Book of Abstracts - AI Summer School 2025

This volume brings together the abstracts of the posters presented at the AI Summer School 2025, hosted by Friedrich Schiller University Jena. As part of the program, every participating student was required to design and present a poster in one of two dedicated poster sessions. These posters represent either the students' own research projects or their critical engagement with existing work in the field of artificial intelligence.
The collection highlights the range of topics and perspectives explored during the summer school. Many contributions review, analyze, and contextualize established methods, while others showcase original approaches and applications. Together, they provide an overview of current research directions in artificial intelligence.

All posters are available online at: `https://ai.uni-jena.de/posters/`

The committee:

Prof. Dr.-Ing. Martin Bücker
Prof. Dr. Gerhard Zumbusch
B.Sc. Tamino Steinert

AI-Driven Financial Portfolio Management

Atifa Tokhi[1]

[1] *Friedrich Schiller University Jena*

Abstract: Artificial intelligence (AI) is reshaping portfolio management by augmenting classical mean-variance optimization with machine learning (ML) for prediction, risk modelling and decision automation. This project surveys state-of-the-art methods (supervised learning, deep learning and reinforcement learning) and implements a compact, reproducible workflow for equity portfolios under realistic frictions (transaction costs, turnover constraints, risk budgets). Preliminary results indicate that ML models can improve risk-adjusted returns relative to equal-weight and mean-variance methods when rigorously regularized and combined with robust risk controls. However, performance remains sensitive to feature design, transaction costs, and overfitting risks.

Real-Time Processing for Consumer BCI

Daria Elagina[1]

[1] *Friedrich Schiller University Jena*

Abstract: Abstract: Someday in the future, neurointerfaces will become as common as smartphones or fitness trackers. While EEG systems used in research require high resolution for better decoding, necessitating many electrodes, consumer EEG headsets are designed with significantly fewer electrodes. Moreover, consumer applications demand minimal training time and near plug-and-play usability. This poster highlights some challenges of consumer BCIs and algorithms designed for such constraints, focusing on lightweight machine learning and deep learning methods that enable adaptive feature extraction and real-time classification on resource-limited devices.

Fingerprint Biometrics: A Comparative Study of Deep Learning Models (CNNs vs. ViTs)

Kemkuini Fouodji Gertrude Ketsia[1]

[1] *Friedrich Schiller University Jena*

Abstract: Fingerprint recognition, while widely adopted for its uniqueness, faces challenges with low-quality or partial prints—especially in pediatric applications. Traditional minutiae-based systems struggle in unconstrained environments, prompting a shift toward deep learning. This study compares Convolutional Neural Networks (CNNs) and Vision Transformers (ViTs) for fingerprint biometrics, evaluating their performance across contact-based and contactless scenarios.

CNNs excel in controlled settings, achieving 98.21% accuracy (Pandya et al.) on contact-based datasets by hierarchically extracting local ridge patterns and minutiae. Conversely, ViTs demonstrate superior robustness in contactless environments, attaining 96% accuracy (Kaplesh et al.) by modeling global patch relationships via self-attention. Key trade-offs emerge:

- CNNs generalize efficiently with small datasets but falter under distortion.
- ViTs handle noise and variability effectively but demand large-scale data.

We conclude that hybrid architectures—combining CNNs' local feature power with ViTs' global context—offer a promising path toward accuracy, efficiency, and robustness. This work underscores the need for context-driven model selection in fingerprint biometrics and charts a course for next-generation solutions.

Prompt your Experiment – The Artificial Scientist

Hinrich Dirksen[1]

[1] *Friedrich Schiller University Jena*

Abstract: Artificial Neural Networks (ANNs) have long been indispensable across physics-from modeling condensed-matter systems to processing particle collider data. This synergy between AI and physics was recently highlighted by the 2024 Nobel Prize. With the emergence of large language models (LLMs), AI can be tailored to create scientific instruments, not just analyze their outputs.

