AI Hub Karlsruhe

5-7 October 2022

Networking Event

Overall Program

	05.10.22	06.10.22	07.10.22
09:15	Opening	KIT-An Overview	KCDS Program
09:30	Keynote Lecture	Keynote Lecture	Keynote Lecture
10:30	Networking breaks		
11:00	Session I	Session IV	Session VI
11:30			
12:00			
12:30	Lunch Breaks		
13:30	Session II	AI @ KIT Poster Session & Networking	Session VII
14:00			
14:30			
15:00	Networking	Session V	Networking
15:30	Session III		Session VIII
16:00			
16:30			

Detailed Program

05 October Wednesday

09:20	AI Hub Karlsruhe – Highlights of the Event
09:30	Keynote: Prof. Michael Feindt, KIT *AI/ML for the Supply Chain*
10:30	Networking Break
11:00	Dr. Sebastian Lerch (KIT) *Artificial Intelligence for Probabilistic Weather Forecasting*
11:30	Renumics: It's the data, stupid! How data-centric AI helps to build robust ML systems for industrial applications - Dr. Stefan Suwelack
12:00	KI Garage: "AI expertise meets entrepreneurial spirit" - Patrick Widmann
12:30	Lunch break
13:30	Dr. Andrea Santamaria Garcia (KIT) *Machine Learning for Particle Accelerators*
14:00	Dr. Vahid Babaei (Max Planck Institute for Informatics) *Data-Driven Functional Fabrication*
14:30	Dr. Noémie Jaquier (KIT) *Learning, adapting, and sequencing skills for robot manipulation*
15:00	Networking & Coffee
15:30	Dr. Rudolf Lioutikov (KIT) *Towards Intelligent and Intuitive Robots*
16:00	Dr. Wieland Brendel (University of Tübingen) *A more principled way towards machines that see the world like humans*
16:30	Maddox: Overcoming four Challenges to solve Visual Quality Inspection with Machine Learning - Dr. Wieland Brendel

06 October, Thursday

09:15	KIT – An Overview - Prof. Oliver Kraft
09:30	Keynote: Prof. A. Aldo Faisal *Towards high performance medical AIs: from drug development to medical interventions*
10:30	Networking Break
11:00	Dr. Georg Steinbuß (Universitätsklinikum Heidelberg) *Histopathology via Deep Learning*
11:30	Neurocat AI: Evaluating Worst-Case Scenarios for Computer Vision Models with Risk Scores - Dr. Holger Trittenbach
12:00	Dr. Ece Özkan Elsen (ETH Zürich) *Machine Learning in Healthcare - A Survival Guide*
12:30	Lunch break
13:30	"AI @ KIT" Poster Session and Networking
15:00	Dr. Charlotte Debus (KIT) *Towards robust and efficient AI at scale*
15:30	Prof. Fabian Sinz (University Göttingen) *Exploring the brain with functional twins*
16:00	How apic.ai uses AI to conserve biodiversity through visual pollinator monitoring – Frederic Tausch
16:30	Dr. Gabriele Schweikert (University of Dundee) *Machine Learning for Epigenomic Data Analysis and individual-specific Imputation*

07 October, Friday

09:15	Prof. Martin Frank, Scientific Speaker of KIT Center MathSEE and KCDS *Meet Graduate School Computational and Data Science | KCDS*
09:30	Keynote: Prof. Debarghya Ghoshdastidar (TUM) *The Myths about Overfitting*
10:30	Networking Break
11:00	TT-Prof. Pascal Friederich (KIT)

	Machine Learning for Accelerated Materials Discovery
11:30	Meet Aimino - Dr.-Ing. Duc Tam Nguyen
12:00	Meet Cyber Valley - Alexander Diehl, Cyber Valley Senior Advisor *Science Entrepreneurship – a European Opportunity*
12:30	Lunch break
13:30	Dr. Nicole Ludwig (University of Tübingen) *Energy Informatics*
14:00	Dr. Mehwish Alam (KIT) *Machine Learning & Knowledge Graphs*
14:30	Dr. Ruben Bach (Uni Mannheim) *When Small Decisions Have Big Impact: The Hidden Consequences of Algorithmic Profiling in Public Service*
15:00	Networking Break
15:30	Dr. James Kahn (Helmholtz AI -KIT) *AI Consulting in Energy Research*
16:00	Validaitor: Towards Quality Assurance in Machine Learning - Yunus Emrah Bulut
16:30	

AI @ KIT Poster Session

1	Jakob Bach - *Leveraging Constraints for User-Centric Selection of Predictive Features*
2	Steffen Schotthöfer - *Low-rank lottery tickets: finding efficient low-rank neural networks via matrix differential equations*
3	Mustafa Demetgül - *Sensorless and Intelligent Machine Monitoring using Motor Current*
4	Tim Ortkamp - *Inverse Radiotherapy Treatment Planning using Machine Learning Outcome Prediction Models*
5	Jan Baumgärtner - *Generative Robot Cell Design*
6	Danni Liu - *Learning Common Representations for Multilingual Neural Machine Translation*

