

Medical Cyber-Physical Systems: IoMT Applications and Challenges

Amanda Watson, Jean Park, Sydney Pugh, Oleg Sokolsky, James Weimer, and Insup Lee

Dept. of Computer and Information Science

University of Pennsylvania

{aawatson,hlpark,sokolsky,weimerj,lee}@seas.upenn.edu

Abstract—The Internet-of-Medical-Things (IoMT) is a network of connected medical devices, hardware infrastructure, and software platforms used to connect healthcare information technology. While the IoMT has expanded access to real-time health information, a side-effect has been inundating clinicians with excessive amounts of information - some of which is critical - but the vast majority is superfluous. This has led to clinician burnout, alarm fatigue, and skepticism when adopting new technologies. Leveraging the IoMT infrastructure, Medical Cyber-Physical Systems (MCPS) seek to address these issues by providing more actionable information, detecting device failures, and improving patient safety and treatment effectiveness. MCPS can provide an effective summary of patient state by fusing data from multiple physiological monitoring devices. Furthermore, MCPS can utilize patient physiology represented as a model to provide automatic closed-loop control of patient treatment, which does not require clinician intervention, under known safe conditions. The need to design such MCPS/IoMT systems that are safe, effective, trustworthy, and helpful to clinicians has presented numerous challenges. In this paper, we discuss the challenges in developing MCPS/IoMT, some of our work in addressing them, and several open research issues.

Index Terms—medical cyber-physical systems, internet of medical things, clinical decision support, medical device integration

I. INTRODUCTION

The Internet of Medical Things (IoMT) is a network of connected medical devices, hardware infrastructure, and software platforms that connect to health care information technology systems using networking technologies [1]. As more connected medical devices come online, healthcare is undergoing a transformation [2]. It is transitioning from standalone devices that can be designed, evaluated, and implemented individually to a network of devices that can simultaneously monitor, evaluate, and manage many aspects of a patient's physiology. The interconnection of the embedded software, network capabilities, and underlying infrastructure combined with the complexities of human physiology have created an environment in which medical cyber-physical systems (MCPS) can be implemented [3].

The growing IoMT has expanded access to real-time health information for patients and clinicians. Traditionally, healthcare has relied on the expertise of clinicians and their experience with their patients. With the addition of new data streams from upgraded medical devices, the burden has fallen on the clinicians to synthesize and interpret it [4]. This has led to clinician burnout [5], [6], alarm fatigue [7], [8], and

skepticism when adopting new technologies [9], [10]. To support clinicians in their efforts and reduce the overall burden the IoMT has placed on them, MCPS are being developed to monitor patient states, provide closed-loop control, and personalized treatment. As this is moving away from the traditional healthcare paradigm, it has created an array of challenges that must be addressed.

In this paper, we discuss the research directions in MCPS and IoMT, the challenges surrounding them, and our approaches to addressing these challenges. These challenges need to be addressed, beginning with the infrastructure and protocols for bringing medical devices online. As new MCPS are being developed, they need to be designed in a manner in which they are safe, effective, and useful to clinicians. As clinicians become more familiar with these technologies, there will be greater opportunity for expansion. This supports a move towards higher levels of autonomy, reducing the burden on clinicians and encouraging advancements in closed-loop control.

The remainder of our paper is structured as follows: In Section II, we introduce IoMT and MCPS research and their open challenges. Section III discusses our work on building infrastructure to support this. Section IV describes our approaches to clinical decision support. Section V details hybrid closed loop control with feedback. Finally, we wrap up with our conclusion and discussion of future work in Section VI.

II. OVERVIEW AND CHALLENGES

MCPS is the lens that provides clarity to the overwhelming volumes of information contained in medical data collected by IoMT systems. As clinicians become increasingly dependent upon the guidance from these systems, their safety and reliability becomes imperative. Research surrounding this encompasses medical device interoperability, clinical decision support, device coordination, and closed-loop control. Figure 1 presents an overview of MCPS/IoMT systems and their organization.