The presented paper pioneers the AI-driven design of Gravitational Wave Detectors (GWDs). It introduces a hybrid framework that transforms the intractable combinatorial search for interferometer topologies into a continuous optimization problem. The pipeline to generate GWD designs includes Boltzmann selection to prioritize high-sensitivity designs, massively parallel BFGS gradient optimization for parameter tuning and automated topology simplification to reduce components.

Validated against physics simulations, AI-generated designs outperform existing baselines for several astronomical targets. Among 50 discovered configurations are presented in a GWD zoo. Studying the generated configurations allows to explore new unconventional optimization possibilities.

The poster investigates this step towards a fully autonomous artificial scientist and the power of AI-driven discovery.

You Only Look Once: Real-Time Object Detection with Neural Networks

Jamie Klöser[1]

[1] *Friedrich Schiller University Jena*

Abstract: Real-time object detection must localize and classify multiple objects fast enough for on-device uses such as assistive robotics, AR, and safety monitoring. This poster explains how YOLO achieves high throughput by predicting boxes, objectness, and classes in a single forward pass (Backbone→Neck→Head) followed by non-maximum suppression. We summarize the box-regression objective (IoU-family losses) and analyze the resolution trade-off—higher input size typically improves AP but reduces FPS—using literature data from recent YOLO variants. We complement this with a practical playbook: where YOLO excels (clear, near objects) and where it fails (tiny, occluded, low-light, motion-blurred scenes), plus quick fixes without retraining (input-size tweaks, threshold/NMS tuning, lightweight tracking). A small, reproducible demo with pretrained weights illustrates these ideas on consumer hardware. Overall, the poster provides an intuitive map from architecture to deployment choices when balancing accuracy and latency in real-time detection.

Impact of AI on the Economy and the Labor market

Jannik Strobel[1]

[1] *Friedrich Schiller University Jena*

Abstract: After the publication of ChatGPT 3.5, Large-Language-Models (LLMs) have gotten great attention by the public and have become synonymous with AI. But LLMs are just the newest addition to the AI transformation of the Economy and traditional Machine Learning, Deep Learning and other Artificial Intelligence Systems still play a big role.

In an increasingly digital world AI becomes increasingly powerful and important and leads to a new era of economic growth by increasing productivity and innovations.

But there are also new risks that come with AI. One of the most significant impacts of AI on the job market is the automation of routine and repetitive tasks. A study by the McKinsey Global Institute estimates that by 2030, up to 800 million jobs worldwide could be lost to automation. [4]

The extent of AIs impact on a particular profession can range from full automation to no change and is determined by a handful of factors that are correlated with the tasks of the profession itself and more general societal factors.

To successfully leverage AIs power, we need proactive policies, standards and regulations.

A Variational Network for MR Image Reconstruction

Jiaxin Zhang[1]

[1] *Friedrich-Alexander-Universität Erlangen-Nürnberg*

Abstract: A Variational Network (VN) is used with compressed sensing to accelerate data acquisition during the examination or enhance the image quality with the same amount of acquired data. By using multiple cascaded gradient decent blocks the VN imitates the iterations in compressed sensing. VN can achieve around 90% of similarity and 35 dB of PSNR.

A Comparison of Two Methodological Approaches for Emotion Recognition from EEG Data

Jonas Engicht[1]

[1] *Friedrich Schiller University Jena*

Abstract: Emotions are essential to everyday human life and may appear trivial. However, numerous questions concerning emotions remain unanswered in psychological research. For example, are emotions categorial or dimensional constructs? Recent studies have used electroencephalogram (EEG) experiments to gain further insight into human emotions. This poster investigates two methodological approaches for recognizing human emotions from EEG data, both based on Russel's emotion theory and using the DEAP dataset. The first approach extracts intrinsic mode functions (IMFs) from the EEG signal using Empirical Mode Decomposition and Variational Mode Decomposition. It calculates the first difference and peak value of the power spectral density to be used as input for a Deep Neural Network. The second approach extracts numerous features from the EEG signal and projects them into three types of feature maps. Extracting key features using deep learning and fusing them together, participants' EEG signals are classified using Super Vector Machine classification. Model evaluations indicate that the IMF approach achieves high training accuracy, but only marginally above chance accuracy for test data. Nevertheless, Variational Mode Decomposing outperforms Empirical Mode Decomposition. However, feature map approaches suggest higher generalizability evidenced by consistently high accuracy with Cross Validation. Although feature maps noticeably outperform IMFs, they are more interpretable. Both approaches, however, only allow binary classification into high/low Arousal/Valence. Future research should investigate models for continuous multidimensional predictions to properly depict more in-depth facets of human emotions.

