Workshop on Machine Learning on Low-Power Devices: Applications and Advanced Topics | (SMR 3926)

06 May 2024 - 10 May 2024
Virtual

01 - ARENAS Brayan Alexis

Development of an algorithm that predicts hand movement in the game rock, paper and scissors with the use of TinyML and arduino nano BLE 33

02 - BABALOLA John Oluwaseun

TinyML: Possibilities, Trends, Prospects, and Challenges in Power Systems

03 - CAPOGROSSO Luigi

A Machine Learning-oriented overview on Tiny Machine Learning: Theory, Methods and Applications

04 - CASTRO ABDALLAH Gustavo Pablo

Development of an embedded system for the prediction of humidity in hydroponic germination phenolic sponge based on RNN/LSTM (Agronomic Industry)

05 - DE SOUZA FARIAS Tiago

Local Feature Alignment for Efficient TinyML Training on Low-Power Devices

06 - QUISPE ALVARADO Jhoel Andres

Development of an animal tracking system using Tiny Machine Learning kits

07 - SHAH Parin Jatinkumar

TinyML Devices are Vulnerable: A Study of Attack kill chain for TinyML Devices

08 - TANG Jialu

Out-of-Distribution Detection in Medical Time-Series Models

R01

Development of an algorithm that predicts hand movement in the game rock, paper and scissors with the use of TinyML and arduino nano BLE 33

Brayan A. Arenas F.[1], **Silvia A. Sotelo-López**[2]

[1]*Faculty of Electrical and Electronic Engineering, Universidad Pontificia Bolivariana, Bucaramanga, Santander, Colombia*

[2]*Department of Basic Sciences, Universidad Pontificia Bolivariana, Bucaramanga, Santander, Colombia*

In the game of rock, paper, scissors, your hands are one of the key elements. This popular game involves hand gestures representing these iconic figures. With the aim to develop a TinyML algorithm capable of predicting three categories (rock, paper or scissors) during the game, we collect kinetic data from an accelerometer, a gyroscope, and a magnetometer with 3-axis using the Arduino Nano BLE with the LSM9DS1 IMU. For the experiment, two opponents wearing a glove equipped with the microcontroller played the hand game[1]. In each round was captured the kinetic signals over a 10-second period, the movement category and the winner motion. Additionally, the Edge Impulse platform was used as a data collection tool for the database [2].

For the prediction of hand gestures, a comparison between two models was proposed: one based on dense neural networks and the other based on 1D convolutional networks. The evaluation of the performance for classification and deployment was considered. With this work it was possible to recognize the advantage to use TinyML technology using low-power devices in a resource-limited environment [3].

[1] Arduino, “Nano 33 BLE - Arduino Documentation”, Arduino.cc, (2024). Disponible en: docs.arduino.cc/hardware/nano-33-ble.

[2] Edge Impulse, “Documentation for ML Professionals - Edge Impulse”, EdgeImpulse.com, (2024). Available at: docs.edgeimpulse.com/docs/readme/for-ml-practitioners.

[3] Khalife, R., Mrad, R., Dabbous, A., Ibrahim, A. (2024). Real-Time Implementation of Tiny Machine Learning Models for Hand Motion Classification. In: Bellotti, F., et al. Applications in Electronics Pervading Industry, Environment and Society. ApplePies 2023. Lecture Notes in Electrical Engineering, vol 1110. Springer, Cham. https://doi.org/10.1007/978-3-031-48121-5-70

R02

TinyML: Possibilities, Trends, Prospects, and Challenges in Power Systems

Babalola, John Oluwaseun [1]

[1] *Electrical/Electronic Engineering Programme, Bowen University, Iwo, Nigeria*

Electric power systems have evolved into more sophisticated and complex structures, becoming increasingly delicate, dynamic, and demanding. The integration of dynamic components such as renewable energy sources (RESs), Smart Grids, distributed energy resources (DERs), and electric vehicles (EVs) necessitates that power systems be more efficient and robust to handle the demands placed on them. To enhance reliability and efficiency, machine learning (ML) algorithms like Artificial Neural Network (ANN), Decision Trees (DT), and Support Vector Machines (SVM) have been employed at various levels in modern power systems [1], [2], [3], [4]. "Internet of Things" (IoT) devices are integrated into power systems to collect comprehensive and complex data, which is then analyzed using machine learning algorithms to achieve goals such as accurate forecasting, fault detection, and protection [5], [6].

