Summer-School TUM-IAS

Artificial Intelligence Meets History and Philosophy of Science and Technology

Location: Institute for Advanced Study, Lichtenbergstraße 2a, 85748 Garching / Munich
Time: September 1-5, 2025
Organization: Prof. Dr. Stefania Centrone (TUM Chair for Philosophy of Science)
in cooperation with TUM Institute for Advanced Studies IAS
& PD Dr. Ulf Hashagen and PD Dr. Rudolf Seising (Deutsches Museum Munich),
& Dr. Andrea Reichenberger (TUM: STS Department, History of Technology)
financially supported by the Udo Keller Foundation.

Contributions by: Stefania Centrone (TUM), Roberto Giuntini (TUM; University of Cagliari), Ulf Hashagen (The Research Institute for the History of Science and Technology, Deutsches Museum Munich), Klaus Mainzer (TUM), Ulrich Marsch (TUM), Andrea Reichenberger (TUM), Rudolf Seising (The Research Institute for the History of Science and Technology, Deutsches Museum Munich), Michael Stöltzner (University of South Carolina), Jörg Wernecke (TUM)

Participation in the school is completely free of charge for students!

All meals and refreshments during the event will be covered by the organizing institution for the benefit of all participants.

Program

Monday - morning:

09:00 - 09:30	Welcome
09:30 - 10:30	Stefania Centrone I: Thinking and Reckoning: From Leibniz's Calculus Ratiocinator to Turing's Universal Machine
10:30 - 11:00	Coffee
11:00 - 12:00	Ulf Hashagen I: A Hardware History of Computational Sciences" (1940-2010)
12:00 - 14:00	Lunch

Monday - afternoon:

14:00 - 15:00	Rudolf Seising I Histories of Artificial Intelligence
15:00 - 15:30	Coffee Break
15:30 - 17:30	Tour to the LRZ
17:30 - 18:30	Discussion time

Tuesday - morning:

09:00 - 09:30	Coffee
09:30 - 10:30	Jörg Wernecke I Ethical Challenges and Ethical Bases in AI
10:30 - 11:00	Coffee
11:00 - 12:00	Andrea Reichenberger I Gender Matters: Revisiting the History and Philosophy of AI
12:00 - 14:00	Lunch

Tuesday - afternoon:

14:00 - 15:00	Reading Time: Technical Determinism
15:00 - 15:30	Coffee Break
15:30 - 16:30	Reading Time: Technical Determinism
16:30 - 17:30	Ulf Hashagen II Seminar: Discussion of Technological Determinism
17:30 - 18:30	Jörg Wernecke II Thematic Focus: Ethics and Responsibility. Overview in the Field of Philosophy and Applied Sciences (AI)

Wednesday - morning:

09:00 - 09:30	Coffee
09:30 - 10:30	Stefania Centrone II Seminar: Alan Turing's Theoretical Breakthrough
10:30 - 11:00	Coffee
11:00 - 12:00	Rudolf Seising II Histories of Quantum Mechanics
12:00 - 14:00	Lunch

Wednesday - afternoon:

14:00 - 16:30	Robotics Exhibition Deutsches Museum
16:30 - 17:30	Cultural event
17:30 - 18:30	Cultural event

Thursday - morning:

09:00 - 09:30	Coffee
09:30 - 10:30	Klaus Mainzer I Introduction into the History and Philosophy of Quantum Computing
10:30 - 11:00	Coffee
11:00 - 12:00	Michael Stöltzner I Particle Physics as an Example of Long-term Integration of AI in a Computer-heaving Data Analysis
12:00 - 14:00	Lunch

Thursday - afternoon:

14:00 - 15:00	Roberto Giuntini I From Concepts to Classical Classification
15:00 - 15:30	Coffee Break
15:30 - 16:30	Michael Stöltzner II Short Lecture
16:30 - 17:30	Exercises time Hands-on experience with Gen AI
17:30 - 18:30	Andrea Reichenberger II Picturing the History and Philosophy of AI: Cartoons and Poetry Slam

Friday - morning:

09:00 - 09:30	Coffee
09:30 - 10:30	Klaus Mainzer II Introduction into the History and Philosophy of Machine Learning
10:30 - 11:00	Coffee
11:00 - 12:00	Roberto Giuntini II Quantum-Inspired Classification and the Helstrom Strategy
12:00 - 14:00	Lunch-Round Table Discussion
14:00	Farewell

Abstracts/Participants:

Stefania Centrone (TUM):

Lecture 1: Thinking and Reckoning: From Leibniz's Calculus Ratiocinator to Turing's Universal Machine

Computation and logical reasoning are closely linked concepts, historically developed through philosophy, logic, and mathematics. This lecture series traces the metaphor of thinking as reckoning from its early philosophical roots (Socrates, Plato, Aristotle) through Leibniz's visionary conception of a lingua characteristica universalis and a calculus ratiocinator, intended as universal tools for precise reasoning and decision-making. The series then highlights key developments in the history of logic, including algebraic logic, formal logical systems (Gottlob Frege), and the emergence of fundamental questions regarding the decidability of deductive problems.

