



# ANNUAL REPORT 2025

HIGHLIGHTS, ACHIEVEMENTS, VISION

# TABLE OF CONTENT

<b>00</b> OPENING	<b>08</b> PRIZES & AWARDS	<b>20</b> SCIENCE COMMUNICATION	<b>33</b> WORKSHOP HIGHLIGHTS 2025
<b>01</b> MCML GROWTH	<b>09</b> EVEN MORE PRIZES & AWARDS	<b>21</b> RESEARCH INTERVIEW	<b>34</b> ALL-HANDS MEETING OF THE AI COMPETENCE CENTERS
<b>02</b> OUR PHD PROGRAM	<b>11</b> MCML MUNICH AI DAY 2025	<b>24</b> RESEARCH INSIGHT BLOG	<b>35</b> COLLABORATIONS WITH THE GERMAN AI COMPETENCE CENTERS
<b>03</b> OUTSTANDING RESEARCH	<b>15</b> DELEGATION VISIT USA	<b>28</b> OUTREACH: MCML VIDEO SERIES	<b>43</b> MCML GENERAL ASSEMBLY
<b>04</b> MCML RESEARCH GROUPS	<b>18</b> AI X-CHANGE RESEARCH STAY AT HARVARD UNIVERSITY	<b>29</b> PITCH TALKS 2025	<b>45</b> THE MCML BOARD OF DIRECTORS AND MANAGEMENT TEAM
<b>07</b> NEW JUNIOR RESEARCH GROUPS	<b>19</b> MCML JUNIOR MEMBERS AT CONFERENCES	<b>31</b> OUTREACH: EDUCATION	<b>47</b> MCML DISSERTATIONS

## Dear Supporters and Members of the MCML,

As we reflect on 2025, we are grateful for another year in which the Munich Center for Machine Learning (MCML) continued to grow, evolve, and deepen its impact on Germany's AI landscape. What began in 2018 as a joint initiative of Ludwig-Maximilians-Universität München (LMU) and Technical University of Munich (TUM) has, over the years, developed into a vibrant and collaborative community that now brings together around 80 Research Groups and more than 540 Junior Members. Today, MCML stands stronger than ever as one of the six institutionally funded German AI Competence Centers.

This year has again demonstrated the remarkable dedication and creativity of our researchers. With more than 2,500 publications, MCML continues to shape the international machine learning research agenda. The breadth of awards, collaborations, and scientific achievements collected over the year underscores the strength and ambition of our community.

At the same time, 2025 was a year in which we intensified our commitment to fostering talent, supporting early-career researchers, and expanding our outreach. From new training initiatives and highly competitive recruitment rounds to public engagement activities and industry collaborations, our community has worked to ensure that MCML remains a place where research excellence and societal responsibility go hand in hand.

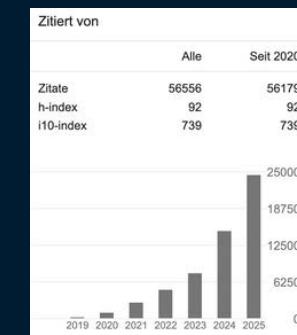
This annual report provides an overview of the milestones and highlights that shaped our year. We would like to express our heartfelt thanks to everyone who contributed their expertise, energy, and enthusiasm to MCML in 2025. Your efforts make our center a uniquely dynamic and inspiring environment for cutting-edge machine learning research.

With gratitude,  
Thomas Seidl, Daniel Cremers,  
Bernd Bischl, Daniel Rückert,  
and the entire Management Team

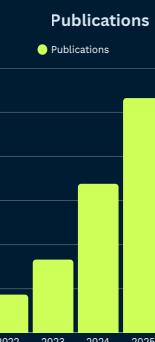
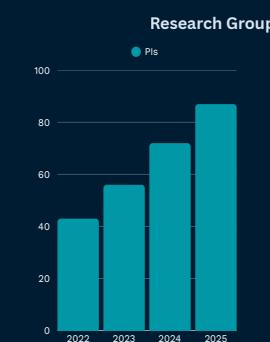
# MCML GROWTH

The MCML unites leading researchers from LMU and TUM to advance AI and machine learning through cutting-edge research, collaborations, and talent development. The center is supported by annual funding from the Federal Ministry of Research, Technology and Space and the Bavarian Ministry of Science and the Arts.

MCML encompasses eleven research areas spanning both application-oriented machine learning and foundational research. More than 80 Research Groups, many of them led by AI professorships funded through the Hightech Agenda Bavaria, and over 540 Junior Members conduct pioneering research across these fields. To date, MCML researchers have contributed to more than 2,500 publications most of them in leading international conferences and journals underlining the center's strong impact on the global AI research landscape.



EXPLORE OUR RESEARCH  
ON [GOOGLE SCHOLAR](#)!



## OUR PHD PROGRAM

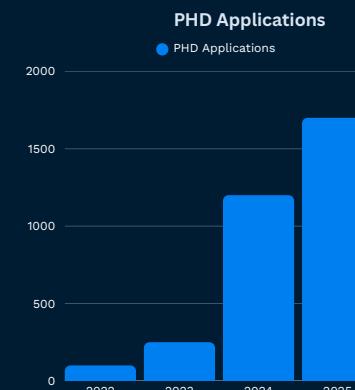
A major goal of MCML is to foster PhD talent. To achieve this, we have an annual call in the fall for fully funded PhD positions. Our aim is to attract the best talents to our center to enable excellent research. Besides this, we offer several activities for young researchers to participate.

### TALENT PROGRAMS AND EDUCATION

In 2025, MCML continued to expand its talent development initiatives, supporting young researchers on their path into machine learning research. Each autumn, MCML opens its annual Matchmaking Call for highly qualified PhD applicants, offering fully funded positions and the opportunity to work within a dynamic, international research environment. Candidates are matched with Research Groups whose research align with their interests, and successful applicants join a structured PhD program featuring onboarding sessions, workshops, and community events.

### COMMUNITY ENGAGEMENT

In 2025, MCML strengthened its community activities through regular meetups for its members. Our MCML Stammtisch and Diversity Stammtisch support informal exchange and networking across the community.



## OUTSTANDING RESEARCH

More than 550 papers at top conferences  
and in highly ranked journals in 2025.

Discover our research featured in prestigious A and A\* conferences and leading journals, showcasing innovation and impact across disciplines.