Smarter Maps for a Warmer World: AI Approaches in Environmental Modelling

Mahshad Forghanifar[1]

[1] *University of Bayreuth*

Abstract: Climate change poses significant threats to agriculture in Sub-Saharan Africa, a region heavily reliant on rainfed cropping systems and highly susceptible to climatic variability. This study examines the use of machine learning techniques to predict changes in maize production suitability across Sub-Saharan Africa under present and future climate scenarios. Maize presence data were combined with 19 bioclimatic variables from the WorldClim dataset to train the model for baseline conditions and for the projected climate scenario SSP3-7.0 by the year 2050. The methodology involved comprehensive data preparation, reclassification of crop distribution datasets, spatial thinning to reduce sampling bias, and extraction of bioclimatic predictors. Model accuracy, evaluated through classification performance metrics and partial response curves, reached approximately 70%. Results indicate a projected reduction in maize suitability in southern parts of Sub-Saharan Africa by mid-century, emphasizing the need for proactive, climate-informed agricultural planning. The findings highlight the effectiveness of spatial ecological modelling as a tool for informing agricultural resilience and adaptation strategies under future climate change.

Using Stochastic Computing with Adjustable Sequence Length to optimize energy consumption of embedded Neural Networks

Niclas Starost[1]

[1] *TU Ilmenau*

Abstract: Using neural networks (NN) on embedded/mobile devices is most of the time infeasible due to energy consumption, performance and the resultant latency.

One way of reducing the computational costs is by exchanging floating point (FP) operations with stochastic computing (SC) operations that can be implemented very efficient in hardware and also tolerate errors by definition.

In this Poster, the methods of stochastic computing is introduced and how they are integrated into NNs. Also one new way of reducing the sequence length after training while preserving sufficient accuracy is shown.

Beyond Memorization: Learning That Generalizes

Niklas Maximilian Becker-Klöser[1]

[1] *Friedrich Schiller University Jena*

Abstract: Deep models often reach high accuracy on held-out data drawn from the independent and identically distributed (IID) test split, but still fail when the data distribution shifts — i.e., on out-of-distribution (OOD) inputs. This points to memorization rather than learning the underlying concept. Contributing factors include double-descent effects, the capacity to fit random labels, and reliance on spurious shortcuts. A compact evaluation setup is to report clean accuracy together with ImageNet-C/P (mean Corruption Error, mCE; perturbation stability) and WILDS (IID → OOD performance gaps). Because of underspecification, models with the same validation score can diverge once the distribution shifts; it is useful to show variability across seeds and minor design changes rather than a single run. In practice, targeted data augmentations, suitable inductive biases, and shift-aware evaluation improve reliability and help transfer apparent fit to the cases that actually matter.

Machine Learning-Based Feature Selection and Classification of NSCLC Using Gene Expression Data

Oleksii Vekha[1]

[1] *Friedrich Schiller University Jena*

Abstract: Non-small cell lung cancer (NSCLC) remains a leading cause of cancer-related mortality causing >1.5 million deaths annually, making accurate early classification crucial. This review examines two computational strategies applied to NSCLC gene expression data: hybrid feature selection with ensemble learning and regularized regression. Their potential for identifying informative gene subsets, improving classification, and supporting biomarker discovery is discussed, highlighting complementary contributions to robust, data-driven diagnostic tools.

