

IoT-PMA: Patient Health Monitoring in Medical IoT Ecosystems

Ariane Ziehn^B, Christian Mandel^B, Kathrin Stich^C, Rolf Dembinski^C,
Karin Hochbaum^C, Steffen Zeuch^{A,B}, Volker Markl^{A,B}

^ATechnische Universität Berlin, Straße des 17. Juni 135, 10623 Berlin, Germany,
{firstname.lastname}@tu-berlin.de

^BDFKI GmbH, Trippstadter Str. 122, 67663 Kaiserslautern, Germany {firstname.lastname}@dfki.de

^CGesundheit Nord gGmbH, Klinikverbund Bremen, 28211 Bremen, Germany.
Kathrin.Stich@klinikum-bremen-mitte.de, {firstname.lastname}@gesundheitnord.de

ABSTRACT

The emergence of the Internet of Things (IoT) and the increasing number of cheap medical devices enable geographically distributed healthcare ecosystems of various stakeholders. Such ecosystems contain different application scenarios, e.g., (mobile) patient monitoring using various vital parameters such as heart rate signals. The increasing number of data producers and the transfer of data between medical stakeholders introduce several challenges to the data processing environment, e.g., heterogeneity and distribution of computing and data, low-latency processing, as well as data security and privacy. Current approaches propose cloud-based solutions introducing latency bottlenecks and high risks for companies dealing with sensitive patient data. In this paper, we address the challenges of medical IoT applications by proposing an end-to-end patient monitoring application that includes NebulaStream as the data processing system, an easy-to-use UI that provides ad-hoc views on the available vital parameters, and the integration of ML models to enable predictions on the patients' health state. Using our end-to-end solution, we implement a real-world patient monitoring scenario for hemodynamic and pulmonary decompensations, which are dynamic and life-threatening deteriorations of lung and cardiovascular functions. Our application provides ad-hoc views of the vital parameters and derived decompensation severity scores with continuous updates on the latest data readings to support timely decision-making by physicians. Furthermore, we envision the infrastructure of an IoT ecosystem for a multi-hospital scenario that enables geo-distributed medical participants to contribute data to the application in a secure, private, and timely manner.

TYPE OF PAPER AND KEYWORDS

Short communication: *IoT data management, patient monitoring, healthcare*

1 INTRODUCTION

The Internet of Things presents a novel computing

architecture for data processing systems: a distributed, highly dynamic, and heterogeneous environment of massive scale [27]. The healthcare sector is one emerging IoT application area with a high potential for economic and social impact [4, 14]. Future IoT scenarios must cope with a dynamic and geographically distributed infrastructure where patients wearing medical sensors

This paper is accepted at the *International Workshop on Very Large Internet of Things (VLIoT 2022)* in conjunction with the VLDB 2022 conference in Sydney, Australia. The proceedings of VLIoT@VLDB 2022 are published in the Open Journal of Internet of Things (OJIOT) as special issue.

are monitored inside and outside hospitals, and data exchange between hospitals and various stakeholders is required [22, 8]. KI-SIGS (*AI space for intelligent health systems*) [22] investigates such a geographically distributed healthcare ecosystem with several application scenarios. In these scenarios, IoT devices drive healthcare systems of massive scale, e.g., medical devices improve monitoring in the medical domain by continuously measuring various vital parameters such as heart rate signals. These potentially mobile IoT devices introduce several challenges to the data processing environment, e.g., heterogeneity and distribution of computing and data, low-latency processing, as well as data security and privacy. The current cloud-based solutions are not suitable for such scenarios: (1) cloud-based solutions collect data centrally before processing which introduces a high risk for companies that deal with sensitive patient data, and (2) the huge number and geo-distributed locations of connected data producers significantly affect the average processing latency [27] and thus violate the critical low-latency requirements of patient monitoring scenarios.

In earlier work, we proposed NebulaStream [27], a novel general-purpose, end-to-end data processing system for IoT applications. NebulaStream addresses the heterogeneity, unreliability, and scalability challenges of the IoT by leveraging an architecture that unifies the advantages of fog and cloud environments. Thus, NebulaStream tackles (1) security and privacy challenges by allowing healthcare organizations to leverage the ownership benefits of private fogs within the ecosystem. In order to overcome (2) the latency challenge, the unified architecture of NebulaStream enables data processing close to the source. It provides in-network processing to scale for an increasing number of data producers and queries. To support a rich set of application scenarios from different fields, NebulaStream provides a user interface (NebulaStream-UI) [20] with various functionalities, such as adding streams or queries and several visualization types for query results. In summary, NebulaStream is a promising solution for both universal patient monitoring applications and the emerging large-scale IoT ecosystem that contains healthcare organizations.

From an application perspective, patient monitoring requires ad-hoc and continuous views on many vital parameters either measured continuously by medical devices, e.g., heart rate signals, or generated in asynchronous processes, e.g., laboratory values. Furthermore, mathematical models, such as scoring systems, exist that define patients' health state indicators on a subset of relevant vital parameters. Several models, e.g., APACHE [28] or TISS [16], have been established for health economics and quality assurance

in daily hospital routines. However, there are no standards for unifying the available vital parameters or the indicators of mathematical models, focusing on ease of use and ad-hoc views on data to support physicians [1]. As a result, physicians manually analyze various data for monitoring, diagnosing, and treating patients. The vast number of vital parameters and their complex dependencies increase the risk of overlooking or misjudging the early stages of critical events, which leads to life-threatening situations. By estimation, hospital errors are the third-leading cause of patient death in the USA, after heart diseases and cancer [11]. Monitoring applications can help reduce hospital errors by supporting physicians in their analysis tasks and providing ad-hoc and continuous views on this valuable information.