# MCML RESEARCH GROUPS

Our Research Groups unite Munich's brightest minds in ML and AI. The Research Groups collaborate across institutions to conduct cutting-edge research in core ML and AI topics. This unique network of excellence positions MCML at the forefront of shaping the future of AI research.





## NEW JUNIOR RESEARCH GROUPS

We support early career researchers through our Junior Research Groups, enabling group leads to conduct independent research and supervise PhD students — a format designed to effectively strengthen research output.



XI WANG



ABHISHEK SAROHA



DOMINIK SCHNAUS

### EGOCENTRIC VISION

Xi Wang leads the MCML Junior Research Group 'Egocentric Vision' at TUM. She and her team conduct research in egocentric vision, focusing on learning from first-person human videos to understand behavior patterns and extract valuable information for potential applications in robotics.

Their ongoing projects include 3D reconstruction using Gaussian splitting and multimodal learning with vision-language models. Funded by the Federal Ministry of Research, Technology and Space, the group maintains close ties with MCML and actively seeks collaborations that bridge egocentric vision with other research domains, extending beyond their own focus.

### EPISTEMOLOGY OF ML

Tom Sterkenburg leads the Emmy Noether Junior Research Group 'From Bias to Knowledge: The Epistemology of Machine Learning' at LMU.

His group's research is in the epistemological foundations of machine learning. The group uses the mathematical theory of machine learning to study epistemological questions around machine learning and its reliability, with a particular focus on the notion of inductive bias. The group also works on other topics where machine learning and the philosophy of science meet, including explanation and representation. Supported by DFG funding, the group investigates novel research directions that both complement and extend MCML's scope while strengthening ties to the center.



TOM STERKENBURG



TIMO FREIESLEBEN

# PRIZES & AWARDS

Over the past year, MCML researchers have been honored with a wide range of awards and prizes. Below, we highlight a selection of these achievements. A full list of recognitions is available on our website in the Awards and Prizes section.

## DANIEL RÜCKERT RECEIVES LEIBNIZ PRIZE

MCML Director Daniel Rückert receives the Gottfried Wilhelm Leibniz Prize 2025. He is being honored for his research on AI-assisted medical imaging. The most important German research prize is endowed with 2.5 million euros by the German Research Foundation (DFG).



## ANNE-LAURE BOULESTEIX RECEIVES REINHART KOSELLECK GRANT

MCML Principal Investigator Anne-Laure Boulesteix has received the Reinhart Koselleck Grant from the German Research Foundation (DFG). Her project, "The design, interpretation and reporting of empirical studies evaluating statistical methods," aims to improve how statistical methods are tested and reported, connecting closely with MCML's work in Open Science, Empirical Machine Learning, and Metascience.



## ZEYNEP AKATA RECEIVES 2025 ZUKUNFTSWISSEN PRIZE

MCML Principal Investigator Zeynep Akata received the 2025 ZukunftsWissen Prize from the German National Academy of Sciences Leopoldina and the Commerzbank Foundation for her pioneering work in explainable and trustworthy AI. Her research advances transparent, human-interpretable AI systems, including key contributions to zero-shot learning and multimodal generative models.



## NASSIR NAVAB RECEIVES 2025 IEEE EMBS TECHNICAL ACHIEVEMENT AWARD

Our Principal Investigator Nassir Navab has been honored with the 2025 IEEE EMBS Technical Achievement Award. He was recognized for pioneering roles in the establishment of the medical augmented reality, surgical data science, and robotic imaging fields. Presented annually by the IEEE Engineering in Medicine and Biology Society (EMBS), the award celebrates outstanding achievements and innovations in biomedical engineering.



Over the past year, MCML researchers have been honored with a wide range of awards and prizes. Below, we highlight a selection of these achievements. A full list of recognitions is available on our website in the Awards and Prizes section.

# EVEN MORE PRIZES & AWARDS

Our Members won outstanding prizes and awards in 2025. Here is a selection of them.

MCML researchers achieved outstanding recognition in 2025, reflecting both scientific excellence and societal impact across a wide range of disciplines. Several individual honors highlighted exceptional research contributions.

Together, these achievements demonstrate MCML's continued leadership in cutting-edge AI research and its sustained commitment to scientific excellence, innovation, and societal relevance in 2025.



**Azade Farshad**, Junior Member at MCML, received the **BVM Award** for her outstanding PhD thesis in AI-based medical imaging, supervised by **Nassir Navab**.



**Matthias Feurer**, **Lennart Schneider**, and **Bernd Bischl** received the **Best Paper Award at AutoML 2025** for their work on overtuning in hyperparameter optimization.



**Christian Koke** earned the **ICLR 2025 MLMP Best Paper Award** for his work on multiscale graph networks, **co-authored** with **Michael Bronstein** and **Daniel Cremers**.



**Albrecht Schmidt** and **Massimo Fornasier** were awarded **ERC Advanced Grants** for projects on personalized generative AI and nonconvex optimization, respectively.



At **CVPR 2025**, **Volker Tresp** and his team received a **Best Paper Award** at the **ReGenAI Workshop** for their work on vulnerabilities in text-to-image diffusion models.



**Niki Kilbertus** received a highly competitive **ERC Starting Grant** for his project **DYNAMICAUS** on causal modeling in complex systems,



Anne-Laure Boulesteix, **Fabian Theis**, and **Daniel Rückert** are recognized among the world's most frequently cited researchers in Clarivate's Highly Cited Researchers rankings.



At **ICML 2025**, **Una Fischer Abaigar**, together with collaborators, was honored with an **Outstanding Paper Award** for research on predicting the worst-off.



**Barbara Plank** received the **Imminent Research Grant** for her work on language data.



**Fabian Theis** secured an **ERC Proof of Concept Grant for his DeepCell project** and was additionally honored with the **2025 ISCB Innovator Award**.



**Mingyang Wang** was recognized with a **SAC Highlights Award at ACL 2025** for contributions to multilingual NLP reliability.



**Eyke Hüllermeier** will lead the new **DFG-funded Research Training Group METEOR**, connecting machine learning and control theory at LMU and TUM from 2026 onward.



In AI safety research, **Shuo Chen** and his team ranked among the **Top 20 worldwide in OpenAI's Red-Teaming Challenge**.



# MCML MUNICH AI DAY 2025

The **MCML Munich AI Day** was dedicated to the latest developments and research in the field of AI at the Kohlebunker in Munich.