A. Medical Device Interoperability

As the IoMT continues to grow and more medical devices come online, the infrastructure to support this also must evolve. This includes the process by which legacy devices are brought online, connection to new medical devices, software

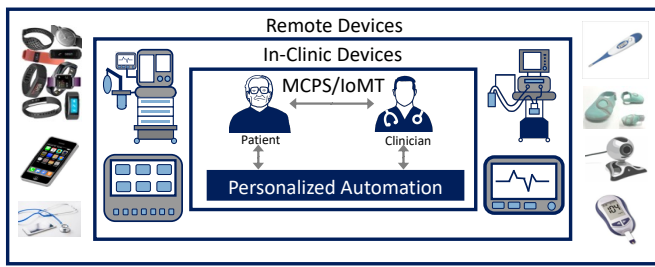


Fig. 1. Overview of IoMT/MCPS

platforms to support access and interconnection among the devices, and hardware infrastructure. Thus far, the many attempts that have begun to build out these protocols and systems have resulted in disconnected, vendor-specific systems. These systems require specialized training and support, and while they offer new functionalities are also placing an additional burden on those using and maintaining them. In the future, standardized methods by which we bring devices online and support their interoperability should be developed. Systems that support access to the data generated and interconnectivity between devices while not requiring repeated specialized training should also be created.

B. Clinical Decision Support

Clinical decision support improves patient care by providing clinicians with patient-specific information that has been processed to present relevant information at appropriate times. This can include pertinent alarms and alerts, diagnostic support, and treatment recommendations. When building these systems, researchers encounter challenges, including not having enough high-quality data, personalizing the support to individual patients, and providing confidence guarantees for safety. Given the influx of medical data from the creation of IoMT systems, clinical decision support can reduce the burden on clinicians to parse and understand all the data they are given. When implemented effectively, these systems can lower costs of care and improve efficiency in providing quality care to patients.

C. Partial Closed-Loop Control with Feedback

Traditionally, healthcare relies on clinician expertise to treat patients. In these scenarios, the clinician(s) evaluate the patient and control the process of care. For example, an anesthesiologist continuously monitors the patient, vital signs and equipment; then makes adjustments to achieve the desired level of sedation while maintaining safety. Even today, medical devices are beginning to assist clinicians with monitoring through smart alarm systems. As these devices become more capable, more tasks can be delegated to them in the future. While the most critical tasks remain at the direction of the expert, automatic controllers can reduce the clinician's overall workload by managing the tedious tasks. Closed-loop control systems are already being utilized in medical devices. Thus far, these systems are limited to organs with simplistic inputs and

outputs, such as pacemakers and defibrillators. In the future, these devices and even combinations of devices can provide support in more complex scenarios.

D. Challenges

When building IoMT/MCPS systems, the following challenges must be addressed:

- 1) *Interoperability*: As more medical devices come online, inconsistent data and systems across vendors and devices must be managed. Additionally, bringing legacy devices online must be managed.
- 2) *Clinician Support*: While the IoMT has expanded access to real-time health information, a side-effect has been inundating clinicians with excessive amounts of information - some of which is critical - but the vast majority is superfluous. IoMT/MCPS systems should be implemented in such a manner to reduce clinician load, not add to it.
- 3) *Autonomy*: MCPS is increasingly utilizing solutions that enable further autonomy. These solutions need to be implemented in a safe manner that can be accepted by clinicians.
- 4) *Security and Privacy*: Medical devices and the systems that handle their data can be vulnerable to security threats. This can impact the safety and effectiveness of the device, its data, and the patient. Thus, the security and privacy of IoMT/MCPS systems is a critical component.

In the following sections, we describe our current work and how we chose to address these challenges.

III. MEDICAL DEVICE INTEROPERABILITY

The IoMT is a complex network of interconnected medical devices, hardware infrastructure, and software platforms. These medical devices have incorporated systems that have been developed to support the storage, transmission, and security of the medical device data produced. As these systems are vendor-specific and thus non-standardized, they require specialized training. These systems are also limited by the amount of data that can be stored, the amount of time it is stored for, and the ability to integrate with devices outside of its vendor-specific system. To address this, we created VitalCore, a medical device integration platform that manages medical devices, proactively keeps them operable, and creates a centralized platform for data collection, analysis, and medical device interoperability. Further, when rapidly prototyping new medical devices and applications, we developed Raproto, an open-source easy-to-use rapid prototyping platform that does not require the time, effort, and expertise needed for custom development.

A. Medical Device Integration

Medical professionals devote considerable resources to collecting medical device data for validation, and analysis. These devices use vendor-specific applications to monitor the device and log the data for a limited time. This leads to a complex

process to acquire medical device data. To remedy this, we created VitalCore [11], a platform that support the integration of medical devices. Further it provides clinical decision support by reducing the need for manual documentation of medical device data and provides access to the data in real time. This platform gives researchers simplified access to a wealth of data collected longitudinally, enabling medical applications. As an added benefit, we created the Medical Device Dashboard to monitor the connected medical devices across Penn Medicine. Using this, we reduced the amount of resources are required of clinical teams including the medical professionals, clinical engineers, and IT analysts. VitalCore simplifies the troubleshooting workflow, thus decreasing downtimes and increasing clinical productivity.




