2026 IEEE/SICE International Symposium on System Integration (SII)
January 11-14, 2026. Cancún, México

Towards Personalized Context-aware Bedside Patient Monitoring

Kai Wu[1,2], Darius Burschka[2,3]

Abstract— **In this work, we propose a bedside monitoring system for ICU patients that integrates patient data from multiple medical devices and collects information about the patient's physical activity. Our proof-of-concept studies demonstrate the potential for more personalized bedside patient monitoring by leveraging context-aware information, thereby enabling a more robust early warning of in-hospital deterioration. Furthermore, we explore visualization techniques to present analysis results to clinical caregivers in a more intuitive and interpretable way.**

I. INTRODUCTION

Intensive care patients are in severe health conditions and require continuous monitoring of their vital signs, so that medical caregivers can react promptly to sudden deterioration. Modern intensive care units (ICUs) are usually equipped with medical devices for monitoring, diagnostics, and therapy. Typically, each patient is connected to a patient monitor that continuously measures vital signs, including heart rate, blood pressure, and oxygen saturation. A ventilator is often used to support breathing and to monitor respiratory functions. Infusion pumps are employed to deliver medications according to the prescribed treatment plan. Depending on the patient's condition, additional devices may be used, such as an enteral feeding pump for nutritional support or a dialysis machine for renal replacement therapy. These systems are typically designed to trigger alarms when measured values exceed predefined thresholds or when treatment settings require adjustment, ensuring timely attention from healthcare caregivers.

While clinically relevant alarms require clinicians to take action, clinically irrelevant alarms create an unnecessary load on the caregivers. Such alarms can be caused by artifacts, patient movements, or transient deviations from normal ranges. Sendelbach et al. stated that the false alarm rate in ICUs ranges from 72-99% depending on different patient cohorts and alarm management settings[1]. In life-critical scenarios, such as ICUs, alarms are designed with a 'better-safe-than-sorry' logic. Clinicians can tolerate a large number of false alarms rather than risk missing a valid alarm. However, with the development and integration of modern medical devices in ICUs, the number of monitored parameters is increasing,

[1]Kai Wu is with the German Heart Center Munich, Lazarettstrasse 36, 80636 Munich, Germany, wuk@dhm.mhn.de

[2]Kai Wu, and Darius Burschka are with the Machine Vision and Perception Group, Department of Computer Engineering, TUM School of Computation, Information and Technology, Technical University of Munich, Parkring 13, 85748 Garching, Germany, {k.wu,burschka}@tum.de

[3]Darius Burschka is with MIRMI - Munich Institute of Robotics and Machine Intelligence, Technical University of Munich, Georg-Brauchle-Ring 60-62, 80992 Munich, Germany

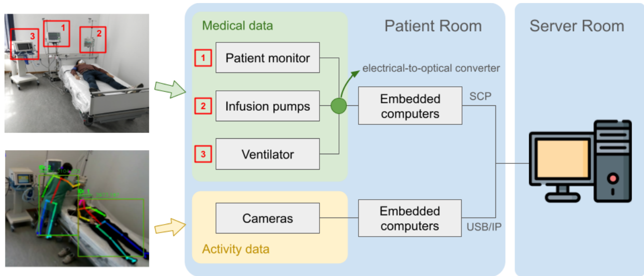

Fig. 1: System for collecting medical and activity data to monitor patients' health status.

thereby increasing the number of alarms. Cho et al. observed 45.5 clinical alarms per patient per hour from a bedside patient monitor, a mechanical ventilator, infusion pumps, and a Continuous Renal Replacement Therapy (CRRT) device[2]. The high false alarm rate leads to noise, stress for nurses, and sleep deprivation for patients. It disrupts patient care and reduces trust in alarms, resulting in a slow or no response to the alarms[3].

