutions. Due to the absence of publicly available datasets suitable for testing the system, as existing datasets only contain top-view images, a synthetic dataset was first generated using a simulation environment. The system was then tested on real data collected during robot operation in a controlled setting designed to simulate clinical conditions. Finally, we evaluated participants' perceptions of the robot, including their acceptance and trust in the robotic system.

II. Related Work

Bed-related fall detection in hospital environments has been the subject of several studies to prevent injuries to patients and, at the same time, reduce the burden on hospital personnel. Different approaches have been proposed, including the use of wearable sensors, bed sensors, and camera-based solutions.

A. Systems for Fall Detection

Fernandez et al. [16] developed an alarm system using Inertial Measurement Units (IMUs) and state machine algorithms to detect falls by analyzing acceleration peaks over one-second windows. Zhao et al. [17] introduced an intelligent bed system that monitors body states to facilitate caregiver collaboration and ease the burden on healthcare staff. To address privacy concerns, Asbjørn et al. [18] used a ceiling-mounted thermal array with an ultrasonic sensor to monitor bedside events. This approach achieved high accuracy while minimizing privacy intrusion.

However, these systems often require extensive calibration and setup, making them costly and difficult to transfer between environments. Moreover, machine learning-based methods rely on large datasets, which are not always available in healthcare contexts.

B. SARs for Fall Detection

Social robots offer new potential for fall detection by enabling autonomous patient monitoring and staff alerting with minimal setup.

Tomoya et al. [19] proposed a household mobile robot system that can follow older people, detect falls, and notify caregivers. The robot exploited the use of a Microsoft Kinect camera combined with a range sensor to detect the skeleton information of the person, identify the depth information to close the gap between the two, and by using a distance calculation between head and knees, was able to determine if the user fell or not. However, such an approach requires an extensive evaluation of the environment with tracking of obstacles and human position updates, which requires high computational performance from the robot. Furthermore, such an approach cannot be shifted for the detection of bed-related falls, which makes adaptation in the clinical environment challenging.

Shifting such a concept to hospital environments, Erzgraber et al. [20] explored the use of a robotic system in healthcare settings to serve as a sentinel to monitor movements in bed. Data were collected at night for approximately 8 hours through a Kinect One camera, providing infrared depth images. Due to the limitations of the robot hardware, a small computer vision model (i.e., YOLOv5S) was used to detect the patient and the bed. A body movement detection algorithm was developed using a k-nearest neighbor algorithm to cluster the information in each image, and according to that, decide whether a warning should be sent to the staff. Although the system was able to detect the patient's movements, it introduced privacy concerns as the robot took pictures of the patient. Furthermore, to allow continuous working of the robot, it had to be attached by a cable to the power outlet, which limits the range of actions of the system and increases safety risks for those (e.g., patients, personnel, family members) moving around it.

Recent efforts have also been redirected to the adoption of a social robot for assessing and evaluating people's health status immediately after a fall [21]. In this work, the authors use computer vision algorithms implemented in OpenCV to assess changes in a fallen person's movements using the cameras of a MiRo robot. While this solution may be effective in evaluating whether people are standing up, it still does not help to prevent falls and prevent people from getting injured, which is the main purpose of this work.

III. Methodology

To address the need for timely and reliable detection of falls from bed, we propose a vision-based, threshold-driven fall detection system for a robotic agent that leverages Human Pose Estimation (HPE) for monitoring patients in bed. The overall system pipeline is shown in Figure 1.

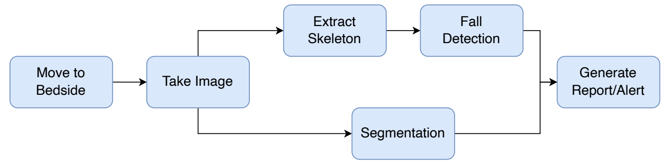

Fig. 1: Overview of the Bed-Fall Detection Pipeline.