Fig. 2. VitalCore System Architecture

The overall architecture of VitalCore is depicted in Figure 2. VitalCore captures the HL7 data feed, a widely supported standard for medical data that streams from medical devices. This data is given into the stream processor to process the data for important segments. Then, the data is either sent to storage or routed directly to our communication protocols that provide the processed data to applications. The applications, which are built on top of VitalCore, can process and analyze the data enabling MCPS and then send any new data back to VitalCore for storage.

VitalCore is currently being leveraged in the Penn Medicine health system. Overall, there are over 3000 integrated medical devices that our VitalCore can support. From VitalCore, we were able to implement three applications: Medical Dashboard, Ventilation Alert, and Anomaly Detection. The Medical Device Dashboard supports clinicians by giving them a user-friendly GUI application where data can be accessed to real-time. It simplifies the troubleshooting for IT staff such as for when devices go offline decreasing expended effort and downtime. Ventilation alert illustrates the applications that can be developed because of VitalCore. Ventilators data is used to generate alerts that are sent to clinicians. Further, we have used the ventilator data to identify patients at high risk of extubation failure [12]. Anomaly detection monitors the usage patterns of medical devices as well as groupings of medical devices allowing for proactive troubleshooting and avoidance of medical device errors.

B. Rapid Prototyping for Research

Research endeavors frequently rely on large amounts high-quality data. To gather this data, a data collection system must be built but this can be a significant drain on resources. These

systems tend to be customized to a specific task; thus, they are not general enough to support other tasks. This is also true for new IoMT devices; evaluating and testing new devices is time-consuming without adding the need to build data collection systems. To solve this problem, we developed Raproto [13], an open-source rapid prototyping platform that does not require the time, effort, and expertise needed for custom development. Further, we built multipurpose and customizable smartwatch applications to give researchers a jump start when leveraging sensors on commercially available smartwatches.




Fig. 3. Raproto Platform

The overall architecture of Raproto is depicted in Figure 5. The Raproto platform consists of the devices, communication protocol, and remote storage. These components make possible the collection, transmission, storage, analysis, and visualization of data. The devices collect data from the available sensors. The communication protocol leverages MQTT to send the data collected on the smartwatches to our remote server. The remote server stores the data, displays it, and allows users to remotely configure devices.

We evaluate our platform in a lab setting as well as in a clinical setting. Overall, we found that we can collect data using our smartwatch applications for greater 24 hours in certain scenarios on a single charge, there is insignificant data loss, and make an ideal tool to preface customized device development for real-world impact and commercialization. We leveraged Raproto in two case studies: In-hospital Stroke Detection and Postpartum Hemorrhage Prediction.

a) *In-hospital Stroke Detection*:: Early detection of stroke leads to early interventions which have been proven to improve clinical outcomes [14], [15]. The in-hospital stroke detection case study included 200 patients over the course of eight months. For this study, we used the Tizen Raproto Application loaded onto twelve Samsung Galaxy Active Smartwatches.

b) *Postpartum Hemorrhage*:: Postpartum hemorrhage is a relatively common complication of childbirth [16], [17]. Their significance can range from mild anemia to heart attack, stroke, and possible death. Timely detection of PPH is crucial to the effective management of the complication. The Postpartum Hemorrhage Prediction Case Study was conducted with 525 patients over the course of 4 months at the Pennsylvania Hospital of Philadelphia.

These case studies provided useful insight into deploying consumer wearables in a hospital environment for data collection. Raproto was successful in allowing clinical researchers to set up and operate smartwatches with minimal engineering

support. Providing a graphical user interface (GUI) interface was essential for the clinical team to monitor battery life and ensure data collection.

IV. CLINICAL DECISION SUPPORT

In the previous section, we discussed our approaches to building the infrastructure for IoMT systems. These systems create an environment in which clinical care can be improved and clinical decision support can be provided. Clinical decision support, when well implemented, has the capability to alleviate the burden of expanding clinical knowledge, the escalating complexity of patient-specific care, and the ever-increasing amount of clinical data. Currently, implementing these systems can be expensive, disruptive, inconsistent, and unvalidated. This leads to distrust and disregard from clinicians. In this section, we discuss our work in developing clinical decision support systems that begin to address these challenges. In this section, we discuss our work in developing clinical decision support systems in the form of smart alarm systems and risk assessment with mitigation guidance.