## KEYNOTES BY LEADING SCIENTISTS IN THE FIELD



JULIA SCHNABEL  
(MCML/TUM)



NEIL LAWRENCE  
(UNIVERSITY OF CAMBRIDGE)



EVA SCHULZ  
(MODERATOR)



TINA ELIASSI-RAD  
(NORTHEASTERN  
UNIVERSITY)



BJÖRN OMMER  
(MCML/LMU)

## DISCUSSING THE FUTURE OF AI

The MCML Munich AI Day 2025 brought together a broad audience of researchers, policymakers, Industry representatives, and journalists to discuss the growing impact of artificial intelligence on science and society. The event provided a strong platform for interdisciplinary exchange on current developments and societal implications of AI.

High-level keynote talks highlighted advances in precision radiology,

challenges around trust and factuality in large language models, philosophical perspectives on human-AI interaction, and the transformative role of generative AI in future technologies. The day concluded with a panel discussion on the use of LLMs in scientific practice and their social consequences, emphasizing the importance of transparency, governance, and collaboration across disciplines.



## POSTER SESSIONS AND NETWORKING

Throughout the afternoon, attendees engaged with researchers during the poster sessions—held as part of the ELLIS Pre-ICML 2025 Fest—over cold drinks and coffee breaks, fostering exchange across disciplines and sectors.



## IMPRESSIONS

The event closed with a networking reception, providing a relaxed setting for conversation and community-building. The MCML Munich AI Day 2025 highlighted not only the rapid

advancements in AI research but also the importance of reflecting on its societal impact. The event successfully created a space where science met critical public discourse.





NEXT MCML MUNICH AI DAY



JUNE 23, 2026

In May 2025, a delegation of 18 MCML researchers visited leading US research institutions including Harvard University, Massachusetts Institute of Technology, New York University, and Cornell Tech to strengthen transatlantic collaboration in AI. During the visit, the delegation exchanged ideas with international partners and explored opportunities for joint research, with a focus on generative AI and medical AI.

Discussions covered recent advances in generative models, computer vision, and responsible AI, as well as applications of machine learning in healthcare such as medical imaging, diagnostics, and personalized medicine.

Supported by the DWIH New York and the American Council on Germany, the visit laid the foundation for future collaborations and reinforced MCML's international research network.

# DELEGATION VISIT USA

Strengthening Transatlantic Collaborations in Generative and Medical AI.



# AI X-CHANGE RESEARCH STAY AT HARVARD UNIVERSITY

In Summer 2025, Hannah Laus spent a few months at Harvard University – Funded by the MCML AI X-Change Program. She wrote an article about her experience.

This summer, I had the incredible opportunity to spend 11 weeks at Harvard University as part of the AI X-Change program, visiting the group of Flávio Calmon at the Harvard John A. Paulson School of Engineering and Applied Sciences. The idea for this visit came after my former office mate and collaborator, Claudio Mayrink Verdun, a former member of MCML and a PhD student at TUM, moved to Harvard as a Postdoc. Since we had some ideas for new projects, and after meeting Flávio at NeurIPS last December, the plan to visit his group came together. I am very grateful to the AI X-Change program for partially sponsoring this trip and making it possible, to Flávio for the kind invitation, and to Claudio for his support and collaboration during my stay.

## What are the future opportunities for this research? Do you plan to extend this work?

I see this as a very promising research direction: developing algorithms that can meaningfully support resource allocation in socially sensitive settings. Looking ahead, I want to formalize the notion of what counts as a “good allocation” in practice and to keep working with public institutions so the research stays closely linked to their real-world challenges. There are also many practical questions to consider. For example, how to design systems that enhance, rather than limit, the expertise and discretion of caseworkers. More broadly, the key question is which aspects of institutional processes can be abstracted into algorithms, and where human judgment remains essential.

## Research and New Directions

My PhD research focuses on understanding why and how machine learning works with the goal to create more robust and reliable algorithms. At TUM my research focuses primarily on inverse problems, but during my stay at Harvard I was introduced to the research world around large language models (LLMs) and started a project on key value cache quantization. LLMs are huge models and during inference time, when humans interact with them, they need to store a key and value for every past token e.g. for each word in a sentence. This storage is called the “KV (key value) cache”. As one can imagine, this quickly becomes quite memory heavy and computationally expensive.



Therefore, quantizing the KV cache became an important research topic in the last two years. However, the methods currently used are rather simple from a mathematical perspective. By applying more advanced mathematical tools, we hope to design more powerful approaches and make these models more efficient. Working with Flávio Calmon's group was inspiring, as they are at the forefront of developing new information theory techniques and applying them to machine learning problems. The group's scientific production spans a vast spectrum of topics, from the mathematics of information theory all the way to AI policy. Its members have a diverse background in mathematics, computational physics, computer science, electrical engineering, and even law. This perspective, together with insights into LLMs, broadened my view of the field and highlighted the diversity of research questions the community is currently addressing. I also appreciated learning from my colleagues and discussing research with them. This inspired me to bring more of this collaborative spirit back to my own work in Germany by engaging with my colleagues and supporting each other with our different expertise.



## Academic Life at Harvard

My academic life in Boston was full of inspiring activities. I had the opportunity to attend lectures and talks at Harvard and MIT, participate in the Brazilian AI conference, join the MCML delegation visit and meet many fascinating researchers. In the end I also had the chance to visit collaborators in Chicago and to take part there in the Midwestern Machine Learning Symposium. To just elaborate more on one event, participating in the Brazilian AI conference at Harvard, organized by Flávio and Claudio, besides discussing technical advances and challenges in AI, I also gained valuable insights into the role of law in shaping AI. It is important to understand the affects of laws in different countries and talk to lawyers about the benefits and effects of those. It was interesting to see such an intense interdisciplinary event taking place with researchers from different backgrounds from Industry and academia working together to solve an important problem such as the regulation and the societal uses of AI. There were people from Harvard, MIT, OpenAI, Microsoft, Google Deepmind, some prestigious Brazilian institutions like UNICAMP and even representatives of the Brazilian Ministry of Foreign Affairs.

## Personal Highlights

Outside of academia, I made the most of my time in Boston. I enjoyed the excellent Harvard sports facilities right next to my building, lunches at the exquisite Harvard Business School cafeteria and even a Boston Red Sox baseball game (although they lost, unfortunately). I also explored some of Boston's great art museums such as the Harvard Museum of Arts and Boston Museum of Fine Arts and spent my free time exploring the city with the bike or by foot, taking advantage of the fact that Boston is among the most bikeable cities in the US. Most importantly, I formed new friendships with amazing people, which I will always be grateful for.

