RoSe: Robot Sentinel as an Alternative for Medicinal or Physical Fixation and for Human Sitting Vigils

Robert Erzgräber, Falko Lischke, Frank Bahrmann and Hans-Joachim Böhme
University of Applied Sciences Dresden, Friedrich-List Platz 1, Dresden, Germany

Keywords: Service Robot, Sitting Vigil, Medical Restraints, Image Classification, Human-robot Interaction.

Abstract: An approach for a Robot Sentinel is described as an alternative to medicinal or physical fixation. The robot offers the opportunity to give the patient some privacy while also offering protection from falling out of bed. This approach is solely based on input data given by a Kinect One. A database with IR data with labels according to the sleep stages of the patient was generated. With given database the presented framework is able to detect the movement of the patient in bed from given input data and therefore warn the staff, if a possible harmful situation occurs. In two different experimental phases the approach could be tested and was able to successfully recognize different sleeping phases of the patient (e.g. unsettled sleep, falling asleep and wakeup phase). An unsettling sleep serves as an indication of waking up and therefore the possible desire of standing up. Recognizing those sleeping phases and counteracting this desire, preserves the patient from falling out of bed and potential injury.

1 INTRODUCTION

A sitting vigil is used in hospital and nursing home environments to observe cognitively impaired patients while sleeping or having to stay in bed, for example after surgery or if the person is endangered by being on their own. Preferably a medical student is hired as sitting vigil but in most cases, an unlearned assistant takes this position. Especially in geriatric cases many patients can loose their sense of their current location and are disoriented when waking up at night. Therefore, several methods include the use of chemicals, which keep the patient in a somnolent condition, or physical restraints to ensure a certain, safe position. To physically restrain a patient, an allowance by the Court of Protection is necessary. If waiting for the allowance would cause immediate harm, the physical restraints can be used, but an allowance has to be requested as soon as possible[1]. Sze et al. have shown in their meta-study (Sze et al., 2012) that the amount of falls, can not be associated with a lesser use of physical restraints in favor for other restriction methods. Alternative methods with their own disadvantages would be the usage of sitting vigils (shortage of staff) or chemical restraints (e.g. misuse, physical and mental side effects). While (Krüger et al., 2013) shows that physical restraints are used in standard care in hospitals of Germany, even though multiple intervention programs are aimed to reduce their usage. Either way, the privacy and the personal freedom of the patient is restricted and it should be one of the last methods to keep the patient in bed. While sitting vigils potentially have a medical background of some sort, they are not allowed to interfere, only to signal the medical staff that an emergency is about to happen or is happening.

In Germany, night-time medical staff of a neurology ward has, per law, at full capacity, 20 patients per nurse (see §6 passage 1.7 in PpUGV (Leber and Vogt, 2020)) to look after. Besides some differences between various wards in a hospital, except for intensive care units, there are more than 10 patients per medical staff to look after. Nevertheless, neither are all patients in one room nor is it possible to look after all of them at the same time. Therefore multiple ward rounds, usually four hours apart, are executed at night, where each room is visited to see whether all patients are in good condition and/or sleeping. Between these ward rounds, the patient's room is only visited if an alarm is signaled by the patient. If no alarm is signaled because, for example, the patient has fallen out of the bed and is lying unconsciously on the floor, the incident will only be discovered on the next

[1]Further informations can be seen in §1906 of the Bürgerliches Gesetzbuch of the Federal Republic of Germany.

Erzgräber, R., Lischke, F., Bahrmann, F. and Böhme, H.
RoSe: Robot Sentinel as an Alternative for Medicinal or Physical Fixation and for Human Sitting Vigils.
DOI: 10.5220/0010827100003123
In *Proceedings of the 15th International Joint Conference on Biomedical Engineering Systems and Technologies (BIOSTEC 2022) - Volume 4: BIOSIGNALS*, pages 184-191
ISBN: 978-989-758-552-4; ISSN: 2184-4305

ward round. This in combination with the rising need for medical staff - an increase of approximately 25% by 2035 is calculated for Germany alone (Sonnenburg and Schröder, 2019) - offers an opportunity to support wards with a robot that can monitor sleeping patients while the staff has the chance to care about the other ones.