Lecture 2: The lecture will focus on Alan Turing's theoretical breakthrough: the concept of the universal Turing machine and special-purpose Turing machines, along with register machines and the Church-Turing thesis. Participants will engage practically by programming simple Turing machines designed to solve specific computational tasks, illustrating how abstract concepts in computability translate into concrete operational procedures.

Roberto Giuntini (TUM; University of Cagliari):

Lecture 1: From Concepts to Classical Classification

In this lecture, we introduce the conceptual and methodological motivations behind the dialogue between quantum theory and machine learning. We begin by discussing the general problem of concept formation and classification in both human and artificial intelligence, drawing from cognitive science, logic, and Gestalt theory. We then transition to the classical framework of supervised machine learning, with a focus on classification tasks.

Using simple examples, we formalize the construction of training and test datasets, the representation of objects as vectors in feature space, and the notion of a classier as an inductive generalization method. Particular attention is devoted to the Nearest Centroid Classier (NCC), which interprets classification in geometric terms by associating classes with their centroid vectors and classifying new inputs based on Euclidean similarity.

- Discussion and Exercises:
- Discuss the philosophical implications of the Gestalt principles in AI.
- Prove that the NCC denes a linear decision boundary.
- Given a 2D dataset, implement (by using Python) the NCC classier and evaluate its accuracy on a test set.
- Reect on the limitations of Euclidean distance as a universal similarity metric.

Lecture 2: Quantum-Inspired Classification and the Helstrom Strategy

Abstract: The second lecture explores how quantum formalism can inspire novel classification strategies, even when operating on classical data. We present two quantum encoding methods amplitude encoding and stereographic encoding that allow real-valued feature vectors to be represented as pure quantum states. From there, we introduce quantum-inspired classifiers: classical datasets are transformed into quantum training sets, where each class is represented not by a centroid vector but by a quantum centroid, i.e., a mixture of quantum states. We then present the Helstrom quantum discrimination problem as a foundational tool for classification: given two quantum centroids, the classier assigns a new input to the class whose centroid is closer in terms of quantum distinguishability. The resulting method, though executable on classical hardware, is derived from deep quantum-theoretic insights and others promising improvements in accuracy.

Discussion and Exercises:

- Encode simple 2D vectors into quantum states using both amplitude and stereographic encodings. Compare the resulting density matrices.
- Derive the Helstrom formula for the optimal discrimination of two quantum states.
- Discuss the epistemological shift: from geometrical similarity (Euclidean distance) to physical distinguishability (trace distance).
- Consider the feasibility and implications of implementing these classifiers on actual quantum hardware.

Ulf Hashagen (The Research Institute for the History of Science and Technology, Deutsches Museum Munich):

Lecture 1: A Hardware History of “Computational Sciences” (1940-2010)

The lecture deals with the question of the extent to which computer hardware development and scientific development have influenced each other since the invention of the computer in the 1940s. Since the advent of the AI boom in the 2010s, this question has assumed a heightened level of urgency in both public and scientific discourse. The prevailing assumption is that the advancement of scientific knowledge is inextricably linked to the development of AI technology. A historical review should help to critically scrutinise this assumption. A historical review should help to critically scrutinise this assumption from a broader historical perspective and to analyse the developments of the last decade.

Seminar 2: Discussion of technological determinism as one of the central historical questions of the history of technology in general

Klaus Mainzer (TUM):

Lecture 1: Introduction into the History and Philosophy of Quantum Computing

Quantum computing represents a fundamental shift in computational theory and practice. This lecture introduces students to its historical evolution, from foundational theoretical insights to current practical implementations. Philosophical implications, including discussions about determinism, reality, and computation in quantum contexts, will also be explored

Lecture 2: Introduction into the History and Philosophy of Machine Learning

This session provides an overview of machine learning from a historical and philosophical perspective. Key developments and methodological debates, from early algorithmic approaches to contemporary deep-learning techniques, will be discussed. Special attention is given to philosophical questions around intelligence, ethics, and the societal impact of machine learning technologies

Andrea Reichenberger (TUM):

Lecture 1: Gender Matters: Revisiting the History and Philosophy of AI

This lecture course offers a critical analysis of how expertise, science communication and public opinion shape and perpetuate biased portrayals of the history and philosophy of AI. We provide an overview of this research by examining two key issues through case studies: firstly, the role of hidden figures as agents of emerging technologies; and secondly, gendered objects and spaces in the context of the rise of computational science and electrical engineering, computational cryptography and GPS technology development.