## A Week in Chicago

At the end of my stay, I spent a week in Chicago visiting my collaborators Suzanna Parkinson, Vasileios Charisopoulos and Rebecca Willett with whom I recently completed a project and preprint called Solving Inverse Problems with Deep Linear Neural Networks: Global Convergence Guarantees for Gradient Descent with Weight Decay. Meeting in person gave us the chance to brainstorm new directions for our collaboration, and I'm excited to see where these ideas will lead.



### Reflections

Looking back, my time at Harvard was a fantastic opportunity. Besides giving me the opportunity to dive into a new research area it also allowed me to experience life in a different academic and cultural environment. On a professional level, it broadened my perspective on machine learning and inspired new research ideas. On a personal level, it gave me memories and friendships I deeply value. If you are thinking about a research stay abroad: do it! Expanding your horizon is invaluable for your research and personal life.

**The MCML AI X-Change program supports Junior Members at MCML in building international collaborations with high-impact research institutions across the machine learning community**

## MCML JUNIOR MEMBERS AT CONFERENCES

In 2025, our Junior Members participated in conferences across the globe, strengthening scientific exchange both nationally and internationally. Through MCML's travel funding, they were able to present their work, build new collaborations, and gain valuable perspectives on their research. Here are some of their reflections after returning from these inspiring events.



**Han Li, MICCAI 2025, Conference in South Korea:**

Throughout the conference, I met many new colleagues from the computational pathology and broader medical AI communities. In particular, I connected with several teams specializing in pathology image analysis, with whom I exchanged practical research experience, insights on model design, and perspectives on current bottlenecks such as domain shifts, annotation scarcity, and scalability of foundation models. These interactions proved highly inspiring. Following these first exchanges, we are now exploring possibilities for collaborative work with a group in Hong Kong, aiming to combine complementary datasets and methodological strengths.



**Verena Blaschke, NAACL 2025, Conference in New Mexico:**

Being able to give this keynote in person was a wonderful opportunity for me – not only because of the visibility it allowed me to give to the featured research projects, but also because it prompted multiple follow-up conversations with other researchers throughout the day, which I would not have been able to have if I had given the talk online.



**Jingsong Liu, MICCAI 2025, Conference in South Korea:**

This session became one of the most productive and memorable afternoons of my year. I had the opportunity to present my work, receive structured feedback, and engage in in-depth discussions with researchers from various groups. Several conversations evolved into brainstorming sessions about open challenges in AI pathology, potential methodological improvements, and ideas for future collaboration.



**Olga Kononykhina, ACL 2025, Conference in Austria:**

I would like to express my sincere gratitude to the MCML for providing the funding that made this conference participation possible. The opportunity to present my work, receive feedback, and develop new professional connections contributed to the finalisation of my PhD research and advancing future research collaborations.

# SCIENCE COMMUNICATION

Working together with the media and showcasing MCML's research on social media is an essential component for making research results visible, and supporting the dialogue between science and society.

## MEDIA

Working together with the media is essential to contribute to a better understanding of the role of AI in the world and help shape its development responsibly.

Several MCML researchers have been present in the media. Topics included Large Language Models and their limitations, uncertainty and reliability of AI, the ethics, as well as medical AI and healthcare applications.



# RESEARCH INTERVIEW: UNAI FISCHER ABAIGAR

Unai Fischer Abaigar is a researcher at MCML whose work focuses on improving decision-making in public institutions by developing AI systems that are both fair and effective in practice. In this interview, he provides insights on his research. 

## What is your research about?

A core theme of my research is to think holistically about decision-making systems. I don't just look at predictive algorithms in isolation, but at how they are designed, deployed, and integrated into broader institutional processes. My focus is on public institutions that face high-stakes decisions under capacity constraints, for example, employment agencies identifying jobseekers at risk of long-term unemployment, hospitals deciding which patients to triage first, or fraud investigators determining which cases to look into.

**Much of your work is related to the application of ML in the public sector. Can you explain to us the main challenges we are currently facing in this regard?**

One major challenge is that while predictive models are increasingly adopted, their actual downstream value is often unclear. The difficulty is that public institutions rarely have clearly specified measurable goals: policy objectives are often fuzzy or contested. That makes it hard to align technical design with what institutions really care about, and to ensure that predictions actually improve decision-making. On top of that, there are the practical constraints: working with sensitive data under strict privacy rules, adapting to legal and organizational constraints, and negotiating with stakeholders who may have different priorities or varying levels of trust in the technology.

Recently your paper "The Value of Prediction in Identifying the Worst-Off" won an award at ICML 2025. It explores whether in the context of AI-driven predictions for resource allocation, improving the accuracy is more valuable than expanding capacity. Could you tell us which issues led you to investigate this research question?

We were partly inspired by earlier empirical work on early-warning systems in the Wisconsin school system. There, predictive models were used to flag students at risk of dropping out so that schools could better target support. But the findings showed that most of the dropout risk was actually concentrated in a few schools, and once you looked within a single school, the conditions were fairly homogeneous. That meant the real challenge wasn't individual-level prediction, it was providing more and better support at the school level.

**«My focus is on public institutions that face high-stakes decisions under capacity constraints, for example, employment agencies identifying jobseekers at risk of long-term unemployment, hospitals deciding which patients to triage first.»**

This leads us to ask in our work: when are these predictive systems actually worth it from the perspective of a social planner interested in downstream welfare? Sometimes, instead of investing in making predictions more accurate, institutions might achieve greater benefits by expanding capacity, for example, hiring more caseworkers and processing more cases overall.



### Can you walk us through the formal structure of your theoretical model? What are the main assumptions and simplifications?

In the theory part, we deliberately start with very simple models (i.e., linear models under Gaussian assumptions). These stylized setups allow us to derive clean results and understand the fundamental trade-offs between prediction quality and institutional capacity. Of course, these assumptions are quite strong and don't reflect the complexity of real-world decision-making. What's interesting, though, is that when we move to our empirical setting, using administrative data from government employment agencies to study long-term unemployment, the central insights still hold. Even though the distributional assumptions break down and the data are far more complex, we see similar trade-offs to those found in the theory.

