Photo: Z. Mesarovic

AI and Society in a Transforming World

Circle U. AI Knowledge Hub

Summer School for PhD students and Postdocs

7 – 11 September 2026

University of Belgrade
BELGRADE • SERBIA

PROGRAMME

Co-funded by the European Union

Circle U.
European University Alliance

Course Information

Artificial intelligence is a transformative technology of our age. This interdisciplinary Summer School explores **the interplay of AI and society** in this transformation. It will examine its technological underpinnings, its social drivers and effects, and how they interact.

It brings together a broad range of participants and speakers from the member universities of the Circle U. Alliance, creating a dynamic and collaborative learning environment.

Participants will engage in lectures, workshops, case studies, and hands-on group projects, gaining both theoretical and practical knowledge to understand the AI systems and the new technological revolution which is built upon and their interplay with human society. Collaborative sessions will encourage interdisciplinary exchange, creative thinking, and critical reflection on the topic.

Key topics:

- Agentic artificial intelligence
- Large language models
- Ethical aspects of artificial intelligence
- AI policy, regulation, and societal impacts
- Human and machine decision-making

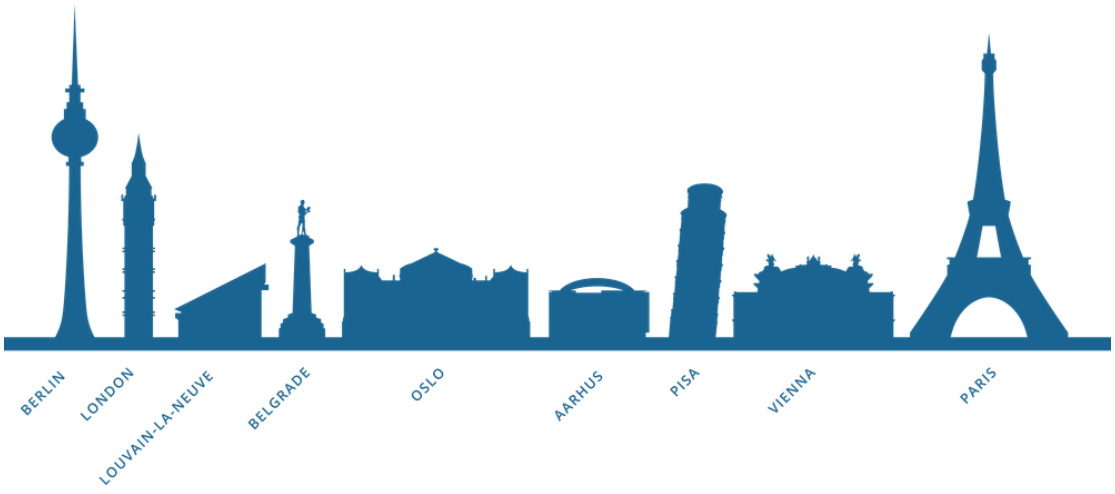

Course Information

Participants will have the opportunity to

- **Explain** the fundamental concepts and technological foundations of artificial intelligence, including key methods, tools, and applications.
- **Analyze** the societal drivers and impacts of AI, including ethical, economic, political, and cultural dimensions.
- **Critically assess** the interaction between AI technologies and human society, identifying both opportunities and risks.
- **Apply interdisciplinary perspectives** to understand complex challenges related to AI and its societal implications.
- **Engage in collaborative problem-solving** through group projects addressing real-world AI-related issues

ECTS: 3
Language of instruction: English

Academic contact:
Professor Miroslav Marić miroslav.maric@matf.bg.ac.rs
Professor Francesco Marcelloni francesco.marcelloni@unipi.it

Administrative contact:
Miloš Bojičić, University of Belgrade milos.bojicic@rect.bg.ac.rs