In this work, we propose a personalized context-aware bedside monitoring system that enhances the interpretation of changes in patient vital signs by incorporating contextual factors, such as circadian rhythms and physical activities. We aim to reduce false alarms and improve the understanding of patient health conditions. Additionally, we explore interpretation and visualization methods to enhance the presentation of patient data and facilitate the explanation of model predictions for clinical use.

II. RELATED WORK

A. *Personalized Adaption for Bedside Alarms*

To reduce false alarms, early research focused on developing more reliable alarm systems by applying advanced signal processing techniques, detecting specific temporal patterns, and leveraging correlations between multiple physiological signals[4]. In addition, various alarm strategies have been explored to mitigate noise pollution in clinical environments. These include annunciation delay and suppression of repeated abnormalities.

Charbonnier and Gentil improved the alarm filtering algorithm further with online adaptive trend extraction[5]. Depending on the physiological state of patients (e.g., asleep or awake, using controlled ventilation or not), they manually defined different sets of alarming thresholds. This relies on prior knowledge about the patient cohort, which may require customization for each hospital, and can be difficult to adapt to individual patient variability. Von Rossum et

978-1-6654-5784-2/26/$31.00 ©2026 IEEE

al. implemented and compared six adaptive threshold-based alarm strategies on heart rate, respiration rate, and axillary temperature. with consideration of individual and situational factors (e.g., surgery related conditions, physiological difference during day and night)[6].

B. Presentation and Interpretation of Patient Data Analysis

Recent advances in medical AI have enabled the development of predictive models that identify patient condition in ICUs, including in-hospital mortality, sepsis, and acute kidney injury (AKI). Leveraging large-scale datasets like MIMIC[7], eICU[8] and HiRID[9], researchers have built machine learning models that outperform traditional scoring systems by learning complex temporal patterns from electronic health records and continuous monitoring data.

However, for life-critical medical AI, although blackbox models have shown great performance in the risk estimation, they need to be understood, interpreted and validated by clinicians. Before an AI system is implemented in clinical setting, apart from technical and clinical evaluations, it is important to transfer the knowledge about what the AI system has focused its attention on through some post hoc explanations[10]. While interpretable models, such as linear models and decision trees, offer transparent insights into their decision-making processes, they have limited performance in handling the complex and multi-modal medical data for comprehensive analysis. Therefore, researchers developed model-agnostic methods such as SHapley Additive exPlanations (SHAP)[11], which interprets machine learning models by attributing the prediction of an instance to its input features. Apart from external explainer, Choi et al. designed a REverse Time AttentIoN model (RETAIN) that utilized the internal attention mechanism to detect influential past visits and significant clinical variables (key diagnoses) within those visits[12]. Instead of considering only discrete diagnosis data, Xu et al. proposed the Recurrent Attentive and Intensive Model (RAIM), which analyzes both continuous monitoring data and discrete clinical events. By incorporating a multi-channel attention layer, their model highlights segments of the temporal data that receive greater attention, which can potentially aid clinicians in verifying and understanding critical patterns[13]. Moving beyond interpreting input contributions to model outputs, Manduchi et al. learned a two-dimensional latent space of patient health state from time-series measurements through deep probabilistic clustering using Variational AutoEncoder (VAE) and Self-Organizing Map (SOM) [14]. In their experiments, they showed reasonable temporal trajectory of the patient in the latent space, offering insight into whether a patient's condition is improving or deteriorating.

State-of-the-art approaches in medical AI primarily rely on purely clinical or physiological data to perform data-driven analyses of patient health conditions. While these models have demonstrated strong predictive performance, they overlook valuable contextual information that can influence patient physiology, such as circadian rhythms, physical activities, and ongoing medication treatments. To the best of our knowledge, no existing research has incorporated these contextual factors to enhance the understanding or modeling of patient physiology, despite their potential to improve personalization and interpretability in clinical decision support.