The proposed framework was built and tested on the self recorded and labelled database, which was formed in two different experimental phases.

1.1 Robot

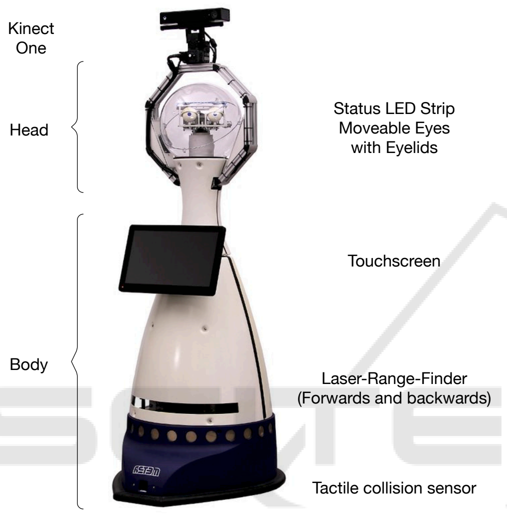

Figure 1: The anthropomorphic robot of the type Scitos G5 features a Kinect One mounted on a pan-tilt unit, a head with led lights and moveable eyes with eyelids, a touchscreen and the body itself. The body contains two laser-range-finder (forwards and backwards), a tactile sensor, and a differential drive.

The robot has a pin-shape design with multiple sensors and features built in (see Figure 1). It is based on the Scitos G5 platform from MetraLabs[2] but has several additions, like the Kinect One on top of the head and two additional speakers which are positioned nearby the head. The camera can be moved and tilted to some degree, which enables the patient to see what direction the robot currently observes. The head features no sensors built into it, but has two eyes which serves directly as interaction point for the patient and are able to keep eye contact while speaking to her or him. Furthermore the robot has a touchscreen monitor built in. During the Robot Sentinel functionality the monitor is usually powered off, but can also be used to display medical data for staff or patient alike. The body has a protective chassis to hide the computing system and sensors. In the lower body of the robot is one front and one back-facing laser-range-finder, and a tactile sensor, which stops the robot instantly as an emergency system in case an unforeseeable collision occurs. The entirety of the robot is 1.8m tall, with the head being at comfortably 1.6m and weighs 80kg. This results in a trustable platform with a human-like appearance.

[2]For further information about the platform please refer to https://www.metralabs.com/mobiler-roboter-scitos-g5/

1.2 State of the Art

Two different approaches to the presented framework that work in a similar way could be found. One is a market-ready product and a camera-based approach with edge-computing.

The camera-based system Ocuvera was created for patient monitoring with the same goal of alarming the hospital staff in advance, if the patient would leave the bed (Bauer et al., 2017). This system is mounted in a docking station at the wall and has therefore a fixed field of view where the patients bed has to be, but can be moved easily between different docking stations. The system has a display and a speaker attached to it, so that music and/or images or videos can be displayed. The difference to our proposed approach is that the robot itself has a human-like appearance with an head, so that the patient has a fixed position to speak to. Moreover, the robot itself is not immobile. The mobility of the robot can react to different positions of the bed in the room without dismounting and mounting the docking station. Even if the bed has to be moved at night the robot can position itself to have an unobstructed view of the patients bed.

Another camera-based fall protection system that uses deep learning methods is described in (Chang et al., 2021). In this case, only a camera, mounted on the wall behind the bed, and a wifi-router with edge-computing is needed. This allows the camera to directly include the birds-eye-view of the patients bed and therefore saves computing time. The router will send an alarm to a mobile device if the patient is sitting up or is in immediate danger of falling out of bed. The advantage of this system over the proposed approach is that the birds-eye-view of the camera does not have to be calculated and therefore offers a better point of view (see Section 2.5). But, like the Ocuvera system, the mobility of the proposed system is a disadvantage while a hidden camera is monitoring the patient.

