PROFESSIONAL EDUCATION

DESIGNING EFFICIENT DEEP LEARNING SYSTEMS

professional.mit.edu/dls | Lead Instructor: Vivienne Sze

Note: All times are US Eastern Daylight Time. Schedule is subject to change.

	DAY 1	DAY 2
9:30–10:30am	• **Introduction and Course Overview** • **Background on Deep Learning:** Learn how deep learning has evolved to where it is today and what are the factors that have enabled it to become popular and widely used. • **Deep Learning Applications:** Discover applications that use deep learning today (e.g., computer vision, speech recognition) and what new applications can benefit from deep learning moving forward (e.g., self-driving cars, medical, internet of things), as well as what are some of the computational challenges of these applications with respect to speed, cost, energy/power consumption, etc. • **Overview of Deep Neural Networks:** Gain an overview of how deep neural networks work, specifically what are the key computations used by deep neural networks.	• **Recap of Day 1** • **Deep Learning on Specialized Hardware (Part 1):** Discuss the key design choices and optimizations that are considered when designing specialized hardware for processing deep neural networks, and the various design tradeoffs to consider.
10:30–11:00am	**BREAK**	**BREAK**
11:00am–12:30pm	• **Popular Deep Neural Network Models:** Explore widely used deep neural networks and what attributes make them effective.	• **Deep Learning on Specialized Hardware (Part 2)** • **Use of Advanced Technologies:** Review emerging technologies used to further improve the speed and efficiency of deep learning systems (e.g., in-memory computing, optical computing).
12:30–1:30pm	**LUNCH BREAK**	**LUNCH BREAK**
1:30–3:00pm	• **Development Resources for Deep Learning:** Review the various resources to support running/training deep neural networks including frameworks (e.g., Tensorflow), and datasets (e.g., ImageNet), and describe how to select the right ones for a given application. • **Training Deep Neural Network Models:** Overview of how deep neural networks are trained.	• **Co-optimization of Algorithms and Hardware (Part 1):** Gain an understanding of how to design efficient deep neural network models in conjunction with hardware for further improvement in processing efficiency, and the various design tradeoffs to consider.
3:00–3:30pm	**BREAK**	**BREAK**
3:30–5:00pm	• **Metrics for Evaluating Deep Learning Systems:** Learn how to evaluate different deep learning systems with benchmarks and comparison metrics, and highlight some of the key tradeoffs and questions one should ask when evaluating which platform is best. • **Deep Learning on Programmable Platforms:** Review the range of deep learning platforms available for different applications, and how deep neural networks are mapped/complied onto these platforms and how that impacts the processing speed. • **Discussion and Day 1 Summary**	• **Co-optimization of Algorithms and Hardware (Part 2)** • **Discussion on Trends in Deep Learning:** Discuss recent implementation trends and opportunities in deep learning research and applications.