Embedded AI for NIME: Challenges and Opportunities

**Teresa Pelinski Victor Shepardson Steve Symons
Franco Santiago Caspe Adan L Benito Temprano Jack Armitage
Chris Kiefer Rebecca Fiebrink Thor Magnusson Andrew McPherson**

Description (up to 750 words)

Cutting edge embedded systems have always been a part of NIME's practices. Low-resource computing hardware, such as microcontrollers or single-board computers, can be embedded into digital musical instruments or interfaces to perform specific functions such as real-time digital signal processing of sensor data and sound [1][2][3][4][5]. Simultaneously, an interest in exploiting the creative potentials of artificial intelligence (AI) for instrument design and musical expression has been growing within the NIME community in the past years [6][7][8][9][10][11].

Recent advancements in embedded computing have allowed for faster and more intensive computation capabilities [12]. However, the deployment of machine learning or symbolic AI techniques still presents several technical challenges (e.g., data bandwidth, memory handling) and higher-level design constraints [13][14][15][16]. Some of these challenges are general to embedded systems, while others are specific to musical interaction, particularly questions regarding real-time performance and latency [17]. With this workshop, we aim to: (1) bring together a body of research practitioners that face such challenges in the context of NIME, (2) articulate these challenges and identify the tools and strategies being currently used to overcome them, (3) forge a community of practitioners of embedded AI for NIME and (4) discuss critical approaches on the use of embedded AI for musical expression.

Deploying AI models on embedded systems is an emerging and fast-changing field. A workshop is an excellent opportunity for practitioners to present works in progress and collaboratively identify shared challenges. We expect this workshop to serve as the starting point for an embedded AI NIME community and as future reference to help newcomers get started with embedded AI.

We will invite any NIME attendees interested in embedded AI to participate in the workshop and send a call for submissions open in terms of format (abstracts / short papers / progress reports / demos / posters) with themes that could include (but not necessarily restricted to):

- Any technical prototype or concrete implementation of resource-constrained systems using AI in the context of NIME
- Design strategies and conceptual frameworks for embedded AI
- Interaction paradigms for systems using embedded AI

- Embedded and real-time neural audio synthesis methods (e.g., neural, artificial life, statistical methods etc.)
- AR/MR/VR systems using AI in the context of NIME
- Mobile computing systems using AI in the context of NIME
- Musical uses of AI in embedded platforms
- Workflows for moving AI implementations between laptop, embedded/real-time and HPC platforms
- Development environments for interactive machine learning in embedded contexts, particularly those targeting non-expert users
- Values, biases, ethical and philosophical issues with embedded AI in musical performance
- Inclusivity and diversity in emerging embedded AI communities

The workshop will have a duration of a half-day split in two cycles spaced by five or six hours to accommodate different time zones. The first half-hour of each cycle will start with 5-minute presentations from each of the organiser labs (Emute Lab, Intelligent Instruments Lab, Augmented Instruments Lab) and 15-minute opening talks by Prof Rebecca Fiebrink and Dr Jack Armitage. The rest of the cycle will consist of brief talks or demos grouped in sessions by topic. After each session, we will have an open discussion with presenters and attendees moderated by the student organisers. We will use collaborative boards (padlets[1]) to share the discussion content among sessions. The last 10 minutes will serve as a wrap-up.

Prior to the event, the submissions will be made available to the workshop participants to familiarise themselves with the contents if they wish to. We will also open a Discord server to facilitate interaction before, during and after the workshop. Moreover, we are aware that embedded AI is not only of interest for researchers in academia and wish to include others, such as makers or hackers, with no institutional affiliations. For this reason, we intend to make the workshop's video recording and submission materials publicly available after the event, with the intention that it will remain for future reference. We expect that this workshop will serve as the preliminary of a series of events that will gather both researchers and makers. We also anticipate that communities such as TinyML, the Julia language, and VR/MR/AR may share many technical concerns and would benefit from mutual exchange with an organised 'embedded AI for NIME' community.

Short Description (up to 70 words)

Despite recent advancements in low-resource computing hardware, such as microcontrollers or single board computers, the deployment of machine learning or symbolic artificial intelligence (AI) techniques still presents several technical challenges and higher-level design constraints. With this workshop, we aim to articulate these challenges in the context of NIME and forge a community of embedded AI for NIME practitioners through a series of talks, demos and collaborative discussions.

Organizers

Teresa Pelinski

Centre for Digital Music, Queen Mary University of London, t.pelinskiramos@qmul.ac.uk

I am a PhD researcher at the Artificial Intelligence and Music CDT and a member of the Augmented Instruments Lab[2] at the Centre for Digital Music, QMUL. I hold a BSc in Physics from Universidad Autónoma de Madrid and a MSc in Sound and Music Computing from Universitat Pompeu Fabra, in Barcelona. My PhD project is in collaboration with Bela[3] and deals with capturing nuanced, high-bandwidth interaction using deep learning techniques in embedded platforms in the context of digital musical instruments.