2 APPROACH

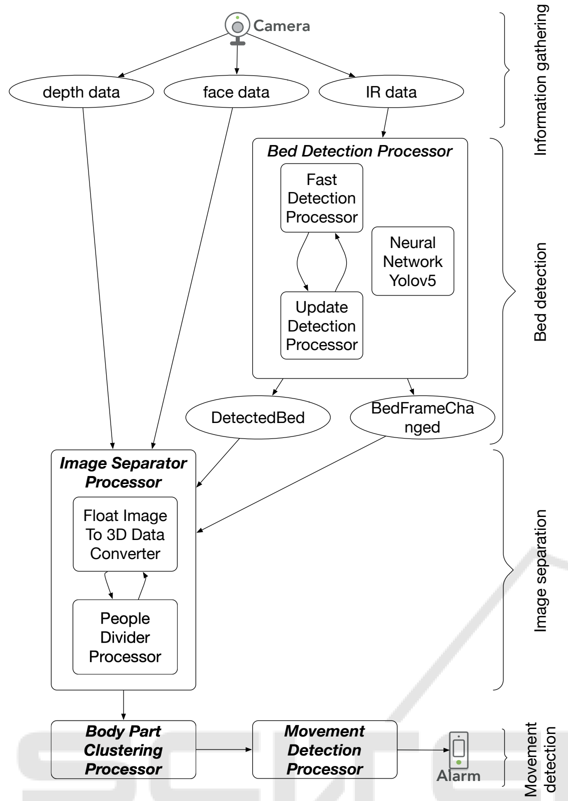

Figure 2: The data flow of the proposed approach from the camera to the warning system, which is enabled if a movement was detected. It is separated into four parts - gathering information, bed detection, image separation and movement detection.

2.1 Information Gathering

In Figure 2 the data flow from the camera to the warning signal is depicted. It starts with the information gathering by the Kinect One. The gathered information is a 640x480px infrared(IR)/depth image and, if spotted, face detections as a rectangle with hypothesis. The RGB image is not usable in this context because the setting will most likely be used at night without any light-sources, so the RGB image will stay at most times black. These information are stored separately into different containers to reduce the amount of information spread throughout the code.

2.2 Bed Detection

As input for the *Bed Detection Processor* the infrared image is used. At the beginning of the program a *Fast Detection Processor* is started, which uses every given image. The information will be sent to a trained YOLOv5S (Jocher et al., 2020) network which therefore returns a rectangle and a hypothesis. The YOLOv5S is the smallest network within the YOLO architecture and consists of 7.3m parameters with a speed of 2ms per image on a V100 GPU. The smallest network was chosen because of hardware restrictions by the robot and the fastest speed for any picture to be calculated by the neural network. An additional training with each possible architecture was implemented but showed that no significant improvement could be achieved by using any bigger YOLO network. The network was trained on a custom dataset composed of infrared data acquired within the experimentation environment and consists of 420 images of beds with different obstructions and from different angles. This dataset includes patients lying in a bed in different positions. Due to these images and their privacy protection the dataset unfortunately cannot be released to the public (see Figure 5 for an impression of the dataset). The neural network pretrained on the COCO data set was used for further training, and within 300 episodes a mean average precision of 95% was reached. For the further training the above mentioned dataset was used, with the included mnemonic data set.

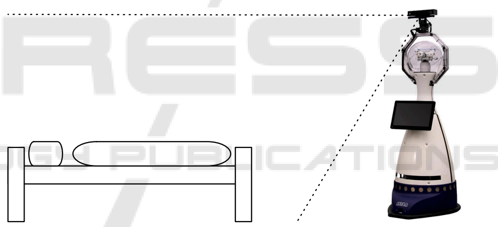

Figure 3: The robot is placed at the bottom of the patient's feet, so that it has a full view of the bed and the patient itself. The angle of aperture of the Kinect One has a wider view so that more than just the bed can be seen.