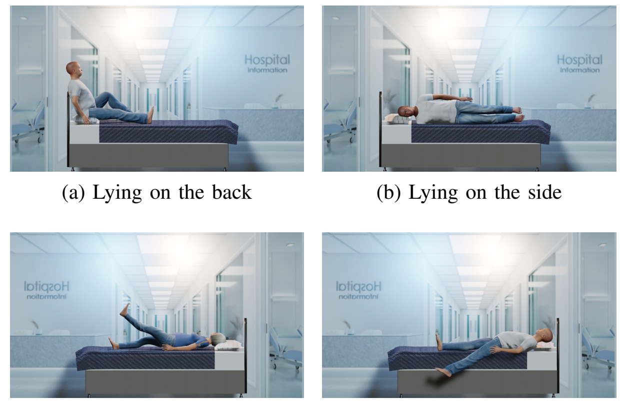

(a) Lying on the back

(b) Lying on the side

(c) With leg up

(d) Possible fall

Fig. 2: Examples of the images generated for the dataset.

A. Dataset Generation

Datasets involving fall activities are diffused in the literature [22], but a dataset focused on images taken at the robot's height is not available. For this reason, a synthetic dataset has been generated for the purpose of this study. The dataset was created using the Blender software[1], which is a free and open-source 3D creation suite, along with the Make Human plugin[2]. The dataset consists of 270 images of a person lying on a bed in different positions, comprising lying on the back, lying on the side, sitting on the bed, with leg up, and falling (see Figure 2). The images comprised different people with different ages, genders, ethnicities, body shapes, clothes, and a hospital background for simulating a real environment.

B. Fall Detection Algorithm

Given the constraints of hospital deployment (limited computational resources, need for real-time inference, and absence of large-scale annotated datasets), a rule-based method is preferable over deep learning models. While learning-based models (e.g., Graph Neural Network, Recurrent Neural Network, spatio-temporal attention mechanisms) have shown promise in activity recognition, they require significant training data and computational power not feasible for bedside deployment on mobile robotic platforms. For this reason, we propose a lightweight rule-based system to detect fall states from patient images that are interpretable, efficient, and robust to variations in patient appearance and pose.

The robot is first guided through the hospital environment using a predefined environmental map. Specific Points of

[1]https://www.blender.org
[2]https://static.makehumancommunity.org

Interest (POIs), the left or right side of the patient bed, are visited, and the robot captures an image of the patient. Human pose keypoints are extracted using the OpenPose framework [23], a bottom-up, multi-person system that employs a Convolutional Neural Network (CNN) to predict confidence maps and part affinity fields for associating body parts. A greedy bipartite matching algorithm assembles full poses, returning 95 keypoints per person—70 facial and 25 full-body—each represented as a tuple $K_i = (x_i, y_i, c_i)$, where (x_i, y_i) are pixel coordinates and c_i is the confidence score.

To ensure consistency in interpreting body landmarks (e.g., left hip vs. right hip), the system identifies which side of the patient is more visible by computing the average confidence of visible keypoints on both sides of the body:

$$S_{\text{left}} = \frac{\sum_{i \in L} c_i}{|L|}, \quad S_{\text{right}} = \frac{\sum_{i \in R} c_i}{|R|} \tag{1}$$

where L and R are the index sets for left and right side keypoints (selected from shoulder, elbow, wrist, hip, knee, and ankle). The side with the higher average score is assumed to face the robot, and subsequent calculations are based on that side to ensure consistency in anatomical reference.

Several critical features are extracted to evaluate the patient's posture and infer their safety state, including:

- *Lying Position*: defined as the vertical displacement between the shoulder and hip keypoints. Empirical evaluation of collected data demonstrated an AUC > 0.9 in distinguishing lying versus sitting/non-sitting postures. Lying position is a key indicator of patient safety.
- *Leg Lifted*: defined as the vertical displacement between the hip and foot keypoints. Such a posture can be imposed by medical conditions and should not be considered dangerous by the system. The foot is assumed to be represented by the most reliably detected among heel and ankle keypoints, based on the confidence score.
- *Leg Outside Bed*: defined in two stages. First, a primary condition checks whether the foot keypoint is vertically below the hip and then the displacement between the knee and foot is computed. A significant increase in the displacement between the knee and foot further supports the hypothesis that the leg is extended downward and potentially off the bed.
- *Head Drop*: defined as the vertical displacement between the ear and hip keypoints. An unusually high value may indicate that the head has dropped significantly, suggesting a loss of postural control or unconsciousness.