Victor Shepardson

Intelligent Instruments Lab, Iceland University of the Arts, victor@lhi.is

I am a doctoral student in the Intelligent Instruments Lab at LHI. Previously I worked on neural models of speech as a machine learning engineer and data scientist. Before that I was an MA student in Digital Musics at Dartmouth College and BA student in Computer Science at the University of Virginia. My interests include machine learning, artificial intelligence, generative art, audiovisual music and improvisation. My current project involves building an AI augmented looping instrument and asking what AI means to people, anyway.

Steve Symons

University of Sussex, Brighton, s.symons@sussex.ac.uk

I am a PhD researcher at the Leverhulme Trust funded be.AI Centre at Sussex University where my research is hosted by the School of Media, Arts and Humanities. I have spent many years as a maker of embedded locative audio systems and of making and improvising music with NIMEs. I am interested in enactive interfaces, woodwork and finding new metaphors for collaborative instruments.

Franco S. Caspe

Centre for Digital Music, Queen Mary University of London, f.s.caspe@qmul.ac.uk

I am a PhD researcher at the Artificial Intelligence and Music CDT and a member of the Augmented Instruments Laboratory of the Centre for Digital Music, QMUL. I hold an Electronic Engineer degree, and MSc in Image Processing and Computer Vision. I worked on R&D of real-time systems for audio, communications, and image classification on diverse platforms from micro-controllers to FPGAs. My PhD project is about musical instrument expression modelling using AI, for informed timbre transfer and instrument retargeting.

Adan L. Benito

Centre for Digital Music, Queen Mary University of London, a.benitotemprano@qmul.ac.uk

I am a PhD researcher from the Augmented Instruments Laboratory at the Centre for Digital Music in QMUL, and a member of the Artificial Intelligence and Music CDT. I am also part of the active development team of the Bela platform and have an active interest in the development of new hardware tools for music-making. I graduated as a Telecommunications Engineer from the University of Cantabria with an MSc in Radio Communications and hold an MSc in Sound and Music Computing from QMUL. My current research focuses on the creation of gestural models fusing representations from sensor and audio domains and their application to instrument augmentation. I also have an interest in all guitar-related technologies and the cultures that surround them.

Jack Armitage

Intelligent Instruments Lab, Iceland University of the Arts, jack@lhi.is

I am a postdoctoral research fellow in the Intelligent Instruments Lab[4]. I have a doctorate in Media and Arts Technologies from Queen Mary University of London, where I studied in Prof. Andrew McPherson's Augmented Instruments Lab. During my

PhD I was a Visiting Scholar at Georgia Tech under Prof. Jason Freeman. Before then, I was a Research Engineer at ROLI after graduating with a BSc in Music, Multimedia & Electronics from the University of Leeds. My research interests include embodied interaction, craft practice and design cognition. I also produce, perform and live code music as Lil Data, as part of the PC Music record label.

Chris Kiefer

Experimental Music Technologies Lab, Department of Music, University of Sussex, UK
c.kiefer@sussex.ac.uk

I am a computer-musician, musical instrument designer and Senior Lecturer in Music Technology at the University of Sussex. As a live-coder I perform under the name 'Luuma'. Recently I have been playing an augmented self-resonating cello as half of improv-duo Feedback Cell, and with the feedback-drone-quartet 'Brain Dead Ensemble'. I co-run the AHRC Feedback Musicianship Network. My research specialises in musician-computer interaction, physical computing, machine learning and complex systems.

Rebecca Fiebrink

Creative Computing Institute, University of the Arts London
r.fiebrink@arts.ac.uk

I am a Professor of Creative Computing at UAL. My research focuses largely on exploring how machine learning can be used to augment human creative practices in music and beyond. My students, collaborators, and I have developed a number of widely used tools for creative end-user machine learning, including Wekinator and InteractML.

Thor Magnusson

Intelligent Instruments Lab, Iceland University of the Arts
thor.magnusson@lhi.is

I am a professor of future music in the Music Department at the University of Sussex and a research professor at the Iceland University of the Arts. I've recently served as an Edgard-Varèse guest professor at the Technische Universität Berlin. My research interests include musical performance, improvisation, new technologies for musical expression, live coding, musical notation, artificial intelligence and computational creativity.

Andrew McPherson

Centre for Digital Music, Queen Mary University of London, a.mcpherson@qmul.ac.uk

I am a Professor of Musical Interaction in QMUL's Centre for Digital Music, where I lead the Augmented Instruments Laboratory. I am also a founder of Bela[5], an embedded platform for rich, low-latency interaction with audio and sensors. With a background in music composition and electronic engineering, my interests include augmented instruments, performer-instrument interaction and foundational technologies for creating new digital musical instruments.

Preferred Length of Workshop

Half-day split into two cycles.