If in two succeeding frames a bed was found and the differences between the outliers of them are within a certain margin, the *Fast Detection Processor* (depicted in Figure 2) is switched to the *Update Detection Processor*. The *Update Detection Processor* will not be toggled by every image of the Kinect One but uses one image every 5 minutes to update the detected bed if needed. If the newly detected bed is outside the margin surrounding the original bed detection, the frame will be updated and then a switch to the *Fast Detection Processor* will be executed. The output of the *Bed Detection Processor* consists of two objects, one being the found quadrangle for the bed with the highest certainty and one being a value if the quadrangle was changed in this time-step. This dual processor design is due to the limited capacities of the CPU of the robot itself as it does not have any GPU. The load on the CPU should be kept low, as it produces less heat and therefore fan noise and could disturb the patients sleep. The face hypothesis is used to decide where the patient is lying. In the current experimental design the robot was placed at the feet of the patient - see Figure 3 - and therefore the face hypothesis was not needed to decide how the patient is lying in bed. In any other context, the robot can move by itself to achieve a good position for the observation and therefore has to decide where the patient is lying.

2.3 Image Separation

The inputs for the *Image Separation Processor* are the outputs of the previous *Bed Detection Processor*:

- Quadrangle of the bed
- Whether or not the quadrangle has changed
- Depth data as a matrix of distance values (float image)
- Face detection hypothesis from the Kinect One

At first the bed detection will be updated if needed and the quadrangle of the bed detection, depth data, and face detection hypothesis are sent to the *FloatImageTo3DDataConverter*. The data converter uses the given depth data image, filters the image according to the detected quadrangle and transforms it to a birds-eye-view. The position for the virtual birds-eye-view camera is in the middle of the detected quadrangle. The equations for the birds-eye-view transformation are:

$$\hat{n} = (n_1, n_2, n_3)^T \tag{1}$$

$$\vec{x} = \begin{pmatrix} depth \\ -depth * azimuthTangent \\ depth * elevationTangent \end{pmatrix} \tag{2}$$

$$\begin{aligned} M\vec{x} &= M \cdot \vec{x} \\ &= \hat{n}(\hat{n} \cdot \vec{x}) + \cos(\alpha)(\hat{x} \times \hat{n} + \sin(\alpha)(\hat{n} \times \vec{x}) \end{aligned} \tag{3}$$

For each point of the point cloud, $\hat{n}$ is the vector to the point in the current basis, α is the rotation, M is the rotation matrix and $depth$ is the length of the vector of the current point. The values for the azimuth and elevation tangents are derived from the Kinect One itself. As result a point cloud is generated with a view from above the bed and enables the robot to determine the possible elevation and movement of the patient.

The birds-eye-view of the detected quadrangle will be sent to the *People Divider Processor*. The purpose of this processor is to divide the bed into several parts matching the limbs of the patient lying in it. For this, two different ways are included. The movement of the head has little to no effect on the movement if the patient wants to stand up or slide out of bed. So, if a face detection hypothesis from the Kinect One is provided within the quadrangle containing the bed, then the image is cut to the lowest point of the detected face as upper line for the detected bed. The resulting image is then split into half, resulting in an upper and a lower body image of the patient. Both images are set as output for the *Image Separation Processor*. Using both the upper and the lower body as separate images results in the opportunity to calculate clusters in each of the images clusters at the same time, thus reducing the calculation time and resulting in multiple clusters in both of them.

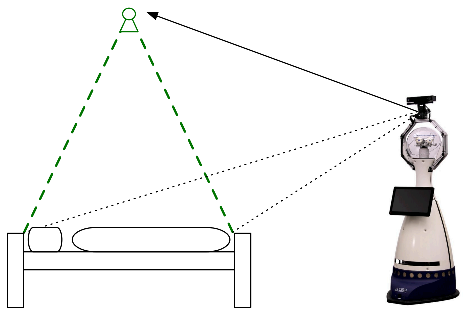

Figure 4: A virtual birds-eye-view camera will be calculated, which enables a better overview and easier gathering of information for the movement detection.