A. Intelligent Alarm Systems

Clinicians are already inundated with excessive amounts of information. While providing the information in its entirety to clinicians does present them with the small subset of critical information, it also requires them to evaluate the non-critical majority. This is highlighted in the Intensive Care Unit, a notoriously noisy environment with an overabundance of alarms. In these environments, clinicians can experience alarm fatigue, desensitization to alarms that can cause missed critical alarms or delayed responses [18], [19]. Today, most medical devices use threshold-based alarms that are set by the clinician to determine when the device should sound an alarm. While this approach allows for customization to the patient and clinician, it has largely been shown to produce a high number of alarms with a high false alarm rate. This leads to the alarms being ignored or turned off entirely. Further, the only information provided is that a threshold was reached, requiring the clinician to evaluate the clinical significance. Leveraging data from multiple medical devices can provide more context, and MCPS systems can be developed to suppress the irrelevant alarms that contribute to clinician alarm fatigue. Further, these alarms can be presented in a manner that supports the current workflow instead of impeding it. This will create an environment in which clinician can respond to and make an informed decision quickly.

As there is an abundance of alarms that can be used to create large datasets for training, obtaining labels for these alarms using traditional methods, such as observational study, is time-consuming and expensive. Thus, we developed an effective method for estimating the performance of a smart alarm system without the need for a large labeled dataset [20]. We leverage data programming [21] to label our dataset by taking as input the unlabeled dataset of alarms and clinician-created labeling functions. It uses a generative model to synthesize the labeling function and determine the confidence the model has

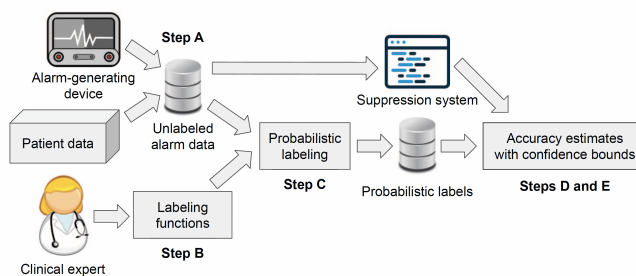


Fig. 4. Our approach to alarm suppression systems

in each generated label. Then, by taking the high-confidence subset of those labels, we evaluate the effectiveness of the suppression system. This process is illustrated in Figure 4. This supports early-stage investigations of suppression systems leading to the prioritization resources. This can select observational studies, which require large amounts of resources, of systems with higher potential to relieve alarm fatigue. This can prioritize where clinician effort to manually label datasets should be placed.

We worked closely with clinicians at the Children’s Hospital of Philadelphia to collect a dataset of over 3,000 low SpO_2 alarms, an alarm that is known to have a high false positive rate [22]. Additionally, they provided us with 62 labeling functions across multiple vital signs (heart rate, SpO_2 , respiratory rate, etc.). We used our method to evaluate these alarms for a suppression system and found that where 81% of the alarms were non-critical. This study demonstrated how to use our system to evaluate a dataset of alarms for a suppression system without investing significant time and effort into labeling the alarm data. To further demonstrate the impact of this system, we repeated this study across four more alarm datasets and found three datasets to have a high percentage of suppressible alarms [23].

B. Risk Assessment and Mitigation Guidance

Assessment and mitigation of risk of disease or complications in healthcare shifts treatment from reactionary to preventative. Moreover, personalized treatment optimizes an individual’s care. One such example is anterior cruciate ligament (ACL) reinjuries. The ACL is a ligament in the knee that stabilizes the joint during motion [24]. ACL injuries of the knee are common, especially among youth athletes [25], [26]. When torn, the injured knee has increasing instability causing a higher risk of future injuries. A common treatment is to surgically repair the ACL and subsequently undergo extensive rehabilitation [27]. Even after fully recovering, ACL reinjury occurs in approximately 20% of patients leading to additional surgical interventions, rehabilitation, and inferior results [28].

