

Descriptions of Courses

[updated Sep. 01, 2022]

▣ Graduate Courses

AI495 Individual Study

This course provides Machine Learning/AI research opportunities to undergraduate students. Students will participate in research projects and learn the latest research topics in the AI/ML field. All Kim Jaechul Graduate School of AI professors can open the course every semester so that students can have various options based on their research interests.

AI501 Machine Learning for AI

In this course, we will learn about introductory materials for machine learning, which is the fundamental and core technology for current generation of artificial intelligence. We will cover the most fundamental ideas and theories of machine learning, and introduce some of the important topics that will be covered in more advanced courses. Specifically, we will cover mathematical backgrounds for machine learning, fundamental concept of machine learning, supervised learning methods (regression & classification), unsupervised learning methods (clustering & dimensionality reduction), ensemble models, Bayesian approaches and models, neural networks, and reinforcement learning.

AI502 Deep Learning

In this course, we will learn about introductory materials for deep learning, which is a machine learning methodology that learns multiple layers of non-linear representations for given prediction tasks, while reviewing some of its applications to computer vision and natural language processing. The course will be mostly focused on understanding deep learning methodology, rather than implementing and using existing deep learning frameworks. We will have three to four lab courses on Tensorflow basics.

AI503 Mathematics for AI

In this lecture, I plan to introduce elementary mathematical concepts frequently used for the area of artificial intelligence. In particular, I will explain some introductory parts of linear algebra, multi-variate calculus, probability(or statistics), algorithms, complexity theory and information theory which are useful to building machine/deep learning models with corresponding applications.

AI504 Programming for AI

Programming for AI aims to introduce several programming languages for deep neural networks and deep probabilistic models. Topic covered includes various deep learning models and probabilistic inference on the programming platform.

AI505 Optimization for AI

Machine learning algorithms in general train their parameters from training data by optimizing their objective functions. This course covers optimization methods with examples of machine learning algorithms.

AI506 Data Mining and Search

Huge amounts of data are being generated everyday, and data-driven decision-making becomes increasingly important. The course covers a variety of topics in data mining, search, exploration, and preprocessing, with a focus on efficient algorithms and tools.

AI599 Special Topics in Machine Learning

This course deals with selected special topics in Machine Learning and related fields that are hard to cover the other courses. It will cover various topics of the fields of Machine Learning and others to keep up with the latest developments and trends.

AI601 Advanced Machine Learning for AI

Machine learning, a sub-field of computer science, has been popular with the era of intelligent softwares and attracted huge attentions from computer vision, natural language processing, healthcare and finance communities to name a few. This course will consider the art of designing good learning algorithms, as well as analyzing an algorithm's computational and statistical properties / performance guarantees. We will also discuss topics such as nonparametric density estimation, nonparametric regression, and Bayesian estimation, and derive the framework for target applications such as privacy, causality, and stochastic learning algorithms.

AI602 Advanced Deep Learning

In this lecture, I plan to cover recent advances in the field of deep learning. Neural networks have been used for many applications in artificial intelligence for more than 30 years. However, due to powerful computing powers and large-scale datasets available nowadays, the field currently made breakthroughs via new techniques, in particular, for last 5 years. I will introduce them as well as their applications.

AI603 Machine Learning Theory

This course covers both classical and recent machine learning theory. In this course we provide fundamental ideas and theoretical frameworks so that students can understand and analyze complexity of algorithms and performance bounds for machine learning algorithms.

AI604 Deep Learning for Computer Vision

This is an introductory course on deep learning for computer vision with emphasis on understanding of convolutional neural networks and their applications to visual recognition tasks such as image classification, localization, and detection. The students will perform term projects, where they implement their own networks using deep learning libraries for their choices of computer vision problems.

AI605 Deep Learning for Natural Language Processing

Natural language processing (NLP), which aims at properly understanding and generating human languages, emerges as a crucial application of artificial intelligence, with the advancements of deep neural networks. This course will cover various deep learning approaches as well as their applications such as document classification, machine translation, question answering, and dialog systems.