2.4 Movement Detection

The *Body Part Clustering Processor* uses both images and a k-nearest-neighbor algorithm (KNN) to cluster the given information in each image. For this part, two lists of images will be stored, one for the upper body and one for the lower body. Both lists will contain a history of 100 images. At every time-step an image (either the upper or the lower body image) is added to their list, so that the list contains 100 consecutive images. These lists have their own KNN worker that tries to find three clusters within this timeframe. The clusters are specified by μ and Σ, where μ represents the mean of the cluster and Σ the covariance matrix. To find a moving cluster between the different timeframes, the Mahalanobis distance is used. Both timeframes of images are now represented in their list of clusters.

The lists of clusters will be used in the *Movement Detection Processor*, which analyzes the given movements and decides whether a warning should be sent. If one or more clusters in either list exceeds a threshold, a warning signal is emitted. If a warning is signaled to the backend, the current infrared stream is displayed at the device for the medical staff and they have the opportunity to either dismiss the current situation as a false positive or react accordingly to save the patient from harm. As devices for the medical staff a tablet, which will be positioned in the nurses room and can therefore only be used by them, and a smartphone is supplied, which can be carried on ward rounds. If a false positive is signaled, the information will be send back to the backend and the *Movement Detection Processor* will get the notice to change the threshold to a higher amount. The threshold for movement detection is lowered over time if no movement is detected. This option is needed as not every patients movements are the same and some patients tend to sleep still while others have a more active sleep.

2.5 Constraints

Figure 5: Example for a situation that could not be reliably analyzed because most of the body is hidden behind the legs. In this situation the patient was not able to move the legs and slept with them standing up. Therefore neither a possible moving upper body could be seen nor would the lower body move.

The constraints of the proposed approach are for a typical set of movements and are not sensitive enough for a more diverse set of movements for either disabled patients or patients with a neurological pain treatment. The movement set for a patient with neurological treatment can be flatten out by medicine and would undergo the threshold that is currently set. A possibility would be that with the decline of the threshold over time a recognition is possible, but further research in such direction is needed to verify this proposition. Also a change of bed linen to a thicker one could introduce some errors as the movement seen through the linen can be obstructed and flatten out.

An unobstructed view of the bed is needed for the proposed approach to work (see Figure 5) and the bed should stand in front of the robot.

3 EXPERIMENTS

3.1 Setup

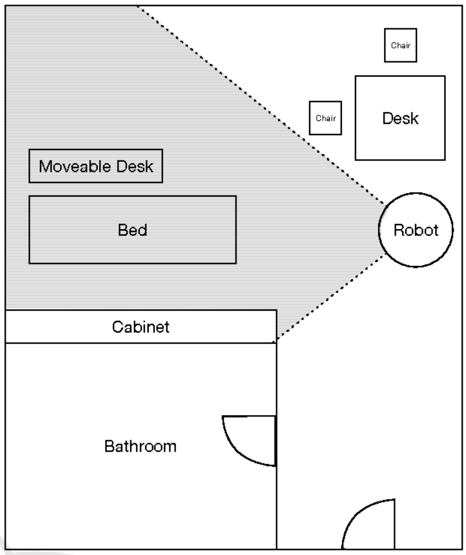

Figure 6: The patients room with the field of vision (striped gray area) of the robot. The robot is positioned to see the bed with some margin around it, but leave several blind spots in the room for the patient. While the robot is visible from every angle of the room the patient is not monitored in every angle.

