

## Homework 9

---

Using the assigned reading listed on the course page, answer the questions below with a short response. Note that we are looking for concise statements that show understanding, not quantity. The total discussion should roughly be a page.

### Hardware for deep learning

1. Do you think that Moore's Law is dead (or dying)? And why?
2. Fig 17 from Sze et al. shows two parallel compute architectures. What are some of the tradeoffs involved in these two strategies? When would you pick one over the other?
3. Section VII from Sze et al. describes co-design – the idea of designing algorithmic architectures that are amenable to cheap implementation in hardware (and vice versa). Propose a modification to any part of standard deep learning algorithms that could make hardware cheaper (in time, memory, energy, or something else). For example, you could propose a nonlinearity that is cheap on-chip, or you could propose a change to how we do SGD that makes things more efficient to run on the cloud. Feel free to be creative. Explain your reasoning and the assumptions it rests on (it's okay if the assumptions turn out to be wrong, we will grade for having a reasonable idea, not for knowing the details of resource consumption on different specific hardware setups).

**Submission:** Upload a PDF of your response through Canvas by **11/16 at 1pm**.