7	Pawel Bielski - *Refining Domain Knowledge for Domain Knowledge Guided Machine Learning*
8	Felix Laufer - *Machine Learning for Perovskite Thin-Film Photovoltaics*
9	Chen Zhou - *Machine Learning for Molecular Dynamic Simulation and Chemical Reaction Prediction*
10	André Orth, Jan Schützke - *Deep Learning for the Analysis of Spectroscopic Data*
11	Matthias Schaufelberger - *CNN-Based Classification of Craniosynostosis Using 2D Distance Maps*
12	Paras Koundal - *Using Graph Neural Networks for Cosmic-Ray Analysis, at IceCube Observatory*
13	Armin Weckmann - *Optimisation in Logistics - Traveling salesman, facility placement and minimum cost flow problems in practice*
14	Dr. Prantik Samanta - *AI in Wastewater Treatment*
15	Pranav Sampathkumar - *Sequential Networks for Cosmic Ray simulations*
16	Paul Mifsud - *Graph Neural Networks applied to Fluid Dynamics*
17	Andrej Rode, Vincent Lauinger - *Optimization of Communication Systems Using Generative Networks and Auto-encoders*
18	Lu Guo - *GANs for generation of synthetic bronchoscopic images*
19	Dennis Gnad - *Physical Side-Channel Attacks and Defenses on Hardware Neural Network Accelerators*
20	Yichen Jia - *Constraints on cloud fraction adjustment to aerosols using explainable machine learning*
21	Joel Arweiler - *Automated Active Learning*

AI Hub @ Karlsruhe 05-07 October 2022

Day 1 - 05.10.22

AI/ML for the Supply Chain

Prof. Michael Feindt

Blue Yonder

I will describe recent developments in data driven AI/ML planning, prediction, decision and automation algorithms along the supply chain from manufacturing via distribution and warehousing to retail. I will show its value in normal times, but especially also in times of disruptions as the Covid 19 pandemic.

In parallel I will tell the Blue Yonder story, from a spin-off of KIT, based on algorithms developed for particle physics at CERN, to a leading global company with 6000 associates recently acquired by Panasonic at a valuation of 8.5 bln US$.

Artificial Intelligence for Probabilistic Weather Forecasting

Dr. Sebastian Lerch

Modern AI methods, in particular deep learning methods based on multi-layered artificial neural networks, provide unprecedented tools for data analysis and prediction. Over the past years, they have transformed many scientific fields, including the environmental sciences, and have been used for weather forecasting in a multitude of ways. In my talk, I will give an overview of applications of AI methods for weather and climate modelling, with a focus on probabilistic weather prediction based on post-processing forecasts from physical weather prediction models. In particular, I will highlight the importance of and opportunities for interdisciplinary collaborations to improve and better understand predictive models.

It's the data, stupid! How data-centric AI helps to build robust ML systems for industrial applications."

Dr. Stefan Suwelack

Data-driven tools will fundamentally change product development and manufacturing; They enable faster development cycles, minimize manual work, and raise quality levels. However, building data-driven processes and tools is hard; it takes a strong collaborative effort, the right infrastructure, and good data science skills to succeed. Based on our extensive experience, we believe that data-driven tools should be built in user-focused, data-centric, and collaborative way. We empower cross-functional teams to build robust ML-enabled solutions based on these paradigms.

We are building the data curation software Renumics Spotlight. The tool allows cross-functional teams to quickly create and iterate training data sets for ML algorithms. We envision Spotlight to be the data curation component in a fully modular ML stack. Together with our customers and partners, we are continuously pushing the envelope how data-driven technology can generate business value. We share these insights with the community through templates such as the "AI-assisted Engineering Canvas" as well as blog and research articles.

KI Garage: AI Expertise meets entrepreneurial spirit

Patrick Widmann

Artificial intelligence and machine learning are changing the economy and society. We have long since become accustomed to cameras recognising our faces and digital assistance systems listening to our words - and yet we are only at the beginning of a revolution. We can delegate more and more decisions to computers, which make them better and, above all, faster. This opens up enormous opportunities for medicine, research, production processes, mobility and many other areas.

Machine Learning for Particle Accelerators

Dr. Andrea Santamaria Garcia

Accelerators are one of the most complex machines in the world, where their manual tuning is a laborious process even for experienced operators. Humans can only act on a few parameters simultaneously and operate on a narrow set of timescales. Additionally, the relationship between parameters change depending on external conditions. Due to the high number of parameters required to operate an accelerator and the nonlinear correlation between them, as well as the need for different operation modes and the presence of many non-controllable variables, accelerators are an excellent candidate to benefit from machine learning methods. Moreover, the increased complexity of future accelerators with higher energy, higher brightness, and higher gradients will require advancing the capabilities of accelerator facilities with technologies like machine learning.

Data-Driven Functional Fabrication

Dr. Vahid Babaei

While digital manufacturing (e.g., 3D printing) has had significant advances on the hardware and material fronts, algorithmic design for digital manufacturing remains one of the biggest barriers to its widespread adoption. Computational Design and Fabrication (CDF) is an emerging research area that tries to address this problem. Functional fabrication is an important paradigm in CDF where given the high-level function or goal, the design is computed, and physically realized by the fabrication hardware. Although the functional fabrication lends itself to powerful abstractions, we still lack a general and practical workflow. I will show how machine learning methods can raise these abstractions closer to a practical level.

Learning, adapting, and sequencing skills for robot manipulation

Dr. Noémie Jaquier

To perform a wide variety of manipulation tasks, humans learn versatile manipulation skills that can efficiently be adapted to novel settings. Moreover, humans are able to smoothly sequence and combine these skills to realize complex motions. In contrast, adapting skills in a fast and data-efficient manner and combining them to generate seamless skill sequences remain key challenges in robotics. In this talk, I will first discuss the use of Bayesian optimization to generalize previously-learned robotic skills to unforeseen settings. In particular, I will discuss how Bayesian optimization can benefit of inductive bias, which can be introduced via task-specific information about the geometry of the search space. Second, I will show how the obtained skills may be combined to generate complex motions. Namely, I will discuss how sequences of skills can be encoded as quadratic programs, in which the relative importance of each skill throughout the task is learned from demonstrations. Finally, I will illustrate the use of the presented learning methods in real robotic applications.