In this paper, we propose IoT-PMA, an end-to-end solution for patient monitoring that includes NebulaStream as data processing system, an easy-to-use UI that provides ad-hoc views on the available vital parameters, and the integration of ML models to enable predictions on the patients' health state. In particular, we leverage NebulaStream's features for the timely and accurate acquisition of relevant health indicators and its rich set of streaming operators to define monitoring queries. We extend the NebulaStream-UI with a patient overview page as entry point for a clinic and a detail page for each patient. The overview page contains general information about the patients combined with a traffic light system to reflect the patients' health state. For each patient on the overview page, we provide a detailed page that visualizes the readings of the patient's vital parameters (time-series data streams) and static demographic information. To evaluate our solution, we implement the *RIDIMP* scenario from the KI-SIGS project to enable a continuous patient monitoring application based on two novel scoring systems, i.e., for hemodynamic and pulmonary decompensation. To this end, we leverage machine learning (ML) techniques to predict the patients' health state, i.e., the score value. Beyond that, NebulaStream's scalability and unique features will allow us to extend our solutions in the future with further monitoring tasks for other diseases and vital parameters to support universal end-to-end patient monitoring applications. The remainder of this paper is structured as follows: We introduce a real-world patient monitoring scenario in Sec. 2. In Sec. 3, we present our end-to-end patient monitoring solution IoT-PMA. In Sec. 4, we outline our vision of an emerging IoT ecosystem focusing on healthcare. We discuss challenges of clinical studies and our lessons learned to verify patient monitoring applications for usage in hospitals in Sec. 5. We present related work in Sec. 6. Finally, we conclude our findings in Sec. 7.

2 PATIENT MONITORING SCENARIO

In this section, we describe the real-world patient monitoring task introduced by the KI-SIGS [22] project *Risk Indicators for cardiopulmonary Decompensation in Intensive care units by Monitoring vital Parameters (RIDIMP)*, its clinical relevance, and the data used.

Clinical Relevance: The project *RIDIMP* focuses on detecting hemodynamic and pulmonary decompensations by monitoring patients' vital parameters. These decompensations are progressive, dynamic lung and cardiovascular function deteriorations. With increasing failure of physiological counter-regulation or delayed medical intervention, organ failure and death are the medical endpoints of these life-threatening events. Intensive care physicians have to deal with a large number of parameters like vital signs, laboratory values, ventilator settings, and applied medication. Additionally, patients in intensive care units are in critically ill conditions, and hemodynamic and pulmonary decompensations can occur suddenly and at any time during intensive care treatment. The etiology and cause of these incidents can be very different and diverse, which introduces an additional challenge for intensive care physicians to detect early decompensation stages in the vast number of parameters. Furthermore, this overwhelming amount of information and its dependencies increases the potential risk that beginning decompensation stages are misjudged, which can cause life-threatening health situations for the patient. Early detection of these events is essential for physicians to intervene therapeutically and improve survival. Common surveillance systems raise alarms if one or more parameters exceed a critical value. However, they neither support the differentiation of vital parameter value ranges nor ad-hoc views on the entire relevant data. A robust monitoring and warning system that visualizes vital parameters and computation results of the developed methods to detect, classify and predict the severity class of decompensation can support the physicians. A suitable user interface would visualize the computed results to everyone involved in the patient's treatment. The bedside physician can be informed directly about a potential risk of hemodynamic or pulmonary decompensation, may intervene therapeutically, and arrange a prolonged observation episode or invasive monitoring and diagnostics.

Patient Data Management System: The data used for developing the scoring systems and training the neural network is stored in the *IntelliSpace Critical Care*

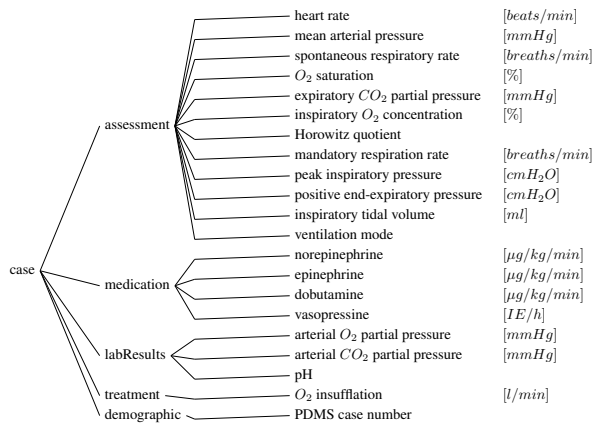
and *Anesthesia Data* management system, developed by Philips [21]. An excerpt of the data containing more than 10.000 cases recorded between 2013 and 2021 had been anonymized. The responsible ethical committee approved the retrospective observational study. Each case, i.e., a patient data record, contains the following data groups: assessment, medication, lab result, treatment, and demographic information.

Data groups provide data entries, such as time series of observations made during the patient's hospital stay. Data are recorded with different sampling rates, e.g., high frequent and regular measurements of vital parameters, e.g., heart rate and arterial pressure. At the same time, the asynchronous data recording process, e.g., laboratory values, depends on the patient's therapy and clinical condition. After the patient is discharged, the data are compressed to a maximal sample frequency of one tuple per hour (per data entry). All methods introduced in Section 3 are developed based on the compressed anonymized historical patient data.

3 IOT-PMA: PATIENT MONITORING APPLICATION IN THE IOT

Patient monitoring is one major emerging IoT application [3]. In order to enable such an application, we propose our end-to-end solution IoT-PMA that includes NebulaStream as data processing system, an easy-to-use UI with ad-hoc views on the available vital parameters, and the integration of ML models to enable predictions on the patients' health state. NebulaStream addresses the challenges of the IoT by leveraging an architecture that unifies the advantages of fog and cloud environments. In particular, NebulaStream tackles security and privacy challenges by allowing healthcare organizations to leverage the exclusive resource usage for private fog owners in contrast to the shared usage of cloud resources. In order to overcome the challenge of low-latency requirements, the unified architecture of NebulaStream enables data processing close to its source. Furthermore, it supports diverse data and programming models to support various monitoring queries and provides flexible scalability for increasing numbers of data producers and queries.

