Real-Time Systems (2023) 59:345–347
https://doi.org/10.1007/s11241-023-09406-0

Guest editorial: special issue on predictable machine learning

Daniel Casini[1] **· Giorgio Buttazzo**[1]

Published online: 17 August 2023

In recent years, deep neural networks have achieved remarkable performance in several tasks related to perception and control. Nevertheless, their usage in safety-critical systems is quite problematic due to a number of reasons.

First, during inference, the execution time of a neural network can be subject to high variations, which may be caused by the specific computing platform, hardware accelerator, or the framework used to manage the execution of the various network nodes. Second, neural models have been shown to be prone to adversarial attacks, which can induce a wrong prediction through imperceptible perturbations applied to the input. Such adversarial attacks have been shown to be also applicable in the real world through properly crafted patches that can be printed and placed on physical objects, so without accessing the vision system. Third, the prediction of a neural model can also be compromised by inputs that are out of the typical distribution of the data samples used during training. Detecting or neutralizing such adversarial attacks or out-of-distribution inputs may have a significant impact on the overall execution time of the neural model.

This special issue includes three selected articles that cover well three sub-topics related to "Predictable Machine Learning": timing vs. accuracy tradeoffs, safe and predictable implementation of machine learning algorithms for embedded systems, and predictability of machine learning algorithms in autonomous driving frameworks.

The first article "Scheduling IDK Classifiers with Arbitrary Dependences to Minimize the Expected Time to Successful Classification" by Abdelzaher et al. considers a model specialized for classification-based machine perception problems to trade off accuracy and execution duration to meet timing constraints while maximizing accuracy.

✉ Daniel Casini
daniel.casini@santannapisa.it

Giorgio Buttazzo
giorgio.buttazzo@santannapisa.it

[1] Scuola Superiore Sant'Anna, Pisa, Italy

The second article "Extending a predictable machine learning framework with efficient GEMM-based convolution routines" by De Albuquerque Silva et al. focuses on the safe real-time implementation of the inference phase of feed-forward deep neural networks on embedded platforms, with the objective of being compliant with avionics requirements.

The third article "Main Sources of Variability and Non-Determinism in AD Software: Taxonomy and Prospects to Handle Them" by Alcon et al. analyzes the source of variability and non-determinism in autonomous driving software, with a focus on the Apollo autonomous driving framework.

We would like to thank all the authors for submitting their excellent work to this special issue and the reviewers for providing their valuable feedback. We thank the Editor-in-Chief of Springer Real-Time Systems, Luis Almeida, and Springer's editorial assistants for their support throughout the organization of the special issue and in managing the reviewing process.

Daniel Casini
Scuola Superiore Sant'Anna, Italy
daniel.casini@santannapisa.it

Giorgio Buttazzo
Scuola Superiore Sant'Anna, Italy
giorgio.buttazzo@santannapisa.it
Guest Editors

Daniel Casini is Assistant Professor at the Real-Time Systems (ReTiS) Laboratory of the Scuola Superiore Sant Anna of Pisa. He graduated (cum laude) in Embedded Computing Systems Engineering, a Master degree jointly offered by the Scuola Superiore Sant Anna of Pisa and University of Pisa, and received a Ph.D. in computer engineering at the Scuola Superiore Sant Anna of Pisa (with honors). In 2019, he has been visiting scholar at the Max Planck Institute for Software Systems (Germany). He is Associate Editor for Elsevier Microprocessors and Microsystems. He has been a technical program committee member in several international conferences including RTSS (IEEE Real Time System Symposium, 2021, 2022, 2023), ECRTS (Euromicro Conference on Real-Time Systems, 2021, 2022, 2023), RTAS (IEEE Real-Time and Embedded Technology and Applications Symposium, 2022, 2023), and ISORC (IEEE International Symposium on Real-Time Distributed Computing, 2020, 2021, 2022).

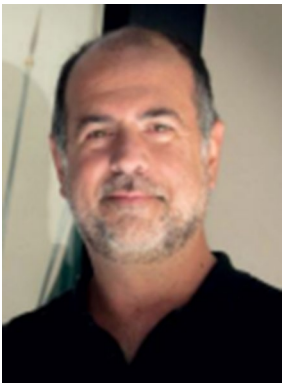

Giorgio Buttazzo is Full Professor of Computer Engineering at the Scuola Superiore Sant Anna of Pisa. He graduated in Electronic Engineering at the University of Pisa in 1985, received a Master in Computer Science at the University of Pennsylvania in 1987, and a Ph.D. in Computer Engineering at the Scuola Superiore Sant Anna of Pisa in 1991. He has been Chair of the IEEE Technical Committee on Real-Time Systems (2010-2012), and Program Chair and General Chair of the major international conferences on real[1]time computing. In 2013, he received the Outstanding Technical Contributions and Leadership Award from the IEEE Technical Committee on Real-Time Systems. He has been Editor-in-Chief of the Journal of Real-Time Systems, Associate Editor of the IEEE Transactions on Industrial Informatics and the ACM Transactions on Cyber-Physical Systems. He is IEEE Fellow since 2012 and has authored 6 books on real-time systems and over 300 papers in the field of real-time systems, robotics, and neural networks, receiving 13 Best Paper Awards.