



Resource provided by



Course Syllabus

Deep Learning in AI

This intermediate course delves into deep learning, the driving force behind transformative breakthroughs in artificial intelligence.




Table of Contents

Course Description	3
Course Competencies	4
Instructional Resources	5
Grading Schema	5
Supplemental Information	6
Course Outline	7
Supplemental Books (Optional)	8
Supplemental Papers (Optional)	9


Course Description

This intermediate course delves into deep learning, the driving force behind transformative breakthroughs in artificial intelligence. Students will explore the core principles of deep learning, mastering the design, training, and evaluation of neural network models to tackle complex, real-world challenges across diverse industries. Through hands-on labs, projects, and competitions, learners will develop practical skills, integrate deep learning with broader AI methodologies, and critically examine its ethical and societal implications. The course equips students to stay at the forefront of this rapidly evolving field, fostering innovation and responsible application of AI.

Prerequisites:

Python programming language

Machine Learning course



Course Competencies

Competency 1
Explain the core principles and historical evolution of deep learning to establish a foundational understanding of its role in AI
Competency 2
Apply deep learning techniques to design and train neural network models for diverse, real-world challenges
Competency 3
Analyze advanced neural network architectures to determine their suitability for specific applications
Competency 4
Integrate deep learning methodologies with other AI approaches to create cohesive solutions
Competency 5
Evaluate the ethical implications and societal impacts of deep learning applications to ensure responsible innovation
Competency 6
Synthesize emerging trends and innovations in deep learning to propose future directions for AI development



Instructional Resources

No prescribed textbook is required due to the rapidly evolving nature of artificial intelligence and Deep Learning. We will curate up-to-date articles, papers, videos, and other resources, make them available via Canvas, to provide the latest advancements and diverse perspectives. This approach ensures flexibility to tailor content to class needs and incorporate emerging topics.

Platforms and Tools:

Kaggle, GitHub, Google Colab, Amazon SageMaker, VS Code, Hugging Face

Recommended Reading:

See [Course Bibliography/Materials](#) section.

Grading Schema

Assignment Type	Percentage of Grade
General Assignments	20%
Labs/Hands-On Activities	25%
Midterm Project	20%
Capstone Project	25%
Course Portfolio	10%

Supplemental Information

General Assignments (20%)

Objective: To reinforce theoretical concepts and develop analytical skills through weekly assignments including case studies, research discussions, problem sets, and design tasks directly related to module content.

Labs/Hands-On Activities (25%)

Objective: To provide students with practical understanding of natural language processing through guided implementation exercises, coding workshops, and interactive demonstrations using industry-standard tools and frameworks.

Midterm Project (20%)

Objective: Individual or group project demonstrating integration and application of knowledge from the first half of the course. Students will design and implement a focused NLP solution addressing a specific problem or use case.

Capstone Project (25%)

Objective: To provide students with comprehensive hands-on experience in creating a practical, real-world application. Students will integrate multiple NLP techniques learned throughout the course to develop an end-to-end solution with documentation and presentation.

Course Portfolio (10%)

Objective: Individual dynamic document continuously updated and maintained on GitHub, serving as a comprehensive record of learning progress, code repositories, reflections, and achievements throughout the course. This portfolio demonstrates professional development and technical growth.

Course Outline

Week starting on / Module	Module Topic
Week 1	Mod 01: Introduction to Deep Learning and Its Applications Mod 02: Environments and Tools
Week 2	Mod 03: Neural Network Optimization and Training Deep Networks
Week 3	Mod 04: Convolutional Neural Networks (CNNs) for Computer Vision
Week 4	Mod 05: Sequence Modeling and Natural Language Processing with RNNs
Week 5	Mod 06: Transformers and Attention Mechanisms for NLP and Beyond
Week 6	Mod 07: Variational Autoencoders (VAEs)
Week 7	Mod 08: Generative Adversarial Networks (GANs)
Week 8	Mod 09: Diffusion Models and Advanced Generative Techniques
Week 9	Mod 10: Deep Reinforcement Learning and AI Agents
Week 10	Mod 11: Reasoning Models and RAG
Week 11	Mod 12: Agent Planning and Orchestration
Week 12	Mod 13: LangChain and Agent Architectures
Week 13	Mod 14: Advanced Agent Systems and Applications
Week 14	Mod 15: Future Trends and Course Review
Week 15	Mod 16: Course Wrap-Up
Week 16	Finals week / Capstone Projects

Supplemental Books (Optional)

Auffarth, B. (2024). *Generative AI with LangChain: Build Large Language Model (LLM) Apps with Python, ChatGPT, and Other LLMs* (2024 Edition). Packt Publishing.
This book provides practical guidance on building LLM applications, including agents and retrieval-augmented generation, suitable for labs and projects on advanced neural networks.

Lanham, M. (2025). *AI Agents in Action*. Manning Publications.
This text offers a comprehensive framework for developing production-ready AI agents, ideal for projects involving complex model architectures.

Scott, M. R. A., & Acero, A. (2025). *Building Agentic AI Systems*. Packt Publishing.
This book focuses on designing and deploying agentic AI systems, providing hands-on examples for labs and projects in neural network applications.

Using Stable Diffusion with Python: Leverage Python to Control and Automate High-Quality AI Image Generation (1st Edition).
This book focuses on implementing diffusion models for image generation, providing hands-on Python examples for generative AI labs.

Supplemental Papers (Optional)

Vaswani, A., et al. (2017). "Attention is All You Need." Advances in Neural Information Processing Systems (NeurIPS). <https://arxiv.org/abs/1706.03762>
This paper introduces the transformer architecture, explaining its self-attention mechanism and dataset requirements, foundational for understanding modern neural networks in Data Science.

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012, 2017 version). "ImageNet Classification with Deep Convolutional Neural Networks." Communications of the ACM. <https://arxiv.org/abs/1409.1556>
This paper details the AlexNet CNN architecture and dataset engineering techniques, key for understanding convolutional networks and data preprocessing.

Goodfellow, I., et al. (2014). "Generative Adversarial Nets." Advances in Neural Information Processing Systems (NeurIPS). <https://arxiv.org/abs/1406.2661>
This paper introduces GANs, explaining their architecture and dataset needs, essential for understanding generative models in Data Science.

Instructors may consider adopting these books and papers to complement curated resources, aligning with course objectives and teaching style.