AI606 Recommender Systems

As people are confronted with unprecedented amounts of information, recommender systems, which provide people with relevant information, become indispensable to support their decision-making process. The course covers a variety of topics in recommender systems, including collaborative filtering, content-based filtering, and scalability issues.

AI607 Graph Mining and Social Network Analysis

Graphs are ubiquitous, representing a variety of information: online social networks, e-commerce purchase history, the World Wide Web, to name a few. This course covers a variety of topics related to understanding, analyzing and utilizing graph data, with a focus on efficient algorithms and AI problems on graphs. The course also introduces related studies in Physics and Social Science.

AI608 AI-based Time Series Analysis

This course provides a survey of the theory and application of time series methods. Topics covered include stationary and non-stationary models, autoregressions, multivariate time series, deep neural models for time series, inference in persistent time series and structural break. Real-world data in finance, manufacturing and healthcare will be provided for practice.

AI609 Parallel and Distributed Computation for AI

In this course, students will learn mathematical theories associated with parallel and distributed computation often arising in modern artificial intelligence. In particular, iterative algorithms and their distributed implementation, convergence, and communication and synchronization among processing nodes, focusing on asynchronous parallel and distributed algorithms. System of equations, nonlinear optimization, variational inequality problem, shortest path problem, dynamic programming, and network flow problem will be addressed as applications.

AI610 Sequential Decision Making under Uncertainty

The subject of this course is sequential decision making under uncertainty in a system whose evolution is influenced by decisions. The decision made at any given time depends on the state of the system and the objective is to select a decision making rule that optimizes a certain performance criterion. Such problems can be solved, in principle, using the classical methods of dynamic programming. In practice, however, the applicability of dynamic programming to many important problems is limited by the enormous size of the underlying state/action spaces as well as uncertainties in the system. "Neuro-dynamic programming" or "Reinforcement Learning" which is the term used in the Artificial Intelligence literature, uses neural networks and other approximation architectures to overcome such bottlenecks to the applicability of dynamic programming, while using Monte Carlo estimation and/or stochastic approximation to learn models or value functions of the system. The methodology allows systems to learn about their behavior through simulation, and to improve their performance through iterative reinforcement. The focus of this course is to understand the mathematical foundations of this methodology in light of the convergence, degree of suboptimality, computational complexity and sample efficiency of different algorithms.

AI611 Deep Reinforcement Learning

This course covers deep learning for reinforcement learning, which is one of the core research areas in machine learning and artificial intelligence. Deep reinforcement learning has various applications that requires intelligent decision and control, and can used as training method for various machine learning models. Students will be able to understand the graduate-level background principles, and capture recent research trends.

AI612 Machine Learning for Healthcare

We introduce machine learning methods, especially deep learning techniques to process large-scale electronic health records and perform various prediction tasks. In addition, we discuss diverse healthcare-related topics (e.g. interpretability, causality) and modalities (e.g. images, text, knowledge graphs).

AI613 Musical Applications of Machine Learning

This course handles various applications of machine learning in music classification, music transcription, music composition and performance, audio signal processing and sound synthesis. The scope includes not only conventional machine learning but also recent advances in deep learning.

AI614 Robot Task and Motion Planning

We study how a robot can perform task and motion planning in an integrated fashion to accomplish high-level goals, such as cooking. We will go over robot kinematics, motion planning, task planning, integrated task-and-motion planning (TAMP), and learning algorithms for TAMP problems.

AI615 Information Theory

This course covers the core concept of information theory, including the fundamental source and channel coding theorems, coding theorem for Gaussian channel, rate distortion theorem, vector quantization, multiple user channel and multiple access channel.

AI616 Deep Learning Theory

This course discusses theoretical foundations and recent research results in deep learning theory. Three central questions in deep learning theory are covered. (1) Approximation: What functions can deep neural networks represent? (2) Optimization: Why can we train deep neural networks to global optimality? (3) Generalization: Why can deep neural networks avoid overfitting?

