

Abstract: The Internet of Things (IoT) heralds the emergence of multitudes of computing-enabled networked everyday devices with sensing capabilities in homes, cars, workplaces, and on our persons, leading to ubiquitous smarter environments and smarter cyber-physical systems. The next natural step in this computing evolution is to develop the infrastructure needed for these computational things to collectively learn. Advances in deep learning offer remarkable results but require significant computing resources. The talk discusses advances in deep learning geared to empower the emerging world of embedded IoT devices. These advances tackle challenges such as resource-efficiency (to fit neural networks on small mobile devices), foundations of quality assurances, and cost of labeling. Recent solutions demonstrate the feasibility of building IoT applications that are powered by effective, efficient, and reliable deep learning models. Evaluation results and experiences presented offer encouraging evidence of viability of deep learning for IoT.