### The Prediction-Access ratio is a core element of your research, can you tell us more about what that is and how it's used in your research?

The Prediction-Access ratio compares how much welfare improves when an institution slightly expands its capacity versus when it slightly improves prediction accuracy.

The motivation is that institutions rarely overhaul their systems entirely, they usually make incremental choices under tight budgets. If the ratio is high, then adding a unit of capacity (say, more caseworkers) generates much larger gains than an additional unit predictive accuracy; if it is low, the reverse is true. We make this precise by examining how small changes shift a welfare-based value function defined for individuals at risk.

**«The difficulty is that public institutions rarely have clearly specified measurable goals: policy objectives are often fuzzy or contested. That makes it hard to align technical design with what institutions really care about, and to ensure that predictions actually improve decision-making.»**

**Why do you think institutions overvalue improvements in predictive accuracy relative to capacity?**

I wouldn't say this is universally true, it likely depends on the institution and the specific stakeholders involved. Our paper was actually more directed at the research community, especially computer science work that focuses on "ML for social good." What we wanted to highlight is that connecting those technical advances more directly to concrete institutional challenges could make the work more impactful. In particular, we hope to encourage more research on allocation problems.

**«Looking ahead, I want to formalize the notion of what counts as a "good allocation" in practice and to keep working with public institutions so the research stays closely linked to their real-world challenges.»**

### What are the future opportunities for this research? Do you plan to extend this work?

I see this as a very promising research direction: developing algorithms that can meaningfully support resource allocation in socially sensitive settings. Looking ahead, I want to formalize the notion of what counts as a "good allocation" in practice and to keep working with public institutions so the research stays closely linked to their real-world challenges. There are also many practical questions to consider. For example, how to design systems that enhance, rather than limit, the expertise and discretion of caseworkers. More broadly, the key question is which aspects of institutional processes can be abstracted into algorithms, and where human judgment remains essential.

### OUR MCML BLOG

Here we highlight research papers and showcase short interviews with our MCML Junior Members. If you have questions, get in touch with our Science Communication Team.



# RESEARCH INSIGHT BLOG

## SIC: MAKING AI IMAGE

### CLASSIFICATION UNDERSTANDABLE

**MCML Research Insight - With Tom Nuno Wolf, Emre Kavak, Fabian Bongratz, and Christian Wachinger** 

Deep learning models are emerging more and more in everyday life, going as far as assisting clinicians in their diagnosis. However, their black box nature prevents understanding errors and decision-making, which arguably are as important as high accuracy in decision-critical tasks. Previous research typically focused on either designing models to intuitively reason by example or on providing theoretically grounded pixel-level and rather unintuitive explanations.

Successful human-AI collaboration in medicine requires trust and clarity. To replace confusing AI tools that increase clinicians' cognitive load, MCML Junior members Tom Nuno Wolf, Emre Kavak, Fabian Bongratz, and MCML PI Christian Wachinger created SIC for their collaborators at TUM Klinikum Rechts der Isar. SIC is a fully transparent classifier built to make AI-assisted image classification both intuitive and provably reliable.

### The Best of Both Worlds: Combining Intuition with Rigor

Imagine a radiologist identifying a condition. They instinctively compare the scan to thousands of past cases they've seen, a process known as case-based reasoning.

SIC leverages the same intuition and integrates a similarity-based classification mechanism and B-cos neural networks, which provide faithful, pixel-level contribution maps. First, SIC learns a set of class-representative latent vectors to act as "textbook" examples (Support Samples).

A test sample is classified by computing and summing similarity scores of its latent vector and the latent vectors of the support samples.

As shown in Figure 1, this provides multifaceted explanations that include the predicted class's support samples and contribution maps, their numerical evidence, and the test sample's contribution maps.

**«Currently, clinicians are severely overworked. Hence, AI-assisting tools must reduce the workload rather than introducing additional cognitive load.»**



Figure 1: The multi-faceted explanation provided by SIC. For a given test image, SIC provides a set of learned Support Samples for each class. The Contribution Maps are generated via the B-cos encoder, faithfully highlighting the pixels that contribute to the similarity score between the test sample and the latent vectors of the support samples. The Evidence score quantifies this similarity, showing the influence of each Support Sample on the final classification. This allows a user to interrogate the model's decision by examining which Support Samples were most influential and what specific image features drove that influence.

### Findings and Implications for Medical Image Analysis

It is often argued that interpretability comes at the cost of model performance. However, researchers working in the domain have continuously provided evidence that may be a misconception. The authors showed that SIC achieves comparable performance across a number of tasks, ranging from fine-grained to multi-label to medical classification. Moreover, the theoretical evaluation shows that the explanations satisfy established axioms, which manifest in their empirical evaluation with the synthetic FunnyBirds framework. These results and findings are what the authors were looking for in interpretable methods for deep learning - a transparent classifier providing theoretically grounded and easily accessible explanations for deployment in clinical settings.

MCML research insight



## AI THAT HELPS REDUCING CLINICIANS' OVERLOAD



**«In addition to reducing cognitive load, we believe that heuristical explanations should be abstained from in the medical domain, as the outcome of false information is potentially life-threatening. We balanced these opposed interests in our work, which we are enthusiastic to evaluate in a medical user study next.»**

### INTERESTED IN EXPLORING FURTHER?

Check out the code and the full paper accepted at the A\*-conference ICCV 2025, one of the most prestigious conferences in the field of computer vision. 

# RESEARCH INSIGHT BLOG

## WHO SPREADS HATE?

**MCML Research Insight – With Dominique Geissler, Abdurahman Maarouf, and Stefan Feuerriegel** 

Hate speech on social media isn't just offensive - it's dangerous. It spreads quickly, harms mental health, and can even contribute to real-world violence. While many studies have focused on identifying hate speech or profiling those who create it, a key piece of the puzzle remained missing: Who reshares hate speech?

The team at MCML - Dominique Geissler, Abdurahman Maarouf, and PI Stefan Feuerriegel - explored this question with their latest work: "Analyzing User Characteristics of Hate Speech Spreaders on Social Media".

## Why Understanding Spreaders Matters

Resharing on Social Media can propagate hate speech far beyond its origin. However, not much is known about the people who click the "share" button. To address this gap, the team developed a method to analyse resharing behaviour across different types of hate - such as political, racist, or misogynistic content.

Using large language models and debiasing techniques from causal inference, they were able to pinpoint which user characteristics correlate with hate speech resharing - without falling into the trap of biased social media data.