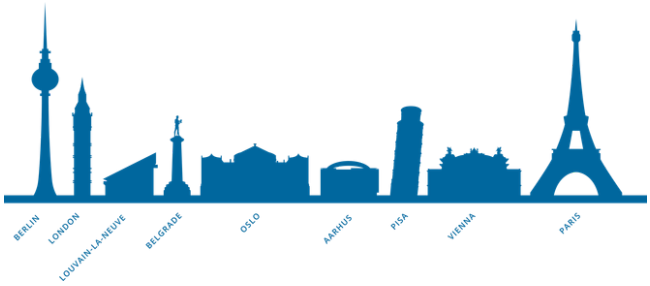

Programme

Day 1 **Trustworthy AI and Research Ethics**	Day 2 **From Language Models to Decision Intelligence**	Day 3 **Reasoning and Explainability in AI**	Day 4 **Intelligent Systems: Swarms, Agents, and LLMs**	Day 5 **AI and Regulations**
10:00-11:30 Prof. Francesco Marcelloni **Human-Centered AI: Balancing Innovation, Trust, and Human Values**	**10:00-11:30** Prof. Stephan Oepen **Open Large Language Models: A European Perspective**	**10:00-12:30** Prof. Guillaume Wisniewski **Reasoning with Large Language Models: Mechanisms, Evaluation, and Critical Use**	**10:00 – 11:30** Dr Trifun Savić **From Low-Power Wireless to Swarm Robotics: Enabling Distributed Intelligent Systems**	**10:00-11:30** Prof. Djordje Krivokapic **One Technology, Many Rules: Understanding the Architecture of AI Regulation**
11:30 – 12:00 Break	**11:30 – 12:00** Break	**12:30-13:00** Break	**11:30-12:00** Break	**11:30-12:00** Break
12:00 – 13:30 Prof. Yves Deville **Responsible use of generative AI in research**	**12:00 – 13:30** Prof. Boris Delibašić **Bridging the Worlds of Machine Learning and Decision Support Systems**	**13:00- 14:30** Dr Vanja Subotić **Are Artificial Neural Networks Explanatory at All?**	**12:00 – 13:30** Prof. Mladen Nikolić **Large Language Models and Intelligent Agents**	**12:00-14:00** **Group projects presentations**
13:30-14:30 Lunch	**13:30-14:30** Lunch	**14:30-15:30** Lunch	**13:30-14:30** Lunch	**14:00-15:00** Lunch
14:30 Group Work	**14:30 Group Work**	**15:30 Group work**	**14:30 Group work**	**Closing words**
18:00 Welcome coctail and buffet for all participants at the courtyard		**17:30** Belgrade walking tour		

УНИВЕРЗИТЕТ У БЕОГРАДУ
UNIVERSITY OF BELGRADE

Introducing our lecturers

Dr Miroslav Marić (Academic Coordinator of the Summer school) is a Full Professor and the Head of the Department of Computer Science and Informatics at the Faculty of Mathematics, University of Belgrade, the President of the Mathematical Society of Serbia, the Head of the GeoGebra Center of Belgrade, and a member of the National Education Council. He received his BSc. degree (2002), MSc (2006) and PhD (2008) in Computer Science from the Faculty of Mathematics, University of Belgrade. His research interests are: Artificial Intelligence, Operational Research, Credit Risk, Mathematical Optimization, Parallel Algorithms, Teaching Methodology in Mathematics and Computer Science, Computer Graphics, etc. Throughout his academic career, he authored over 120 research papers, mostly in the field of Artificial Intelligence. Also, he supervised over 100 MSc and 5 PhD thesis.

Miroslav was the Team leader for four national projects, participated at several national and international projects and was a member of program and organizational committees of several national and international conferences. Notably, he was the leader of the project Final Exam – A Platform for Equal Inclusion of Students in the Educational Process which won the Award Innovative Practice 2020 given by the Zero project founded by the Essl foundation and United Nations in Vienna. Also, he lead the project of the Ministry of Education, eGovernment, cooperated by the Faculty of Mathematics and the Faculty of Organizational Sciences at the University of Belgrade, which resulted in the creation of the „eVežbaonica" platform which was later awarded the DIS Plaque for its outstanding contribution to the development of Informatics in Serbia in 2022.

He won the Saint Sava Award for the year 2020, which is the most significant award in the field of education given by the Ministry of Education of the Republic of Serbia. He received the EdTech Western Balkans Person of the Year award.