Towards Intelligent and Intuitive Robots

Dr. Rudolf Lioutikov

Artificial intelligence approaches have produced impressive results across a wide spectrum of fields and applications in recent years. These successes in combination with an increasing demand for assisted living, elderly care and local production have caused the expectation of an imminent deployment of intelligent autonomous robots in our everyday life. These future agents will be expected to work in close interaction with non-expert users in both general everyday situations and professional tasks. A new generation of intelligent robots will be required that is capable of communicating intent to non-expert users as well as understanding intent from the action of the user in a natural way. These robots will appear more intuitive to non-expert users as well as be able to deduct valuable information through more intuitive interaction with the non-expert user. This talk will present my past, current, and future research on such Intuitive Robots.

A more principled way towards machines that see the world like humans

Dr. Wieland Brendel

Machines perceive the world very differently from humans. This gap between humans and machines has narrowed little in the past years, despite even the recent emergence of impressive large-scale foundation models. In this talk, I discuss our recent systematic work to unravel the hidden assumptions underlying existing representation learning techniques and to develop new techniques that leverage the compositional 3D nature of our visual world in a more principled way.

Maddox: Overcoming four Challenges to solve Visual Quality Inspection with Machine Learning

Dr. Wieland Brendel

In many factories, ensuring the quality of manufactured products is still a laborious manual task performed by legions of humans. That's surprising given that modern machine learning algorithms should be well adapted to detect the scratches, deformations or missing parts that mark defunct or substandard products. In this talk, I will discuss four barriers that prevent the application of machine learning in visual inspection and how we overcome them at Maddox.

Day 2 - 06.10.22

Short Talk 2: Karlsruhe Institute of Technology – an Overview

Prof. Oliver Kraft

Vice-President Research

Towards high performance medical AIs: from drug development to medical interventions

Prof. A. Aldo Faisal

Imperial College London

AI has tremendous potential for addressing the unmet need for healthcare globally, yet many applications are mainly aimed at reproducing existing human-performed methodologies, instead of enabling entirely new strategies that can only be enabled with advanced machine learning methodologies. We will discuss here based on our recent work: We will show how Deep Probabilistic Models (Kadirvelu et Faisal, Nature Medicine, in Press; Ricotti et Faisal, Nature Medicine, in Press) on rapidly accelerating drug development and clinical trial using real-world human wearable data can systematically outperform traditional clinical trials in neurodegenerative diseases. We then show how we developed off-policy distributional reinforcement learning methods that allow us to learn from routine hospital operations data how to treat patients better, we illustrate this with the AI Clinician system (Komorowski et al, Nature Med, 2018; Gottessman eta al. Nature Med, 2019) which is being trialled in 4 British hospitals to treat intensive care patients with sepsis.

The capability of these system highlights how AI for healthcare needs to develop human-AI interaction technologies, as management of explainability and trust is becoming key for regulatory approval and general adoption of such systems, and we will complete this keynote with some considerations and novel insights that we gained here.

Histopathology via Deep Learning

Dr. Georg Steinbuß

Histopathology describes the inspection of carefully prepared tissue slices for traces of a disease. Framed in microscopic slides, the tissue samples are usually evaluated by a pathologist via microscope. Since such manual evaluation is tedious and there is a lack of experienced pathologists, deep learning has been suggested to complement and improve histopathology. Our goal is to render deep learning in histopathology from a purely academic activity to a tool helpful in the diagnosis of tumor patients. To accomplish this, we train convolutional neural networks on a broad spectrum of different tissue and disease types but also make these networks available to other pathologists.

Evaluating Worst-Case Scenarios for Computer Vision Models with Risk Scores

Holger Trittenbach

Neurocat

Computer vision models can fail in various and unexpected ways when images differ only slightly from the training distribution. For example, variations through weather or lighting conditions, and slight, imperceptible changes to the image can cause the prediction of detected objects to change or disappear. Such failures are a significant threat to operating machine learning models in safety and security-critical systems such as autonomous driving. Although the threat is well-understood, it is an open question how to systematically assess and quantify the associated risks.

In this talk, we frame the problem of risk estimation for computer vision models. We first propose a method to quantify the probability of observing the worst-case, i.e., a variation to the image that causes the model to fail. We then show how to use our method to decide on the deployment of models in safety and security-critical contexts. Finally, we show how to scale risk estimation to production-scale ML-Ops environments.

Machine Learning in Healthcare - A Survival Guide

Dr. Ece Özkan Elsen

ETH zürich

In recent years, modern machine learning (ML) algorithms have broken records achieving impressive performance. They have been used in various fields, such as gaming, protein structure prediction, creating new images, translation, and autonomous driving. Because computer-based systems have become an integral part of modern hospitals, numerous machine learning-based methods have also been developed for healthcare. In this talk, we will get to know the types of methods commonly used in healthcare and the typical tasks they can solve with concrete examples. We will further discuss the limitations, challenges, and opportunities of ML for healthcare and showcase the current research examples.

Towards robust and efficient AI at scale

Dr. Charlotte Debus

KIT Karlsruher Institut für Technologie

Next to everyday life applications like smart phones, autonomous driving or voice assistants, AI methods have also revolutionized data analysis, system monitoring and control optimization in scientific research and engineering. In these domains, data is often acquired over time, adding an extra dimension and as such more complexity to the problem. The introduction of transformer architectures has opened up new ways to accurately predict behaviours in such dynamic systems. However, the sheer size of transformer models poses entirely new challenges, pushing AI research towards large scale models that run on multiple accelerators and supercomputers. The increased demand in compute resources comes at the price of growing energy consumption, which now raises the question regarding sustainability and environmental friendliness of AI applications. In the talk, we will look at a few examples of scientific applications of AI-based time series forecasting as well as their scalability and energy efficiency, and how we can balance these apparently competing trends of Scalable AI and Green AI.