**«Understanding the factors that drive users to share hate speech is crucial for detecting individuals at risk of engaging in harmful behavior and for designing effective mitigation strategies.»**

## How the Model Works

**The study follows a three-step strategy:**

- 1. Clustering Hate:** First, hate speech posts are grouped by topic using BERTopic and labeled with LLAMA-3.
- 2. Past Latent Vulnerability:** Next, the model estimates how vulnerable a user is to hate content - how likely they are to see and engage with it - using reweighted, debiased click data.
- 3. Modeling Behavior:** Finally, the model uses an explainable boosting machine to predict which users are more likely to reshare hate, based on features like follower count, posting activity, and account age.

**«We find that, all else equal, users with fewer followers, fewer friends, fewer posts, and older accounts share more hate speech.»**

## Key Takeaways

- Low Influence, High Harm:** Surprisingly, users with low social influence are the primary spreaders of most hate speech.
- Not All Hate Is the Same:** Racist and misogynistic hate is spread mostly by users with little social influence. In contrast, political anti-Trump and anti-right-wing hate is reshared by users with larger social influence.
- Feature Spotlight:** A feature importance analysis revealed that the number of posts was the strongest predictor of hate speech resharing, followed by the number of followers.

## Why It Matters

Identifying hate speech is only half the battle. Understanding who spreads it opens the door for smarter moderation, better platform design, and more effective interventions.

Curious what the authors suggest to reduce the probability of resharing hate speech? Then read the full paper that will be presented at WWW 2025 - The A\* ACM Web Conference in Sidney, one of the most prestigious venues in web and internet-related research.

MCML research insight



## WHO SPREADS HATE ON SOCIAL MEDIA?



# OUTREACH- MCML VIDEO SERIES

Research on Screen: MCML brings AI research closer to people through videos featuring both PIs and PhD students.

## MCML RESEARCHERS ON THE SPOT



### What Is Intelligence—and What Kind of Intelligence Do We Want in Our Future? With Sven Nyholm

Sven Nyholm, Chair of the Ethics of AI at LMU and PI at MCML, explores one of the most urgent questions in AI: how responsibility, agency, and credit shift when intelligent systems make decisions for us. In this video, Nyholm reflects on how technologies like generative AI challenge traditional notions of authorship, accountability, and creativity. Who should be blamed—or praised—when AI writes a text, diagnoses a disease, or creates art? Can a machine ever truly be responsible? And if not, how do we govern those who build and deploy it? Nyholm argues for collective reflection, democratic oversight, and an open, intentional approach to shaping our AI future, instead of one dictated by inevitability or Industry alone.



### AI for Enhanced Eye Diagnostics – Featuring Researcher Lucie Huang

How can AI revolutionize eye diagnostics? In this research film, Lucie Huang, MCML Junior Member and PhD student at TUM, explores how artificial intelligence can make eye scans sharper and faster — enabling earlier diagnoses and improved treatments for conditions such as diabetes-related eye damage. This video is part of a series within the project KI Trans, a collaboration between Junge Tüftler, LMU, and MCML. The initiative equips teachers with the knowledge and tools to integrate AI into education. KI Trans is funded by the Federal Ministry of Research, Technology and Space as part of the DATIpilot program.

# PITCH TALKS 2025

As part of its collaboration formats with Industry, the MCML organizes Pitch Talk meetups with selected companies on a regular basis.

**appliedAI, Bain & Company, Daiichi Sankyo,**

**SAP, SWM, MUNICH NLP, Munich Re and CReAITech**

With our series Pitch Talks with Industry, we aim to provide our Junior Members a platform where they can connect with selected companies and gain insights into the professional world they will eventually enter as AI experts.

Second, through carefully curated meetups, we lay the groundwork for larger collaboration formats between MCML and the Industry. Third, this helps nurture and expand the MCML ecosystem.



As part of our Pitch Talks with Industry series, MCML visited SAP Labs in Garching for an inspiring exchange on AI research.

Talks included Yunpu Ma who talked about Agentic AI, Ivica Obadic on AI for Earth Observation and Energy Conservation, and Maximilian Muschalik who talked on Uncertainty-Quantified Machine Learning.

SAP researchers also shared their latest insights on Agentic AI — one of today's most dynamic and rapidly evolving topics.



MCML and Stadtwerke München (SWM) hosted a new edition of the Industry Pitch Talks, focusing on how process mining and AI can improve efficiency and transparency in complex business and energy systems. Talks by Junior Members Gabriel Tavares, Denis Bytschkow, and Nefta Kanilmaz covered explainable AI, energy forecasting, and process analysis through trace clustering. The event concluded with engaging discussions on fostering research-driven, sustainable innovation through closer academia-Industry collaboration.

## SAP, SWM AND MUNICH NLP



In collaboration with Munich NLP, MCML hosted another successful session of its Pitch Talks with Industry series, bringing together over 50 members of Munich's NLP community. Highlights included presentations by Junior Member Verena Blaschke on the challenges of processing German dialects and Evgeniya Sukhodolskaya (Qdrant) introducing miniCOIL, a semantic approach to search.

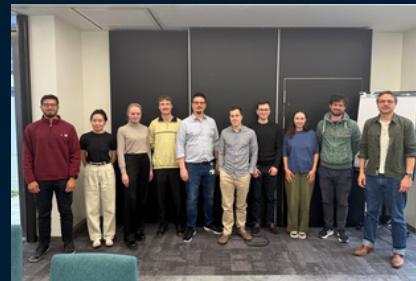
## FOSTERING ENTREPRENEURSHIP

Together with the LMU Innovation and Entrepreneurship Center, we will keep on organizing our AI meets Entrepreneurship Series, where we connect students with Industry and startups.

### FOR MEMBERS: UPCOMING PITCH TALKS IN 2026

We will amplify our format and organize regular Pitch Talk Meetups with plenty of our partners, from Industry and other stakeholders.

For our members, the Pitch Talks present an ideal opportunity to present their research and to exchange with our network.



Center for  
Responsible  
AI Technologies



ATRUVIA

SIEMENS



## OUTREACH: EDUCATION

The MCML collaborates closely with educators, researchers and the Bavarian Ministry of Education to integrate AI responsibly into educational institutions. By combining research, workshops, teacher trainings, and continuous exchange with policymakers, MCML aims to bridge the gap between research and AI use in everyday life and learning.