Introducing our lecturers

Prof. Yves Deville is a Full Professor at the Engineering School of UCLouvain, Belgium. He obtained an undergraduate degree in Computer Science from the University of Namur in 1983, an M.S. degree in Computer Science from Syracuse University (USA) in 1986, and a Ph.D. from the University of Namur in 1987. He joined UCLouvain as an Associate Professor in 1991.

His research interests include Artificial Intelligence, constraint satisfaction, and optimization. He is the author or co-author of more than 150 scientific articles and book chapters.

From 2014 to 2024, Yves Deville was Senior Advisor to the President of UCLouvain for Digital University and Open Science. Since 2014, he is also chairing the university's Information Systems Steering Committee. He is currently Special Advisor to the President for Artificial Intelligence.

He led the development of UCLouvain's digital strategy, with a focus on Open Education and Open Science. He is now responsible for defining and implementing the university's strategy for Artificial Intelligence.

Introducing our lecturers

Prof. Francesco Marcelloni is Full Professor at the **University of Pisa**, Director of the Knowledge Hub on AI of the European Alliance Circle U., and coordinator of the multidisciplinary Good AI Lab and the public–private IT2PAO Lab at the University of Pisa. He is also a member of the University's Board of Directors. His research interests include federated learning, explainable artificial intelligence (XAI), fake news detection and debunking, and clustering algorithms. He is the (co-)author of one book and more than 300 scientific publications. He has edited five volumes and four special issues of international journals. He received the IEEE Transactions on Fuzzy Systems Outstanding Paper Award (2021) and the IEEE Computational Intelligence Magazine Outstanding Paper Award (2022). He currently serves as Associate Editor of international journals and as General Chair of international conferences.
He has coordinated numerous national and international research projects.

Introducing our lecturers

Prof. Stephan Oepen studied Linguistics, German and Russian Philology, Computer Science, and Computational Linguistics at Berlin, Volgograd, and Saarbrücken. He first joined the **University of Oslo** in 2002, as a researcher at the Department of Linguistics and Scandinavian Studies, and was appointed full professor at the Department of Informatics in 2007. Between 2011 and 2017, he headed the Research Group for Language Technology at the University of Oslo; in 2019–20 he served as the manager of the Section for Machine Learning at the Department of Informatics; and for the period 2021–24 he was elected as department head. He co-leds two European research projects on open Large Language Models, HPLT and OpenEuroLLM.

Prior to his current appointment, he worked, among others, at the German Research Center for Artificial Intelligence and Saarland University (both Germany), YY Technologies (Mountain View), and Stanford University (both USA). His research revolves around the integration of linguistics and computing, where he has published some hundred peer-reviewed research articles and conference papers and has co-founded the world-wide DELPH-IN network and the Nordic Language Processing Laboratory (NLPL). Oepen has supervised to completion ten doctoral projects and about twenty-five Master of Science degrees.

Introducing our lecturers

Guillaume Wisniewski is a Senior Lecturer (maître de conférences) in Information Science and Computational Linguistics at **Université Paris Cité**, where he is a member of the Laboratoire de Linguistique Formelle (LLF). His research in natural language processing centres on the emergent capabilities of large language models, investigating which mechanisms allow LLMs to speak and reason so effectively, and how these abilities arise over the course of training, as well as on audio and speech models. Alongside this work, he develops tools for under-resourced languages and contributes to computational documentary linguistics, supporting the description and preservation of languages that lack substantial digital resources.

Dr Trifun Savić is a research scientist and innovator working at the intersection of low-power wireless systems, embedded intelligence, and swarm robotics. He earned his PhD from Sorbonne University, where he conducted research at Inria Institute on advanced ICT solutions deployed across three continents. He previously completed his Bachelor's and Master's degrees at University of Montenegro. Dr. Savić is the inventor of three patents and has published over 20 research papers in leading journals and conferences. His career spans academia, industry, and public service, including leadership roles in IoT innovation and government. In 2024, he received Montenegro's Award for Scientific Achievements for the most successful invention.