Clinicians have determined many risk factors that lead to a higher rate of ACL re-tear [29]–[33]. Using these risk factors, individual weaknesses can be identified and then addressed during the recovery, potentially leading to better outcomes. Thus, we developed a system that evaluates the re-tear risk

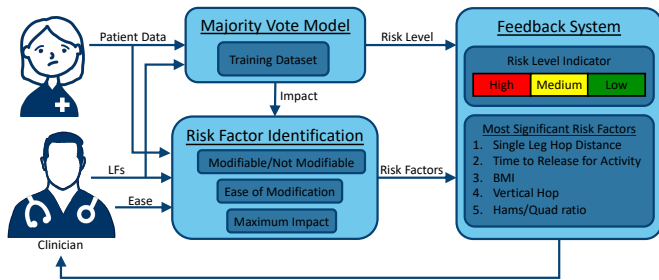


Fig. 5. System Architecture

of each individual patient and identifies the risk factors that have the highest impact [34]. From this, clinicians can amend rehabilitation protocols to address the deficiencies, potentially leading to better outcomes. We leverage clinical knowledge and expertise in a simplistic ensemble-based algorithm to create a re-learn risk prediction for each patient. Then, the highest impact risk factors can be determined. This process is shown in Figure 5. Using the results from this system, it is possible that clinicians can tailor treatment plans for each patient and optimize factors such as when a patient can return to sports.

Our approach involves working closely with clinicians at the Children’s Hospital of Philadelphia. We collected a dataset with 442 youth patients who have undergone ACL reconstruction surgery. This dataset includes surgical information, demographics, family history, rehabilitation information, and data surrounding the initial injury. Then we asked clinicians to incorporate their expert knowledge into labeling functions. These labeling functions are evaluated for their predictive capabilities and are assigned an impact score. Patients are evaluated for their risk level of re-learn, and this is presented to the clinicians with a list of recommendations for risk factors to be focused on during rehabilitation. Clinicians can then use this to personalize each patient’s rehabilitation in an effort to improve overall outcomes. Similar systems [12] have since been developed for other contexts, such as determining patients at high risk for extubation failure.

V. HYBRID CLOSED LOOP WITH FEEDBACK

Clinical decision support uses patient data integrated from a number of sources, processed to gain insights, and presents it to clinicians to interpret and make decisions. This data can also be used to directly control intervention devices creating a closed-loop system. These closed-loop systems are already being used for implantable devices [35] by monitoring patient state and automatically adjusting the intervention provided. These systems also generally include the ability to alert clinicians if the monitored parameter(s) moves outside of the normal range. Moving forward, these systems can be used to reduce clinician workload and incorporate behavioral feedback from the clinician and patient.

A. PCA Infusion Pump

Patient-controlled analgesia (PCA) is a method of pain control that allows patients to manage their analgesia via a patient-controlled analgesia pump. These PCA pumps can be configured to provide a scheduled dose of analgesia and/or allow patients to press a button when they require more pain management. However, the medications used to provide pain management can cause side effects, including respiratory depression leading to hypoxia or even death [36]. PCA pumps have safeguards that can be configured for each specific patient to prevent these side effects from occurring. [37] For some patients, these safeguards are not sufficient and adverse events do occur. Thus, PCA pumps could benefit from the implementation of a closed loop control mechanism [38], [39].

To provide an additional level of safety, patients can be monitored for signs of respiratory depression, an early signal that more significant side effects could occur. Respiratory depression can be monitored through a pulse oximeter that measures a patient’s SpO_2 levels [40]. SpO_2 levels measure the oxygen saturation of the blood, which decreases when respiratory depression occurs. In a closed-loop scenario, the PCA pump monitors the patients SpO_2 and if it detects respiratory depression, it will stop the pump from delivering more pain management regardless of patient request [41]. To verify the safety and effectiveness of a control mechanism for a PCA pump, we leveraged a set of safety requirements [42], categorized them, and formalized the model and requirements in UPAAL. We performed a formal verification of that model and performed conformance testing. We found no violations occurred during this testing [43].

B. Social Isolation and Loneliness

Social isolation is a lack of social contacts, connections, and/or interactions. Loneliness is the feeling of social isolation that one experiences generally due to a disparity between expected and actual social contacts [44]. Social isolation and loneliness are concerns in the elderly community as those who are socially isolated or lonely experience shortened lifespan, depression, anxiety, and dementia, etc. [44]. As this is a growing concern in the elderly community, we have begun work to support the identification and monitoring of social isolation and loneliness in the elderly community using a system of smart devices. For example, we plan to monitor an individual’s activity levels, sleep quality, time out of residence, etc., with pervasive sensing techniques that do not require large amounts of direct interactions with the sensing technology. This is essential as the population in which we are interested has lower familiarity and patience when dealing with smart device errors [45].