Exploring the brain with functional twins

Prof. Fabian Sinz

Deep neural networks have set new standards in modelling the responses of large-scale populations of neurons to natural stimuli, yielding models that can accurately predict the response of thousands of neurons to novel stimuli. This allows us to treat the model as a functional digital twin of the neural circuit and probe neurons in ways that would not be feasible experimentally. With that, we can derive new hypotheses about the neural circuits that can then be verified in subsequent experiments. In this talk, I will give an overview over the models and their application in understanding the computational properties of visual cortex.

"How apic.ai uses AI to conserve biodiversity through visual pollinator monitoring"

Frederic Tausch

apic.ai was founded in Karlsruhe, Germany in 2018. Our company combines expertise in software and hardware development in the field of image processing. We use it in a science-based way for the purpose of behavioural analysis of insects. We are united by the conviction that technology can provide the basis for better decisions by creating transparency over complex interdependencies. We have decided to use our skills to meet the social challenge of the loss of biodiversity.

There is no doubt that insect extinction is primarily due to human influence. We believe, however, that in most cases this damage is not caused intentionally, due to greed or malice. Instead, we are convinced that it stems from a lack of information about the impact of pesticide use, intensive agriculture, expansive urban development and other factors. Through our contribution, we want to ensure that ignorance is no longer a reason for the loss of biodiversity. Like humans and nature, we believe that prosperity and sustainability can go hand in hand.

Machine Learning for Epigenomic Data Analysis and individual-specific Imputation

Dr. Gabriele Schweikert

Epigenomic Modifications are reversible chemical marks on top of the DNA, that do not change the underlying sequence itself. Personal, cell-type-specific epigenomes result from a combination of genetic variants and a cellular memory of past cellular events. Epigenetic mechanisms are therefore essential mediators of gene–environment interactions. Functionally, they contribute to the control of current and future transcription and thus play important roles during development, disease progression and ageing.

Recently, efforts to record personal epigenomes across tissues have become feasible. However, the large number of assays required for a complete epigenomic map continues to be a limiting factor for personalized epigenomics.

Machine Learning approaches are poised to fill this gap. In this talk I will present eDICE, which is based on the transformer architecture and is capable of predicting individual-specific epigenomic landscapes. We achieve high prediction accuracy by learning factorised representations. At the same time, eDICE has unprecedented generalisation capabilities. The complete model fits into GPU memory and does not require complicated training schemes as the number of parameters is several orders of magnitude smaller than in previous models. These are essential preconditions to apply computational imputation for personalised epigenomics and to use these methods at scale.

Day 3 - 07.10.22

Meet KIT Graduate School Computational and Data Science | KCDS

Prof. Martin Frank, Scientific Speaker of KIT Center MathSEE and KCDS

KIT Graduate School Computational and Data Science (KCDS) is a new graduate school at KIT Center MathSEE starting in winter semester 2022/23 that offers an interdisciplinary training program for doctoral researchers in the field of model-driven and data-driven computational science. We are open for doctoral researchers and PIs who are interested in interdisciplinary research projects that revolve around computational methods such as mathematical models, simulation methods and data science techniques, all the while building bridges between mathematical sciences and an applied SEE discipline (science, economics and engineering). If you are ready to conquer the data-driven challenges of tomorrow, we encourage you to join us!

More info on KCDS: https://www.kcds.kit.edu/

The Myths about Overfitting

Prof. Debarghya Ghoshdastidar

Overfitting is the practice of using a complex machine learning model that perfectly fits the training data. Historically, overfitting has been considered a “bad practice” that is expected to produce predictors which perform poorly on new data. However, the recommendation has flipped in recent times, where overfitted neural networks perform surprisingly well in computer vision, natural language processing among others. For instance, in the ImageNet image classification benchmark (with 14 million images), the best architectures have more than a billion parameters and yet achieve 90% accuracy. This naturally raises the question whether overfitting is a good practice or a bad practice.

In this talk, I will discuss the mathematical foundations behind the classical and modern views about overfitting. I will start with a brief introduction to the statistical theory that leads to the conclusion “overfitting is a bad practice”. I will then discuss some recent theoretical results that debunk the myths:

1. large models with too many parameters always overfit on the training data;

2. models that perfectly fit the training data cannot predict well on unseen data.

The above results are the basis of two promising research directions in machine learning theory: Neural Tangent Kernels -- that capture the training dynamics of wide neural networks -- and Double-Descent phenomenon -- a precise characterisation of the performance of overfitted models. We will finally see why the classical and the modern theories of overfitting are not at odds with each other.

Machine Learning for Accelerated Materials Discovery

T.T.-Prof Dr. Pascal Friederich

Machine learning can enable and accelerate the design of new molecules and materials in multiple ways, e.g. by learning from large amounts of (simulated or experimental) data to predict molecular or materials properties faster, or even by interfacing machine learning algorithms for autonomous decision making directly with automated high-throughput experiments. In this talk I will give a brief overview of our research activities on graph neural networks for materials property prediction [1,2], machine learning accelerated atomistic simulations [3,4], as well as using machine learning methods for decision making processes in automated materials science and chemistry labs [5].

[1] Reiser et al., Software Impacts 2021,
https://www.sciencedirect.com/science/article/pii/S266596382100035X

[2] Reiser et al., arXiv:2208.09481

[3] Friederich et al., Nature Materials 2021, https://www.nature.com/articles/s41563-020-0777-6

[4] Li et al., Chemical Science 2021, https://pubs.rsc.org/en/content/articlehtml/2021/sc/d0sc05610c

[5] Luo et al., Angewandte Chemie 2022,
https://onlinelibrary.wiley.com/doi/full/10.1002/anie.202200242

Meet Aimino

Dr.-Ing. Duc Tam Nguyen

Aimono helps users create smart applications that are reliable, secure and deployment ready. Aimono empower real products with machine learning for different industries including aerial imagery, agriculture, healthcare and smart building with various applications such as workflow automation, predictive maintenance, generating synthetic human data etc.