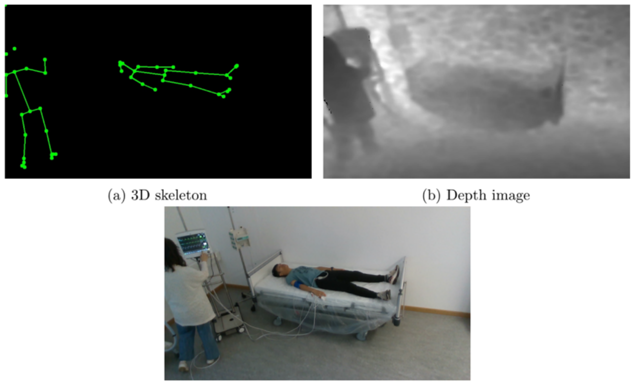

(a) 3D skeleton (b) Depth image (c) Color image

Fig. 2: Anonymized image data derived from raw camera recordings storing only depth images and extracted skeleton poses for subsequent analysis.

III. SYSTEM OVERVIEW

A. Data Collection

Setting up the data collection system in the ICU faces multiple challenges. Workstations are usually not allowed to be put in the ICU. On the one hand, such workstations may not be designed with materials that are easy to clean or resistant to harsh disinfectants used in ICUs. On the other hand, they can disrupt temperature regulation, airflow, or noise levels, which prevents patients from resting. Therefore, in our recording system, the workstation is located in the server room, and data is transferred through embedded computers mounted in the ceiling of the ICU room. Scripts on the embedded computers acquire patient measurements and device settings from bedside medical devices in the ICU and send the data, along with camera data streams, to the workstation via a secure clinic Ethernet connection (see Fig. 1).

To ensure the safety of intensive care patients, signals in the medical devices must not be interfered with while acquiring measurements from them. During our data collection procedure, we acquire data from medical devices, including a GE patient monitor, an Evita ventilator, and a B. Braun infusion station, via the serial ports. Serial ports are directly connected to electrical-to-optical converters, so that no electrical signals are passed into the medical devices. The serial data streams are requested and decoded according to the device protocol in the embedded computers, and then sent to the workstation via Secure Copy Protocol (SCP).

To observe patient physical activity, three RGBD cameras (Realsense D435) are mounted on the ceiling, filming the patient from different viewpoints to avoid occlusion. Cameras are connected to embedded computers via USB3-A ports, streaming RGB and depth frames to the workstation with USB/IP service. A bandwidth of approximately 900

Mbps is required to stream data for each camera in real-time. After receiving the image streams, pose estimation is performed online on the workstation, where the human pose is extracted as a skeleton (comprising joints and connecting bones) from the RGB images. To preserve patient privacy, only the recognized poses and depth image are saved for further analysis (see Fig. 2).

B. *Skeleton-based Activity Recognition*

We use OpenPose[15] to extract human pose as skeleton-structured information, which is represented by joints and bones. A two-stream attention-based Graph Convolutional Network[16] is adopted to recognize actions. It takes the joint and bone streams as two separate pathways and processes them with graph neural networks. Instead of setting a fixed skeleton graph, the model learns to adapt the graph structure during training. This helps the network focus more on important joints and their relationships for a given action.

C. *Context-aware Personalized Patient Monitoring*

From our previous study, different states of physical activity cause notable transient changes in vital signs. For example, lying down from a sitting position can cause a decrease in the heart rate of 20 bpm[17]. It can drop further when the person falls into a deep sleep state, which may trigger a false alarm in the patient monitor. With patient activity recognized by the camera systems, we provide the possibility of performing activity-aware bedside patient monitoring, allowing the alarm threshold to be adjusted based on different activities and thereby reducing false alarms caused by patient activities.