Links to Supporting Media

https://embedded-ai-for-nime.github.io/

Footnotes

1. https://padlet.com/ ↩

2. https://instrumentslab.org ↩

3. https://bela.io ↩

4. https://iil.is ↩

5. https://bela.io ↩

Citations

1. Sullivan, J., Vanasse, J., Guastavino, C., & Wanderley, M. (2020). Reinventing the Noisebox: Designing Embedded Instruments for Active Musicians. In R. Michon & F. Schroeder (Eds.), *Proceedings of the International Conference on New Interfaces for Musical Expression* (pp. 5–10). Birmingham City University. https://doi.org/10.5281/zenodo.4813166 ↩
2. Meneses, E., Wang, J., Freire, S., & Wanderley, M. (2019). A Comparison of Open-Source Linux Frameworks for an Augmented Musical Instrument Implementation. In M. Queiroz & A. X. Sedó (Eds.), *Proceedings of the International Conference on New Interfaces for Musical Expression* (pp. 222–227). UFRGS. https://doi.org/10.5281/zenodo.3672934 ↩

3. Momeni, A., McNamara, D., & Stiles, J. (2018). MOM: an Extensible Platform for Rapid Prototyping and Design of Electroacoustic Instruments. In T. M. Luke Dahl Douglas Bowman (Ed.), *Proceedings of the International Conference on New Interfaces for Musical Expression* (pp. 65–71). Virginia Tech. https://doi.org/10.5281/zenodo.1302681 ↵

4. Berdahl, E., Salazar, S., & Borins, M. (2013). Embedded Networking and Hardware-Accelerated Graphics with Satellite CCRMA. *Proceedings of the International Conference on New Interfaces for Musical Expression*, 325–330. https://doi.org/10.5281/zenodo.1178476 ↵

5. Schmeder, A., & Freed, A. (2008). uOSC : The Open Sound Control Reference Platform for Embedded Devices. *Proceedings of the International Conference on New Interfaces for Musical Expression*, 175–180. https://doi.org/10.5281/zenodo.1179627 ↵

6. Fiebrink, R., & Sonami, L. (2020). Reflections on Eight Years of Instrument Creation with Machine Learning. *Proceedings of the International Conference on New Interfaces for Musical Expression*, 237–242. https://doi.org/10.5281/zenodo.4813334 ↵

7. Martin, C., Morreale, F., Benedikte, W., & Hugo, S. (2021). Workshop in Critical Perspectives on AI/ML in Musical Interfaces. *Proceedings of the International Conference on New Interfaces for Musical Expression*. https://critical-ml-music-interfaces.github.io/ ↵

8. Tahiroğlu, K., Kastemaa, M., & Koli, O. (2020). Al-terity: Non-Rigid Musical Instrument with Artificial Intelligence Applied to Real-Time Audio Synthesis. *Proceedings of the International Conference on New Interfaces for Musical Expression*, 337–342. https://doi.org/10.5281/zenodo.4813402 ↵

9. Yaremchuk, V., Medeiros, C. B., & Wanderley, M. (2019). Small Dynamic Neural Networks for Gesture Classification with The Rulers (a Digital Musical Instrument). *Proceedings of the International Conference on New Interfaces for Musical Expression*, 150–155. https://doi.org/10.5281/zenodo.3672904 ↵

10. Hantrakul, L. (2018). GestureRNN: A neural gesture system for the Roli Lightpad Block. *Proceedings of the International Conference on New Interfaces for Musical Expression*, 132–137. https://doi.org/10.5281/zenodo.1302703 ↵

11. Caramiaux, B., & Tanaka, A. (2013). Machine Learning of Musical Gestures. *Proceedings of the International Conference on New Interfaces for Musical Expression*, 513–518. https://doi.org/10.5281/zenodo.1178490 ↵

12. Wang, F., Zhang, M., Wang, X., Ma, X., & Liu, J. (2020). Deep learning for edge computing applications: A state-of-the-art survey. *IEEE Access*, *8*, 58322–58336. ↵

13. Situnayake, D., & Plunkett, J. (2022). *AI at the Edge: Solving Real World Problems with Embedded Machine Learning*. O'Reilly Media, Inc. https://learning.oreilly.com/library/view/ai-at-the/9781098120191/ ↵

14. Blasch, E., Pham, T., Chong, C.-Y., Koch, W., Leung, H., Braines, D., & Abdelzaher, T. (2021). Machine learning/artificial intelligence for sensor data fusion– opportunities and challenges. *IEEE Aerospace and Electronic Systems Magazine*, *36*(7), 80–93. ↵

15. Shafique, M., Theocharides, T., Reddy, V. J., & Murmann, B. (2021). TinyML: Current Progress, Research Challenges, and Future Roadmap. *2021 58th ACM/IEEE Design Automation Conference (DAC)*, 1303–1306. ↵

16. David, R., Duke, J., Jain, A., Reddi, V. J., Jeffries, N., Li, J., Kreeger, N., Nappier, I., Natraj, M., Regev, S., Rhodes, R., Wang, T., & Warden, P. (2020). TensorFlow Lite Micro: Embedded Machine Learning on TinyML Systems. *CoRR*, *abs/2010.08678*. https://arxiv.org/abs/2010.08678 ↵

17. McPherson, A., Jack, R., & Moro, G. (2016). Action-Sound Latency: Are Our Tools Fast Enough? *Proceedings of the International Conference on New Interfaces for Musical Expression*, 20–25. https://doi.org/10.5281/zenodo.3964611 ↵