Introducing our lecturers

Prof. Boris Delibašić is a full professor at the **University of Belgrade** – Faculty of Organizational Sciences. His research interests are related to data science, machine learning, business intelligence, multi-criteria decision analysis, and decision support systems. He is the deputy coordinator of the EURO (Association of European Operational Research Societies) working group on Decision Support Systems. He earned his PhD in 2007 at the University of Belgrade. He was a visiting lecturer at Friedrich Schiller University in Jena, Germany, from 2006 to 2011. He received the Fulbright visiting program scholarship for the 2011/2012 academic year at the Center for Data Analytics and Biomedicine, Temple University, Philadelphia, Pennsylvania, USA. He is often engaged as a consultant, helping companies utilize their data and decision support system solutions.

Prof. Mladen Nikolić is an associate professor at the Faculty of Mathematics, **University of Belgrade** where he teaches artificial intelligence and machine learning courses. His work in AI spans two decades and various subfields, primarily in machine learning, but also in symbolic AI. He conducted research mainly in graph-based learning, computer vision, and various applications of machine learning. In the last years, he is increasingly involved in large language models and their applications.

Introducing our lecturers

Prof. Dr. Đorđe Krivokapić LL.M. is an Associate Professor of Information and Communication Technology Law at the Faculty of Organizational Sciences, **University of Belgrade**. He is the responsible professor for courses in the areas of information and communication technology law, data protection, business law, data & business ethics, business systems organisation, and human resource management within the curricula of the Faculty of Organizational Sciences, the Faculty of Electrical Engineering, the Faculty of Medicine, and multidisciplinary studies at the University of Belgrade, as well as the Faculty of Drama Arts in Belgrade.

He holds a law degree from the University of Belgrade and a master's degree from the University of Pittsburgh, USA. In 2016, he defended his PhD on internet defamation at the University of Belgrade. Prof. Krivokapić spent time furthering his education and preparing his doctoral dissertation at the Berkman Center for Internet & Society, Harvard Law School, USA, in 2010 and 2012.

Introducing our lecturers

Dr Vanja Subotić earned her BA (2017), MA (2018), and PhD in Philosophy (2023) from the Faculty of Philosophy at the University of Belgrade, as well as MSc in Computational Methods in Humanities and Social Sciences from the Center for Multidisciplinary Studies at the **University of Belgrade** (2021). Her PhD thesis is the first to explore the philosophy of AI in Serbia, with a specific focus on the methodological aspects of large language models. Subotić presently works as a Research Associate at the Faculty of Philosophy. Previously, she was an EPSA Postdoctoral Research Fellow at the University of Turin and an ERASMUS Visiting Teaching Staff member at the University of Rijeka. From 2022, she served as an expert consultant for AI and emerging technology for the Horizon 2020 project TechEthos. To date, Subotić has published 18 papers in peer-reviewed journals and edited volumes (categories M21a+, M21, M22, M23, M24, M14 as per Serbian quantitative regulations), co-edited a monograph, and participated in 40 international and national conferences and panels, as well as 17 workshops, winter schools, and summer schools.

Lecture apstracts

Prof. Francesco Marcelloni, University of Pisa

Human-Centered AI: Balancing Innovation, Trust, and Human Values

Artificial Intelligence is undergoing a profound transformation driven by the rapid evolution of foundation and generative models, which are increasingly integrated into scientific, industrial, and societal processes. While these systems provide remarkable capabilities in language understanding, content generation, reasoning, and decision support, they also raise significant technical and socio-technical challenges. Issues such as hallucinations, bias, opacity, privacy risks, misinformation, and reduced human oversight raise fundamental questions about trust, accountability, and the role of humans in AI-driven environments.

Human-Centered AI (HCAI) has emerged as a key paradigm for addressing these challenges by promoting AI systems aligned with human values, societal goals, and ethical-legal principles. Beyond performance, HCAI emphasizes transparency, robustness, fairness, human agency, and the augmentation of human capabilities. These principles are increasingly reflected in international regulations, including the EU AI Act.