On the other hand, how vital signs respond to physical activity reveals, to some extent, the health condition of patients. Our system offers the opportunity to conduct continuous stress tests for patient movement. Advanced bedside ventilator provides the possibility to measure oxygen consumption (VO_2), with which the Metabolic Equivalent of Task (MET) can be estimated by

$$\text{MET} = \frac{VO_2[ml/kg/min]}{3.5}. \quad (1)$$

It measures the level of exertion required to perform an activity. Using the proposed system, we can analyze transient changes in patients' physiological signals, for example, how rapidly and to what extent the heart rate increases at the onset of activity, whether it stabilizes during the activity, and how quickly it returns to baseline levels afterward. Such dynamic responses have been traditionally leveraged in cardiac stress testing and may offer valuable insights into a patient's recovery progress.

In addition to the transient physiological changes, long-term patterns can also provide valuable insights into the patient's overall health. By monitoring the sleep–wake cycle of a patient, we can analyze physiological signals over longer timescales, capturing rhythmic patterns. Ben-Dov found that sleeping heart rate and the heart rate dip possessed a linear relationship to hazard ratios for all-cause mortality[18]. Our previous work utilized Gaussian Process model to learn personalized circadian rhythm patterns and detect rhythm-aware anomalies. With the extracted information, we observed a noticeable improvement in mortality prediction[19].

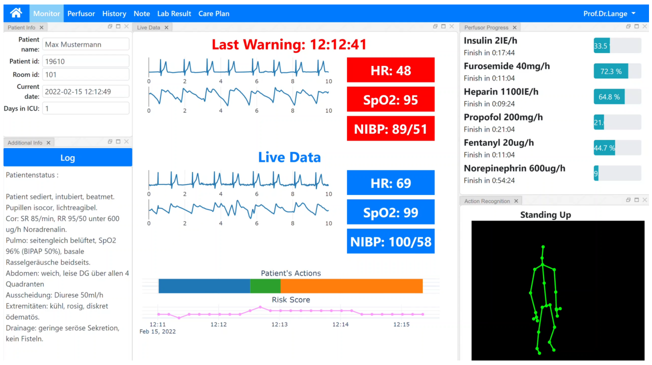

Fig. 3: Web-based visualization providing an overview of patient data. Data from multiple sources is synchronized and summarized to extract essential information from the overwhelming data streaming.

D. *Visualization and Interpretation of patient data analysis*

With the collected patient data and model analyses, we intend to provide a clear and intuitive overview of patient data for clinical caregivers. In current ICU settings, bedside measurements are often distributed across various devices. An overview of the patient's complete profile can be accessed from computers with a Patient Data Management System (PDMS), which is usually located in the corresponding nurse stations. To enable more flexible and mobile access, we developed a web-based platform, allowing clinicians to access complete patient data on the move. In addition, we explore ways to incorporate contextual information into existing clinical data to enable a more comprehensive understanding of patient health conditions. Fig. 3 shows an example of an overview dashboard of the data we want to present. Besides medical data, we visualize the current activity of the patient (see the skeleton plot in the lower right corner), the activity history, and the estimated risk score from risk estimation models (see lower middle plots).

Alongside the website dashboard, we developed an application on the HoloLens Augmented Reality (AR) headset to explore alternative methods for patient data presentation (see Sec. IV-B).

IV. Results

Initial versions of the various modules have been implemented and tested as a proof of concept and will be further refined for readiness in field studies. In this section, we demonstrate the functionality of the individual modules that have been implemented so far.

A. *Context-aware Physiological Modeling*

In our preliminary studies, we investigated how vital signs are influenced by physical activities, circadian rhythm.

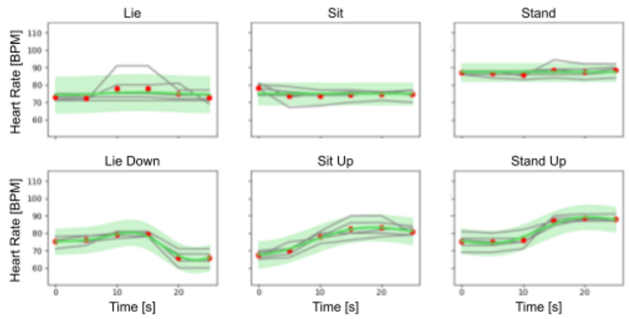

(a) Example of physiological patterns from a single volunteer during different types of movements.