TinyML, an ecosystem of hardware, software, and algorithms supported by a growing community, enables ML models to run on resource-constrained IoT devices [7]. Within this framework, TinyML can be defined as "machine learning-aware architectures, frameworks, techniques, tools, and approaches capable of performing on-device analytics at mW (or below) power range settings for various sensing modalities (vision, audio, speech, motion, chemical, physical, textual, cognitive), primarily targeting battery-operated embedded edge devices suitable for large-scale implementation in the IoT or wireless sensor network domain" [8].

IoT is projected to expand rapidly and surpass 24.1 billion devices worldwide by 2030, with nearly four devices per person, owing to its numerous potential applications. In 2019, devices collected less than 20 zettabytes of data, and by 2025, it is estimated that IoT device data volume will reach 79.4 zettabytes [9]. Given this massive data volume and variety, known as Big Data, customized platforms and methods are required to handle data in power systems [10]. Compared to traditional machine learning, TinyML offers low latency and real-time predictions/inferences by shifting computations from servers to edge devices, enabling them to make independent choices without relying on cloud servers. Edge computing is facilitated by combining the data collection and data analytics layers, avoiding long-distance transmissions and enabling real-time decision-making [7], [11].

However, IoT devices are often engineered with limited processing capabilities and resources to save costs, prioritizing economic concerns over performance. Consequently, IoT devices or edge servers are initially limited to basic data processing tasks due to their constrained capacity and may struggle with computationally demanding tasks [12]. Currently, TinyML does not support training models directly on edge devices. Instead, models are developed using computational resources and downscaled to fit embedded devices. The compacted models, represented as C arrays, are then deployed on devices for inference using data collected by sensors [7]. Researchers are focused on creating better algorithms to optimize resources for deployment on small devices.

Machine learning-based approaches in power systems have advantages over traditional model-based approaches due to their independence from system models and parameters. Machine learning techniques use statistically driven exploratory data analysis to develop solutions/inferences. However, the computational requirements for this process cannot be met by low-power devices, greatly limiting the adoption of TinyML in power systems applications [4]. Nevertheless, recent successes with on-board training of edge devices using federated learning (FL) and transfer learning (TL) suggest that the application of TinyML in the power system industry is likely to garner much interest from researchers [7].

Typical ML applications in renewable generation forecasting, for instance, can benefit greatly from FL in TinyML. ML algorithms are used for more accurate forecasting to ensure system stability amidst variable energy supply mixes in power systems today [4]. Federated learning of TinyML devices can be deployed across various sources to enable more accurate and real-time predictions, thus ensuring system stability. Similarly, load forecasting can be significantly improved with the use of TinyML for real-time and more accurate forecasting, using FL across TinyML devices and smart meters deployed at the consumer end. Smart meters deployed with TinyML can also be utilized for key tasks such as load analysis and management [5], [13].

The interaction between energy users, power providers, and system operators becomes increasingly complex due to the growing adoption of smart grid applications and the global restructuring of the electricity market. The dynamic interactions between suppliers and customers contribute to the unpredictability of power prices, emphasizing the importance of accurate energy price forecasting for stakeholders in the electricity market employing optimal bidding methods and risk management [4], [14]. Federated learning of TinyML devices can be used to implement deep learning frameworks for day-ahead forecasting of electricity prices as demonstrated by Zhang, Li, and Ma (2020) [14]. Similarly, the study conducted by Pourdaryaei et al. (2019) [15] employing a two-stage feature selection approach and an optimized adaptive neuro-fuzzy inference system methodology to develop a forecasting engine for energy price forecasting could greatly benefit from TinyML by training the devices on feature selection and implementing FL of the TinyML framework for the neuro-fuzzy system.

Additionally, due to the complex and variable nature of power grids, they are more susceptible to issues like storms, fires, earthquakes, and cyberattacks due to increased integration of RES and DERs. Early identification and detection of faults are essential to ensure the power system operates efficiently and securely. TinyML, if deployed along the power system, can be used to achieve accurate and fast fault detection and classification on transmission lines [16]. TinyML could also be used in microgrid island detection by extracting and processing valuable information or features and employing deep federated learning to identify if the microgrid has been islanded [17].

