Narayan Srinivasa, IEEE Fellow (2024 – 2026)

Intel Corporation

2200 Mission College Boulevard, Santa Clara, CA 95054

Email: narayan.srinivasa@intel.com

Lecture Topic 1: On Solving Hard Optimization Problems in an Energy Efficient Way

Abstract: The core philosophy in solving complex optimization problems is to describe system energy as a cost function encoded in the variables of these problems. As time evolves, the dynamics of interaction between the system components that encode these variables move it closer to the optimal solution. Even if a massively parallel system can be implemented to solve NP-hard optimization problems with traditional computers, exponential resources would be needed in space, time, or precision. These types of problem are ubiquitous, and can show up, for example, when deciding a route for package deliveries, parsing DNA sequence data, training an artificial neural network or disease control. To address this challenge, many highly differentiated approaches remain viable, from quantum computing and annealing to optical, analog, and neuromorphic computing. In this talk, I will provide an overview of these approaches and highlight some recent results obtained using Intel's Loihi neuromorphic chip for enabling a new class of solutions. These solutions have yielded compelling outcomes for some NP-hard optimization problems in a manner that is substantially more energy efficient and scalable than current state-of-the-art solvers. We believe that exploring these concepts further could have an impact for a broad range of applications from SWaP constrained edge devices to very large-scale cloud computing problems including complex diagnosis, DNA sequencing, logistics, and 5G MIMO problems.

Lecture Topic 2: Opportunities, Challenges, and Threats Using Synthetic Data for AI

Abstract: Industry and governments are making massive investments in AI to boost productivity and automation. Forward-thinking businesses are already ahead of the game. Research by Forrester in 2019 found that just over half of decision-makers in the global data and analytics space have already implemented, are currently implementing, or plan to expand/upgrade their AI investments. IDC predicts that three-quarters of enterprises will embed AI into their technology and process development by 2022 and that by 2024, AI will be integral to all parts of these businesses. A key requirement for AI is data. You can't do AI without data. And not just any data. You need huge quantities of high-quality, accurate data to build useful models of AI for the real world. Although data-gathering behemoths like Google, Facebook, and Amazon have no problems in obtaining this data from their products and services, other businesses often lack access to the datasets they need. In this talk, I will first discuss synthetic data generation approaches using AI to mitigate this data access problem. I will then discuss the opportunities, challenges, and threats arising from these approaches to the democratization of AI and in accelerating AI adoption.

Lecture Topic 3: Machine Learning & Neuromorphic Computing

Abstract: A historical perspective on Machine Learning that underpins most of the models today will be provided followed by the reasons for why machine learning and its applications have taken the world by storm today. This will be followed by an introduction to neuromorphic computing followed by the various applications that are driving the need for such technology in the future. The talk will contrast the two from the perspective of biological plausibility and the nature of applications that are well suited for them.