This lecture will discuss the main technical and societal risks associated with contemporary AI systems, with particular attention to generative AI. It will then introduce the foundations and design principles of Human-Centered AI, examining how AI systems can remain trustworthy, controllable, and beneficial for society. Finally, the lecture will present our ongoing research on explainable AI and federated learning.

Lecture apstracts

Prof. Yves Deville, UCLouvain

Responsible use of generative AI in research

Generative AI is changing the nature and methods of scientific research. Research is one of the sectors that could be most profoundly transformed by generative AI. In May 2026, the European Union published the third version of its Living guidelines on the responsible use of generative AI in research [1]. This lecture will discuss the challenges posed by generative AI, particularly in research, as well as the importance of proposing guidelines for its use. It also described different generative AI tools that could be used for research.This talk will be interactive as the participants will be asked to share their own experience with generative AI in their research.

[1] Living guidelines on the responsible use of generative AI in research. An ERA Forum stakeholders' document. May 2026. https://research-and-innovation.ec.europa.eu/document/download/2b6cf7e5-36ac-41cb-aab5-0d32050143dc_en?filename=ec_rtd_ai-guidelines.pdf

Lecture apstracts

Prof. Stephan Oepen, University of Oslo

Open Large Language Models: A European Perspective

Large Language Models (LLMs) are the core of much current generative AI; they have become a critical societal infrastructure. Recent breakthroughs have largely been obtained through scaling, i.e. ever increasing model sizes, training data volumes, and computational cycles. This development is dominated by a handful of large players, predominantly in the USA and China, while for university researchers working on LLM is anything but a level playing field. This lecture will review the fundamentals of LLM development, including data selection, architecture variation, different training phases, and evaluation, with particular emphasis on multilingual perspective and the languages of Europe, as well as on transparency and replicability in terms of training data, software, and results.

Lecture apstracts

Prof. Guillaume Wisniewski, Universite Paris-Cite

Reasoning with Large Language Models: Mechanisms, Evaluation, and Critical Use

This lecture explores how large language models produce reasoning, and how to draw on them with discernment. It opens with mechanisms, examining how next-token prediction, in-context learning, and techniques such as chain-of-thought give rise to apparently reasoned behaviour, and why grasping these foundations is essential to using such models well. It then addresses evaluation: the benchmarks and probing methods used to assess reasoning, what they reveal, and how contamination and shortcut learning can render them misleading. Practical strategies follow, covering prompting, task decomposition, and verification for applying LLMs to reasoning problems. The course closes with the limits such as encoded biases, brittleness, and the ethical questions raised by delegating inference to opaque systems, cultivating informed, critical use throughout.

Lecture apstracts

Prof. Boris Delibašić, University of Belgrade

Bridging the Worlds of Machine Learning and Decision Support Systems

ML has become the dominant paradigm for predictive modeling across domains, delivering impressive performance in tasks such as customer churn prediction, fraud detection, and risk assessment. However, real-world decision-making rarely depends on prediction accuracy alone. Decision makers must consider constraints, policies, interpretability, and the ability to justify and adapt decisions under changing conditions. These requirements are traditionally addressed by Decision Support Systems (DSS), which emphasize transparency, consistency, and actionable recommendations. This talk explores how to bridge the gap between ML and DSS, arguing that predictive models must be transformed into decision-aware systems to be truly useful in practice. We will introduce the concept of model transformation, where standard ML models are augmented or restructured to incorporate domain knowledge, monotonicity constraints, and decision logic. A central focus will be on concept-based representations, where intermediate, human-understandable concepts (e.g., usage, engagement, risk) mediate between raw data and final decisions. Using a telecom customer churn case study, we will demonstrate how DSS models can be used to convert predictive insights into structured, interpretable decision rules that support counterfactual reasoning (“what should change to alter the decision?”). We will show how such models naturally support key DSS properties, including comprehensibility, consistency, and sensitivity analysis, while maintaining competitive predictive performance.

The talk will also discuss recent developments in explainable AI, concept learning, and algorithmic recourse, highlighting how these areas converge toward a common goal: enabling systems that not only predict outcomes but also support robust, transparent, and actionable decisions. The session will conclude with open challenges and research directions at the intersection of ML and DSS.