Meet Cyber Valley: Science Entrepreneurship – a European Opportunity

Alexander Diehl, Cyber Valley Senior Advisor

Cyber Valley is Europe's largest research consortium in the field of artificial intelligence. Academic and private sector partners are building bridges between curiosity-driven basic research and applied research. The state of Baden-Württemberg, the Max Planck Society with the Max Planck Institute for Intelligent Systems, the Universities of Stuttgart and Tübingen, as well as Amazon, BMW AG, IAV GmbH, Mercedes-Benz Group AG, Dr. Ing. h.c. F. Porsche AG, Robert Bosch GmbH, and ZF Friedrichshafen AG are Cyber Valley's founding partners of this initiative. In 2019, Fraunhofer-Gesellschaft also joined Cyber Valley as a partner. Moreover, Cyber Valley receives support from the Christian Bürkert Foundation, the Gips-Schüle Foundation, the Vector Foundation, and the Carl Zeiss Foundation.

Forecasting Renewable Energy using Machine Learning

Dr. Nicole Ludwig

Most mitigation strategies to face climate change involve the electrification of our grid and a shift towards clean energy resources, thus no fossil fuels. However, without fossil fuels, we have to rely on more volatile renewable energy sources, and the energy system is facing unprecedented challenges which call for new methods and perspectives. These challenges include, among many others, more uncertainty in energy

generation through complex interactions with the weather and climate and active participation of consumers through, for example, PV panels on household rooftops. This talk will introduce challenges when forecasting weather-driven renewable energy time series and introduce solutions focusing on (primarily probabilistic) machine learning.

Machine Learning & Knowledge Graphs

Dr. Mehwish Alam

Knowledge Graphs (KGs) constitute a large network of real-world entities and relationships between these entities. KGs have recently gained attention in many tasks such as recommender systems, question answering, etc. Due to automated generation and open-world assumption, these KGs are never complete. Recent years have witnessed many studies on link prediction using KG embeddings which is one of the mainstream tasks in KG completion. To do so, most of the existing methods learn the representations of the entities and relations whereas only a few of them consider contextual information as well as the textual descriptions of the entities. This talk will give an overview of the methods and benchmark datasets proposed for KG Completion by taking into account multimodality. Applications of such methods to the real-world problems associated with scholarly data will also be discussed.

When Small Decisions Have Big Impact: The Hidden Consequences of Algorithmic Profiling in Public Service

Dr. Ruben Bach

Algorithmic profiling is increasingly used in the public sector to support the allocation of limited public resources. For example, in criminal justice systems algorithms inform the allocation of intervention and supervision resources, child protection services use algorithms to target risky cases and to allocate resources such as home inspections to identify and control health hazards, immigration and border control use algorithms to filter and sort applicants seeking residence in the country, and Public Employment Services use algorithms to identify job-seekers who may find it difficult to resume work and to allocate support programs to them. However, concerns are raised that profiling tools may suggest unfair decisions

and thereby cause (unintended) discrimination. To date, empirical evaluations of such potential side-effects are rare. Using algorithm-driven profiling of jobseekers as an empirical example, we illustrate how different modelling decisions in a typical data science pipeline may have very different fairness implications. We highlight how fairness audits, statistical techniques as well as social science methodology can help to identify and mitigate biases and argue that a joint effort is needed to promote fairness in algorithmic profiling.

AI Consulting in Energy Research

Dr. James Kahn

HELMHOLTZ AI
Energy

The Helmholtz AI research platform was launched in 2019 and aims to empower scientists to use artificial intelligence methods in scientific domain problems. One of the integral parts in achieving this goal is AI consulting. Using so-called free-of-charge consulting vouchers, collaboration requests for small to medium-sized projects, scientists may receive help in leveraging the capabilities of data-driven modelling as well as effectively improving the quality and scalability of existing AI pipelines. In this talk, the Helmholtz Local Unit Energy @ KIT will showcase a set of completed vouchers enabling the automated detection of thermal leakages in building rooftop insulation from drone footage. The implemented approach improved the IoU@50 recall by a factor of three compared to the state-of-the-art. In line with the Green AI movement the Local Unit Energy also quantified the electricity consumed during modelling utilising the Helmholtz' HAICORE AI supercomputer.

Validaitor: Towards Quality Assurance in Machine Learning

Yunus Emrah Bulut

VALIDAITOR

Machine Learning (ML) is disruptive and transformative but with its weaknesses and vulnerabilities. In order to get most out of the ML, a wholistic quality assurance should be put in place. That's the required condition of an Artificial Intelligence (AI) system to be trustworthy and robust. Upcoming regulations on AI across the globe try to foster the adoption of quality and risk management practices in AI system development.

This talk is about how to assure quality in ML and how Validaitor (a spin-off from KIT) enables a wholistic quality and risk management in ML development. We discuss how we should test ML models before putting them in place and how to monitor them when they're in use. We also talk about how Validaitor approaches to the ML quality management by bringing testing, monitoring and interpretability into a single platform.