When social isolation and/or loneliness are identified, we determine personalized interventions that change as a patient progresses. For example, our recommendations could suggest that a person call a friend or family member in the next week or go on a walk at a local park. These recommendations can be updated at a time interval or as goals are achieved. As the recommendations are updated, they should reflect the

new patient's condition, whether that patient has progressed or regressed. Based on how the patient's state changed, there is the feedback that can provide assistance in selecting future recommendations for that patient. For example, a patient may see more impact on their loneliness level from calling their friend than taking a walk in the local park. In future iterations of recommendations, this should be accounted for. Thus, we recommend accounting for patient-specific feedback in systems that recommend personalized intervention. This could also be applicable in physical therapy and rehabilitation plans [34] or drug selection and dosing. Individualizing feedback could improve utility and increase clinician and patient uptake.

VI. CONCLUSION AND FUTURE WORK

In this paper, we considered the challenges created when applying MCPS to underlying IoMT infrastructure. We discussed the challenges brought by medical device interoperability, creating systems that support clinicians instead of adding to their workload, and the autonomy, security, and privacy issues surrounding hybrid closed-loop control with feedback. As the capabilities of MCPS expand, we must transform the methods by which we develop and evaluate them. In the present, evolving research in this field presents opportunities for immediate impact and real-world functionality. We are pushing towards a vision of the future in which technology autonomously provides comprehensive medical care. As we strive towards this reality, we have developed the IoMT and MCPS, but we still have many more challenges to surpass.

ACKNOWLEDGEMENTS

The authors would like to thank all of our collaborators for their contributions to the works presented in this paper. This work was supported in part by NSF 1915398, NSF 2125561, NIH, R01EB029767, and NIH, R01EB029363.

