

Split up these questions (front and back) among your neighbors. Each person writes a very brief response to their questions. Discuss.

Neural Computation

- How do neural nets compute? (How does that differ from traditional programming?)
- What are the “data structures” of neural computing and efficient operations we can do with them?
- How can we update parameters to optimize an objective function?

ML Systems

- What are the inputs to and outputs of AI systems?
- What abstractions do systems provide, and how can they compose? (ML APIs)
- How do we evaluate ML solutions?

Learning Machines

- How can systems improve from experience?
- What can be learned from data vs interaction?
- How can we evaluate learning: does it generalize?

Context and Implications

- What problems can we use AI to solve?
- What **should** we use AI for?
- What are the limits of AI systems?
- What might happen socially when AI systems are deployed broadly? (effects on work, education, creativity, ...)
- How can AI systems be aligned with human values? What are the risks if they aren't?
- How might we design AI systems to align with human values? to honor each other and our neighbors?
- How do privacy and copyright relate with AI?
- What is creativity? Agency? Truth?

Course Reflections

- What's the point of college in an AI age?
- In what ways could this class be a picture of the future? How is it a relic of the past?