Despite its potential, TinyML faces challenges, primarily limited memory and power as low-power edge devices, greatly restricting their computing capability. Another major challenge is the difficulty in updating TinyML devices since training is typically done elsewhere, and the model is only compressed into a light model to run and make inferences on the device. This significantly limits the ability of devices to update their learning process and stay up to date [12].

[1] T. Chen and C. Liu, ‘Soft computing based smart grid fault detection using computerised data analysis with fuzzy machine learning model’, *Sustain. Comput. Inform. Syst.*, vol. 41, p. 100945, Jan. 2024, doi: 10.1016/j.suscom.2023.100945.

[2] C. Ying, W. Wang, J. Yu, Q. Li, D. Yu, and J. Liu, ‘Deep learning for renewable energy forecasting: A taxonomy, and systematic literature review’, *J. Clean. Prod.*, vol. 384, p. 135414, Jan. 2023, doi: 10.1016/j.jclepro.2022.135414.

[3] O. A. Alimi, K. Ouahada, and A. M. Abu-Mahfouz, ‘A Review of Machine Learning Approaches to Power System Security and Stability’, *IEEE Access*, vol. 8, pp. 113512–113531, 2020, doi: 10.1109/ACCESS.2020.3003568.

[4] M. Farhoumandi, Q. Zhou, and M. Shahidehpour, ‘A review of machine learning applications in IoT-integrated modern power systems’, *Electr. J.*, vol. 34, no. 1, p. 106879, Jan. 2021, doi: 10.1016/j.tej.2020.106879.

[5] E. Hossain, I. Khan, F. Un-Noor, S. S. Sikander, and Md. S. H. Sunny, ‘Application of Big Data and Machine Learning in Smart Grid, and Associated Security Concerns: A Review’, *IEEE Access*, vol. 7, pp. 13960–13988, 2019, doi: 10.1109/ACCESS.2019.2894819.

[6] S. E. Collier, ‘The Emerging Enernet: Convergence of the Smart Grid with the Internet of Things’, *IEEE Ind. Appl. Mag.*, vol. 23, no. 2, pp. 12–16, Mar. 2017, doi: 10.1109/MIAS.2016.2600737.

[7] M. Ficco, A. Guerriero, E. Milite, F. Palmieri, R. Pietrantuono, and S. Russo, ‘Federated learning for IoT devices: Enhancing TinyML with on-board training’, *Inf. Fusion*, vol. 104, p. 102189, Apr. 2024, doi: 10.1016/j.inffus.2023.102189.

[8] P. P. Ray, ‘A review on TinyML: State-of-the-art and prospects’, *J. King Saud Univ. - Comput. Inf. Sci.*, vol. 34, no. 4, pp. 1595–1623, Apr. 2022, doi: 10.1016/j.jksuci.2021.11.019.

[9] S. F. Ahmed, Md. S. B. Alam, S. Afrin, S. J. Rafa, N. Rafa, and A. H. Gandomi, ‘Insights into Internet of Medical Things (IoMT): Data fusion, security issues and potential solutions’, *Inf. Fusion*, vol. 102, p. 102060, Feb. 2024, doi: 10.1016/j.inffus.2023.102060.

[10] H. Yang, X. Liu, D. Zhang, T. Chen, C. Li, and W. Huang, ‘Machine learning for power system protection and control’, *Electr. J.*, vol. 34, no. 1, p. 106881, Jan. 2021, doi: 10.1016/j.tej.2020.106881.

[11] M. Mohammadi, A. Al-Fuqaha, S. Sorour, and M. Guizani, ‘Deep Learning for IoT Big Data and Streaming Analytics: A Survey’, *IEEE Commun. Surv. Tutor.*, vol. 20, no. 4, pp. 2923–2960, 2018, doi: 10.1109/COMST.2018.2844341.

[12] L. Yang and A. Shami, ‘IoT data analytics in dynamic environments: From an automated machine learning perspective’, *Eng. Appl. Artif. Intell.*, vol. 116, p. 105366, Nov. 2022, doi: 10.1016/j.engappai.2022.105366.