The patients are lying alone in the room, where the sessions are recorded. At each session the robot will be placed at the bottom of the patient's feet facing him - depicted in Figure 3. To ensure the goal of monitoring the patient and keep her or him within the field of view, the camera setting is facing the bed so that the angle of aperture of the Kinect One aligns with a margin to it (Figure 6). There is also a gap between the bed and the robot so that the medical staff or the patient itself can go safely in front of the robot without the risk of tripping or falling over any cable. To ensure operation over a period of eight to twelve hours, the robot is attached to a power outlet, and to not further disturb the patient the cable will be attached at the start of the session. The robot is placed in the room at dinner so that the patient has the opportunity to ask questions about the robot or the recording session and to familiarize her/himself with the robot. The recording session will start at the signal of the patient when nighttime arrives. Therefore the amount of video data is reduced to the needed timespan and the privacy of the patient can be extended as long as she or he needs. For the recording only the Kinect One is active and the robot will not move at all. That ensures an undisturbed sleep and minimizes the noise that is emitted by the robot. The proposed approach can also be used as an all-day monitoring system if needed, but was currently only tested at night.

3.2 Execution

The chosen hyperparameter configuration for the YOLOv5S network utilized for the bed detection in the infrared image is listed in Table 1. This set of hyperparameters was automatically determined and optimized by utilizing the evolutionary algorithm provided by the YOLOv5 framework.

All executions of the experiments were done in two separate instances with the resulting of five nights being recorded. Each recorded session is in median eight hours and eleven minutes long, with the minimum value of six hours and forty minutes and maximum value of nine hours and fifteen minutes. Each of the recordings were seen by the authors and each movement was labeled. The possible labels were:

- sleeping still
- fall asleep phase
- unsettled sleep
- bathroom visit
- wakeup phase

These phases were chosen according to the different characteristics the robot has to chose from. The standard case would be sleeping still, where no movement could be seen and the patient is sleeping. While the patient is either falling asleep or is in a wakeup phase, the robot has to be at the utmost surveillance mode as in these phases the patient tends to move the most and has as such the highest possibility to fall out of bed. If the patient is in an unsettled sleep she or he can move but most likely will not fall out of bed, because the movements tend to switch the sleeping position. As for the last possible phase, the patient would be out of bed and the robot has a phase where it does not alter the threshold at all. In this phase the robot should only emit an emergency signal if the patient is lying on the floor, but this case was not observed in both experimental phases.

Each possible phase had a margin of maximum of five minutes to include possible not detected movements. For the training the first label - sleeping still - was removed, as it does not contain any information for the robot to learn and for the other labels a margin was included, so that a no movement phase is in them also included. This resulted in training data of sixteen hours and fifteen minutes with a typical length of about twenty to thirty minutes for each patient (depicted in Table 2). In this dataset several bathroom visits could be recorded but luckily no falling incidents were recorded.

In the first experimental phase, three patients (1-3) could be recorded with a moderate set of movements, which resulted in a basic set of information. The second phase had two patients (4 & 5) with a more unique set of movements including disability in the lower limbs and a neurological pain treatment. Both of these patients resulted in constraints for the algorithm which are described in Section 2.5.

4 EVALUATION

The recorded data are analyzed in Table 2 and are showing that a significant amount of data could be logged. The experimental phases showed that the amount of bathroom visits increased with the age of the patients and that the patient with the neurological treatment had the most unsettled sleep. It also shows that the patients tend to sleep about 8 hours and were up before breakfast was served. Both visits at night resulted in unsettling sleeps for the patients, with number 1 and 2 not being asleep when the first visit had come. In most cases just a quick checkup on the patient had been done, but in some cases a medical treatment had to be given, which resulted in a longer phase in which the patient was awake and tried to get to sleep again. The mean time per scene tended to be around 30 to 35 minutes, which included about 20 to 25 minutes off the to be recorded movements due to keep the uprising and the laying down phase within the record. This shows that in most cases the patients will have some troubled sleep before rising up and/or a bathroom visit. Therefore, for the standard visits without the need of the patient waking up, a virtual visit by the robot would significantly increase the sleep quality overall.