Poster Presentations

Leveraging Constraints for User-Centric Selection of Predictive Features

Jakob Bach

Feature selection identifies the most valuable predictors in a dataset. Thus, feature-selection techniques are popular for obtaining small, interpretable, yet highly accurate prediction models. While optimizing technical quality metrics, standard feature-selection techniques might not satisfy user needs for two reasons. First, existing methods do not consider domain knowledge. Such domain knowledge can restrict which feature combinations make sense to users. Second, traditional feature-selection techniques typically yield only one feature set, which might not suffice in some scenarios. For example, users might be interested in finding different feature sets with similar prediction quality, offering alternative explanations of the data. Constraints on feature sets alleviate both these shortcomings. First, constraints allow users to express domain knowledge, e.g., known physical laws, novel scientific hypotheses, etc. Second, constraints can formalize the notion of alternative feature sets. Our research studies different types of constraints that make feature selection more user-centric. We investigate how to formulate and integrate such constraints into existing feature-selection techniques. Further, we study the impact of constraints on feature-selection results, e.g., if prediction quality remains stable under constraints. Our experiments show that it often is possible to find high-quality feature sets adhering to user constraints.

Low-rank lottery tickets: finding efficient low-rank neural networks via matrix differential equations

Steffen Schotthöfer

Neural networks have achieved tremendous success in a large variety of applications. However, their memory footprint and computational demand can render them impractical in application settings with limited hardware or energy resources. In this work, we propose a novel algorithm to find efficient low-rank subnetworks. Remarkably, these subnetworks are determined and adapted already during the training phase

and the overall time and memory resources required by both training and evaluating them is significantly reduced. The main idea is to restrict the weight matrices to a low-rank manifold and to update the low-rank factors rather than the full matrix during training. To derive training updates that are restricted to the prescribed manifold, we employ techniques from dynamic model order reduction for matrix differential equations. Moreover, our method automatically and dynamically adapts the ranks during training to achieve a desired approximation accuracy. The efficiency of the proposed method is demonstrated through a variety of numerical experiments on fully-connected and convolutional networks.

Sensorless and Intelligent Machine Monitoring using Motor Current

Mustafa Demetgül

With the rise of Industry 4.0, the concept of intelligent manufacturing has been further developed, and human-machine interaction has become significant. In order to further improve the efficiency and flexibility of industrial manufacturing, artificial intelligence-assisted manufacturing, and intelligent manufacturing have become a topic of long-term research. Therefore, to extend the life cycle of the entire industrial system and reduce the probability of failure, the quality of the workpiece must be inspected. Therefore, to detect the presence of defects early and solve the problem before a failure occurs, a more comprehensive maintenance strategy must be specified, especially preventive maintenance will be mentioned more often. Diagnosis is made with many sensors. This is shown in the figure below. In addition, this study has many contributions to machine monitoring. These are shown in the figure below. In this study, a new approach for monitoring is implemented. With the help of the PLC, the motor current data obtained from the motor can detect the small axis misalignments that occur in the machines.

Inverse Radiotherapy Treatment Planning using Machine Learning Outcome Prediction Models

Tim Ortkamp

Half of all cancer patients receive radiotherapy as part of their treatment schedule. Being a mostly non-invasive medical intervention, radiotherapy delivers high-energetic, ionizing radiation to target cancerous tissue while sparing healthy tissue. To control the underlying trade-off between tumor control probability (TCP) and normal tissue complication probability (NTCP), radiotherapy dosage needs to be simulated pre-treatment on the patient's imaging data, i.e., computed tomography images, and respectively optimized. Existing attempts to directly optimize the treatment plan for TCP and NTCP rely on simple low-parametric, often univariate models, like, for example, the Lyman-Kutcher-Burman (LKB) model, and have been consistently questioned with regards to accuracy and usability. Therefore, this research project aims at

superseding these approaches by incorporating state-of-the-art machine learning models for TCP and NTCP into the radiotherapy treatment plan optimization process.

Generative Robot Cell Design

Jan Baumgärtner

The commissioning of robot cells is divided between mechanical engineers who design the structure of the cell and robotics engineers who optimize the robot's performance in terms of energy, cycle time, or precision. This workflow neglects that the structure of the cell often has much more influence on these properties than the robots programming. My work investigates tools and algorithms that automatically design robot cells integrating structural design and programming into an end-to-end solution. This is done using techniques borrowed from continuous dynamic systems optimizations and discrete search problems.

Learning Common Representations for Multilingual Neural Machine Translation

Danni Liu

Neural machine translation (NMT) has become the backbone of many state-of-the-art translation tools nowadays. Given the over 7000 languages in the world, multilingual NMT models have become an attractive research direction for the: 1) ease of deployment, 2) potential of cross-lingual knowledge-sharing. The latter is especially useful in low-resource conditions where training data (translation pairs) is limited. To improve the translation quality on such low-resource languages, we focus on learning common features between different languages. We especially focus on an extreme of low-resource conditions: zero-shot translation, i.e. translation between languages that do not have translated sentence pairs.

We show that the models commonly adopted in NMT contain an inherent bias towards learning language-specific representations, thereby inhibiting cross-lingual knowledge-sharing. In response, we propose improved architectures that facilitate the learning of common representations across languages. We also present multiple methods to analyze the otherwise blackbox hidden representations learned by the translation model. Based on these analyses, we show that our proposed models lead to more similar representations of related languages.

Refining Domain Knowledge for Domain Knowledge Guided Machine Learning

Pawel Bielski

Conventional data-driven machine learning approaches learn relevant patterns solely from data. In some fields, where there is not enough data, learning only from data may not be sufficient. Domain Knowledge Guided Machine Learning integrates domain knowledge from taxonomies and proved to be especially beneficial in predicting rare events in data. However, recent approaches assume that domain knowledge quality is always high and has only a positive impact. It is unclear whether this assumption is valid or whether some parts of domain knowledge have low quality and negatively impact the prediction. It is further unclear how to identify low-quality domain knowledge and potentially improve it in the context of Domain Knowledge Guided Machine Learning. In this work, we describe the problem of low-quality domain knowledge in the context of Domain Knowledge Guided Machine Learning. We propose methods to automatically identify and refine low-quality domain knowledge from a medical taxonomy on a next visit prediction task. Our results show that the refinement of domain knowledge is especially beneficial for rare data, with up to 1.5 percentage points improvement in accuracy over the methods with unrefined domain knowledge and up to 2.5 percentage points over the methods with no domain knowledge.