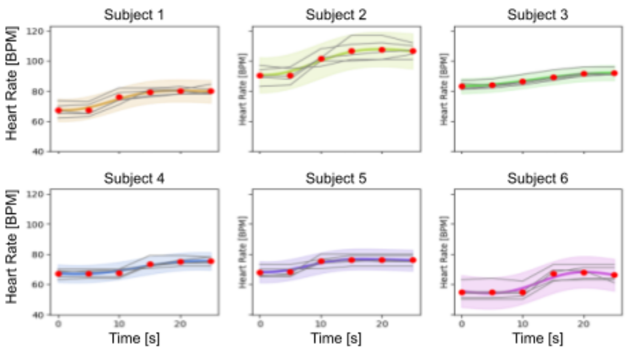

(b) Example of physiological patterns from six volunteers (subjects) during the standing-up movement.

Fig. 4: Correlation between physical activity and heart rate. The gray lines are the raw measurements, red dots indicate the mean values, and the colored regions show the data variance.

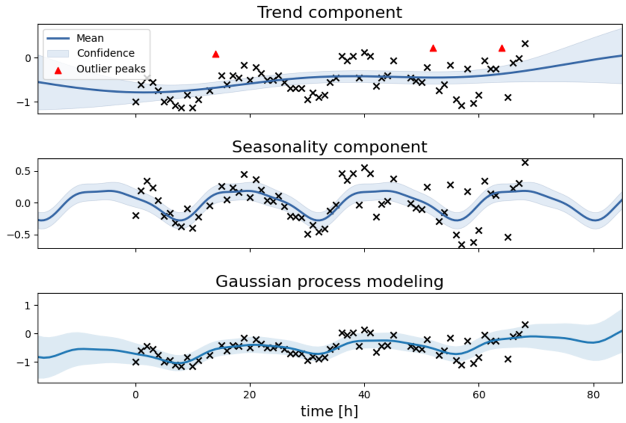

Fig. 5: Example of circadian pattern in heart rate measurements from one patient. From the raw measurements, day-night fluctuation and an overall upward trend are extracted with Gaussian Process models.

1) Activity-aware Physiological Modeling: For the development of activity-aware physiological models, we collected data from 10 volunteers in a mock-up ICU room. Each participant was instructed to perform a set of 10 different actions, including lying, sitting, standing, walking, sitting up, and standing up. Each action was repeated five times. As they are performing the actions, two data streams were recorded simultaneously. First, an Intel RealSense D435 camera was mounted on the ceiling to capture video footage of the participants performing the actions. These recordings were used to train and evaluate the action recognition model. Second, each participant was connected to a commercially available GE CARESCAPE B650 monitor, which is commonly used in hospitals to track patient vital signs. We recorded three key vital signs at a sampling interval of 5 seconds: heart rate, oxygen saturation, and non-invasive blood pressure.

From the recorded data, we observed fast responses in the vital signs, with signals typically stabilizing within approximately 30 seconds across all activities. These observations highlight two key aspects: (1) the distinct temporal patterns induced by different physical activities (see Figure 4a), and (2) the variability in baseline physiological values across individuals (see Figure 4b). As demonstrated in our previous work[17], by modeling these patterns using Gaussian Process models, we show the possibility to personalize the model to a new patient robustly with only a few measurements, using a Kalman filter with dynamic time warping metrics. Furthermore, this approach enables the prediction of vital sign responses to specific activities, given baseline measurements during lying, sitting, and standing.