In the remainder of this section, we summarize the major methods used to implement IoT-PMA, i.e., the scoring systems to classify decompositions (Section 3.1) and their prediction model (Section 3.2). Furthermore, we explain the integration of the scoring systems and the prediction model into NebulaStream in Section 3.3, the monitoring extension on the NebulaStream-UI in Section 3.4, and the application deployment for the retrospective observational study in Section 3.5.


Figure 1: Monitoring Parameters by Data Groups.
Table 1: Severity Class of Decompensation and Score Intervals.

severity class of decompensation	hemodynamic score interval	pulmonary score interval
none	0 - 3	0 - 4
moderate	4 - 5	5 - 20
severe	>5	>20

3.1 Scoring Systems for Hemodynamic and Pulmonary Decompensation

The fundamental first step to overcoming the medical challenge of identifying vital parameters and parameter dependencies requires defining new scoring systems for hemodynamic and pulmonary decompensation. In order to capture the early stages of decompensation, avoid focusing on special etiology, and minimize data gaps, we chose routinely and frequently used parameters and severity thresholds and had to renounce special, invasive laboratory results, vital signs, and procedures. In total, 20 relevant parameters have been identified as vital risk parameters for the two decompensation types. Figure 1 [13] summarizes all parameters by their data group and provides their measurement unit.

The final medical definition is described in two novel scoring systems proposed by Stich et al. [24]. These systems define various value ranges for the patient's monitoring parameters which are punished with different point values, i.e., 0 points (pts) for normal conditions and up to 3 pts (4 pts) for pulmonary (hemodynamic) decompensation for massively deviant behavior. The hemodynamic decompensation score is based on six identified monitoring parameters, including the two vital parameters, mean arterial pressure and heart rate, combined with circulatory drug support. The pulmonary decompensation score takes 14 monitoring parameters into account, inter alia, it includes the vital parameters

Table 2: Pulmonary Decompensation Scores for Identified Parameter and their Value Ranges.

parameter:	0	1	2	3
spontaneous respiratory rate	10-25	26-30	31-35	>35
O ₂ saturation	96-100	95-90	85-89	<85
expiratory CO ₂ partial pressure	35-45	30-34 46-49	25-29 50-58	<25 >59
arterial O ₂ partial pressure	70-100	69-65	64-60	<60
arterial CO ₂ partial pressure	35-45	30-34 46-49	25-29 50-58	<29 >59
pH	7.35-7.45	7.46-7.49	7.5-7.55	>7.55
inspiratory O ₂ concentration	30-35	36-49	50-60	61-100
O ₂ insufflation	0	2-5	6-8	>8
Horowitz quotient	400-600	399-200	199-100	<100
mandatory respiration rate	10-20	21-23	24-26	>26
peak inspiratory pressure	10-25	26-28	29-30	>31
positive end-expiratory pressure	5-8	9-11	12-15	16-25
inspiratory tidal volume	401-500	301-400	201-300	<200
ventilation mode	spontaneous breathing	oxygen insufflation	assisted spontaneous breathing	bivent

Table 3: Hemodynamic Decompensation Scores for Identified Parameter and their Value Ranges.

parameter:	0	1	2	3	4
heart rate	50-90	45-49 91-100	40-44 101-110	40-44 101-110	<40 >110
mean arterial pressure	65-80	64-60	59-50	59-50	<50
catecholamine therapy	none	singular	singular	combined	singular/combined in high dose
norepinephrine	0	0.01-0.09	0.1-0.39	0.1-0.39	>0.4
epinephrine	0	0.01-0.09	0.1-0.39	0.1-0.39	>0.4
dobutamine	0	1-3	3.1-5	3.1-5	>5
vasopressin	0	0	0	0	>0.01

spontaneous respiratory rate and peripheral oxygen saturation, as well as different laboratory values taken from blood gas analysis and oxygenation or support by a ventilator. Tables 2 and 3 provide the complete list of identified parameters, their value ranges, and assigned scores for the two decompensation types [13].

The final score is the sum of all points from one patient's information at a specific point in time and indicates the patient's severity class of decompensation, as given in Table 1.

3.2 Prediction Model for Severity Classes

In order to enable not only monitoring but also warnings for evolving decompensations, a gated recurrent unit-based (GRU) neural network has been developed by Mandel et al. [13]. The TensorFlow [25] network has been trained to predict the maximal severity class of the decompensation within a 24-hour prediction timeframe with 0.85 AUROC for hemodynamic decompensation and 0.9 AUROC for pulmonary decompensation. The network can estimate the underlying decompensation score for prediction times of up to 24 hours with mean errors of 6.3% of the maximal possible pulmonary, and 9.6% of the hemodynamic score based on 60h observation period.

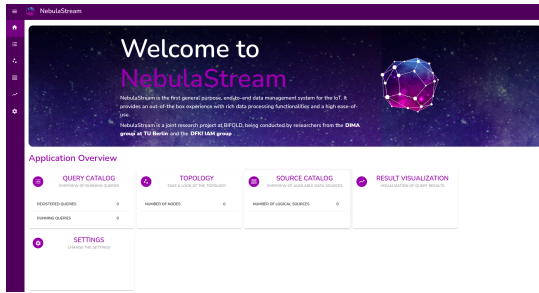


Figure 2: Overview of NebulaStream-UI.




Figure 3: Patient Overview Page.