To eliminate dependence on camera distance and patient size, all features were normalized using the shoulder-hip distance. This normalization ensures scale invariance, improving robustness to variations in patient positioning or camera perspective.

Each posture-related condition contributes a weighted score toward two competing states: safe and fall. The score functions are defined as:

$$\begin{aligned} S_{\text{safe}} &= w_1(\text{lying}) + w_2(\text{leg-lifted}) + w_3(\neg\text{leg-out}) \\ &\quad + w_4(\neg\text{head-drop}) \\ S_{\text{fall}} &= w_5(\neg\text{lying}) + w_6(\text{leg-out}) + w_7(\text{head-drop}) \end{aligned} \tag{2}$$

where w_i are empirically tuned weights reflecting the significance of each cue in predicting falls. Notably, leg-lifted contributes only to the safe score, serving as a mitigating factor for false predictions when specific postures are required by the medical conditions but are not dangerous. The final class probabilities for the safe and fall states are computed by applying the softmax function over the respective scores. This transformation converts the raw scores into normalized probability values that sum to one, ensuring a probabilistic interpretation of the model output:

$$\begin{aligned} P_i &= \frac{e^{S_i - \max(S_{\text{safe}}, S_{\text{fall}})}}{e^{S_{\text{safe}} - \max(S_{\text{safe}}, S_{\text{fall}})} + e^{S_{\text{fall}} - \max(S_{\text{safe}}, S_{\text{fall}})}}, \\ &\text{for } i \in \{\text{safe}, \text{fall}\} \end{aligned} \tag{3}$$

where P_i denotes the probability of the i-th class, and $\max(S_{\text{safe}}, S_{\text{fall}})$ is the maximum score between the two classes included for numerical stability.

C. Optimization of Classification Thresholds

To optimize the classification thresholds associated with each postural feature, the system was evaluated on the generated dataset. A Bayesian Optimization strategy [24] was adopted to identify the optimal threshold values for each decision boundary. Threshold tuning was considered a hyperparameter optimization problem, and by building Gaussian process surrogate models, the system was able to explore the parameter space and converge to the optimal decision boundary configuration. The objective was to maximize the system's F1-score, a harmonic mean of precision and recall. By iteratively refining the thresholds set based on performance feedback, the system converged toward an optimal decision boundary configuration that maximizes fall detection accuracy while minimizing false predictions.

D. Privacy-Preserving Approach

To ensure patient privacy and comply with ethical data handling standards, the system avoids storing raw images containing identifiable visual information. Instead, a segmentation-based anonymization strategy has been adopted to remove sensitive details while preserving structural and contextual information. Specifically, the Segment Anything Model (SAM) [25] was employed to segment each image into semantically distinct regions, where each segment was assigned a unique color. This produces a stylized, privacy-preserving representation in which facial features and other identifying characteristics are no longer discernible.

Figure 3 illustrates the anonymization process applied to a sample image. The resulting segmented images retain sufficient spatial and anatomical context to support clinical interpretation, enabling their use in generating postural reports for healthcare personnel. In scenarios where a fall is detected or imminent, these anonymized visual summaries

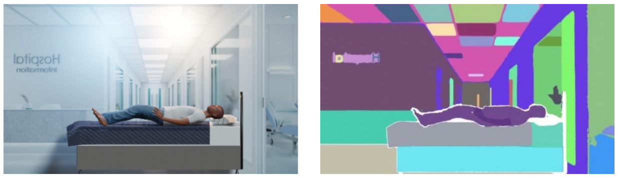

(a) (b)

Fig. 3: Privacy Preserving Approach. (a) Original Image. (b) Segmented Image.

(a) (b)

Fig. 4: Examples of the experimental setup. (a) Participant in a *SAFE* condition. (b) Participant in a *FALL* condition.

can be included in automated alerts without compromising the patient's identity. Furthermore, the use of these segmented images opens a path toward the creation of a labeled dataset that respects patient privacy. Healthcare professionals can annotate these images with posture labels (e.g., "safe", "falling"), facilitating future model development and validation without requiring access to original visual data.

IV. Use Case

To support hospital staff unable to continuously monitor patients, we developed a robot that autonomously navigates hospital environments, visiting predefined points of interest (i.e., beds in hospital rooms) to assess posture, detect potential falls, and interact with patients.