[13] Y. Wang, Q. Chen, T. Hong, and C. Kang, ‘Review of Smart Meter Data Analytics: Applications, Methodologies, and Challenges’, *IEEE Trans. Smart Grid*, vol. 10, no. 3, pp. 3125–3148, May 2019, doi: 10.1109/TSG.2018.2818167.

[14] R. Zhang, G. Li, and Z. Ma, ‘A Deep Learning Based Hybrid Framework for Day-Ahead Electricity Price Forecasting’, *IEEE Access*, vol. 8, pp. 143423–143436, 2020, doi: 10.1109/ACCESS.2020.3014241.

[15] A. Pourdaryaei, H. Mokhlis, H. A. Illias, S. Hr. A. Kaboli, and S. Ahmad, ‘Short-Term Electricity Price Forecasting via Hybrid Backtracking Search Algorithm and ANFIS Approach’, *IEEE Access*, vol. 7, pp. 77674–77691, 2019, doi: 10.1109/ACCESS.2019.2922420.

[16] T. S. Abdelgayed, W. G. Morsi, and T. S. Sidhu, ‘A New Approach for Fault Classification in Microgrids Using Optimal Wavelet Functions Matching Pursuit’, *IEEE Trans. Smart Grid*, vol. 9, no. 5, pp. 4838–4846, Sep. 2018, doi: 10.1109/TSG.2017.2672881.

[17] A. A. Abdelsalam, A. A. Salem, E. S. Oda, and A. A. Eldesouky, ‘Islanding Detection of Microgrid Incorporating Inverter Based DGs Using Long Short-Term Memory Network’, *IEEE Access*, vol. 8, pp. 106471–106486, 2020, doi: 10.1109/ACCESS.2020.3000872.

A Machine Learning-oriented overview on Tiny Machine Learning: Theory, Methods and Applications

Luigi Capogrosso[1], **Franco Fummi**[1], **and Marco Cristani**[1,2]

[1] *(Presenting author underlined) Department of Engineering for Innovation Medicine, University of Verona, Verona, Italy*

[2] *QUALYCO S.r.l, Spin-off of the University of Verona, Verona, Italy*

Over the past decades, a prodigious amount of research has been invested in improving embedded technologies to enable real-time solutions for many complex and safety-critical applications. In this regard, hardware-specific (e.g., Edge TPUs) and Micro-Controller Unit (MCU)-based embedded systems have earned a lot of attention, primarily due to their low power requirements and high performance, and secondarily for their maintainability, adaptability, and reliability.

From these premises, since 2018, **Tiny Machine Learning (TinyML)** has positively revolutionized the field of Artificial Intelligence by promoting the joint design of resource-constrained IoT hardware devices and their learning-based software architectures. TinyML carries an essential role within the fourth and fifth industrial revolutions in helping societies, economies, and individuals employ effective AI-infused computing technologies (e.g., smart cities, automotive, and medical robotics).

The challenges for TinyML practitioners are formidable, e.g., in modern neural networks, among the best currently available technologies, the number of required parameters has skyrocketed to the order of billions, with larger networks having better results and broader applicability. Unfortunately, the energy required to run these networks is proportional to their size, making this trend of scaling up neural networks energetically unsustainable at large scales. This is another reason why TinyML has to be considered as a necessary, other than promising, research direction.

Given its multidisciplinary nature, the field of TinyML has been approached from many different angles: we aim to present and discuss with you this comprehensive survey [1], in which we provide an up-to-date overview focused on all the **learning algorithms within TinyML-based solutions**. The survey is based on the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) methodological flow, allowing for a systematic and complete literature survey.

In particular, we examine the three **different workflows for implementing a TinyML-based system**, i.e., ML-oriented, HW-oriented, and co-design. As a further and unique contribution, the survey emphasizes the Machine Learning point of view, not only reporting the latest trends in **TinyML frameworks** but also suggesting recent variations and advancements in the Machine Learning technologies that a TinyML practitioner may want to explore to improve on the state-of-the-art. In this regard, we also covered the principles around designing **TinyML model architectures, hardware-aware training strategies, effective inference optimizations, and benchmarking methodologies**. This unique combination equips readers in both academic and industrial spheres with universal concepts essential for implementing TinyML in production settings. In this review we also present the distinct features of **hardware devices and software tools that represent the current state-of-the-art for TinyML edge applications**. Finally, we discuss the challenges and future directions.