REFERENCES

- [1] G. J. Joyia, R. M. Liaqat, A. Farooq, and S. Rehman, "Internet of medical things (iomt): Applications, benefits and future challenges in healthcare domain." *J. Commun.*, vol. 12, no. 4, pp. 240–247, 2017.
- [2] A. Ghubaish, T. Salman, M. Zolanvari, D. Unal, A. Al-Ali, and R. Jain, "Recent advances in the internet-of-medical-things (iomt) systems security," *IEEE Internet of Things Journal*, vol. 8, no. 11, pp. 8707–8718, 2020.
- [3] I. Lee, O. Sokolsky, S. Chen, J. Hatcliff, E. Jee, B. Kim, A. King, M. Mullen-Fortino, S. Park, A. Roederer *et al.*, "Challenges and research directions in medical cyber-physical systems," *Proceedings of the IEEE*, vol. 100, no. 1, pp. 75–90, 2011.
- [4] I. Lee, O. Sokolsky, and J. Weimer, "Medical cyber-physical systems: Challenges and future directions," in *Encyclopedia of Systems and Control*. Springer, 2021, pp. 1213–1220.
- [5] V. J. Dzau, D. G. Kirch, T. J. Nasca *et al.*, "To care is human—collectively confronting the clinician-burnout crisis," *N Engl J Med*, vol. 378, no. 4, pp. 312–314, 2018.
- [6] M. Linzer, "Clinician burnout and the quality of care," *JAMA Internal Medicine*, vol. 178, no. 10, pp. 1331–1332, 2018.
- [7] K. J. Ruskin and D. Hueske-Kraus, "Alarm fatigue: impacts on patient safety," *Current Opinion in Anesthesiology*, vol. 28, no. 6, pp. 685–690, 2015.
- [8] J. Welch, "Alarm fatigue hazards: the sirens are calling," *JAMA*, vol. 307, no. 15, pp. 1591–2, 2012.
- [9] C. Robinson, "Clinician adoption of healthcare information technology," *Canadian Nursing Informatics J*, vol. 2, no. 1, pp. 4–21, 2007.
- [10] A. S. McAlearney, D. J. Chisolm, S. Schweikhart, M. A. Medow, and K. Kelleher, "The story behind the story: physician skepticism about relying on clinical information technologies to reduce medical errors," *International Journal of Medical Informatics*, vol. 76, no. 11–12, pp. 836–842, 2007.
- [11] H. Choi, A. Lor, M. Megonegal, X. Ji, A. Watson, J. Weimer, and I. Lee, "Vitalcore: Analytics and support dashboard for medical device integration," in *2021 IEEE/ACM Conference on Connected Health: Applications, Systems and Engineering Technologies (CHASE 2021)*.
- [12] J. Park, A. Watson, X. Ji, K. C. Quinn, J. Weimer, and I. Lee, "Autowean: Extubation failure risk estimation for critically ill patients," in *2022 IEEE/ACM Conference on Connected Health: Applications, Systems and Engineering Technologies (CHASE 2022)*.
- [13] A. Watson, H. Choi, I. Lee, and J. Weimer, "Rapproto: an open source platform for rapid prototyping of wearable medical devices," in *Proceedings of the Workshop on Medical Cyber Physical Systems and Internet of Medical Things*, 2021, pp. 1–6.
- [14] D. Dulli and E. A. Samaniego, "Inpatient and community ischemic strokes in a university hospital," *Neuroepidemiology*, vol. 28, no. 2, pp. 86–92, 2007.
- [15] H. Park, H. Cho, Y.-D. Kim, D. Lee, H.-Y. Choi, S. Kim, and J. Heo, "Comparison of the characteristics for in-hospital and out-of-hospital ischaemic strokes," *European journal of neurology*, vol. 16, no. 5, pp. 582–588, 2009.
- [16] C. Burke, "Active versus expectant management of the third stage of labor and implementation of a protocol," *The journal of perinatal & neonatal nursing*, vol. 24, no. 3, pp. 215–228, 2010.
- [17] B. E. Hamilton, J. A. Martin, and M. J. Osterman, "Births: provisional data for 2020," 2021.
- [18] C. W. Paine, V. V. Goel, E. Ely, C. D. Stave, S. Stemler, M. Zander, and C. P. Bonafide, "Systematic review of physiologic monitor alarm characteristics and pragmatic interventions to reduce alarm frequency," *Journal of hospital medicine*, vol. 11, no. 2, pp. 136–144, 2016.
- [19] B. J. Drew, P. Harris, J. K. Zegre-Hemsey, T. Mammone, D. Schindler, R. Salas-Boni, Y. Bai, A. Tinoco, Q. Ding, and X. Hu, "Insights into the problem of alarm fatigue with physiologic monitor devices: a comprehensive observational study of consecutive intensive care unit patients," *PLoS one*, vol. 9, no. 10, p. e110274, 2014.
- [20] S. Pugh, I. Ruchkin, C. P. Bonafide, S. B. DeMauro, O. Sokolsky, I. Lee, and J. Weimer, "High-confidence data programming for evaluating suppression of physiological alarms," in *2021 IEEE/ACM Conference on Connected Health: Applications, Systems and Engineering Technologies (CHASE)*. IEEE, 2021, pp. 70–81.
- [21] A. J. Ratner, C. M. De Sa, S. Wu, D. Selsam, and C. Ré, "Data programming: Creating large training sets, quickly," *Advances in neural information processing systems*, vol. 29, 2016.
- [22] S. T. Lawless, "Crying wolf: false alarms in a pediatric intensive care unit," *Critical care medicine*, vol. 22, no. 6, pp. 981–985, 1994.
- [23] S. Pugh, I. Ruchkin, C. Bonafide, S. Demauro, O. Sokolsky, I. Lee, and J. Weimer, "Evaluating alarm classifiers with high-confidence data programming," *ACM Transactions on Computing for Healthcare*, vol. 3, no. 4, pp. 1–24, 2022.
- [24] F. R. Noyes, "The function of the human anterior cruciate ligament and analysis of single-and double-bundle graft reconstructions," *Sports Health*, vol. 1, no. 1, pp. 66–75, 2009.
- [25] N. A. Beck, J. T. R. Lawrence, J. D. Nordin, T. A. DeFor, and M. Tompkins, "Acl tears in school-aged children and adolescents over 20 years," *Pediatrics*, vol. 139, no. 3, 2017.
- [26] J. S. Frank and P. L. Gambacorta, "Anterior cruciate ligament injuries in the skeletally immature athlete: diagnosis and management," *JAAOS-Journal of the American Academy of Orthopaedic Surgeons*, vol. 21, no. 2, pp. 78–87, 2013.
- [27] N. A. Bates, A. L. McPherson, M. B. Rao, G. D. Myer, and T. E. Hewett, "Characteristics of inpatient anterior cruciate ligament reconstructions and concomitant injuries," *Knee Surgery, Sports Traumatology, Arthroscopy*, vol. 24, no. 9, pp. 2778–2786, 2016.
- [28] N. D. Schilaty, C. Nagelli, N. A. Bates, T. L. Sanders, A. J. Krych, M. J. Stuart, and T. E. Hewett, "Incidence of second anterior cruciate ligament tears and identification of associated risk factors from 2001 to 2010 using a geographic database," *Orthopaedic journal of sports medicine*, vol. 5, no. 8, p. 2325967117724196, 2017.
- [29] M. Buckthorpe, "Recommendations for movement re-training after acl reconstruction," *Sports Medicine*, pp. 1–18, 2021.