End-to-End AI-based Drug Formulation Optimization

Saba Seifi[1]

[1] *Friedrich Schiller University Jena*

Abstract: AI technology and machine learning present a transformative opportunity in the drug discovery, formulation, and testing of pharmaceutical dosage forms. By utilizing AI algorithms that analyze extensive biological data, including genomics and proteomics, researchers can identify disease associated targets and predict their interactions with potential drug candidates. This enables a more efficient and targeted approach to drug discovery, thereby increasing the likelihood of successful drug. The adoption of new technologies is helpful to address global healthcare challenges and medical emergencies, such as the recent COVID pandemic with minimal face-to-face contact and reducing the need for extensive and costly animal testing. AI algorithms can analyse and optimize drug candidates by considering various factors, including efficacy and safety.

Automatic Differentiation in AI; Why Julia & Enzyme?

Salim Alkhaddoor[1]

[1] *Friedrich Schiller University Jena*

Abstract: Automatic differentiation (AD) is a foundational technique for computing exact derivatives of numerical programs with machine precision and efficiency, especially in artificial Intelligence applications. Unlike numerical or symbolic differentiation, AD applies the chain rule directly to the executed program, enabling scalable gradient computation for machine learning, scientific computing, and high-performance simulations. This poster examines the principles of forward- and reverse-mode AD, their computational trade-offs, and hybrid approaches for higher-order derivatives. We focus on Julia's AD ecosystem, including ForwardDiff.jl, ReverseDiff.jl, Zygote.jl, and the LLVM-based Enzyme.jl, which enables language-agnostic, compiler-level differentiation. Enzyme's integration as an LLVM pass allows aggressive post-differentiation optimizations, efficient memory management, and application to legacy HPC codes. Benchmark comparisons demonstrate competitive performance in both time and memory usage against existing Julia-based AD tools. The presented workflow highlights practical considerations such as checkpointing, custom gradient rules, and mixed-precision computation, underscoring the potential of compiler-driven AD to unify productivity, performance, and portability in modern computational workflows.

Semi-Automated 2D Material Flakes Detection

Ting-Chun Ou[1]

[1] *Friedrich Schiller University Jena*

Abstract: Identifying monolayer flakes of 2D materials remains a significant challenge due to their slightly optical differences from bilayer and multilayer regions. Manual identification is time-consuming and inconsistent. In this study, we propose an automatic flake detection method based on image processing techniques that can reliably locate monolayer flakes under controlled imaging conditions. By fixing parameters such as illumination intensity and light source angle, we created a model that characterizes the optical features of monolayer regions. We then used this model to annotate monolayer flakes on the original sample and extended it to new samples of the same material. This approach enables the rapid and consistent identification of monolayers, thereby enhancing the efficiency and objectivity of preparing and studying 2D materials. Furthermore, this framework gives an idea for optical machine learning, which could lead to more scalable flake detection across various imaging conditions and material systems.

Bildung, Erziehung und KI: Algorithmus trifft Aufklärung

Tobias Kresse[1]

[1] *Friedrich Schiller University Jena*

Abstract: Das Poster „Bildung, Erziehung und KI: Algorithmus trifft Aufklärung“ postuliert die Pädagogik als einen entscheidenden Analyserahmen für die Gestaltung von KI-Systemen. Hierfür werden die Grundbegriffe „Erziehung“ (gesteuertes Training) und „Bildung“ (kritische Selbstorganisation) nach Koerrenz auf die Lernmechanismen von KI übertragen. Die Analyse zeigt, dass KI-Systeme primär Produkte einer „Erziehung“ sind, deren inhärente Probleme (wie Bias oder mangelnde Transparenz) als Erziehungsprobleme interpretiert werden können. KI kann das Ziel Bildung und Erziehung zu simulieren attestiert werden. Auf dieser Grundlage werden Prinzipien für ein bildungsförderliches Alignment entwickelt, die auf Transparenz, den Umgang mit Unsicherheit und Dialog abzielen. Das Poster schließt mit der These, dass menschliche Bildung der normative Horizont für die KI-Entwicklung sein muss. Die pädagogische Theorie liefert somit den Ordnungsrahmen, um die Fähigkeiten und Grenzen von KI präzise zu erklären.