Machine Learning for Perovskite Thin-Film Photovoltaics

Felix Laufer

Hybrid metal-halide perovskites are a promising candidate as absorber material for the next generation of thin-film solar cells. At laboratory scale, they already enable high-performance perovskite solar cells (PSCs). However, transferring the fabrication process from lab-size solar cells to large-areas remains a key challenge. This is where machine learning (ML) methods can be applied to improve the understanding and reproducibility of the scalable thin-film formation process. We explore a unique photoluminescence (PL) dataset acquired during the perovskite layer fabrication with unsupervised k-means clustering. We show that k-means generates clusters that correlate with the performance of the final solar cell and identify disadvantages of the used fabrication method. Also, performance of PSCs are predicted using supervised k-nearest neighbors. This is a first step towards ML-based inline process monitoring and therefore is part of the path towards an autonomous laboratory driven by high throughput experiments and ML combined with high performance computing.

Automated Active Learning

Joel Arweiler

Labelling huge amounts of unlabelled data is one of the major challenges associated with “Big Data”. One approach to produce so-called pseudo labels is called “Active Learning”, where a model is trained with a small amount of labelled training data and subsequently is used to predict pseudo labels for the remaining dataset. A major problem here is, that the training data often does not represent the entire dataset, which leads to incorrect pseudo label predictions. For the classic Active learning algorithm, pseudo labels with a low confidence score will be passed to a human for correcting especially the uncertain results. This process is repeated until the model performance saturates. In our work, we propose a methodology for replacing the time-consuming labelling process by humans with an automated selection approach, for which t-SNE dimensionality reduction is used as a preprocessing step. This allows only images with a high similarity value to be added to the initial training data, increasing the probability of correct predictions for these samples, even when the initial training data does not represent the entire unlabelled data. Iteratively, the model will explore the “data space” by itself, resulting in an increasing prediction accuracy on the unlabelled dataset.

Deep Learning for the Analysis of Spectroscopic Data

André Orth, Jan Schützke

Different measurement techniques in the field of material science generate a large amount of data to evaluate, but still require manual intervention for analysis of the measured signals. For example, in-situ X-ray diffraction (XRD) measures a crystalline powder sample under varying conditions and generates hundreds of one-dimensional signals. Currently, the evaluation process involves a two-stage approach of identifying different phase variants before a refinement model fits the exact parameters. Alternatively, the use of neural networks for the analysis of powder XRD scans proves highly accurate and applicable for measurement techniques that require high-throughput evaluation methods. Similarly, neural networks have been shown to remove artifacts from the spectroscopic signals and diffractograms and work even for two-dimensional diffraction data. We demonstrate that neural networks are well suited to be used with spectroscopic and diffractometric data from different techniques within the material science domain and are looking for projects that generate and provide such data.

CNN-Based Classification of Craniosynostosis Using 2D Distance Maps

Matthias Schaufelberger

Craniosynostosis affects infants and leads to irregular head growth. 3D surface scans are a radiation-free alternative to traditional diagnosis using computed tomography. We propose a CNN-based classification approach on 2D images obtained from the 3D scans. We propose mapping approaches to visualize the 3D head shape in 2D images using coordinate transforms, ray casting, and distance extraction. We train multiple CNNs and compare them to competing classification approaches on a dataset of 496 patients. Resnet18 on spherical mapping outperformed the competing classifiers with an accuracy of 98.4%. Image attribution using Integrated Gradients reveals typical head features on the forehead and back of the head which contributed to the prediction. We demonstrated a versatile mapping approach to obtain 2D images from the 3D head shape and employ the first CNN-based classification of craniosynostosis, outperforming existant approaches. By releasing a Python package we enable other groups to easily contribute to our methodology.

Using Graph Neural Networks for Cosmic-Ray Analysis, at IceCube Observatory

Paras Koundal

The IceCube Neutrino Observatory, located near the South Pole, is a multi-component detector that detects high-energy particles from astrophysical sources. These astrophysical accelerators generate charged particles called cosmic rays (CRs). The ability to determine the underlying features and behavior of such sources is provided by CRs and CR-induced air-showers. When combined with the IceTop surface array, IceCube provides unique three-dimensional detection capabilities and cosmic-ray analysis in the transition region from galactic to extragalactic sources. The poster will detail a ongoing study using the entire in-ice shower footprint and extra composition-sensitive air-shower variables for doing an improved approximation of cosmic-ray composition.

Optimisation in Logistics - Traveling salesman, facility placement and minimum cost flow problems in practice

Armin Weckmann

Logistics-related problems are among the oldest optimisation problems in the world. Algorithms to solve them are well established in the literature and many software packages. In practice, optimisation issues of different kinds arise simultaneously and are coupled. This coupling leads to a huge amplification in problem

complexity and poses serious problems on designing an overall optimal network at large scale. We outline several optimisation problem combinations that we face in everyday business.

AI in Wastewater Treatment

Dr. Prantik Samanta

Human population is growing significantly over the past decades, so does the stress on the existing wastewater treatment plants (WWTPs). The WWTPs demand high energy to perform operations such as aeration, chemical addition, cleaning etc. The WWTPs in the US account roughly 3 to 4% of the total consumed energy per year. The global wastewater treatment market is projected to rise from US$ 281 billion in 2022 to US$ 489 billion by the end of 2029. Hence, the application of AI would not only help to optimize WWTP operations/processes and save substantial amount of energy (at the end US$) in future but also it has the potential to pre-eliminate any sort of cross-contamination that may lead to health hazard. In addition, the WWTPs generally restore historical data about the plant operations. This certainly create the platform to start off with the AI tools to optimize the energy requirements of the existing WWTP and leading towards a better future.