[1] Capogrosso, Luigi, et al. “A Machine Learning-oriented Survey on Tiny Machine Learning.” IEEE Access (2024).

R04

Case study: Development of an embedded system for the prediction of humidity in hydroponic germination phenolic sponge based on RNN/LSTM (Agronomic Industry).

Gustavo P. Castro Abdallah

Universidad Nacional de Rafaela (Argentina)

In this presentation, I will present an innovative project addressing a crucial challenge in hydroponic agronomy: the precise prediction of substrate moisture during the germination stage. Our proposal focuses on constructing an intelligent sensor integrating Machine Learning technologies to optimize this crucial phase of the cultivation process.

The project spans multiple stages, from the design and development of equipment for acquiring environmental and substrate data to creating a robust dataset reflecting real conditions in the growing environment. We will highlight the process of developing a Recurrent Neural Network (RNN) or Long Short-Term Memory (LSTM) network, suitable for modeling the sequential nature of temporal data.

I will delve into the training and evaluation process of the neural network, emphasizing strategies to optimize model accuracy and generalization. Additionally, we will discuss the effective deployment of the trained model on ESP32 microcontrollers, enabling its practical integration into real-time monitoring and control systems.

By the end of this presentation, participants will gain a deep understanding of how the integration of artificial intelligence and sensor engineering can significantly enhance efficiency and productivity in the germination stage of hydroponic agronomy.

R05

Local Feature Alignment for Efficient TinyML Training on Low-Power Devices

Tiago de Souza Farias[1], Amanda G. Valério[1]

[1]*Physics Departament, Federal University of São Carlos, São Carlos, SP, Brazil*

The prevalence of artificial intelligence (AI) has increased markedly in recent years, becoming a pervasive aspect of contemporary life. According to a survey in [1], 50% of respondents claim AI-powered products and services have significantly impacted their lives. Additionally, 57% of participants anticipate that AI will alter their future work practices. Furthermore, this study indicates that the cost of training and the size of AI models are on the rise annually.

Training many machine models are computationally intensive primarily because they utilize the backpropagation algorithm for parameter optimization. This method constructs a computational graph, demanding substantial memory and energy resources. A notable example of this is the training of GPT-3, which required approximately 175 billion parameters and 1,287 MWh of energy consumption. [1]

One potential solution is to identify neural networks that are more resource-efficient, allowing machine learning development on low-power devices [2, 3]. One approach to this end involves the modification of the training rule, whereby the use of backpropagation could be replaced with local training algorithms [4, 5]. These algorithms focus on restricting the differentiable graph to local operations within the network. By segmenting the network into autonomous blocks that can be dynamically loaded and unloaded from memory as needed, the overall memory demand is significantly reduced.

In this work, we investigate local feature alignment [5, 6] as a strategy to reduce memory consumption during neural network training. This method, designed for local training approaches, predicts the input of a specific local region within the network based on its output. This predictive mechanism allows for the training of network parameters using only local data, thereby enabling the network to invert its outputs.

We demonstrate the feasibility of this technique through its application to a regression and a classification problem in computer vision. The neural network is optimized to reconstruct images from the latent vector, which is represented as the last layer of the network. This is achieved by employing an encoder-only network configuration. Furthermore, local training is utilized to successfully learn to predict the images, effectively reducing memory consumption. This practical application not only demonstrates the effectiveness of the technique for solving a problem but also illustrates its potential to run machine learning on low-powered devices.

[1] N. Maslej et al.,“The AI Index 2024 Annual Report,” AI Index Steering Committee, Institute for Human-Centered AI, Stanford University, Stanford, CA, April 2024.

[2] A. Goel, C. Tung, Y. -H. Lu, G. K. Thiruvathukal, 2020 IEEE 6th WF-IoT, pp. 1-6 (2020)

[3] S. Kuninti, S. Rooban, Journal of Physics: Conference Series, 012169 (2021).