3.3 NebulaStream Extensions

In order to specify the monitoring queries to derive the patients' current score continuously, we use NebulaStream's rich set of streaming operators. This operator set includes all traditional streaming ETL operators, temporal aggregations, window joins in event-time, and Complex Event Processing operators, such as the temporal sequence operator and conjunctions. To derive a prediction of the severity class of decompensation, the trained TensorFlow model needs to be integrated into NebulaStream. Ongoing work integrates a new operator, including a placement strategy for inferring ML models. This feature enables the usage of TensorFlow models inside NebulaStream and allows users to add trained ML models to a query. The specified data stream in the query sends its tuples as input to the model. The resulting stream contains the models' output, i.e., the estimated maximal decompensation score. The predicated score can then be visualized for the physicians in the NebulaStream-UI.

3.4 NebulaStream-UI Extensions

The NebulaStream-UI provides an intuitive and visual way to interact with data from NebulaStream and is publicly available [19, 20]. Figure 2 introduces the standard features of the NebulaStream-UI. The *Query Catalog* allows the user to submit new queries and provides an overview of all queries and their current status. The *Topology* page visualizes the current tree-like network topology containing all connected workers and the coordinator, including node-specific details such as resources or IP addresses of the nodes. The *Source Catalog* allows registering new logical streams by providing a stream name and schema. A logical stream represents a logical view over a set of IoT devices, e.g., a logical stream *HeartRate* combines all heart rate readings from physical devices into one consistent stream. Furthermore, it provides an overview of all registered streams and their schemata. The *Result Visualization* page enables the user to select a graph type,

e.g., a line chart, to visualize results. Finally, on the *Setting* page, the user defines the IP address and port to connect to a running NebulaStream instance.

In IoT-PMA, physicians interact with NebulaStream through the NebulaStream-UI that has been extended by the following three components for our health care scenario, in accordance with the physicians of GeNo: (1) the patient overview page, (2) the patient detail page, and (3) the TensorFlow model page.

Patient Overview Page: We extend the navigation bar with a hospital symbol (white cross) that helps the physicians navigate to the patients' overview page. The overview page shows each patient in the clinic on a card grid. Figure 3 shows the overview page of one clinic with two patients (1). The card grid component constantly listens to the assigned request (query) and updates the page every time it receives a message, i.e., adding a new patient or updating a patient's information on its card. Each patient card includes the patient's name and room number. Furthermore, the heart color in the right upper corner of each patient card (2) represents the currently more critical decompensation severity class (either pulmonary or hemodynamic) given the introduced scoring model in Section 2. We display the severity class as a traffic light system, i.e., green for no decompensation, yellow for moderate, and red for severe decompensation. As an additional feature, the page provides a search field (3) that allows medical staff to filter for specific patient details, e.g., room number or name. The physicians will be redirected to (2) the patient detail page by clicking on a patient.

Patient Detail Page: The patient detail page (see Figure 4) provides comprehensive information about the patient, particularly the patient's biosignals, all available demographic information, and health history. The page is divided into three components: Patient Details (1), Patient BioSignals (2), and BioSignal Time Series (3).

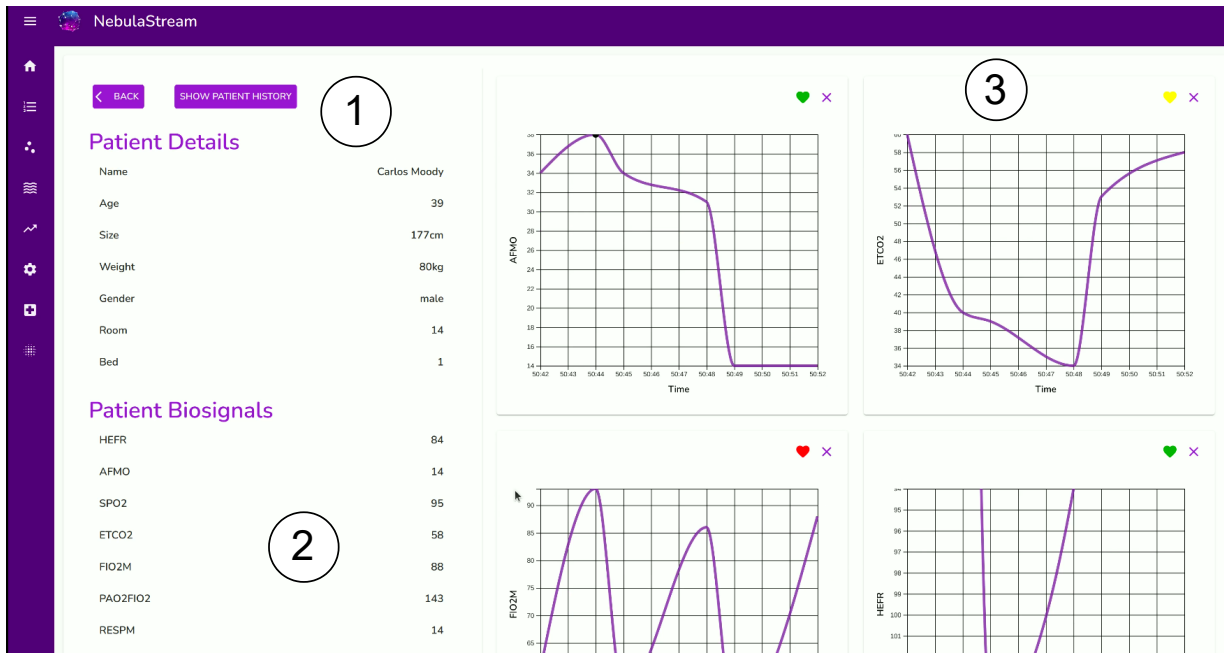


Figure 4: Patient Detail Page.