Session 2: Picturing the History and Philosophy of AI: Cartoons and Poetry Slam

In this session, all participants will have the opportunity to showcase their talents as word acrobats, creative illustrators and hobby artists. Everyone will have the opportunity to contribute, either individually or in groups, to the discussion of the controversial topic of gender in AI. Participation is optional. Listening, watching, discussing and participating are all encouraged!

Rudolf Seising (The Research Institute for the History of Science and Technology, Deutsches Museum Munich):

Lecture 1: Histories of Artificial Intelligence

Several histories of research on AI are presented and discussed. Ideas, methods and theories from logic and mathematics, computer technology, brain research and psychology led to a new research goal, namely to describe those characteristics that stood for human intelligence and learning so precisely that machines could be constructed that simulate these characteristics. Differentiations were made in the 1980s. Approaches of artificial neural networks, decision and classification trees complement Good Old Fashioned AI (GOFAI). Aspects of intelligence are not imitated here by manipulating symbols, but rather by algorithms that search through large amounts of data for patterns - in a way that humans are poorly or not at all able to do.

Lecture 2: Histories of Quantum Mechanics

Since the 1920s, quantum mechanics has developed from the contradictions that emerged in classical physics (wave-particle duality, photoelectric effect, Compton effect). It paved the way for new interpretations of the perceived world (indeterminacy relation, probability density function, Copenhagen vs. realistic interpretation) and established itself as a new fundamental theory of natural science. The results of this new fundamental theory have had an impact on many scientific fields, including plans for quantum computing and quantum machine learning.

Michael Stöltzner (University of South Carolina):

Lecture 1: Particle physics as an example of long-term integration of AI in a computer-heaving data analysis

Elementary particle physicists sometimes claim that many aspects of artificial intelligence have been around for a while in the sense that cutting-edge machine learning models take the place of earlier computer-based methods of data-analysis, such as Monte-Carlo simulations of the detector background. Interestingly, the rise of LLM coincides with a shift of the field from testing theoretical particle models to rather open model-independent explorations of “new physics”. The use of LLM in such searches to identify interesting signatures has prompted discussions about the appropriate degree of supervision, the control or explainability of the models used, and the danger of hallucination in the sense that alleged signals of new physics may lead to a refocusing of research activities.

Lecture 2 (includes practical exercise): Gen-AI in Engineering Ethics and Related Classes

Motivated by my experiences as a long-time teacher of Engineering Ethics and on the basis of a grant integrating genAI methods into the class, I will discuss the various strategies to employ how to make students get acquainted with these technologies by experiment and thus obtain a sense of the dangers and opportunities, and the ethical consequences on their studies. The lecture will have a practical part in which students let Chat-GPT rewrite and grade previous work and let it argue for unethical and even illegal behavior. Students’ experiences provide the basis for ethical reflections on the gen-AI in the classroom, based on the previous lectures on AI ethics.

Jörg Wernecke (TUM):

Lecture 1: Ethical Challenges in AI and Their Ethical Bases: an Introduction and Overview

AI is becoming an emerging technology that is permeating almost all areas of life. However, the many advantages of this new technology also raise questions regarding its normative, i.e. ethically appropriate or necessary, application. Ethics of AI analyzes and discusses its possibilities and limitations within different fields of application. Against this background, this course examines the question of what specific ethical challenges exist (e.g. bias, manipulation, moral errors of judgment, etc.) and what analytical tools ethics provides to enable the development of responsible innovation. The course is therefore also an introduction to applied ethics.

- Short overview of AI and their normative ethical challenges.
- Ethical decision making by AI – Example Autonomous Driving (Moral machines)
- Toolbox of Ethics: Ethical Reasoning and Justification of moral decisions and actions.

Lecture 2: Thematic focus: Ethics and Responsibility. Overview in the Field of Philosophy and Applied Sciences (AI)

- Concrete Applications – testing by students
- Different applications of AI and their ethical challenges:
- AI & Medicine
- AI & Robotics
- AI & Society
- AI & Moral Machines
- Concrete Applications – testing by students