[4] P. Baldi,P. Sadowski, Neural networks : the official journal of the International Neural Network Society. **83**, 51-74 (2015).

[5] T. de S. Farias, PhD’s thesis, UFSM (2023).

[6] T. de S. Farias, J. Maziero, Frontiers in Artificial Inteligence. Vol. **5** (2023).

R06

Development of an animal tracking system using Tiny Machine Learning kits

Jhoel Quispe Alvarado[1], **Nicolás Catalano**[1,2], **Luis H. Arnaldi**[2] **and Laila Kazimierski**[3]

[1] *Instituto Balseiro, Universidad Nacional de Cuyo, Comisión Nacional de Energía Atómica (CNEA),*
San Carlos de Bariloche R8402AGP, Argentina
[2] *Centro Atómico Bariloche, Comisión Nacional de Energía Atómica (CNEA),*
San Carlos de Bariloche R8402AGP, Argentina
[3] *Consejo Nacional de Investigaciones Científicas y Técnicas, Centro Atómico Bariloche (CONICET)*

The objective of this project is to design and develop a small, energy-efficient monitoring system (hardware) for tracking animals in their natural habitat using TinyML kits. We utilized data from monitoring campaigns on the movement of Chelonoidis chilensis, a species in a vulnerable state, to design various classification models for animal behavior. These kits allow embedding a pre-trained neural network that enables real-time behavior classification. Additionally, we equipped each kit with GPS, battery, and memory to facilitate their use in animal monitoring. It should be noted that, in the case of species in vulnerable states, understanding how individuals move daily and seasonally helps establish guidelines that contribute to their conservation.

It's important to note that, for species in vulnerable states, understanding their daily and seasonal movement patterns helps establish conservation guidelines.

In conclusion, a concrete application of a study like this will allow us to classify different behaviors such as eating, walking, mating, digging nests to lay eggs, etc. In particular, finding the deposited eggs by the female to help in their conservation. The talk is part of the final project I am carrying out at the Instituto Balseiro to obtain a degree in Telecommunications Engineering. It is also part of an ongoing interdisciplinary project involving engineers, physicists, and biologists from the Centro Atómico Bariloche and other institutions, Argentina.

R07

TinyML Devices are Vulnerable: A Study of Attack kill chain for TinyML Devices

Parin Shah, Yuvaraj Govindarajulu, Pavan Kulkarni and Manojkumar Parmar

AIShield, Bosch Global Software Technology, Bengaluru, India

Recent advancements in artificial intelligence (AI) and machine learning (ML) have led to the development of TinyML [1], enabling AI computations on resource constrained devices without relying on cloud connections. With improved bandwidth and reduced latency TinyML, has the potential to decentralize cloud applications, marking the beginning of a new era of distributed intelligence. These deployment of AI/ML software on hardware throughout the real world, have witnessed significant adoption across industries such as plant automation, factory robots and edge computing. Despite offering rapid data analysis and real-time responses crucial for various applications, TinyML devices face security risks. This is mainly due to their vulnerable deployment environments and the shortcomings of the embedded security mechanisms to protect against AI attack such as adversarial attacks[2], where the attacks potentially compromise the integrity and confidentiality of sensitive data or disrupt critical operations.

In our research, we demonstrate how adversaries can transfer adversarial attacks from powerful host machines to smaller, less secure devices like the used in industrial and IoT applications, highlighting an extension of adversarial threats to tiny devices. Considering an attack kill chain from MITRE ATLAS[4], we show how the attacker can compromise functionality on the tiny device. Here adversary conducts reconnaissance to identify target AI models and analyze their architecture, framework, and deployment specifics which exploits vulnerabilities or intercept communication to gain unauthorized access. Then, through a powerful attack such as 'Model Extraction [3]', which is aimed at reverse-engineering and replicating the AI model, the adversary rebuilds an AI model. Finally, the adversary make use of this stolen model to craft adversarial samples from a powerful system to compromise the low-power device.