As no falling incident should happen under the supervision of the robot, we concentrated on recognizing the uprising movements. In most cases, if the patient wants to stand up the movement of the lower body part indicates a sliding process towards either side of the bed. It is then followed by the upper body sitting up and therefore creating a movement cluster, or in one case by tuck up one's leg. If the patient has an unsettled sleep the movement of the lower and upper body part could indicate a change of sleeping position. In these cases, the upper body will not move upward but a shoulder is raised if a sidewise sleeping position is reached. In either case the movement process is smaller and happens over a longer period of time, which results in a smaller covariance matrix Σ in the clusters. Both movement possibilities are rec-

Table 1: Hyperparameter configuration for the bed detection. The listed parameters were automatically optimized by utilizing the evolutionary algorithm provided by the YOLOv5 framework. A default YOLOv5S network was configured with these values.

param	value	param	value	param	value
lr_0	0.00855	lrf	0.193	momentum	0.88
weight_decay	0.00049	warmup_epochs	4.51	warmup_momentum	0.95
warmup_bias_lr	0.193	box	0.0541	cls	0.386
cls_pw	0.974	obj	2.23	obj_pw	1.42
iou_t	0.2	anchor_t	5.1	fl_gamma	0.0
hsv_h	0.00888	hsv_s	0.727	hsv_v	0.454
degrees	0.0	translate	0.056	scale	0.604
shear	0.0	perspective	0.0	flipud	0.0
fliplr	0.5	mosaic	0.919	mixup	0.0

Table 2: All scenes that were labeled with the exception of fall asleep and wakeup, because there were for each patient just one scene. For mean time per scene the data given are in minutes and total time asleep are given in hours.

patient	# scenes total	# unsettled sleep	# bathroom visit	mean time per scene	total time asleep
1	4	1	1	35	6:40
2	4	1	1	40	9:15
3	6	3	1	32,5	8:00
4	9	4	3	24	8:40
5	9	5	2	27	8:20

ognized by the *Movement Detection Processor* using the amount of found clusters and their mean μ and covariance matrix Σ. If either part reaches a certain threshold a warning is indicated. The threshold can be adjusted if a warning was called and dismissed by the medical staff. In this case, the threshold would be moderately increased.

With the proposed approach we were able to recognize the anticipated greater movements prior to the uprising of the patient for their bathroom visit and smaller movements which indicated an unsettled sleep. Due to finetuning, which had to be done for each patient, a threshold between both of these states could be established. On the second experiment one edge case was introduced that could not yet be reliably detected. The edge case consists of a patient who is paralyzed downwards from the hip (see Figure 5). Furthermore, we gathered data on unsettled sleep induced by neurological pain.

Additionally, differences in the amount of time taken for a patient to sit up or slide to the side of the bed could be observed. Despite those differences, in every case, at least as the patient rises, an unsettled sleep was detected. According to the proposed approach the video-stream would have been established while the patient rises. In this time the medical staff can intervene through the robot by either speaking with the patient or playing soothing music and probably slowing down the process of standing up. This would give the medical staff some time to intervene personally if needed.

After each experimental phase information from each site were gathered. This included the medical staff - in this case, nurses and doctors - and the patients. As mentioned in Figure 1, the robot has a LED strip at its head, which was only dimmable to low emission mode but could not be turned off, and in some cases, the fans of the robot could be heard. This was not a point to end the recording - which was at every point of the night possible - but was only mentioned in the talk at the morning. Each patient that mentioned either point of the above, identified it as a minor nuisance. To further reduce the fan speed of the computer, the amount of the clusters to be found by the KNN was reduced to three. It was empirically identified to be the sweet spot between noise and recognition.

The amount of images saved in the list for the *Body Part Clustering Processor* was also empirically identified to be 100, as to be in the middle of time taken to calculate the KNN and emitted noise by the fans.

5 CONCLUSION AND OUTLOOK

We have generated a database, which could not be found prior, containing different sleeping data, recorded in IR images and labelled different sleeping phases of several patients. In return the robot was able to successfully detect different sleeping phases and could monitor the patient, if she or he has an unsettled sleep or a wakeup phase was reached. With these calculated information the robot is able to warn the medical staff prior or while the intention of standing up arises and can intervene in a situation that can be possibly harmful for the patient. The robot is especially useful as it is an embodied interaction partner and can therefore be easily recognized as the one currently speaking to the patient, opposing a solely camera-based approach with attached speakers. The mobility of the robot will come in handy, as to broaden the service in observing multiple patients at once in a multi-bed room with a single device. Additionally, the robot can also be used at daytime as enhancement and assistance in the context of MAKS therapy in hospital wards, proposed in (Bahrmann et al., 2020).