Patient Details (1): This component provides a grid structure, resulting in a two-column layout for the patient detail page, containing all patient demographic information. Additionally, by pressing the *Show Patient History* bottom, the physician can view the patient’s medical history and add further details. Finally, the *Back*-bottom navigates the physician to the patient overview page.

Patient BioSignals (2): This component follows the two-column grid layout of the Patient Detail component (1) and presents the latest values of the patient’s biosignals. A separate view of concrete values is required to support documentation processes in the daily clinical routine.

BioSignal Time Series (3): This component displays each patient biosignal in a line chart. Each graph represents the time on the x-axis and the biosignal value on the y-axis. Furthermore, the physicians can remove and add new graphs using the *Add Graph* button. Note that the NebulaStream-UI provides several graph types besides line charts. Both features enable physicians to flexibly create ad-hoc views on the required vital parameters and support their fast decision-making.

TensorFlow Model: An additional menu point is added for the TensorFlow Model page (see Figure 5). The component contains the graphical representation of the trained TensorFlow neuronal network (see Section 2). The image was generated with netron [23] and added to




Figure 5: TensorFlow Page.

the NebulaStream-UI to explain the forecasting process and present the insides of the model layers. Furthermore, model updates outside of NebulaStream are considered after observing significant variance in more recent data. The page will be updated in case of updates on the TensorFlow model. In the future, this page can also contain further insides on the monitoring task and its model, e.g., feature detection results.

3.5 Application Deployment

As our target scenario is a retrospective observational study, the data access of our application is limited to the patients’ historical data (see Section 2). Furthermore, given the ethical restrictions, we can only use an anonymized view of the data from ten patients. The data contain the required vital parameters (parameter identifier and value), relative timestamps (hours in the

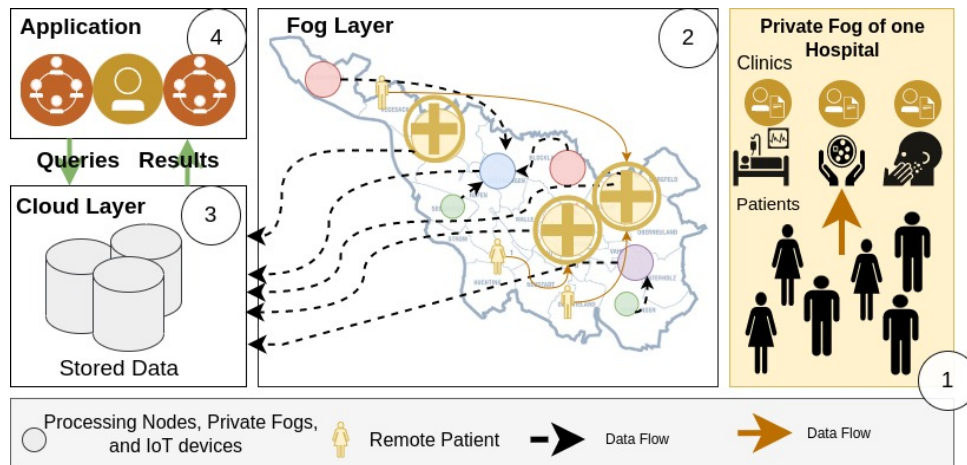


Figure 6: IoT Health Care Scenario.

hospital), and a surrogate key per patient. The data excerpts have file sizes ranging from 100 kB to 1 MB, depending on the time the patient spends in the clinic. Additionally, we can use demographic information and anonymized patient histories of archetype cases provided by the medical project team members.

We use a Raspberry Pi cluster and install our data generator on each Raspberry Pi to represent a patient. The generated data are converted into a JSON format and delivered to the message broker Mosquitto [18], which implements the MQTT protocol. The NebulaStream-UI allows for reading data from this lightweight broker as it fits the heterogeneous hardware requirements of the IoT and can be used for low-end devices as well as servers. The NebulaStream workers listen to the topic of the broker defined in the running query and receive the generated data. The data are displayed using respective IDs, e.g., patientID, and subsequently supplied to the visual components of the UI.

4 IOT ECOSYSTEM FOR HEALTH MONITORING

For our vision of a future IoT health monitoring ecosystem, we extend the introduced intensive care monitoring scenario (see Section 2) to several monitoring tasks running in parallel inside one clinic and other clinics inside a clinic network. In Figure 6, we present such a large-scale IoT scenario focusing on healthcare monitoring with the clinic network *Gesundheit Nord* (GeNo) [6] as a representative in the German city of Bremen. The network consists of four hospitals in Bremen and 62 specialist clinics. Each hospital with several clinics (1) has a private fog. A private fog is owned by an organization, e.g., GeNo or a third-party provider, and its resources are exclusively

used. Thus, the patient data are only available to the clinics of the fog owner, which mitigates issues of privacy and security [1]. The number of clinics depends on the hospital size. In the case of the GeNo, the number of clinics per hospital ranges from 12 to 24. Each clinic cares about up to 16 patients and measures different vital parameters. NebulaStream can handle the collection and low-latency processing of the generated data within the private fog. It leverages the available computation resources of connected devices in the private fog and collects data in a local data center. We add monitoring queries in NebulaStream when the task maps to the appropriate streaming operator set inside NebulaStream. The NebulaStream-UI visualizes the query results of each clinic and allows the specialist to add specific graphs or adjusted queries for individual patients. Furthermore, additional resources can be provided in the private fog to include ML tasks, e.g., for predictions that indicate the patients' health development and enable proactive treatment. Such an IoT health monitoring system also allows fast information transfer, e.g., if a patient moves to another clinic or laboratory values can be added to the patient and are immediately accessible.