The system uses the Temi robot [26], equipped with a 13.3-inch tablet, 13MP camera with tracking and facial recognition, microphone, audio system, LiDAR, 3D and depth sensors for obstacle avoidance and mapping. This robot was chosen for its capabilities of moving autonomously and fluidly, with a maximum speed of 1 m/s [27], in complex environments, recognizing people and objects.

Temi navigates using a preloaded map and detects dynamic obstacles in real time. At each point of interest, it uses its camera to assess the patient's posture. If a possible fall position is recognized, the robot asks the patient if they need any help. In case the patient answers positively or does not answer after multiple prompts, the robot informs the hospital personnel to provide immediate assistance. In case the robot does not recognize any fall position, it proceeds to dialogue with the patient about their general health status, including the quality of their sleep and their mood, and checks whether they have any pain or discomfort. The robot then moves to the next POI (i.e., the next patient). The robot continues its monitoring process until all the POIs have been visited.

From the surveillance and interaction, the robot creates reports for the healthcare personnel. The system is designed to be real-time and person-agnostic to ensure greater accuracy and reduce costs compared to sensor-based solutions.

A. Experimental Settings

An experimental setup was created to test the system in a real controlled environment, and a POI was set at the center of a living area facing a sofa. A warning zone was set at the end of the room to simulate the position of the hospital personnel if the users showed any sign of falling or if they interacted with the robot asking for help. In order to collect participants' acceptance and perception of the robot, we designed an in-between-subject study, where participants were randomly assigned to only one of two experiment categories (*FALL*, *SAFE* - see Figure 4). Participants assigned to the *FALL* condition were asked to simulate at their best without compromising their security, a fall pose on the sofa; 2) participants assigned to the *SAFE* condition were instructed to completely lay on the sofa.

At the end of their study trial, participants were asked to complete a questionnaire including information about their demographics and their experience in interacting with the robot. The questionnaire was structured in two parts, with the first focusing on evaluating the user experience in interacting with the robot, while the second part focused on the user's perception of the robot's social skills and intelligence (Perceived Social Intelligence (PSI) scales test [28]). In particular, we collected participants' responses to evaluate the following dimensions: RB (Recognizes Human Behaviours), IH (Identifies Humans), SOC (Social Competence), FRD (Friendly), HLP (Helpful), TRU (Trustworthy).

A total of 27 participants were recruited for the study, with 74% of the participants stating to be males and 25.9% females, with an age ranging from 19 to 34 years old ($\mu = 25, \sigma = 3.36$).

V. Results

To measure the feasibility of the deployment of the proposed system in hospital environments, we evaluated its performance with both synthetic and real-world data and assessed people's perceptions and trust in the robot.

A. Fall Detection Performance Evaluation

The proposed threshold-based fall detection system was evaluated on the synthetic dataset generated using the Blender software (*Test Set*) and on the data collected during the use of the robot in a controlled environment (*Real Set*). The system achieved an accuracy of 91.7% on the *Test set* and 90.9% on the Real set. The precision, recall, and F1-score for each class are reported in Table I. The system demonstrated high performance in detecting both safe and fall states, with F1-scores exceeding 90% for safe states and 88% for fall states. Notably, the system's robustness was further confirmed by the real-world evaluation, which showed consistent performance across different participants' demographics and postures.

TABLE I: Evaluation of the Fall Detection System on the Test and Real Sets

Set	Class	Accuracy	Precision	Recall	F1-score
Test	SAFE	0.917	0.906	0.967	0.935
	FALL		0.938	0.833	0.882
Real	SAFE	0.909	0.867	0.929	0.897
	FALL		0.944	0.895	0.919

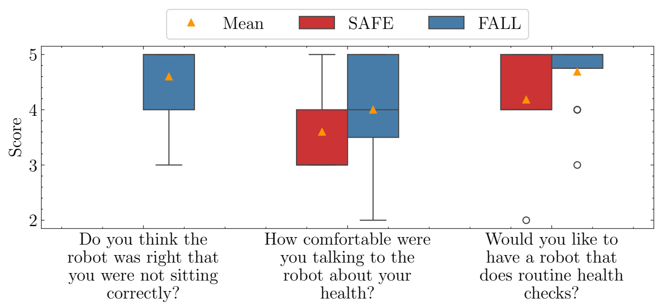

Fig. 5: Participant responses to questions related to their interaction with the robot, segmented by experimental condition (*SAFE*, *FALL*)

