



# Call for Papers

## tinyML Research Symposium 2021

Monday, March 22, 2021

[OpenReview](#)

New academic & commercial research symposium as part of the [tinyML Summit 2021](#)

### Program Chairs

Vijay Janapa-Reddi, Harvard Univ.

Boris Murmann, Stanford Univ.

### Program Committee

Edith Beigne, Facebook

Vikas Chandra, Facebook

Yiran Chen, Duke Univ.

Hiroshi Doyu, Ericsson

Adam Fuks, NXP

Wolfgang Furtner, Infineon

Song Han, MIT

Jeremy Holleman, Syntiant

Prateek Jain, Microsoft

Kurt Keutzer, Berkeley

H. T. Kung, Harvard

Matthew Mattina, ARM

Tinoosh Mohsenin, Univ. of Maryland

Edwin Park, Qualcomm

Priyanka Raina, Stanford Univ.

Jae-sun Seo, ASU

Mingoo Seok, Columbia Univ.

Dennis Sylvester, Univ. of Michigan

Jonathan Tapson, GrAI Matter Labs

Marian Verhelst, KU Leuven

Pete Warden, Google

Hoi-Jun Yoo, KAIST

**Helen Lee**

### Publications Chair

Zain Asgar, Stanford Univ.

### Publicity Chair

Theocharis Theocharides, Univ.

Cyprus

### Important Deadlines

~~Date: Nov 23 2020 11:59pm AOE~~

New Date: **Nov 30** 2020 11:59pm AOE

Author Notification: **Feb 5** Jan 15, 2021

Camera Ready: Feb 15, 2021

### Submission Page Limit

6 - 8 pages, including references.

Submissions must be blind (i.e., no author information allowed)

### Submission Website

Tiny machine learning (tinyML) is a fast-growing field of machine learning technologies and applications including algorithms, hardware, and software capable of performing on-device sensor (vision, audio, IMU, biomedical, etc.) data analytics at extremely low power, typically in the mW range and below, and hence enabling a variety of always-on use-cases and targeting battery-operated devices. tinyML systems are becoming “good enough” for (i) many commercial applications and new systems on the horizon; (ii) significant progress is being made on algorithms, networks, and models down to 100 kB and below; and (iii) initial low power applications in vision and audio are becoming mainstream and commercially available. There is growing momentum demonstrated by technical progress and ecosystem development. The first annual tinyML research symposium serves as a flagship venue for research at the intersection of machine learning applications, algorithms, software, and hardware in deeply embedded machine learning systems. We solicit papers from academia and industry combining cross-layer innovations across two or more topics. Submissions must describe tinyML innovations that intersect and leverage synergy between at least two of the following subject areas:

#### tinyML Datasets

- Public release of new datasets to tinyML
- Frameworks that automate dataset development
- Survey and analysis of existing tiny datasets that can be used for research

#### tinyML Applications

- Novel applications across all fields and emerging use cases
- Discussions about real-world use cases
- User behavior and system-user interaction
- Survey on practical experiences

#### tinyML Algorithms

- Federated learning or stream-based active learning methods
- Deep learning and traditional machine learning algorithms
- Pruning, quantization, optimization methods
- Security and privacy implications

#### tinyML Systems

- Profiling tools for measuring and characterizing performance and power
- Solutions that involve hardware and software co-design
- Characterization of tiny real-world embedded systems
- In-sensor processing, design, and implementation

#### tinyML Software

- Interpreters and code generator frameworks for tiny systems
- Optimizations for efficient execution
- Software memory optimizations
- Neural architecture search methods

#### tinyML Hardware

- Power management, reliability, security, performance
- Circuit and architecture design
- Ultra-low-power memory system design
- MCU and accelerator architecture design and evaluation

#### tinyML Evaluation

- Measurement tools and techniques
- Benchmark creation, assessment and validation
- Evaluation and measurement of real production systems

All submitted papers must be anonymized; no author information. The review process is double-blind. Accepted papers will be published as peer-reviewed online proceedings. An author of an accepted paper must attend the symposium to give a presentation.

Accepted full papers will be invited to submit an extended manuscript to a “Title:

T I N Y



[OpenReview](#)

**Paper**

**Template:** [ACM](#)  
([SIGCONF](#))

TinyML" Special Issue in ACM Transactions on Embedded Computing Systems (TECS).