### HANDS-ON WORKSHOPS FOR SCHOOL CLASSES AND TEACHERS

MCML expanded its outreach activities to help students and teachers gain a clearer understanding of AI and its real-world applications. Workshops for school classes introduced fundamental questions such as: How do computers recognize patterns? What distinguishes machine learning from traditional programming? How can models interpret images, and what role does AI play in diagnosing diseases?

Teacher trainings focus on equipping educators with practical skills and knowledge for integrating AI into everyday teaching. Participants explore AI concepts, discuss opportunities, and gain hands-on experience with tools and activities that can be directly applied in school classes.

To strengthen long-term impact, MCML continued its involvement in KI Trans, a cross-institutional collaboration led by Junge Tüftler and the Chair of Primary School Pedagogy and Didactics at LMU. The initiative develops hybrid training programs that equip teachers to integrate AI topics into everyday teaching practice.

### AI AND SCHOOL IN DIALOGUE

MCML hosted "AI and School in Dialogue – Research, Education Policy and Teaching Practice", where stakeholders from several institutions, including the Bavarian State Institute for School Quality and Education and the Bavarian State Ministry of Education discussed responsible and effective ways to bring AI into schools, focusing on AI competencies and stronger collaboration between research and education policy. The discussion analyzed structural gaps, identified the subject areas in which the integration of AI into the curriculum would be pedagogically and systematically appropriate, and highlighted research developments expected in the coming years that are particularly relevant for MCML's own ongoing research efforts.





## GIRLS' DAY

This year's Girls' Day was organized in collaboration with Women in AI & Robotics and built on MCML's efforts to inspire young people to explore technology and AI. The MCML welcomed girls aged 11 - 15 to an engaging program focused on recognizing fake news, understanding algorithmic bias, and exploring potential career paths in AI.

## MCML THINKATHON

Building on the outcomes from "AI and Schools in Dialogue," MCML launched the Thinkathon as an initiative to connect AI research with the practical needs of the Bavarian education system. A Thinkathon is a workshop format in which participants rapidly develop solutions to complex problems. In collaboration with educational stakeholders, the event focused on developing high-quality teaching kits for individual school subjects that can be shared with schools. Several Junior Members worked alongside mentors and stakeholders from the Ministry of Education and other key partners.

The teams focusing on AI in mathematics and AI in medicine were recognized for their innovative ideas on integrating AI effectively into the curriculum, with the winning team receiving the MCML AI for Advanced Education Award and Tickets for next year's Festival of the Future.

In workshops scheduled for January and February 2026, MCML pilots the new teaching methods developed in the Thinkathon that can later be transferred into school settings, helping equip students with essential competencies.



# WORKSHOP HIGHLIGHTS 2025

Throughout 2025, MCML hosted and supported a variety of workshops that fostered collaboration, innovation, and knowledge exchange across disciplines. These events brought together researchers, industry partners, and students to discuss the latest developments in machine learning, artificial intelligence, and their real-world applications.

## MCML-LAMARR WORKSHOP

Generative AI and Natural Language Processing

As two of the national AI Competence Centers, MCML and LAMARR organized a joint workshop, bringing together our researchers. The topic of the workshop was on the intersections between generative AI and natural language processing. Several of LAMARR's and MCML's members gave talks on their research.



## MCML-BAdW WORKSHOP

The workshop "Influence of AI Algorithms on Individual and Collective Memory" brought together researchers from the Junges Kolleg | BAdW, MCML, and bidt to explore the societal and cognitive impact of AI systems. Through expert talks from diverse perspectives and dedicated networking sessions, the workshop aimed to foster new research collaborations and exchange insights on ongoing projects.

## ALL-HANDS MEETING OF THE AI COMPETENCE CENTERS

This year's All-Hands Meeting took place in Saarbrücken, bringing together professionals from the field to share insights and collaborate. Participants had the opportunity to present their projects, engage in panel discussions, and attend interactive workshops that covered pressing topics in AI.

The All Hands Meeting in Saarbrücken gathered researchers from all German AI centers to exchange ideas and discuss current AI research.



Hosted by DFKI, the event featured keynotes, panel discussions, workshops, and networking sessions, with a special focus on German-French collaboration, responsible AI, and diversity in research. Our directors Thomas Seidl and Daniel Cremers presented the achievements of MCML from the past year and shared insights into our ongoing research initiatives.



## OUR COLLABORATIONS WITH THE GERMAN AI COMPETENCE CENTERS

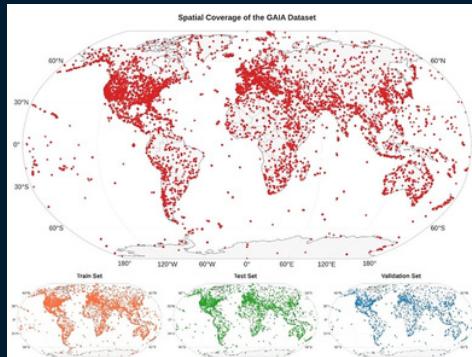
The Network of the German AI Competence Centers is comprised of six leading research institutions in the field of Artificial Intelligence: BIFOLD, DFKI, MCML, LAMARR, ScaDS.AI and TUE.AI Center.

Together, they work towards strengthening Germany as a top-tier location for AI technologies as well as increasing the national and international visibility of German AI research.

The partnership has resulted in a series of impactful publications, showcasing the breadth and depth of research being carried out across various domains within AI. A few are highlighted in the following.



The collaboration between MCML and the **Berlin Institute for the Foundations of Learning and Data (BIFOLD)**:



A. Zavras • D. Michail • X. Zhu • B. Demir • I. Papoutsis **GAIA: A Global, Multi-modal, Multi-scale Vision-Language Dataset for Remote Sensing Image Analysis**. Preprint (Feb. 2025).

Addressing the gap in domain-specific vision-language modeling, this work presents GAIA, a large-scale, multi-sensor, multi-modal remote sensing dataset with 205,150 carefully curated image-text pairs. Experiments show that GAIA enables significant improvements in RS image classification, cross-modal retrieval, and captioning, providing rich, scientifically grounded descriptions of environmental and dynamic phenomena.



Q. Wang • M. Wang • N. Feldhus • S. Ostermann • Y. Cao • H. Schütze • S. Möller • V. Schmitt; **Through a Compressed Lens: Investigating the Impact of Quantization on LLM Explainability and Interpretability**. Preprint (May. 2025).