Sequential Networks for Cosmic Ray simulations

Pranav Sampathkumar

A hybrid model of generating cosmic ray showers based on neural networks is presented. We show that the neural network learns the solution to the governing cascade equation in one dimension. We then use the neural network to generate the energy spectra at every height slice. Pitfalls of training to generate a single height slice is discussed, and we present a sequential model which can generate the entire shower from an initial table. Errors associated with the model and the potential to generate the full three dimensional distribution of the shower is discussed.

Graph Neural Networks applied to Fluid Dynamics

Paul Mifsud

Graph Neural Networks train on data represented by graphs, for example the user profiles of social networks. The neural network learns the interactions between a node and its graph neighbourhood. It achieves this through what is called message passing: a message is constructed between a node and nodes within its graph neighbourhood. These messages are aggregated and passed through an update function that updates the state of the relevant node. This message passing can be repeated many times, per training step, allowing the change in state of one node to propagate throughout the graph. Because this message passing is done between nodes, the size of the graph of each training example can vary. In applications to fluid dynamics, we build the graph by considering the particles or fluid packets as the nodes, and create edges between nearby particles. The GNN is then trained to learn the physical interaction between these particles. The node features we use are the previous velocities and the particle type. The edge features are the displacement and distance between the particles. The concept has been proven to work but training is challenging, and many question regarding training difficulty, convergence, error bounds, and stability remain open.

Optimization of Communication Systems Using Generative Networks and Auto-encoders

Andrej Rode, Vincent Lauinger

In this poster we present novel approaches to optimize the physical layer of communication systems by exploiting various techniques from machine learning. Specifically, we optimize the geometric constellation shaping by end-to-end learning via auto-encoders and apply generative approaches like generative adversarial networks (GANs) and variational auto-encoders (VAEs) to equalize channel distortions.

GANs for generation of synthetic bronchoscopic images

Lu Guo

Many bronchoscopic image-processing tasks aiming to improve vision-based bronchoscopic navigation and assistance system, for example the depth estimation task, have been supported by deep learning-based

algorithms due to their outstanding performance compared to traditional methods. However, these algorithms often need a large amount of training data with sufficient quality. For bronchoscopic domain, the access to real intra-operative images and associated data is limited. A typical way to circumvent this obstacle is to use the synthetic data, which can be generated by Generative adversarial networks (GANs). In this work, with help of virtual bronchoscopy, which is a non-invasive technology allowing the creation of bronchoscope-like inner views of human bronchus and produces virtual bronchoscopic images, two GANs-based approches are applied to generate synthetic realistic bronchoscopic images. One of them tries to transfer bronchoscopic images between virtual and realistic image domain directly, while the other one focuses on generating realistic bronchoscopic image textures. Both of them only require a small number of training data, especially real bronchoscopic images. The results of these two approaches are compared and evaluated accordingly.

Physical Side-Channel Attacks and Defenses on Hardware Neural Network Accelerators

Dennis Gnad

Neural network (NN) computations are one of the fastest growing application domains to require increasingly more computing performance. To increase their efficiency, computing power for NNs can come from specialized hardware accelerators, such as Graphical Processing Units (GPUs), Field-Programmable Gate Arrays (FPGAs), or even dedicated Tensor Processing Units (TPUs). Like other hardware, there exists an inherent risk of physical and architectural side-channel and fault attacks on these accelerators. Furthermore, many applications in which neural networks are used are privacy-sensitive applications, include company secrets in the NN model itself, or have safety implications such as in medicine or autonomous driving. Specifically when dedicated hardware is deployed in edge computing, the risk of physical attacks is high. In this line of work, recently supported by a DFG project, we aim to increase the security of neural network accelerators with novel ideas that combine methods from side-channel security with those from machine learning. This poster presents related work, including our own, and gives a generic overview of the topic.

Constraints on cloud fraction adjustment to aerosols using explainable machine learning

Yichen Jia

Aerosol-cloud interactions (ACI), especially cloud adjustments including cloud fraction (CLF) adjustment, remain major uncertainties in the climate system. Separating the role of aerosols from meteorological confounders in such a complex buffered aerosol-cloud-climate system remains a challenge. The parameterizations of ACI in Earth System Models (ESMs) are likely inadequately representing adjustments in CLF. This ongoing work analyzes the CLF adjustment by isolating aerosol effects on CLF with machine-learning techniques in global observational data sets with the ultimate aim of evaluating ESM parameterizations.

Machine Learning for Molecular Dynamic Simulation and Chemical Reaction Prediction

Chen Zhou

Accuracy and efficiency are key factors of chemical space exploration and chemical reaction evaluation, which are essential for rational compound design and synthesis in the field of chemical, material and pharmaceutical industries. The demand of accelerating the process of target molecule discovery and reaction design has prompt the development of virtual screen (VS) that can save both expense and human effort for performing numerous chemical experiments. Among various of VS techniques, machine learning (ML) has emerged as an efficient tool as the computation costs of classic quantum calculation methods (e.g. density function theory) are often prohibitive for large molecules or long time scale. In our work, we focus on an aryl sulfone oxide that can serve as a component of host materials of organic light-emitting diodes (OLED). We have demonstrate that our fine-tuned feedforward fully connected neural network can predict energies (MAE = 0.0376 - 0.0379 eV) and forces (MAE = 0.0449 - 0.0450 eV/A) of the target molecule within "chemical accuracy" threshold (1 kcal/mol = 0.043 eV). With the ML model and active learning framework, we hope to understand the potential molecular transformation in a detailed level.