In the fog layer (2), outside the private clinic network, remotely monitored patients moving around in the city and sent their data only to the private fog clinic network. In the private fog, their data is analyzed and stored. Both patients and physicians are informed in case of critical events. In order to keep ML models for the prediction of critical cases updated, federated learning among different hospitals is required. Federated learning on sensitive patient data is supported by leveraging the ownership rights of private fogs instead cloud-based AI solution where sensitive data is collected on a central server. Moreover, various other IoT devices, processing nodes, or private fogs are distributed within the city to

gather, (pre-)process, and store data of the surroundings en-route to the cloud. The cloud layer (3) is the endpoint for the data of an entire geographical area and allows for a centralized, global view of the data, e.g., for smart city applications (4). The applications outside the clinic network also require aggregated and up-to-date hospital information, e.g., bed coverage or COVID cases. The clinic network can share that coarse granular information with the public cloud by defining new streams that provide the requested information using NebulaStream, while the detailed patient data stay in the private fog.

Using NebulaStream as the backbone for IoT data processing, we can enable universal low-latency patient monitoring in the private fog of a clinic network. Our scalable approach focuses on ad-hoc views on the relevant data and its analysis results to support physicians. To this end, it enables the flexible addition of new medical devices and monitoring tasks without affecting the latency performance of the UI client. Furthermore, the private fog can be part of the NebulaStream fog layer of a more general IoT ecosystem of various stakeholders, technologies, and companies, e.g., as proposed by KI-SIGS or a smart city. Hospitals can contribute essential statistics and information for such an application and leverage the private fog to secure sensitive patient data.

In order to enable our large-scale and geo-distributed IoT health monitoring vision, NebulaStream is under heavy development by an active team of researchers at TU Berlin and DFKI GmbH, led by Prof. Volker Markl. A closed-beta release (NebulaStream 0.2.0.) is available under the Apache License v2.0.

5 CHALLENGES OF CLINICAL STUDIES FOR PATIENT MONITORING

A prospective observational clinical study is currently under preparation to perform validation on real-time data using the developed methods and our proposed application (see Section 3). Moving the setting from retrospective to prospective introduces challenges not limited to our solution. In the remainder of this section, we describe relevant challenges of verifying the usage of patient monitoring applications in smart hospitals and discuss our lessons learned.

Ethical Challenges: Healthcare scenarios introduce a broad range of ethical challenges. Therefore, KI-SIGS offers workshops to present rules and discuss the resulting challenges. In the following, we describe two significant ethical challenges.

Privacy and Security: To obey data protection rules and ensure the privacy of patient data, we had to ensure

that only involved medical team members could re-identify patients and have access to non-anonymized data. Therefore, personal data and parameters of the patient that could lead to the identification of patients were partially manually excluded. Furthermore, the patient data are not allowed to leave the clinic network. Thus, an access-protected research server accessible via VPN containing the selected, cleaned, and anonymized patient cases were set up.

Method Evaluation: During the *Clinical Study*, the responsible physicians are not allowed to consider the support of the system and the developed methods for patients' treatments. In particular, allowing the bedside physician to access the results of the developed methods, i.e., the potential risk of decompensation, would directly influence the physician's decisions and medical actions. For example, the physician could intervene therapeutically, arrange a prolonged observation episode, or arrange invasive monitoring and diagnostics. However, a non-validated warning system implies risks of unnecessary or potentially harmful intervention and prolonged stay in the intensive care unit. We consider avoiding these issues by only allowing members of the research team to access the results and reports of the system. Bedside physicians will not get notifications about a possible upcoming decompensation, and thus the warnings system will not influence the course of treatment.

Sample Rates of Patient Data: The retrospective study and its developed methods are based on hourly compressed historical data. In particular, no aggregation of detailed patient information is applied, and only a specific data entry at a single timestamp is stored. Live data will appear in different and more granular frequencies and must be processed to match the hourly frequency of the developed methods to guarantee similar results and accuracy. A retraining of the neural network on raw data is not considered as the training sample is too small. This challenge demonstrates that current patient data management systems require updates on their storing procedures to train ML models on the available data and apply them to real-time streams. Advanced solutions to compress vital signals (time-series) and enable their reconstruction exists. Allowing for the reconstruction of the data can help the training of accurate ML and other prediction models usable in real-world health monitoring applications.

Device Integration: The usage of medical monitoring devices increased with the IoT. However, their integration into a unified application to support physicians and provide a central view of all taken

measurements is still challenging. Communication standards for the healthcare sector are under heavy development to enable smart hospitals and data migration in healthcare ecosystems. Fast Healthcare Interoperability Resources (FHIR) [5] is the latest standard of the Health Level Seven International (HL7) [9] invented to enable data exchange between healthcare institutions and device integration. Adopting standard technologies widely used by other industries is essential for universal health applications. Integrating FHIR in NebulaStream is crucial for future work since it enables the envisioned patient monitoring application and the participation of health organizations in general IoT scenarios such as smart cities. The devices that gather the different pieces of information will help medical practitioners to make better-informed and timely decisions while reducing human-made errors in the field of medical science.

6 RELATED WORK

IoT Healthcare Monitoring System: Mamdiwar et al. [12] discuss different architectures, data processing, and transfer methods, as well as several computing paradigms suitable for the IoT and healthcare monitoring. They introduce an overview of applications that use IoT devices for healthcare monitoring. Most of these applications run on personal mobile phones and use wearable sensors to monitor, e.g., activity or glucose monitoring. For the monitoring of illnesses, e.g., diabetes, applications envision push-based communication with physicians, e.g., when a critical event is detected or infectious diseases, e.g., COVID, are remotely diagnosed [2, 8, 15]. More critical illnesses, e.g., cardiac patients [7], are continuously monitored inside the hospital using various vital parameters. In contrast to our solution, these monitoring scenarios propose cloud solutions for processing and storing the data. Hospital applications introduce scientific methods to detect a specific disease pattern but do not provide a visual presentation or the possibility for the physicians to interact with the application [10].

