



COURSE DESCRIPTION

2026 ELITE UNDERGRADUATE SUMMER CAMP EUSC × AI FOR SCIENCE AND ENGINEERING

Course Schedule	Each course meets twice a week, with each session lasting 2.5 hours
Course Load	Participants may choose 2–3 courses to attend during the program
Program Period	July 13 – July 27, 2026

Instructor	Course Title
Prof. Sanjib GHOSH	AI in Quantum Science and Technology 量子科技中的人工智能
Prof. ZHOU Kai & Prof. YIN Yi	Physics Awakening - When Nature teaches AI Models 物理觉醒-当AI读懂自然的法则
Prof. KE Zhihai & Prof. LI Weitang	AI for Chemistry and Materials Science 人工智能在化学与材料科学中的应用
Prof. Kenneth SHUM	AI for Math - Autoformalization of Mathematical Theorems 人工智能与数学证明：从自动形式化到AI辅助证明
Prof. ZHANG Gongqiu	AI for Finance and Investments 人工智能在金融与投资中的应用
Prof. CHEN Yi	Foundations of Machine Learning: From Linear Models to Deep Learning 机器学习基础：从线性模型到深度学习

Course 1

Course Title	AI in Quantum Science and Technology 量子科技中的人工智能
Instructor	Prof. Sanjib GHOSH
Course Highlights	<p>This short course explores the emerging interplay between quantum mechanics and artificial intelligence (AI). On one hand, modern AI techniques are becoming powerful tools for advancing quantum science, enabling efficient data analysis in quantum mechanics, quantum state reconstruction, and intelligent control of quantum experiments. On the other hand, quantum mechanical systems themselves offer new paradigms for information processing, inspiring the development of AI models based on quantum mechanical with the potential for significant computational advantages. Through a combination of conceptual discussions and practical demonstrations, the course introduces participants to how AI can enrich quantum physics and how quantum principles can enhance the capabilities of next-generation AI technologies.</p>
Indicative Syllabus	<p>Session 1: Introduction</p> <ul style="list-style-type: none">• Foundational Principles of AI & ML• Motivation for AI in quantum science• Scope of AI in quantum: simulation, experiments, and data• Quantum mechanics for better AI• New concepts and the promise
	<p>Session 2: AI for Quantum</p> <ul style="list-style-type: none">• Nature of Quantum Data• The Reconstruction Bottleneck: quantum tomography• ML-Driven Quantum data analysis• AI-Assisted Quantum Experiments• ML for adaptive control of quantum systems• ML techniques for analyzing quantum measurements
	<p>Session 3: Quantum for AI</p> <ul style="list-style-type: none">• The Quantum Advantage• Quantum systems for information-processing• Quantum-inspired machine learning models• Data Mapping: in phase, amplitudes, Hamiltonian• Quantum Neural Networks (QNNs)• Architecture of quantum neural networks
	<p>Session 4: Training, Realities, and the Road Ahead</p> <ul style="list-style-type: none">• Training and Optimization of QNN• Advantages and Limitations• Case study: MNIST dataset• Implementing a simple quantum neural network• Training and evaluating the model• Current challenges and future directions

Course 2

Course Title	Physics Awakening - When Nature teaches AI Models 物理觉醒-当AI读懂自然的法则
Instructor	Prof. ZHOU Kai & Prof. YIN Yi
Course Highlights	<p>This course explores the interplay between physics and AI – from physics to AI, and back. It examines how concepts like entropy, variational principles, and energy landscapes shape modern machine learning, and how AI in turn helps solve physical inverse problems. Through the lens of physics, it dives deep into energy-based models (EBMs) as a unifying framework, connecting Boltzmann machines, diffusion models, and stochastic quantization. The course also covers Physics-Informed Neural Networks (PINNs), teaching how to embed differential equations into neural networks for solving forward and inverse problems in science and engineering. At the core are inverse problems – understanding how both PINNs and generative models can be used to infer hidden parameters, reconstruct fields, and discover physical laws from sparse or noisy data.</p>
Indicative Syllabus	<p>Session 1: Energy, Entropy, and Variational Principles Key topics: Entropy, free energy, partition functions, and the principle of least action, their connection to AI; From Boltzmann Machines to modern Energy based models</p>
	<p>Session 2: From Brownian motion to Diffusion Key topics: Emergent Phenomenon and Self-Organization, Brownian Motion and Diffusion Model, Stochastic Dynamics</p>
	<p>Session 3: Advanced Generative Models & Frontiers Key topics: MLE and variational free energy for generative AI, Advanced Generative models from physics perspective; Diffusion models as energy-based models and the stochastic quantization perspective</p>
	<p>Session 4: Inverse Problems and Physics-Informed Neural Networks (PINNs) Key topics: Inverse Problems solving with machine learning, PINNs fundamentals: bridging physics-driven prior/constraints and AI</p>