Despite its performance and interpretability, the proposed system presents different limitations. The generalization of the system is limited by the quality and visibility of the extracted keypoints, and even though the system assumes that part of the body is not always visible, the identification relies strongly on the OpenPose framework, which cannot estimate keypoints if the user is covered by blankets or other objects. The use of the generated dataset allowed the identification of optimal threshold values, which showed reliance while testing on the real-world data. However, we asked participants not to lay in any dangerous position, and complex poses were not considered in the evaluation, which may possibly lower the system's performance.

B. Post-Experiment Questionnaire Results

To assess how participants perceived the robot during the experiment, we analyzed questionnaire responses related to their interaction experiences. Figure 5 illustrates the distribution of participant scores across key questions, segmented by experimental condition (*FALL* vs. *SAFE*).

Participants in the *FALL* condition, who experienced a simulated fall scenario, received proactive engagement from the robot, which initiated a dialogue to check on their status and offer assistance. When asked whether the robot was correct in identifying that they were not sitting properly, participants in the *FALL* group reported a high level of agreement ($\mu = 4.60, \sigma = 0.63$). This result suggests that the robot's situational awareness was perceived as accurate and relevant. In terms of comfort when discussing their health with the robot, *FALL* participants again rated the interaction more positively ($\mu = 4.00, \sigma = 1.20$) compared to those in the *SAFE* condition ($\mu = 3.60, \sigma = 0.89$). When asked whether they would be willing to have a robot perform routine health checks, participants expressed strong agreement except for two participants. These findings indicate that participants who experienced contextually meaningful interaction with the robot (i.e., in the *FALL* scenario) not only validated the robot's assessments but also reported higher comfort in communication and greater openness to robotic health monitoring in the future. However, by using an independent samples t-test analysis, no statistical difference was discovered across the two conditions.

TABLE II: Summary of the PSI Scales according to the type of session participants were assigned (FALL, SAFE)

	FALL $\mu(\sigma)$	SAFE $\mu(\sigma)$	p-value
RB	3.859 (0.890)	3.159 (0.735)	0.041*
IH	4.250 (0.987)	4.159 (0.605)	0.788
SOC	3.562 (1.063)	2.886 (0.990)	0.108
FRD	3.438 (1.059)	2.955 (0.941)	0.235
HLP	4.328 (0.604)	3.773 (0.817)	0.053
TRU	4.203 (0.579)	3.636 (0.832)	0.047*

μ stands for the mean, σ for the standard deviation, and * indicates statistical significance with $\alpha = 0.05$

To evaluate the participants' perceptions of the robot's social intelligence, a statistical comparison using the independent samples t-tests was conducted on the scales of the PSI. Results revealed differences between the two experimental conditions (see Table II). The robot was perceived as significantly more capable of recognizing human behaviors in the *FALL* condition ($\mu = 3.859$, $\sigma = 0.890$) compared to the *SAFE* condition ($\mu = 3.159$, $\sigma = 0.735$), with a p-value of 0.041. Similarly, trustworthiness scores were significantly higher in the *FALL* group ($\mu = 4.203$, $\sigma = 0.579$) than in the *SAFE* group ($\mu = 3.636$, $\sigma = 0.832$), yielding a p-value of 0.047. While people's perception of helpfulness (HLP), social competence (SOC), friendliness (FRD), and human identification (IH), did not show statistically significant differences between conditions, their mean scores were consistently higher in the *FALL* condition, suggesting that the robot was perceived more favorably when it was actively involved in a critical assistance scenario (i.e., responding to a potential fall).

VI. Conclusion

This study presents a robotic-based solution designed to detect bed-related falls in hospital environments. The system utilizes an HPE approach to extract patient posture data and determine fall risk through a threshold-based decision model. A distinguishing feature of the solution is its privacy-preserving architecture that does not store original images, but it uses segmented and anonymized visual data that retains essential structural information while removing identifiable details. This design enables effective alert generation while safeguarding patient privacy. In addition to its privacy-aware design, the system emphasizes interpretability and ease of deployment. Unlike complex machine learning models that require extensive training and tuning, this rule-based method offers plug-and-play adaptability, allowing for rapid implementation in diverse environments with minimal data requirements. The use of HPE modules also ensures scalability, enabling future enhancements to address challenges such as partial occlusions, variable lighting conditions, and different room configurations.