Scoring Systems: Medical scoring systems are used for medical research and education but are also essential in health economics and quality assurance. They assess clinical conditions and procedures at particular, mostly retrospective points in time. There are already established and important, non-AI-based scoring systems applied in the field of intensive care medicine: The Acute Physiology and Chronic Health Evaluation (APACHE) [28] is a complex scoring model for evaluating disease severity and forecasting the

probability of survival. The commonly used Simplified Acute Physiology Score (SAPS) [17] describes patients' physical condition and permits comparison of disease severity, while the Sequential Organ Failure Assessment (SOFA) [26] focuses on particular diseases and organ failure. The Therapeutic Intervention Scoring System Score (TISS) [16] depicts therapy and care efforts. In contrast to these scoring systems, we developed two AI-based systems that focus on pulmonary and hemodynamic decompensation events and allow the prediction of these incidents in a timely manner.

7 CONCLUSION

We propose IoT-PMA implemented on top of the IoT data processing system NebulaStream for a real-world scenario. Our application implements two scoring systems defined by the medical project members to indicate hemodynamic and pulmonary decompensation severity classes. Furthermore, we designed extensions of the NebulaStream-UI with the medical project members that visualize all relevant information with continuous updates on new data. Thus, our application supports intensive care physicians in their decision-making using ad-hoc views on relevant vital parameters and monitoring the patients' severity class. Work in progress integrates the trained neural network for predicting the severity class of decomposition in time. We present the current challenges of our solution for a prospective clinical study, which indicates hurdles of emerging patient monitoring applications. Finally, we envision the infrastructure of larger, geographically distributed IoT application scenarios such as the intelligent health ecosystem proposed by healthcare ecosystems or smart cities using NebulaStream. Those applications require the latest statistics or capacity information from hospitals and other stakeholders for their users. NebulaStreams' architecture enables stakeholders to participate in such application scenarios without exposing sensitive data to the public.

ACKNOWLEDGEMENTS

This work is supported by the German Ministry for Education and Research as BIFOLD (01IS18025A, 01IS18037A) and the German Federal Ministry for Economic Affairs and Climate Action (BMWK) through the KI-SIGS – KI-Space for intelligent health systems (grant no. 01MK20012P). Furthermore, the authors would like to thank Serge Autexier, Christoph Lüth as well as Christoph Int-Veen for their comments, feedback, and constructive criticism in the context of the *RIDIMP* project. Additionally, we thank our amazing students

Aljoscha Lepping, Hoang Mi, Laura Mons, Richard Bendler, Malte Hoberg, David Wenzel, and Mark Sumegi for their work on the NebulaStream project.

REFERENCES

- [1] A. Ahmed, H. Arkian, D. Battulga *et al.*, “Fog computing applications: Taxonomy and requirements,” *arXiv preprint arXiv:1907.11621*, 2019.
- [2] M. Ahmid and O. Kazar, “A cloud-iot health monitoring system based on smart agent for cardiovascular patients,” in *2021 International Conference on Information Technology (ICIT)*, 2021.
- [3] M. A. Akkaş, R. Sokullu, and H. E. Cetin, “Healthcare and patient monitoring using iot,” *Internet of Things*, vol. 11, p. 100173, 2020.
- [4] S. H. Almotiri, M. A. Khan, and M. A. Alghamdi, “Mobile health (m-health) system in the context of iot,” in *2016 IEEE 4th international conference on future internet of things and cloud workshops (FiCloudW)*. IEEE, 2016, pp. 39–42.
- [5] M. L. Braunstein, “Patient — physician collaboration on fhir (fast healthcare interoperability resources),” in *2015 International Conference on Collaboration Technologies and Systems (CTS)*, 2015.
- [6] Gesundheit Nord, “Gesundheit Nord – Klinikverbund Bremen,” <https://www.gesundheitnord.de/kbn.html>, accessed 23th February 2022.
- [7] U. Gogate and J. Bakal, “Healthcare monitoring system based on wireless sensor network for cardiac patients,” *Biomedical & Pharmacology Journal*, 2018.
- [8] T. Hafsiya and B. Rose, “An iot-cloud based health monitoring wearable device for covid patients,” in *2021 7th International Conference on Advanced Computing and Communication Systems (ICACCS)*, 2021.
- [9] HL7 Deutschland e. V., “HL7 (Health Level 7) - digital standards for healthcare,” <https://hl7.de>, accessed 22nd April 2022.
- [10] L. Y. Kent and I. F. B. Kamsin, “Implementation of iot in patient health monitoring and healthcare for hospitals,” in *Proceedings of the 3rd International Conference on Integrated Intelligent Computing Communication & Security (ICIIC 2021)*, Bangalore, India, 2021.
- [11] M. A. Makary and M. Daniel, “Medical error—the third leading cause of death in the us,” *Bmj*, vol. 353, 2016.
- [12] S. D. Mamdiwar, A. R. Z. Shakruwala *et al.*, “Recent advances on iot-assisted wearable sensor systems for healthcare monitoring,” *Biosensors*, 2021. [Online]. Available: <https://www.mdpi.com/2079-6374/11/10/372>
- [13] C. Mandel, K. Stich, S. Autexier *et al.*, “Using gated recurrent unit networks for the prediction of hemodynamic and pulmonary decompensation,” in *Proceedings of the 44th. International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE Engineering in Medicine and Biology Society, 445 Hoes Lane, Piscataway, NJ 08854 USA, 2022, p. n.a., to appear.
- [14] J. Manyika, M. Chui, P. Bisson, J. Woetzel, R. Dobbs, J. Bughin, and D. Aharon, “Unlocking the potential of the internet of things,” *McKinsey Global Institute*, vol. 1, 2015.
- [15] P. Mhatre, A. Shaikh, and S. Khanvilkar, “Non invasive e-health care monitoring system using iot,” 2020.
- [16] D. R. Miranda, A. de Rijk, and W. Schaufeli, “Simplified therapeutic intervention scoring system: the tiss-28 items—results from a multicenter study,” *Critical care medicine*, 1996.
- [17] R. P. Moreno, P. G. Metnitz, E. Almeida *et al.*, “Saps 3—from evaluation of the patient to evaluation of the intensive care unit. part 2: Development of a prognostic model for hospital mortality at icu admission,” *Intensive care medicine*, 2005.
- [18] Mosquitto, “Eclipse Mosquitto™ -An open source MQTT broker,” <https://mosquitto.org>, accessed 25th March 2022.
- [19] NebulaStream Team, “Nebulastream - data management for the internet of things,” <https://www.nebula.stream>, accessed 22nd July 2022.
- [20] NebulaStream Team, “Welcome to Nebulastream UI,” <https://github.com/nebulastream/nebulastream-ui>, accessed 31st March 2022.
- [21] Philips Healthcare, “Intellispace critical care and anesthesia,” <https://www.philips.de/healthcare/product/HCNOCTN332/intellispace-critical-care-and-anesthesia>, accessed 18th February 2022.