2) Rhythm-aware Physiological Modeling: To analyze the circadian rhythm pattern, there is less need for high time resolution. Therefore, we chose the MIMIC-IV dataset[7] to evaluate our idea, where vital signs are recorded every hour. After cohort selection and data cleaning, we end up with 33778 unique ICU admissions and 27952 unique patient subjects. The number of ICU survivors is 24483 (87.6%). For each patient, we first fit a Gaussian Process model to each vital sign signal. We designed the kernel function as a combination of a periodic kernel and a radius basis function. The former is constrained to have a repetitive pattern with a period of 22 to 26 hours (see seasonality component in Fig. 5), whereas the latter aims to capture the long-term trend in the vital signs (see trend component in Fig. 5). From the temporal signals, we extract three key features: the learned periodic pattern, the long-term trend, and abnormal measurements that fall outside the confidence bounds of the Gaussian Process model. These features, combined with static patient information, are then used as inputs to mortality prediction models. As shown in our previous work[19], incorporating rhythm-aware information improves the performance of mortality and length-of-stay prediction models obviously.

B. Patient Data Presentation

In this section, we introduce the visualization and interpretation methods we developed for the AR headset. Our preliminary survey revealed that medical professionals in different roles have distinct priorities regarding the patient data they focus on. To address this, we designed two role-specific dashboards that offer a comprehensive yet targeted

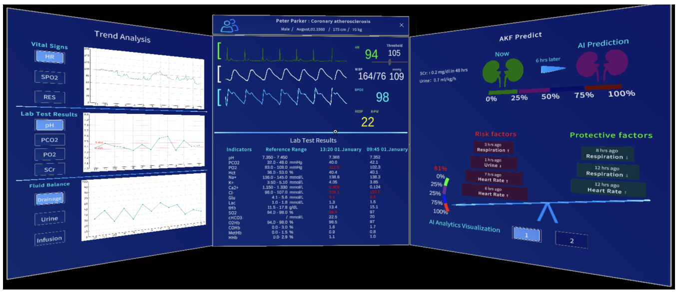

Fig. 6: Role-specific patient data dashboard customized for doctors[20], providing an overview of live monitoring data, past physiological trends, lab results, and warnings from risk prediction models.

Fig. 7: Flexible layout of different data modules allowing users to call up past measurements and obtain a comprehensive overview to support diagnostic decisions.

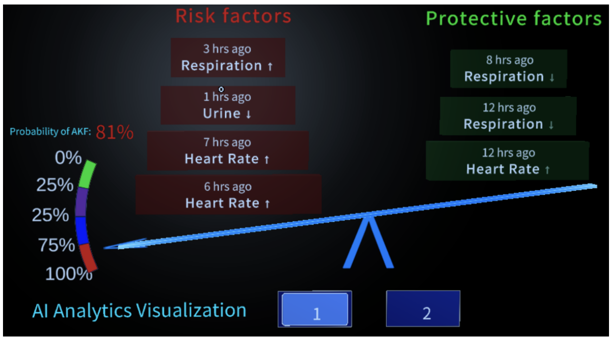

(a) Physiological trends that contribute to a positive prediction of acute kidney failure.

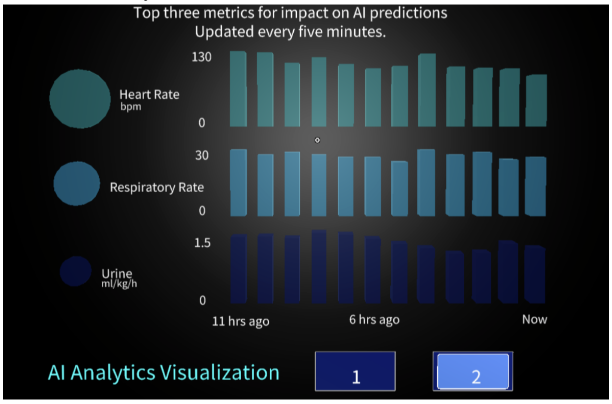

(b) Detailed measurements of the top three signals that support a positive prediction.

Fig. 8: Visual interpretation for acute kidney failure prediction[20].

overview of relevant information.