AI617 Machine Learning for Robotics

Robots and autonomous systems have been playing a significant role in the modern economy. Our goal is to guide students to get familiar with these recent cutting-edge learning advances in robotics. Through this course, students will gain a basic understanding of learning algorithms (including deep learning) for robotics, such as (1) how agents perceive the unstructured environments, (2) how agents make decisions, and (3) how agents learn and adapt actively and continuously. It will help the audience with their research in learning for robotics or related topics.

AI618 Generative Model and Unsupervised Learning

The goal of this course is to provide in-depth discussions on generative models and unsupervised learning. Students are expected to learn not only the necessary mathematical tools such as probability theory, optimal transport, and stochastic differential equations but also the specific implementation of the algorithms from classical GANs to latest models.

AI619 AI for medical imaging and signals

The goal of this course is to learn and implement artificial intelligence technologies used in modern medical imaging and signal fields. Students learn the principles of medical image and signal acquisition, learn medical diagnosis based on this, and learn the latest technology on how to implement it with artificial intelligence.

AI620 Bias and Ethics in Natural Language Processing

This course will prepare students for real-world research in Natural Language Processing, considering ethical requirements from society as well as research communities. The course will explore how biases in data and models can be both identified (using statistical analysis, adversarial attacks, and model explanations) and mitigated (with modified training, fine-tuning, and dataset adaptation)

AI621 Computational Image Generation and Manipulation

Computational image generation and manipulation is an emerging field created by the convergence of computer graphics, computer vision and photography. Its role is to overcome the limitations of the traditional imaging system and related applications by using computational techniques, which produce a richer, more vivid, and more perceptually meaningful representation of our visual world.

AI701 Bayesian Machine Learning

Bayesian Learning conducts model selection and predictive inference based on Bayesian principles. Topics covered include variational Bayesian inference, Bayesian hierarchical models, Bayesian optimization and Bayesian deep learning.

AI702 Interpretability and Interactivity in AI

Interpretability and interactivity of artificial intelligence techniques emerge as important issues. This course will cover various model interpretation approaches and interactive user interfaces applicable to deep neural networks, the core techniques in artificial intelligence.

AI703 Systems and Applications of Artificial Intelligence and Machine Learning

This course covers advanced research topics in Systems research for Artificial Intelligence and Machine Learning. The course is designed to cover system software, distributed machine learning frameworks, and AI applications in the context of Cloud and Mobile computing. We will carry on an in-depth study on the environment in which AI applications run including NPU, GPU, CPU and mobile APUs as well as software systems design to run AI applications on various platforms.

AI704 Embodied Intelligence

We will study both classical and recent research papers to address the question of "how to design intelligent robots". This will include papers on, but not limited to, perception, planning, learning, and partial observability.

AI705 Large Language Models

In-depth study on data-driven tokenization (e.g. Byte Pair Encoding) and Transformer architecture, which are the foundation of large language models. Learning the characteristics of Encoder-only LMs, Decoder-only LMs, and Encoder-Decoder LMs and being able to implement them. Multi-GPU, Multi-node GPU training and Cloud Computing for large language models. Exploring recent variants of language models, including multimodal LMs and Mixture-of-Experts language models.

AI810 Special Topics in Artificial Intelligence

This course deals with selected special topics in Artificial Intelligence and related fields that are hard to cover the other courses. It will cover various topics of the fields of AI and others to keep up with the latest developments and trends.

AI960 M.S. Thesis Research

Discussions with academic advisor, checking of research progress, and presentation of the current status of thesis progress are made for improved research content of the dissertation.

AI966 M.S. Seminar

This course will provide an exclusive opportunity to meet with professionals who are working in forefront of various fields. With this course, the students will keep up with the latest developments and trends in the fields of AI and others. .

AI980 Ph.D. Thesis Research

Discussions with academic advisor, checking of research progress, and presentation of the current status of thesis progress are made for improved research content of the dissertation.

AI986 Ph.D. Seminar

This course will provide an exclusive opportunity to meet with professionals who are working in forefront of various fields. With this course, the students will keep up with the latest developments and trends in the fields of AI and others.