IEEE Design&Test

Call for Contributions to Special Issue on

Robust Resource-Constrained Systems for Machine Learning

Motivation: Machine Learning is nowadays embedded in several computing devices, consumer electronics and cyber-physical systems. Smart sensors are deployed everywhere, in applications such as wearables and perceptual computing devices, and intelligent algorithms power our connected world. These devices collect and aggregate volumes of data, and in doing so, they augment our society in multiple ways; from healthcare, to social networks, to consumer electronics and many more. To process these immense volumes of data, machine learning, and in particular, deep learning via deep neural networks, is emerging as the de facto analysis tool, that powers several aspects of our Big Data society. Applications spanning from infrastructure (smart cities, intelligent transportation systems, smart grids, to name a few), to social networks and content delivery, to e-commerce and smart factories, and emerging concepts such as self-driving cars and autonomous robots, are powered by machine learning technologies. While some applications require analytics and trends identification that can run on large-scale data centers and the cloud; an abundance of emerging systems, however, require real-time inference and decision support, based on the acquired data. Such scenarios may use customized hardware accelerators, are typically bound by limited resources, and limited connectivity and bandwidth. As such, this complex computation shifts from the datacenters and the cloud, to the fog and the edge; near-sensor computation and near-sensor intelligence are starting to emerge as necessities, in order to continue supporting the paradigm shift of our connected world. The need for real-time intelligent data analytics (especially in the era of Big Data) for decision support near the data acquisition points, emphasizes the need of revolutionizing the way we design, build, test and verify processors, accelerators and systems that facilitate machine learning (and deep learning in particular) implemented in resource-constrained environments for use at the edge and the fog. As such, traditional Von Neumann architectures are no longer sufficient and suitable, primarily because of limitations in both performance and energy efficiency caused especially by large amounts of data movement. Furthermore, due to the connected nature of such systems, security and reliability are also critically important. Robustness, therefore, in the form of reliability and operational capability under the presence of faults, whether malicious or accidental, is a critical need for such systems. Moreover, the operating nature of these systems relies on input data that is characterized by the four “V’s”: Velocity (speed of data generation), Variability (variable forms and types), Veracity (unreliable and unpredictable) and Volume (i.e. large amounts of data). Thus, the robustness of such systems needs to consider this issue as well. Furthermore, robustness in terms of security, and in terms of reliability to hardware and software faults, in particular, besides their importance when it comes to safety-critical applications, are also a positive factor in building trustworthiness towards these disrupting technologies from our society. To achieve this envisioned robustness, we need to re-focus on problems such as design, verification, architecture, scheduling and allocation policies, optimization, and many more, for determining the most efficient, secure and reliable way to implement these novel applications within a robust, resource-constrained system, which may or may not be connected.

Special Issue Scope:, This special issue, aims to address a key aspect of fog and edge based ML algorithms; **robustness (as defined above) under resource-constraint scenarios**. **The special issue attempts to present emerging works in how we design robust systems, both in terms of reliability, fault-tolerance and security, while operating with a limited number of resources, and possibly in the presence of harsh environments which may eliminate connectivity and pollute the input data.** The SI aims to bring together views from academia and industry in order to exchange ideas, positions and research directions. The special issue is expected to cover emerging challenges, as well as new research results in the field. Further, the special issue aims to explain how we can take advantage of existing and emerging hardware and software technologies, as well as ML algorithms (such as adversarial machine learning), in addressing the associated design and verification challenges of robust ML systems.

Topics of Interest

This special issue will feature works that address issues related to accelerators, memory, and their coordination, and will discuss issues spanning across the hardware and software stacks to build a holistic cross-layer spectrum. We expect that contributions will highlight advances in technology, design methodologies, design automation, test and verification, etc. Submissions in the following topics are requested, though not restricted to:

- Design and test issues of robust, resource-constrained ML-based systems
- Security and reliability aspects of such systems (e.g. attack mitigation, fault tolerance, etc.)
- Adversarial machine learning: from algorithms to architectures for resource-constrained systems
- Energy-efficient architectures and accelerators for learning algorithms
- Optimization techniques (quantization, approximate computing, etc.) for deep neural networks while maintaining robustness
- Specialized memory architectures and memory optimizations for ML-systems (like weight compression, memory design for DNNs, etc.)
- In-memory and in-storage computation for robust ML-systems
- Neuromorphic and bio-inspired approaches for resource-constrained and robust ML architectures
- Design automation tools and platforms for resource-constrained robust ML systems.
- Robust machine/deep learning systems using Emerging technologies (Memristors, post-CMOS, etc.)
- Integration into existing systems (like smart CPS and smart IoT systems)
- Case studies (edge devices such as drones, wearables, etc.)

Important Dates

- Manuscript submission: March 15th, 2019
- First round of reviews: April 30th, 2019
- Revised Manuscripts Due: May 30th 2019
- Second round of reviews: June 1st , 2019
- Final manuscripts due: June 30th, 2018

Submission Guidelines

Submission guidelines for IEEE D&T papers can be found here:

https://www.ieee-ceda.org/publication/ieee-design-test-dt/paper-submission-instructions

Prospective authors should follow the submission guidelines for IEEE Design & Test. All manuscripts must be submitted electronically to IEEE Manuscript Central at https://mc.manuscriptcentral.com/dandt. *Indicate that you are submitting your article to the special issue on Robust Resource-Constrained Systems for Machine Learning.*

Manuscripts must not exceed 5,000 words, including figures (with each average-size figure counting as 200 words) and a maximum of 12 references (50 for surveys). This amounts to about 4,000 words of text and a maximum of five small to medium figures. Accepted articles will be edited for clarity, structure, conciseness, grammar, passive to active voice, logical organization, readability, and adherence to style. Please see IEEE Design & Test Author Resources for links to Submission Guidelines Basics and Electronic Submission Guidelines and requirements.

Guest Editors

Muhammad Shafique	muhammad.shafique@tuwien.ac.at	Technical University of Vienna
Theocharis Theocharides	ttheocharides@ucy.ac.cy	University of Cyprus
Onur Mutlu	onur.mutlu@inf.ethz.ch	ETH Zurich
Jungwook Choi	choij@us.ibm.com	IBM Research, US