The MITRE ATLAS[4] kill chain, when applied to TinyML low-power devices, highlights potential vulnerabilities and attack stages for small-scale AI models operating on constrained hardware. These devices, often found in IoT and edge computing, are susceptible to reconnaissance and exploitation due to limited resources and security controls. In TinyML devices, an attacker could manipulate sensor inputs (delivery) to influence model outcomes (exploitation), compromising the device's reliability. The kill chain underscores the importance of securing data pathways, implementing robust anomaly detection, and maintaining secure device firmware to mitigate risks.

Finally, we propose defense mechanisms as migitation strategy to strengthen the security posture on TinyML devices.

[1] Ray, P.P.: A review on tinyml: State-of-the-art and prospects. Journal of King Saud University - Computer and Information Sciences 34(4), 1595–1623 (2022).

[2] Biggio, B., Roli, F.: Wild patterns: Ten years after the rise of adversarial machine learning. In: Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security. pp. 2154–2156 (2018).

[3] Tram'er, F., Zhang, F., Juels, A., Reiter, M.K., Ristenpart, T.: Stealing machine learning models via prediction apis (2016)

[4] MITRE ATLAS (2024), https://atlas.mitre.org/, [Accessed 25-04-2024]

R08

Out-of-Distribution Detection in Medical Time-Series Models

Jialu Tang[1], Yuan Lu[1], Jungwoo Oh[2], Edward Choi[2], and Aaqib Saeed[1,3]

[1] *Eindhoven University of Technology, The Netherlands*
[2] *KAIST, South Korea*
[3] *Eindhoven Artificial Intelligence Systems Institute, Eindhoven University of Technology, The Netherlands*

The advent of TinyML has led to significant attention for its application in personalized health monitoring due to its real-time data processing capability on local devices [1]. In the clinical domain, embedded sensors and data-driven models based on TinyML can identify abnormalities in electrocardiogram (ECG) signals, leading to timely and appropriate treatments. These models have demonstrated exceptional performance in detecting arrhythmic abnormalities such as paroxysmal Atrial Fibrillation [2]. However, ensuring the reliability and safety of data-driven systems is critical in high-stakes clinical settings [3]. This necessitates not only exceptional generalizability on known data distributions but also reliable identification of samples that fall outside of these distributions.

To address this challenge, we propose ARMOR, a novel framework that leverages robust representation learning and multi-task adversarial training to enhance the out-of-distribution (OOD) detection performance in medical time-series data. ARMOR is designed to operate efficiently on resource-constrained devices, making it suitable for broad range of applications. By learning shared representations that capture the essential characteristics of in-distribution (ID) samples and effectively discriminate them from OOD samples, ARMOR enables reliable anomaly detection in real-time on low-cost IoT devices with k-nearest neighbors (kNN).

We evaluate the effectiveness of ARMOR through extensive experiments on ECG and EEG datasets. Our results demonstrate that ARMOR significantly improves the OOD detection performance compared to existing methods. Specifically, ARMOR achieves a 15% and 11% reduction in the FPR95 (False Positive Rate at 95% True Positive Rate) for ECG and EEG datasets, respectively, when using a distance-based kNN detector. These improvements highlight the ability of ARMOR to learn robust representations that capture the essential characteristics of ID samples while effectively distinguishing them from OOD samples.

Furthermore, our work establishes a benchmark for achieving OOD detection methods for medical time-series data. This paves the way for future research in developing robust and reliable models that can operate within the constraints of low-cost compute. By addressing the challenges of OOD detection in resource-constrained environments, ARMOR will contribute to the advancement of trustworthy and efficient personalized health monitoring systems.

[1] Zhang, Angela and Xing, Lei and Zou, James and Wu, Joseph C, **Shifting machine learning for healthcare from development to deployment and from models to data**, Nature Biomedical Engineering, 6, (2022).

[2] Baek, Yong-Soo and Lee, Sang-Chul and Choi, Wonik and Kim, Dae-Hyeok, **A new deep learning algorithm of 12-lead electrocardiogram for identifying atrial fibrillation during sinus rhythm**, Scientific reports, 11,(2021).

[3] Feng, Jean and Phillips, Rachael V and Malenica, Ivana and Bishara, Andrew and Hubbard, Alan E and Celi, Leo A and Pirracchio, Romain, **Clinical artificial intelligence quality improvement: towards continual monitoring and updating of AI algorithms in healthcare**, NPJ digital medicine, 5,(2022).