A succeeding study will research on how the Robot Sentinel performs and how the patients react when the robot directly intervenes, if a possibly harmful situation is discovered. This includes a direct intervention triggered by the medical staff or automatically playing music to soothe the patient back to sleep. For this situation, the medical staff will be handed a tablet or smartphone with an application that displays the current situation and emits an alarm. The staff will have the possibility to dismiss the current situation, which will be recorded for further adjustment of the algorithm, monitoring the current situation.

As it could be seen in the second experimental phase, a possible detection of pain even while sleeping, is a possibility for the proposed approach and could serve as an early-warning system for the medical staff to intervene prior to the occurrence of an incident. Also the monitoring and recording of atypical sleeping behaviors can be useful for further diagnostics.

It was seen that the typical sleep circle of about 1.5 hours is an indicator for movements in between each cycle and was mostly used for a bathroom break. It could be possible to determine a wider range of vital signs from the patient to describe the sleep stage that she or he is currently in. This would also improve the warning process in a way to differentiate between a sleeping or an awake person.

ETHICAL STATEMENT

All human studies described have been conducted with the approval of the responsible Ethics Committee, in accordance with national law and in accordance with the Helsinki Declaration of 1975 (as amended). A declaration of consent has been obtained from all persons involved.

ACKNOWLEDGEMENT

The presented work was funded by the 'European Regional Development Funds (ERDF)' (ERDF-100293747 & ERDF-100346119). The support is gratefully acknowledged. We also want to thank all participants of this project that let us monitor them throughout their nights at the hospital and the medical staff that provided many informations to further improve the approach for the hospital use.

REFERENCES

Bahrmann, F., Vogt, S., Wasic, C., Graessel, E., and Boehme, H.-J. (2020). Towards an all-day assignment of a mobile service robot for elderly care homes. *American Journal of Nursing*, 9(5):324–332.

Bauer, P., Kramer, J. B., Rush, B., and Sabalka, L. (2017). Modeling bed exit likelihood in a camera-based automated video monitoring application. In *2017 IEEE International Conference on Electro Information Technology (EIT)*, pages 056–061.

Chang, W.-J., Chen, L.-B., Ou, Y.-K., See, A. R., and Yang, T.-C. (2021). A bed-exit and bedside fall warning system based on deep learning edge computing techniques. In *2021 IEEE 3rd Global Conference on Life Sciences and Technologies (LifeTech)*, pages 142–143.

Jocher, G., Stoken, A., Borovec, J., NanoCode012, ChristopherSTAN, Changyu, L., and et. al. (2020). ultralytics/yolov5: v3.0.

Krüger, C., Mayer, H., Haastert, B., and Meyer, G. (2013). Use of physical restraints in acute hospitals in germany: a multi-centre cross-sectional study. *International Journal of Nursing Studies*, 50(12):1599–1606.

Leber, W.-D. and Vogt, C. (2020). Reformschwerpunkt pflege: Pflegepersonaluntergrenzen und drg-pflege-split. In *Krankenhaus-Report 2020*, pages 111–144. Springer.

Sonnenburg, A. and Schröder, A. (2019). *Pflegewirtschaft in Deutschland: Entwicklung der Pflegebedürftigkeit und des Bedarfs an Pflegepersonal bis 2035*. Gesellschaft für Wirtschaftliche Strukturforschung (GWS) mbH.

Sze, T. W., Leng, C. Y., and Lin, S. K. S. (2012). The effectiveness of physical restraints in reducing falls among adults in acute care hospitals and nursing homes: a systematic review. *JBI Evidence Synthesis*, 10(5):307–351.