



Electronics

an Open Access Journal by MDPI

CiteScore: 6.1

Impact Factor: 2.6

Special Issue Reprint

Hardware Acceleration for Machine Learning

Edited by: Sergio Spanò, Gian Carlo Cardarilli and Luca Di Nunzio

This Reprint, titled “Hardware Acceleration for Machine Learning”, provides a coherent and updated snapshot of the rapidly evolving ecosystem for machine learning acceleration. As hardware platforms transition from general-purpose solutions to specialized heterogeneous architectures, this collection synthesizes cutting-edge developments essential for researchers and designers. The eight contributions within this Reprint cover a full spectrum of advancements, from low-power IoT applications and embedded FPGA and SoC implementations to ASIC edge solutions and efficient data-processing algorithms. Key research presented includes multi-dimensional benchmarking of deep learning inference frameworks, novel power-optimized RTL designs for CNN accelerators, unified FPGA/CGRA pipelines for time-critical edge AI, and new adaptive optimizers designed for training efficiency. Further studies address the performance impacts of data preprocessing techniques, general-purpose hyperdimensional computing accelerators, and innovative approaches to overcoming resource constraints in edge AI through DSP-free architectures and integer-state dynamics in quantized spiking neural networks. By combining these diverse strategies, this collection offers a comprehensive reference for those interested in building implementable, scalable, and energy-sustainable hardware platforms for next-generation intelligent systems.

<https://www.mdpi.com/books/reprint/12880>