Overview of Reinforcement learning

Uday Kaipa[1]

[1] *TU Ilmenau*

Abstract: This poster presents an introductory overview of Reinforcement Learning (RL), a branch of machine learning in which agents learn optimal decision-making policies through interaction with their environments to maximize cumulative rewards. It contrasts RL with supervised learning, highlighting key differences in feedback, data requirements, and learning dynamics. Core RL concepts—states, actions, rewards, and policies—are explained along with the exploration-exploitation trade-off. The taxonomy of RL algorithms is outlined, covering value-based, policy-based, and actor–critic methods, as well as distinctions between model-free and model-based approaches, discrete and continuous action spaces, and episodic versus continuous tasks. Real-world applications are showcased in domains such as game playing, robotics, autonomous systems, recommendation systems, and language model fine-tuning (RLHF). The poster summarizes landmark successes like AlphaGo, DQN, and robotic control, discusses key challenges including reward design and sample efficiency, and points to future directions such as multi-agent cooperation, hierarchical learning, and safer, more efficient training methods.

Bridging Physics and Learning: A Hybrid Deep Model for Structural Deformation Prediction

Yahuan Shi[1]

[1] *TU Berlin*

Abstract: Accurate and efficient simulation of structural deformation is essential in computational mechanics. Traditional methods such as the Finite Element Method (FEM) offer high fidelity but suffer from high computational cost and meshing complexity. In this work, we present a physics-informed neural framework based on the Deep Energy Method (DEM) for solving 2D elasticity problems in a mesh-free and label-free manner. The model learns the displacement field by minimizing a physically motivated energy functional composed of internal strain energy, boundary traction energy, divergence residual, and soft boundary constraints. We further investigate the influence of sampling strategies, numerical integration schemes, and propose a lightweight Kolmogorov-Arnold Network (KAN) as an efficient alternative to traditional dense networks. Experimental results demonstrate that our approach achieves competitive accuracy with significantly reduced inference time and improved training stability. This hybrid modeling technique offers a scalable and interpretable alternative for real-time structural simulation and digital twin applications.

Federated Learning on mobile devices

Yamdjeu Sitcheu Mariza[1]

[1] *Friedrich Schiller University Jena*

Abstract: Federated Learning (FL) trains models collaboratively across mobile devices without sharing raw data, preserving privacy and meeting regulatory requirements. This poster outlines typical FL architecture, key challenges such as limited resources, heterogeneity in devices or data and strategies like model compression, secure aggregation. Real-word use cases including predictive text, voice assistants demonstrate FL's potential for scalable, privacy-preserving AI at the edge.

Retrieval-Augmented Generation (RAG)

Zhe Wang[1]

[1] *LMU Munich*

Abstract: This poster explores Retrieval-Augmented Generation (RAG), a technique developed to enhance knowledge-intensive Natural Language Processing (NLP) tasks by integrating external, up-to-date information into the generation process of large language models. RAG addresses critical limitations of pre-trained models, such as static knowledge, a lack of provenance for generated answers, and the tendency to produce factually incorrect "hallucinations". By grounding responses in retrieved documents, RAG improves factual accuracy and provides interpretability.

The poster details key stages of the RAG pipeline, focusing on advancements in query translation and indexing methods. To counter ambiguous queries that lead to poor document retrieval, techniques like RAG-Fusion generate multiple query variations and use reciprocal rank fusion to identify the most relevant documents. For indexing, the Recursive Abstractive Processing for Tree-Organized Retrieval (RAPTOR) method is presented. This approach constructs a hierarchical tree of document summaries, enabling efficient and contextually relevant retrieval by traversing or collapsing the tree based on the query's specificity. Additionally, the query decomposition method is highlighted as a strategy to break down complex problems into sequential, solvable sub-problems.