- [22] Project management AI-SIGS, “KI-SIGS-spaces for intelligent healthcare systems,” <https://ki-sigs.de>, accessed 18th February 2022.
- [23] L. Roeder, “Lutz Roeder’s Netron,” <https://netron.app>, accessed 18th March 2022.
- [24] K. Stich, C. Mandel, K. Hochbaum *et al.*, “Neudefinierte Scoring-Systeme zur retrospektiven Klassifikation hämodynamischer und pulmonaler Dekompensationen im Rahmen eines KI (Künstliche Intelligenz)-Trainingsprozesses,” in *Berichtsband zum ”DIVI21Virtuell - 21. Kongress der Deutschen Interdisziplinären Vereinigung für Intensiv- und Notfallmedizin e.V.”*, 2021. [Online]. Available: <https://kongress.divi.de/divi-21>
- [25] TensorFlow, “An end-to-end open source machine learning platform,” <https://www.tensorflow.org>, accessed 15th February 2022.
- [26] J.-L. Vincent, R. Moreno, J. Takala, S. Willatts *et al.*, “The sofa (sepsis-related organ failure assessment) score to describe organ dysfunction/failure,” 1996.
- [27] S. Zeuch, A. Chaudhary, B. D. Monte, H. Gavriilidis, D. Giouroukis, P. M. Grulich, S. Breß, J. Traub, and V. Markl, “The nebulastream platform for data and application management in the internet of things,” in *10th Conference on Innovative Data Systems Research, CIDR 2020, Amsterdam, The Netherlands, January 12-15, 2020, Online Proceedings*. [www.cidrdb.org](http://cidrdb.org), 2020. [Online]. Available: <http://cidrdb.org/cidr2020/papers/p7-zeuch-cidr20.pdf>
- [28] J. E. Zimmerman, A. A. Kramer, D. S. McNair, and F. M. Malila, “Acute physiology and chronic health evaluation (apache) iv: hospital mortality assessment for today’s critically ill patients,” *Critical care medicine*, 2006.



investigation of algorithms for time series prediction and classification tasks.

Christian Mandel is a Senior Researcher at the research department for Cyber-Physical Systems of the German Research Center for Artificial Intelligence (DFKI). His scientific work comprises the development of AI-based software for rehabilitation robotic devices, as well as the



Kathrin Stich is a specialist for anaesthesiology and member of the medical team of the *RIDIMP* project. She is working as a senior physician at the clinic for intensive care and emergency medicine at Klinikum Bremen-Mitte.



Nord GmbH, focusing on strategic planning, AI research projects, and teaching structuring in medicine. She is a KI SIGGS steering committee member and the specialist societies DGSMP, BDC, DHV, and DGAM.

Dr. med. Karin Hochbaum is a specialist in surgery with additional qualifications in medical quality management and social medicine. She worked in various management positions in hospital management and is the head of corporate development and medical strategy at Gesundheit

AUTHOR BIOGRAPHIES




a focus on distributed systems as well as data and software engineering.

Ariane Ziehn is a Ph.D. candidate at the IAM group (DFKI). Ariane’s research interests include distributed systems, complex event processing, and IoT environments. She received her M.Sc. in Information Systems Management at TU Berlin with



an adjunct professorship from the medical faculty of the RWTH Aachen University in 2011.

Prof. Dr. Rolf Dembinski is the Director of the Clinic for Intensive Care Medicine and Emergency Medicine at the Klinikum Bremen Mitte. He is a specialist in anaesthesiology with a degree in human medicine at the FU Berlin and additional training in emergency medicine and intensive care medicine. He received his Ph.D. in 1999, habilitated in 2006, and earned



Steffen Zeuch is a Senior Researcher at the DIMA group (TU Berlin) and IAM group (DFKI). His research interests are modern hardware and the IoT. He published research papers on query optimization and execution as well as novel system architectures. He did his Ph.D. in Computer Science at Humboldt University Berlin.



Volker Markl is a Full Professor and Chair of the Database Systems and Information Management (DIMA) Group at TU Berlin, Chief Scientist and Head of the Intelligent Analytics for Massive Data Research in DFKI, and Director of the Berlin Institute for the Foundations of Learning and Data (BIFOLD). He has published numerous research papers on indexing, query optimization, lightweight information integration, and scalable data processing. He is a Fellow of the ACM as of 2021.