Figure 6 presents an example dashboard tailored for doctors. In addition to displaying real-time bedside vital sign measurements, it visualizes historical trends in physiological signals and lab test results, with abnormal values clearly highlighted. The dashboard also integrates early warnings generated by AI models. As shown on the right side of the figure, the system indicates that the patient is at risk of developing acute kidney failure. In comparison, the nurse dashboard presents less information on physiological trends and lab results, but includes an intuitive interface for recording patient assessments using gesture or voice control, such as the Glasgow Coma Scale (GCS) and the Richmond Agitation-Sedation Scale (RASS).

In addition to the pre-designed dashboards, users can display several data modules and arrange them freely using gesture control. This flexible layout allows clinicians to place relevant data side by side for easier comparison and to support their diagnostic decisions (see Fig. 7).

To enhance the interpretability of the AI model outputs, we designed graphical explanation of SHAP outputs, allowing us to quantify and visualize the contribution of each feature to the model's decision-making process. Using the SHAP importance scores, we show the vital sign changes that most strongly influence a positive prediction of acute kidney failure (see Fig. 8a), along with detailed trends in the patient's vital signs over the past twelve hours (see Fig. 8b). This information is intended to guide clinicians in identifying abnormal patterns and applying their medical expertise to assess whether the patient is truly at risk.

V. Conclusion

In this work, we introduce a pipeline from data collection systems in the ICU to patient data analysis and visualization using modern techniques. We presented a bedside patient monitoring system designed to capture both physiological measurements and patient activity. By modeling physiological patterns influenced by physical activity and circadian rhythms, we proposed algorithms for the personalized, context-aware adaptation of normal physiological signal ranges. This approach has the potential to reduce false alarms caused by external activities in conventional threshold-based alarm systems. Furthermore, the characterization of transient physiological responses to exertion levels and circadian rhythms offers valuable insights into a patient's overall health status. In our previous work, we demonstrated that incorporating context-aware physiological changes and adaptive normal ranges can enhance predictive models for patient health conditions[19]. Nevertheless, due to the limitations of existing medical datasets, further validation and evaluation of activity-aware physiological models are necessary in future research. Finally, we deployed a website interface and an application on the HoloLens headset, aiming to provide a more intuitive and user-friendly interface for clinicians to access and analyze the complex patient data. We explored visualization methods tailored to the needs of different clinical roles and scenarios, enhancing the interpretability of the analysis results from our early warning models.

ACKNOWLEDGMENT

The project was funded by the Bavarian State Ministry of Science and Arts within the framework of the "Digitaler Herz-OP" project under the grant number 1530/891 02. We gratefully acknowledge the close collaboration and support from the German Heart Center. We also thank Prof. Gudrun Klinker for her valuable contributions to the development of augmented reality visualization concepts.

REFERENCES

[1] Sue Sendelbach and Marjorie Funk. Alarm fatigue: a patient safety concern. *AACN advanced critical care*, 24(4):378–386, 2013.

[2] Ok Min Cho, Hwasoon Kim, Young Whee Lee, and Insook Cho. Clinical alarms in intensive care units: Perceived obstacles of alarm management and alarm fatigue in nurses. *Healthcare informatics research*, 22(1):46–53, 2016.

[3] Kelly Creighton Graham and Maria Cvach. Monitor alarm fatigue: standardizing use of physiological monitors and decreasing nuisance alarms. *American Journal of Critical Care*, 19(1):28–34, 2010.

[4] Michael Imhoff and Silvia Kuhls. Alarm algorithms in critical care monitoring. *Anesthesia & Analgesia*, 102(5):1525–1537, 2006.

[5] Sylvie Charbonnier and Sylviane Gentil. On-line adaptive trend extraction of multiple physiological signals for alarm filtering in intensive care units. *International Journal of Adaptive Control and Signal Processing*, 24(5):382–408, 2010.