Course 3

Course Title	AI for Chemistry and Materials Science 人工智能在化学与材料科学中的应用
Instructor	Prof. KE Zihai & Prof. LI Weitang
Course Highlights	<p>This course takes an interdisciplinary approach, combining AI with chemistry and materials science for molecular design, drug/material discovery, and chemical synthesis. It features hands-on learning with practical coding sessions using Python and PyTorch, covering cutting-edge applications such as drug discovery, molecular property prediction, catalyst design, and energy materials. Students will gain industry-relevant skills in data-driven modeling and generative AI for chemical research through expert-led lectures by specialists in AI, chemistry, and materials science.</p>
Indicative Syllabus	Session 1: Introduction to AI in Chemistry and Materials Science
	Session 2: AI-Driven Applications: Synthesis, Drug/Material Discovery & Structure/Property Prediction
	Session 3: AI-Driven Simulation of Molecular Dynamics & Neural Network Potentials
	Session 4: Generative Models, Robotic Chemists, and Future Direction

Course 4

Course Title	AI for Math -- Autoformalization of Mathematical Theorems 人工智能与数学证明：从自动形式化到AI辅助证明
Instructor	Prof. Kenneth SHUM
Course Highlights	<p>AI is beginning to do mathematics. This course focuses on the recent convergence of formal proof assistants (especially Lean) with large language models and automated reasoning, including case studies inspired by systems such as Harmonic's Aristotle. Students will see how LLMs can suggest proof ideas in natural mathematical language, how those ideas can be translated into formal statements, and how Lean's kernel provides a rigorous correctness guarantee. This approach has helped accelerate progress on difficult problems, including work connected to Erdos conjectures. The course emphasizes hands-on capability: learning to read and write Lean proofs, building intuition for proof search and tactic-based automation, and understanding what "verified mathematics" really buys us (and where it still struggles). We will demonstrate practical workflows for AI-assisted theorem proving. By the end, students will be able to formalize and prove nontrivial problems, and to critically assess both the promise and the limitations of AI systems as collaborators in mathematical discovery. The landscape of doing mathematics is changing, and we should begin adapting to this paradigm shift.</p>
Indicative Syllabus	<p>Session 1:</p> <ol style="list-style-type: none">1. Recent developments in AI for math2. Lean fundamentals I: expressions and definitions, types, function, Simple propositions
	<p>Session 2:</p> <ol style="list-style-type: none">1. Lean fundamentals II: propositions-as-types, proof terms vs tactic mode2. Autoformalization using Harmonic's Aristotle
	<p>Session 3:</p> <ol style="list-style-type: none">1. Common mathlib patterns: rewriting, typeclass inference, coercions2. Mini-project: Formalize 1 to 2 short textbook-style theorems, students draft a formalization plan.
	<p>Session 4:</p> <ol style="list-style-type: none">1. Project presentations + reflection2. Wrap-up: where the field is going; next steps/resources

Course 5

Course Title	AI for Finance and Investments 人工智能在金融与投资中的应用
Instructor	Prof. ZHANG Gongqiu
Course Highlights	This course provides a strong link to cutting-edge practice (quantitative investment, derivatives pricing, and risk management) and active academic research. It offers a structured map of AI in finance with a clear entry point and pointers to go deeper, integrating theory and practice so that theory guides practice while real-world constraints motivate theory.
Indicative Syllabus	Session 1: Mathematical foundations of AI for quantitative finance: financial backgrounds; mathematical formulations of financial models and problems; deep learning and reinforcement learnings basics; overview of AI in finance and investment.
	Session 2: Portfolio selection, optimal control and reinforcement learning: mean-variance/risk formulation of portfolio optimization problems; model-based analytical solutions; reinforcement learning for portfolio optimization.
	Session 3: Derivatives pricing and hedging with deep learning: learning pricing map of options with deep neural networks; mathematical foundations of approximation capacity of deep neural network; deep hedging of financial derivatives.
	Session 4: Generative financial models: generative models and how they can capture stylized features of financial data; tail and dependence modeling of financial data; generative approaches for financial computation.

Course 6

Course Title	Foundations of Machine Learning: From Linear Models to Deep Learning 机器学习基础：从线性模型到深度学习
Instructor	Prof. CHEN Yi
Course Highlights	<p>This course provides a progressive introduction to machine learning, beginning with core concepts and practical workflows. It emphasizes linear regression and logistic regression as foundational building blocks—covering loss functions, optimization, and evaluation—that naturally lead into deep learning. Along the way, students complete simple hands-on exercises using scikit-learn to connect theory to real-world data problems, and conclude with a high-level overview of modern language models and Transformers.</p>
Indicative Syllabus	<p>Session 1: Introduction to Machine Learning</p> <ol style="list-style-type: none">1. ML definition and significance2. Supervised, unsupervised, and reinforcement learning3. Model evaluation metrics4. Introduction to linear regression and least squares
	<p>Session 2: Linear Regression and Regularization</p> <ol style="list-style-type: none">1. Batch and stochastic gradient descent2. Scikit-learn hands-on example3. Underfitting, overfitting, and bias-variance trade-off4. Ridge and Lasso regularization
	<p>Session 3: Classification and Model Selection</p> <ol style="list-style-type: none">1. Classification formulation2. Logistic regression and sigmoid function3. Maximum likelihood estimation4. Softmax regression and cross-entropy loss5. Introduction to model selection
	<p>Session 4: Deep Learning and Transformers</p> <ol style="list-style-type: none">1. Deep learning concepts and use cases2. Neural network components and training3. Language models in NLP4. Transformer concepts and model types: BERT, GPT, and T5