- [30] K. E. Webster, C. V. Nagelli, T. E. Hewett, and J. A. Feller, "Factors associated with psychological readiness to return to sport after anterior cruciate ligament reconstruction surgery," *The American journal of sports medicine*, vol. 46, no. 7, pp. 1545–1550, 2018.
- [31] A. Kotsifaki, R. Whiteley, S. Van Rossom, V. Korakakis, R. Bahr, V. Sideris, P. Graham-Smith, and I. Jonkers, "Single leg hop for distance symmetry masks lower limb biomechanics: time to discuss hop distance as decision criterion for return to sport after acl reconstruction?" *British Journal of Sports Medicine*, 2021.
- [32] M. V. Paterno, B. Huang, S. Thomas, T. E. Hewett, and L. C. Schmitt, "Clinical factors that predict a second acl injury after acl reconstruction and return to sport: preliminary development of a clinical decision algorithm," *Orthopaedic journal of sports medicine*, vol. 5, no. 12, p. 2325967117745279, 2017.
- [33] H. Grindem, L. Snyder-Mackler, H. Moksnes, L. Engebretsen, and M. A. Risberg, "Simple decision rules can reduce reinjury risk by 84% after acl reconstruction: the delaware-oslo acl cohort study," *British journal of sports medicine*, vol. 50, no. 13, pp. 804–808, 2016.
- [34] A. Watson, P. Lu, E. Greenberg, J. T. R. Lawrence, T. J. Ganley, I. Lee, and J. Weimer, "Rt-acl: Identification of high-risk youth patients and their most significant risk factors to reduce anterior cruciate ligament reinjury risk," in *2021 IEEE/ACM Conference on Connected Health: Applications, Systems and Engineering Technologies (CHASE)*. IEEE, 2021, pp. 35–45.
- [35] M. Mirowski, "The automatic implantable cardioverter-defibrillator: an overview," *Journal of the American College of Cardiology*, vol. 6, no. 2, pp. 461–466, 1985.
- [36] H.-C. Dinges, S. Otto, D. K. Stay, S. Bäumlein, S. Waldmann, P. Kranke, H. F. Wulf, and L. H. Eberhart, "Side effect rates of opioids in equianalgesic doses via intravenous patient-controlled analgesia: a systematic review and network meta-analysis," *Anesthesia & Analgesia*, vol. 129, no. 4, pp. 1153–1162, 2019.
- [37] M. R. Cohen and J. Smetzer, "Patient-controlled analgesia safety issues," *Journal of pain & palliative care pharmacotherapy*, vol. 19, no. 1, pp. 45–50, 2005.
- [38] D. Arney, R. Jetley, P. Jones, I. Lee, and O. Sokolsky, "Formal methods based development of a pca infusion pump reference model: Generic infusion pump (gip) project," in *2007 Joint Workshop on High Confidence Medical Devices, Software, and Systems and Medical Device Plug-and-Play Interoperability (HCMDSS-MDPnP 2007)*. IEEE, 2007, pp. 23–33.
- [39] P. Masci, A. Ayoub, P. Curzon, I. Lee, O. Sokolsky, and H. Thimbleby, "Model-based development of the generic pca infusion pump user interface prototype in pvs," in *International Conference on Computer Safety, Reliability, and Security*. Springer, 2013, pp. 228–240.
- [40] T. Lam, M. Nagappa, J. Wong, M. Singh, D. Wong, and F. Chung, "Continuous pulse oximetry and capnography monitoring for postoperative respiratory depression and adverse events: a systematic review and meta-analysis," *Anesthesia & Analgesia*, vol. 125, no. 6, pp. 2019–2029, 2017.
- [41] Y. D'Arcy, "Eyeing capnography to improve pca safety," *Nursing2022*, vol. 37, no. 9, pp. 18–19, 2007.
- [42] R. Jetley and P. Jones, "Safety requirements based analysis of infusion pump software," *IEEE RTSS/SMDS*, pp. 310–325, 2007.
- [43] B. Kim, A. Ayoub, O. Sokolsky, I. Lee, P. Jones, Y. Zhang, and R. Jetley, "Safety-assured development of the gpca infusion pump software," in *2011 Proceedings of the Ninth ACM International Conference on Embedded Software (EMSOFT)*. IEEE, 2011, pp. 155–164.
- [44] E. National Academies of Sciences, Medicine *et al.*, *Social isolation and loneliness in older adults: Opportunities for the health care system*. National Academies Press, 2020.
- [45] B. B. Neves, R. L. Franz, C. Munteanu, R. Baecker, and M. Ngo, "'my hand doesn't listen to me!' adoption and evaluation of a communication technology for the 'oldest old,'" in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 2015, pp. 1593–1602.