Lecture apstracts

Prof. Mladen Nikolić, University of Belgrade

Large Language Models and Intelligent Agents

Large language models (LLMs) have propelled artificial intelligence from a steadily advancing field with important but relatively limited applications into a pervasive, general-purpose technology. Today, most technology companies strive to integrate LLMs into their products, while individuals increasingly incorporate them into their daily work. At the same time, their rapid adoption raises concerns about reliability, control, and broader societal impact. In the context of human–AI interaction, LLMs are best understood as the source of intelligence within intelligent agents. Intelligent agents are autonomous or semi-autonomous, goal-driven systems capable of perceiving, communicating, reasoning, and taking actions that affect real-world processes. They can also be seen as a new software design paradigm, in which developers define goals, communication protocols, and structural constraints, while delegating execution to adaptive, learning-based components.

In this lecture, we will introduce the machine-learning foundations of LLMs, including their basic architecture (transformers), training approaches (pretraining and fine-tuning), and evaluation. We will examine their current capabilities and limitations, such as strong performance in reasoning and generation alongside issues like hallucinations and sensitivity to prompting, as well as challenges in real-world deployment. We will then focus on the design of intelligent agents as the primary means of deploying LLM-based intelligence in practical applications. We will present key building blocks and design patterns, including tool use and retrieval-augmented approaches, and discuss their advantages and limitations. Finally, we will briefly consider practical and ethical aspects of deploying such systems, including robustness, human oversight, and their broader impact on work and society.

Lecture apstracts

Dr Trifun Savić, University of Montenegro

From Low-Power Wireless to Swarm Robotics: Enabling Distributed Intelligent Systems

Today's intelligent systems are increasingly decentralized, mobile, and collaborative. Rather than relying on a single powerful machine, emerging applications run on networks of constrained, AI-enabled autonomous devices. This lecture examines the technological foundations enabling this shift and explores how advances across communications, localization, robotics, and artificial intelligence are reshaping modern autonomous platforms. We begin with low-power wireless communication, focusing on how resource-constrained devices exchange information efficiently under strict energy and bandwidth limits. Next, we address localization, including indoor positioning, cooperative ranging, and high-precision outdoor methods, which provide mobile agents the spatial awareness essential for reliable navigation and coordinated movement. The lecture then discusses how artificial intelligence extends these capabilities through on-device perception and local decision-making on embedded hardware. Finally, we integrate these layers within swarm robotics, where groups of simple robots cooperate through local interactions rather than centralized commands. Inspired by collective behaviors in nature, robotic swarms unlock powerful applications in smart agriculture, warehouse automation, disaster response, and environmental monitoring—illustrating a vital interdisciplinary frontier in modern ICT engineering.

Lecture apstracts

Dr Vanja Subotić, University of Belgrade

Are Artificial Neural Networks Explanatory At All?

The last fifty years of philosophy of science witnessed the mushrooming of many definitions and types of scientific models and explanations (Van Frassen 1980, Giere 1988, 2004). Most often, however, models were understood to have a representational function: they were supposed to accurately represent and describe the target phenomenon (or its relevant aspects) to yield an explanation and understanding of it. The explanation in question was either deductive or inductive, depending on the specific scientific branch or method used therein (Hempel 1958, Simon 1961, Salmon 1971).

On the other hand, most methodological accounts of artificial neural networks (ANNs), including large language models (LLMs) and convolutional neural networks (CNNs), have stressed that the connection between predictive and explanatory power is severed in such models because the similarity between them and their target systems significantly deviates from the ideal of representative models. This has led many to claim that ANNs are not explanatory at all (Boge 2021, Greif 2021). In the lecture, we will discuss what kind of explanation we cannot reasonably expect from ANN-based models. We will introduce mechanistic explanations (Subotić 2024) and assess whether they align with the actual modeling practice across several scientific domains, e.g., cognitive neuroscience, computational linguistics, or biology. Finally, we will link this kind of explanation to recent attempts to increase the transparency of ANNs through mechanistic interpretability.

Photo: Z. Mesarović

www.circle-u.eu