

# Artificial Intelligence: Principle and Practice – Schedule

---

⚠️ Warning: Under construction ⚠️

Day 0. (Mon. Oct. 4th / Tue. Oct. 5th) **Initialization:** overview basic concepts of artificial intelligence, symbolic computation, machine learning, deep learning, and neural networks, and multi-agent systems as well as AI ethics and ML ops. In **lab** we'll write some basic decision trees to control simulated and real robots and tap into knowledge bases to make a trivia answering system.

Day 1. (Wed. Oct. 6th / Thur. Oct. 7th) **Machines that Learn:** fundamental concepts of machine learning, including optimization and gradient descent, probabilistic modeling and factor graph representation, regression, classification, clustering, regularization, and ensembles. In **lab** we'll parse demographic data and build a shallow machine learning model to predict the likelihood of a user's income.

Day 2. (Mon. Oct. 11th / Tue. Oct. 12th) **Deep Dive:** basic concepts of deep learning, common architectures, layer types, and activation functions as well as backpropagation, optimization, loss functions, and training paradigms. In **lab**, we'll build a neural network from custom-built layers to solve an image classification problem.

Day 3. (Wed. Oct. 13th / Thur. Oct. 14th) **The Zoo:** overview of specific deep learning architectures including DNNs, AEs, CNNs, GANs, RNNs, attention-based models, GNNs, and multimodal systems. In **lab** we'll do neural network 'surgery' on several popular models and try to improve their performance.

Day 4. (Mon. Oct. 18th / Tue. Oct. 19th) **Reinforcement Learning:** the classical utility agent, various extrinsic and intrinsic rewards, training setups, and environments, cooperative/competitive/mixed-mode multi-agent reinforcement learning systems, and advantages and pitfalls of RL. In **lab** we'll train a mixed-mode multi-agent reinforcement learning system in various simulated environments.

Day 5. (Wed. Oct. 20th / Thur. Oct. 21st) **Intelligence:** transfer learning, generalization, zero/one/few-shot learning, distillation, information theory, corellators, criticality, generalization, transferability, the relationship between intelligence and learning, and the limits of AI generality. In **lab** we'll play with large language models including gpt-3 ada, babbage, curie, davinci, davinci-instruct, davinci-codex. We'll also attempt small and large examples of transfer learning all the way to transferring a pretrained language model to perform computer vision tasks.

Day 6. (Mon. Oct. 25th / Tue. Oct. 26th) **Hello World:** AI safety, security, privacy, federated learning, and ethics, AI in the news, in acedemia, and in the workplace, ML ops, human-AI interaction, environmental impact, and the impact of AI on human life. In **lab** students will put MLops to work on a small self-guided group project of any size. Groups with reasonable metric performance will receive additional compute resource credits.

Day 7. (Wed. Oct. 27th / Thur. Oct. 28th) **Project Day!** Students present their projects to the class. Students will also have a chance to meet with other professors and other students to discuss their projects and the course.