The system has been evaluated on both synthetic datasets and real-world data collected during robot-assisted monitoring in a controlled hospital-like environment. The results showed high and consistent performance, with F1-scores exceeding 90% for identifying safe states and 88% for detecting falls. Furthermore, user feedback collected via a post-experiment questionnaire highlighted that the robot has been perceived as more capable of understanding human behavior and more trustworthy when actively engaged in critical scenarios, such as fall detection and response. An added value of the robotic component is its capacity to foster companionship and build trust with users, contributing to overall system acceptance.

As a patrolling robot, it cannot monitor all patient rooms simultaneously, however, future developments could explore hybrid solutions by integrating environmental sensors or wearable devices to enable continuous and comprehensive monitoring. Our next phase of research will focus on generating labeled datasets for the training and validation of advanced models, and improving the system's performance by incorporating temporal features from video streams. Ultimately, the system will undergo extensive evaluation in real-world hospital settings to assess its feasibility, robustness, and effectiveness in long-term clinical deployment.

References

[1] W. Chen, Z. Jiang, H. Guo, and X. Ni, "Fall detection based on key points of human-skeleton using openpose," *Symmetry*, vol. 12, no. 5, p. 744, 2020.

[2] B. Lindgren, "The rise in life expectancy, health trends among the elderly, and the demand for health and social care," tech. rep., 2016.

[3] E. Nowossadeck, "Population aging and hospitalization for chronic disease in germany," *Deutsches Ärzteblatt International*, vol. 109, no. 9, p. 151, 2012.

[4] D. Evans, B. Hodgkinson, L. Lambert, and J. Wood, "Falls risk factors in the hospital setting: a systematic review," *International journal of nursing practice*, vol. 7, no. 1, pp. 38–45, 2001.

[5] WHO., *WHO Global Report on Falls: Prevention in Older Age*. World Health Organization, 2007.

[6] W.-C. Chou, W.-Y. Lin, M.-Y. Lee, and K. F. Lei, "Design and assessment of a real-time accelerometer-based lying-to-sit sensing system for bed fall prevention," in *2013 IEEE International Conference on Systems, Man, and Cybernetics*, pp. 1471–1475, IEEE, 2013.

[7] A. Härmä, W. Ten Kate, and J. Espina, "Bed exit prediction based on movement and posture data," in *IEEE-EMBS International Conference on Biomedical and Health Informatics (BHI)*, pp. 165–168, IEEE, 2014.

[8] Y. Masuda, T. Yoshimura, K. Nakajima, M. Nambu, T. Hayakawa, and T. Tamura, "Unconstrained monitoring of prevention of wandering the elderly," in *Proceedings of the Second Joint 24th Annual Conference and the Annual Fall Meeting of the Biomedical Engineering Society][Engineering in Medicine and Biology*, vol. 3, pp. 1906–1907, IEEE, 2002.

[9] C. Sri-Ngernyuang, P. Youngkong, D. Lasuka, K. Thamrongaphichartkul, and W. Pingmuang, "Neural network for on-bed movement pattern recognition," in *2018 11th Biomedical Engineering International Conference (BMEiCON)*, pp. 1–4, IEEE, 2018.

[10] M. Inoue, R. Taguchi, and T. Umezaki, "Vision-based bed detection for hospital patient monitoring system," in *2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 5006–5009, IEEE, 2018.

[11] K.-S. Song, Y.-H. Nho, and D.-S. Kwon, "Histogram based fall prediction of patients using a thermal imagery camera," in *2017 14th International Conference on Ubiquitous Robots and Ambient Intelligence (URAI)*, pp. 161–164, IEEE, 2017.

[12] R. Igual, C. Medrano, and I. Plaza, "Challenges, issues and trends in fall detection systems," *Biomedical engineering online*, vol. 12, no. 1, p. 66, 2013.

