Course number and name: **CS 07559: Advanced Models of Deep Learning**

Credits and contact hours: 3 credits / 3 contact hours

Instructor's or course coordinator's name: Silvija Kokalj-Filipovic

Instructional materials:

1. Fundamentals of Deep Learning. Nithin Buduma, Nikhil Buduma and Joe Papa. O'Reilly Media. 2022. 2nd Edition.
2. Natural Language Processing with Transformers. Lewis Tunstall, Leandro von Werra and Thomas Wolf. O'Reilly Media. 2022. 1st Edition.
3. Online resources.

Specific course information

Catalog description: This course will teach students the comprehensive landscape of deep learning including theoretical foundations and mechanics of training neural nets designed to perform various tasks, assessing the data and computational needs for training and deploying various types of deep neural nets. The emphasis will be on the latest breakthroughs in algorithms and models, while leveraging a popular programming platform to implement variations and combinations of those algorithms and make them deployable and efficient.

About this course: This hands-on course will teach students the comprehensive landscape of deep learning including theoretical foundations and mechanics of training deep neural nets (models) designed to perform diverse tasks, while assessing the data and computational needs for training and deploying different types of models. The emphasis will be on the latest breakthroughs in algorithms and models, while leveraging a popular programming platform to implement variations and combinations of those algorithms and make them deployable and efficient. The final part of the course will apply the learned principles and methods to generative AI, including transformers. Important note: The goal of this course is not only a deep understanding of the essential components of deep learning models and processes, but also a contextualization of the state-of-the-art LLMs and the AI chatbots and agents derived from them. Students will become more efficient, creative and conscientious in acquiring the skills for LLMs, agentic AI and beyond.

Prerequisites: Linear Algebra, Probability, Python Proficiency

Specific goals for the course

1. **Deep Learning in General.** Students will develop advanced knowledge on the practical aspects of deep learning (DL) model design, training and deployment, including programming skills, algorithm selection and modification.

2. **Neural Net Algorithms and Structure**. Students will understand the concepts and technology behind deep neural net algorithms and structure by starting an advanced project immediately and demonstrating the impact of changing the structural and algorithmic components as the unit is being taught.

3. **Convolutions**. Students will analyze convolutions in depth, including comparison with attention mechanisms in transformers, and as elements of more complex architectures such as ResNets. Students will also implement an advanced convolutional network and a transformer.

4. **Comparative Analyses.** Students will know the concepts behind advanced topics and perform comparative analysis between some advanced DL models and applications.

Required list of topics to be covered:

1. Basic optimization algorithms for DL training
   a. Loss functions
   b. Backpropagation: Stochastic gradient descent and derived algorithms
   c. Computational graphs
   d. Convergence issues, overfitting, gradient exploding and saturation

2. Basic architectural elements of DL models
   a. Activation Functions
   b. Linear layers
   c. Convolutions
   d. Regularization layers
   e. Self-attention and cross-attention mechanisms
   f. Autoencoders

3. Generative AI approaches
   a. Variational inference (VAE)
   b. Game-theoretic approach (GAN)
   c. Learning attention probabilities (transformers)
   d. Diffusion
   e. Combinations

4. Pytorch
   a. Batching and tensor methods

b. Basic training and evaluation framework: forward and backward passes of the iterative loop
c. Pytorch lightning framework
d. Distributed training
e. Higher-level frameworks (e.g., Hugging Face)

5. LLM basics
   a. Packaging, API, repositories
   b. Prompting
   c. Customization: transfer-learning and retrieval augmented generating (RAG)
   d. Multimodal LLMs: language, vision, motion

Illustrative Sample of Possible Lectures

Lecture 1:

Introduction, Course preview (Definitions, Motivation, State-of-the art); Introduction to programming environments: Pytorch, Hugging Face, Running code in Google Colaboratory vs. Jupyter Hub to the class GPU server.

Lecture 2:

Review of Linear Algebra with DL tensors, and Probability w/Deep Learning (DL) examples

Lecture 3:

Simple Neural Net w/Regression: Demonstration of DL layers, models and graphs; DL training as Optimization / Stochastic gradient; Development Notebook (Pytorch training template)

Lecture 4:

Stochastic gradient descent deep dive; Data sets and Data Loaders for Pytorch training and inference

Lecture 5:

Convolutional Neural Net (Feature extractor vs classifier); Example architectures (auto encoder and resnet), Representation learning

Lecture 6:

Autoencoder demo; Learning temporal data - Recurrent Neural Nets to Transformers (exercise with unfolded images)

Lecture 7:

Transformers: introduction and comparative analysis with recurrent neural net

Lecture 8:

Visual transformer (ViT); Architectural types of transformers

Lecture 9:

Deep dive into DL activations and loss functions, Generative adversarial network (GAN); PyTorch Lightning demonstration

Lecture 10:

Basic models of generative AI (GenAI): GAN review, Variational Autoencoders (VAE, VQ-VAE), generative loss functions and generative model evaluations

Lecture 11:

Deep dive how to fine-tune DL training: Tools and metrics, regularizations, scheduling, visualization and debugging (Tensorboard API, MLFlow, MLOps, Hyperparameter tuning with Ray Tune)

Lecture 12:

Advanced GenAI Models (LLMs, GPT transformers, Stable Diffusion)

Lecture 13:

Foundation Models, Memory Requirements, Distributed Training, Explainable AI

Lecture 14:

Model robustness, adversarial attacks, hallucinations, prompt injections, AI system management of the future