Title

Deep Generative Models

Course Description

Generative models are widely used in many subfields of AI and Machine Learning. Recent advances in parameterizing these models using deep neural networks, combined with progress in stochastic optimization methods, have enabled scalable modeling of complex, high-dimensional data including images, text, and speech. In this course, we will study the probabilistic foundations and learning algorithms for deep generative models, including variational autoencoders, generative adversarial networks, autoregressive models, normalizing flow models, and energy-based models. The course will also discuss application areas that have benefitted from deep generative models, including computer vision, speech and natural language processing, and reinforcement learning.

Related Courses

There are a few graduate courses on machine and deep learning offered by the department. The most relevant course is CS 260: Deep Learning which focuses on deep neural networks broadly — given the growth of the field, the overlap between this course and the proposed course is about 15% where the instructor typically covers 1-2 generative models at a high level. Another relevant course is CS 263: Natural Language Processing which covers generative models for text. In contrast, the proposed course is not focussing on one data modality alone and touches upon probabilistic foundations that can be applied more broadly, with illustrative examples covering text, images, audio, videos. Finally, there are other special topic courses offered occasionally covering graph neural nets and reinforcement learning that briefly touch upon a few generative models, but again a comprehensive course covering the probabilistic foundations of these models, intersections with deep neural nets, applications across various data modalities and societal concerns with real-world deployments is lacking.

Prerequisites

Students should have done introductory course in machine learning (eg, CS M146) and/or data science (eg, CS M148). A course in deep learning is preferred but not required (eg, CS 260).

Learning Outcomes

Knowledge outcomes: Learn probabilistic foundations, algorithms, and applications for deep generative models

Skills outcomes: Learn to implement these models in Python, learn deep learning libraries, learn to derive algorithms for generative modeling, learn to read papers

Attitudes and values outcomes: learn societal challenges and opportunities of deep generative models (e.g., deepfakes)

Grade Breakdown

- 2 HWs: 20%
- Midterm: 15%
- Final Exam: 25%
- Course Project: 40%

Syllabus (with links to lecture slides)

Week		**Lecture Topics**	
1		Introduction [slides 1, 2]	
2		Autoregressive Models [slides 1, 2]	

3		Variational Autoencoders [slides 1, 2]	
4		Normalizing Flows [slides 1, 2]	
5		Generative Adversarial Networks [slides 1, 2]	
6		Evaluating generative models, In-class mid term on Feb 10 [slides 1]	
7		Energy-based models [slides 1, 2]	
8		Model combination [slides 1]	
9		Score-based generative models, diffusion models	
10		Societal considerations, In-class Presentations for final project	

Additional Reading: Surveys and Tutorials

1. Tutorial on Deep Generative Models. Aditya Grover. Deep Learning for Science Summer School, July 2020.
2. How to Train Your Energy-Based Models. Yang Song and Diederik P. Kingma. February 2021.
3. Tutorial on Deep Generative Models. Aditya Grover and Stefano Ermon. International Joint Conference on Artificial Intelligence, July 2018.
4. Tutorial on Generative Adversarial Networks. Computer Vision and Pattern Recognition, June 2018.
5. Tutorial on Deep Generative Models. Shakir Mohamed and Danilo Rezende. Uncertainty in Artificial Intelligence, July 2017.
6. Tutorial on Generative Adversarial Networks. Ian Goodfellow. Neural Information Processing Systems, December 2016.
7. Learning deep generative models. Ruslan Salakhutdinov. Annual Review of Statistics and Its Application, April 2015.