[13] G. Maggi, L. Raggioli, A. Rossi, and S. Rossi, "Assessment of Distraction and the Impact on Technology Acceptance of Robot Monitoring Behaviour in Older Adults Care," *IEEE Transactions on Affective Computing*, no. 01, pp. 1–14, 2025.

[14] E. Martinez-Martin and A. Costa, "Assistive technology for elderly care: an overview," *IEEE access*, vol. 9, pp. 92420–92430, 2021.

[15] E. G. Christoforou, S. Avgousti, N. Ramdani, C. Novales, and A. S. Panayides, "The upcoming role for nursing and assistive robotics: Opportunities and challenges ahead," *Frontiers in Digital Health*, vol. 2, p. 585656, 2020.

[16] J. Fernández-Bermejo Ruiz, J. Dorado Chaparro, M. J. Santofimia Romero, F. J. Villanueva Molina, X. Del Toro García, C. Bolaños Peño, H. Llumiguano Solano, S. Colantonio, F. Flórez-Revuelta, and J. C. López, "Bedtime Monitoring for Fall Detection and Prevention in Older Adults," *International Journal of Environmental Research and Public Health*, vol. 19, p. 7139, June 2022.

[17] D. Zhao, Y. Wu, C. Yang, J. Yang, H. Liu, S. Wang, Y. Jiang, and Y. Hiroshi, "A novel multifunctional intelligent bed integrated with multimodal human–robot interaction approach and safe nursing methods," *IET Cyber-Systems and Robotics*, vol. 5, no. 3, p. e12097, 2023.

[18] D. Asbjørn and T. Jim, "Recognizing bedside events using thermal and ultrasonic readings," *Sensors*, vol. 17, no. 6, p. 1342, 2017.

[19] A. Tomoya, S. Nakayama, A. Hoshina, and M. Sugaya, "A mobile robot for following, watching and detecting falls for elderly care," *Procedia computer science*, vol. 112, pp. 1994–2003, 2017.

[20] R. Erzgräber, F. Lischke, F. Bahrmann, and H.-J. Böhme, "RoSe: Robot Sentinel as an Alternative for Medicinal or Physical Fixation and for Human Sitting Vigils:," in *Proceedings of the 15th International Joint Conference on Biomedical Engineering Systems and Technologies*, pp. 184–191, Science and Technology Publications, 2022.

[21] T. Georgiou, K. Singh, L. Baillie, and F. Broz, "Small robots with big tasks: A proof of concept implementation using a miro for fall alert," in *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, HRI '20, (New York, NY, USA), p. 206–208, Association for Computing Machinery, 2020.

[22] S. S. Khan and J. Hoey, "Review of fall detection techniques: A data availability perspective," *Medical engineering & physics*, vol. 39, pp. 12–22, 2017.

[23] Z. Cao, G. Hidalgo, T. Simon, S.-E. Wei, and Y. Sheikh, "Openpose: Realtime multi-person 2d pose estimation using part affinity fields," *IEEE transactions on pattern analysis and machine intelligence*, vol. 43, no. 1, pp. 172–186, 2019.

[24] M. Binois and N. Wycoff, "A survey on high-dimensional gaussian process modeling with application to bayesian optimization," *ACM Transactions on Evolutionary Learning and Optimization*, vol. 2, no. 2, pp. 1–26, 2022.

[25] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, *et al.*, "Segment anything," in *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 4015–4026, 2023.

[26] Robotemi, "Temi." `https://www.robotemi.com/`, 2025. [Online; accessed 06-Mar-2025].

[27] A. Rossi, F. Garcia, A. C. Maya, K. Dautenhahn, K. L. Koay, M. L. Walters, and A. K. Pandey, "Investigating the effects of social interactive behaviours of a robot on people's trust during a navigation task," in *Towards Autonomous Robotic Systems* (K. Althoefer, J. Konstantinova, and K. Zhang, eds.), (Cham), pp. 349–361, Springer International Publishing, 2019.

[28] K. A. Barchard, L. Lapping-Carr, R. S. Westfall, A. Fink-Armold, S. B. Banisetty, and D. Feil-Seifer, "Measuring the perceived social intelligence of robots," *ACM Transactions on Human-Robot Interaction (THRI)*, vol. 9, no. 4, pp. 1–29, 2020.