

6th International Workshop on Embedded and Mobile Deep Learning

WORKSHOP co-located with ACM MobiSys 2022
Portland, Oregon – Friday, 1 July 2022
<https://emdl22.github.io>



In recent years, breakthroughs from the field of deep learning (DL) have transformed how sensor data (*e.g.* images, audio, ambient light, and even accelerometers and GPS) can be interpreted to extract the high-level information needed by bleeding-edge sensor-driven systems like smartphone apps, wearable devices, and self-driving cars. Today, the state-of-the-art in computational models that, for example, recognize a face, track user emotions, or monitor physical activities, are increasingly based on deep learning principles and algorithms. Unfortunately, deep models typically exert severe demands on local device resources, and these conventionally limit their adoption within mobile and embedded platforms. As a result, in far too many cases existing systems process sensor data with machine learning methods that have been superseded by deep learning years ago.

At the same time, two new families of applications have emerged that aim to push the boundaries of user experience when using mobile devices in everyday life. On the one hand, the hardware and software progress of augmented and virtual reality (AR/VR) technologies have sparked a promise towards unified mixed-reality environments, such as the recently coined concepts of “meta-universe” and “telepresence”, where users will be able to interact through new forms of digital communication. On the other hand, cyber-physical systems have started gaining traction with the gradual maturation of robot platforms and DL-based perceptual models. As such, an increasing interest around embodied intelligence in the form of home robots, autonomous vehicles and drones has been formed. Despite the application-level opportunities, both types of systems pose significant engineering challenges. Robust operation, responsiveness, energy efficiency and user-oriented design form a multi-objective design and engineering problem that calls for radically novel solutions, beyond the conventional approaches of today.

In this context, the mobile computing community is in a unique position to begin the careful study of two core technical questions: 1) how to design robust and energy-efficient DL models, or adapt existing ones, to meet the stringent needs of mixed-reality mobile devices and cyber-physical embedded systems, and 2) how to exploit the capabilities of embedded DL through software and hardware innovations to enable the emerging applications of AR/VR and robotics. As such, we particularly encourage submissions on these two topics. More specific topics of interest include, but are not limited to:


- Applications of DNNs with real-time requirements
- Compression of DNNs for mobile devices
- Systems and networking techniques for DL-driven AR/VR applications
- Resource-efficient DNNs for AR/VR systems
- Neural models or sensors for modeling user activities and behavior
- Continuous vision systems supported by DNNs
- Embedded hardware accelerators for DNNs
- Distributed training, including Federated Learning
- Applications of DNNs with real-time requirements
- Optimizing commodity processors (GPUs/NPUs) for DL
- Networking optimizations for edge & cloud offloading
- Resource & energy management for embedded DL

This year we particularly encourage submissions describing systems and applications of DL-driven AR/VR and mobile robots.

Workshop Organizers

PC CHAIRS

Stefanos Laskaridis (*Samsung AI, UK*)
Stylianos I. Venieris (*Samsung AI, UK*)
Qing Wang (*TU Delft, NL*)



FULL PAPER SUBMISSIONS

Solicited submissions include both full technical workshop papers and white position papers. Maximum length of such submissions is 6 pages, and if accepted they will be published by ACM and appear in the ACM Digital Library.

Submission Deadline: 8 April 2022 AOE

WORK-IN-PROGRESS AND DEMO SUBMISSIONS

Abstracts describing work-in-progress and demonstrations are also welcome and warmly encouraged. Submissions are limited to 2 pages, and if accepted included in the program as a short oral presentation – but will only be published on the workshop website (not the ACM DL). Deadlines for this informal track remain open even past the early registration deadline of MobiSys 2022; author notifications will be rolling (*i.e.* max of 4 days after submission) to enable early authors to take advantage of available discounts.