[6] Mathilde C van Rossum, Lyan B Vlaskamp, Linda M Posthuma, Maarten J Visscher, Martine JM Breteler, Hermie J Hermens, Cor J Kalkman, and Benedikt Preckel. Adaptive threshold-based alarm strategies for continuous vital signs monitoring. *Journal of clinical monitoring and computing*, pages 1–11, 2022.

[7] Alistair Johnson, Lucas Bulgarelli, Tom Pollard, Steven Horng, Leo Anthony Celi, and Roger Mark. Mimic-iv. *PhysioNet. Available online at: https://physionet. org/content/mimiciv/1.0/(accessed August 23, 2021)*, 2020.

[8] Tom J Pollard, Alistair EW Johnson, Jesse D Raffa, Leo Anthony Celi, Roger G Mark, and Omar Badawi. The eicu collaborative research database, a freely available multi-center database for critical care research. *Scientific Data*, 5:180178, 2018.

[9] Stephanie L Hyland, Martin Faltys, Matthias Hüser, Xinrui Lyu, Thomas Gumbsch, Cristóbal Esteban, Christian Bock, Max Horn, Michael Moor, Bastian Rieck, et al. Early prediction of circulatory failure in the intensive care unit using machine learning. *Nature medicine*, 26(3):364–373, 2020.

[10] Sobhan Moazemi, Sahar Vahdati, Jason Li, Sebastian Kalkhoff, Luis JV Castano, Bastian Dewitz, Roman Bibo, Parisa Sabouniaghdam, Mohammad S Tootooni, Ralph A Bundschuh, et al. Artificial intelligence for clinical decision support for monitoring patients in cardiovascular icus: a systematic review, 2023.

[11] Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30, 2017.

[12] Edward Choi, Mohammad Taha Bahadori, Jimeng Sun, Joshua Kulas, Andy Schuetz, and Walter Stewart. Retain: An interpretable predictive model for healthcare using reverse time attention mechanism. *Advances in neural information processing systems*, 29, 2016.

[13] Yanbo Xu, Siddharth Biswal, Shriprasad R Deshpande, Kevin O Maher, and Jimeng Sun. Raim: Recurrent attentive and intensive model of multimodal patient monitoring data. In *Proceedings of the 24th ACM SIGKDD international conference on Knowledge Discovery & Data Mining*, pages 2565–2573, 2018.

[14] Laura Manduchi, Matthias Hüser, Martin Faltys, Julia Vogt, Gunnar Rätsch, and Vincent Fortuin. T-dpsom: An interpretable clustering method for unsupervised learning of patient health states. In *Proceedings of the Conference on Health, Inference, and Learning*, pages 236–245, 2021.

[15] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7291–7299, 2017.

[16] Lei Shi, Yifan Zhang, Jian Cheng, and Hanqing Lu. Two-stream adaptive graph convolutional networks for skeleton-based action recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12026–12035, 2019.

[17] Kai Wu, Ee Heng Chen, Xing Hao, Felix Wirth, Keti Vitanova, Rüdiger Lange, and Darius Burschka. Adaptable action-aware vital models for personalized intelligent patient monitoring. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 826–832. IEEE, 2022.

[18] Iddo Z Ben-Dov, Jeremy D Kark, Drori Ben-Ishay, Judith Mekler, Liora Ben-Arie, and Michael Bursztyn. Blunted heart rate dip during sleep and all-cause mortality. *Archives of internal medicine*, 167(19):2116–2121, 2007.

[19] Kai Wu, Ee Heng Chen, Felix Wirth, Keti Vitanova, Rüdiger Lange, and Darius Burschka. Risk estimation for icu patients with personalized anomaly-encoded bedside patient data. In *2023 45th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, pages 1–5. IEEE, 2023.

[20] H. S. Wang. Utilizing augmented reality headset for task-oriented patient monitoring in intensive care units. Master's thesis, Technical University of Munich, Germany, 2024.