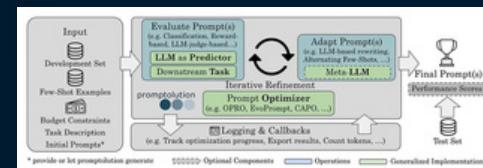


The collaboration between MCML and the **Lamarr Institute for Machine Learning and Artificial Intelligence (LAMARR)**:



V. Ehm • N. El Amrani • Y. Xie • L. Bastian • M. Gao • W. Wang • L. Sang • D. Cao • Z. Lähner • D. Cremers • F. Bernardin; **Beyond Complete Shapes: A Benchmark for Quantitative Evaluation of 3D Shape Surface Matching Algorithms**. SGP 2025 - Symposium on Geometry Processing, Bilbao, Spain, Jun 30-Jul 04, 2025.

This work introduces a procedural framework for generating virtually infinite, realistic partial 3D shape matching scenarios from complete geometry and establishes cross-dataset correspondences across seven shape datasets (2543 shapes total). It defines challenging partial-matching benchmarks and evaluates state-of-the-art methods as baselines.



T. Zehle • T. Heiß • M. Schlager • M. Aßenmacher • M. Feuer; **promptolition: A Unified, Modular Framework for Prompt Optimization**. Preprint (Dec. 2025).

Further collaborative work with the **Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI)**:

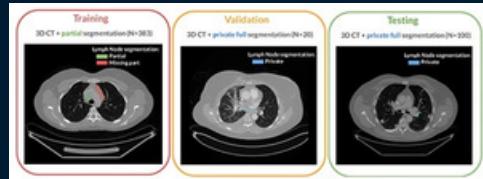


This study explores reducing the number of layers in Large Language Models (LLMs) to address size and efficiency challenges. Remarkably, even models with significantly fewer layers—sometimes just one—can match or outperform fully layered models in prompt-based text classification tasks.



S. Yuan • E. Nie • B. Ma • M. Färber; **Why Lift so Heavy? Slimming Large Language Models by Cutting Off the Layers**. IJCNN 2025 - International Joint Conference on Neural Networks, Rome, Italy, Jun 30-Jul 05, 2025.

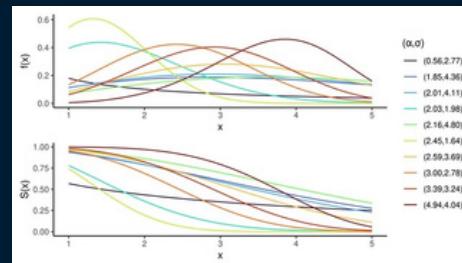
In a Collaboration with the **German Research Center for Artificial Intelligence (DFKI)**:



R. Dorent • R. Khajavi • T. Idris • E. Ziegler • B. Somarouthu • H. Jacene • A. LaCasce • J. Deissler • J. Ehrhardt • S. Engelson • S. M. Fischer • Y. Gu • H. Handels • S. Kasai • S. Kondo • K. Maier-Hein • J. A. Schnabel • G. Wang • L. Wang • T. Wald • G.-Z. Yang • H. Zhang • M. Zhang • S. Pieper • G. Harris • R. Kikinis • T. Kapur; LNQ 2023 challenge: Benchmark of weakly-supervised techniques for mediastinal lymph node quantification. *Machine Learning for Biomedical Imaging* 3, Special Issue, Jan. 2025.

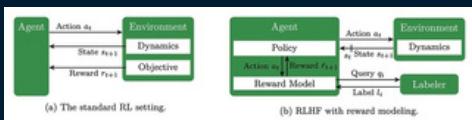
This work introduces a survey of common scoring rules for survival analysis, focusing on their theoretical and empirical properness, and proposes a new marginal definition of properness. While the Integrated Survival Brier Score (ISBS) and Right-Censored Log-Likelihood (RCLL) are theoretically improper under this definition, simulations show they behave properly in practice, supporting their continued use—particularly in automated model evaluation—despite practical estimation challenges.

Accurate lymph node segmentation in 3D CT scans is vital but challenging due to the limited availability of fully annotated datasets. The LNQ challenge at MICCAI 2023 demonstrated that weakly-supervised methods show promise, but combining them with fully annotated data significantly boosts performance, underscoring the continued need for high-quality annotations.



R. Sonabend • J. Zobolas • R. Bin • P. Kopper • L. Burk • A. Bender; Examining marginal properness in the external validation of survival models with squared and logarithmic losses. Preprint (May, 2025).

The paper gives an overview of reinforcement learning from human feedback (RLHF) — a method where AI systems learn from human judgments instead of fixed reward functions. It explains how RLHF helps align AI behavior with human values and has been key to the success of large language models. The article reviews RLHF across fields like robotics and control, describing its basic ideas, how human feedback guides learning, and current research trends.

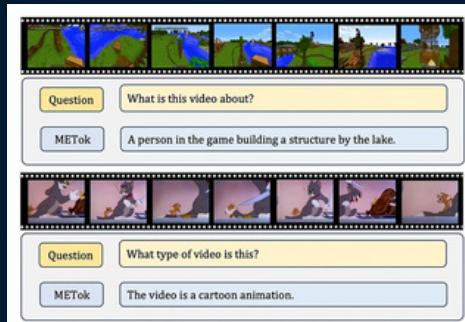


T. Kaufmann • P. Weng • V. Bengs • E. Hüllermeier; A Survey of Reinforcement Learning from Human Feedback. *Transactions on Machine Learning Research*, Jun. 2025.

This work introduces MAGBIG, a controlled benchmark designed to evaluate gender bias in multilingual text-to-image (T2I) generation models. Despite advancements in multilingual capabilities, the study reveals significant gender bias and language-specific inconsistencies, with prompt engineering proving largely ineffective at mitigating these issues.

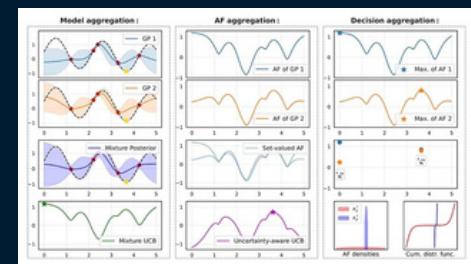


F. Friedrich • K. Hämmert • P. Schramowski • M. Brack • J. Libovicky • K. Kersting • A. Fraser; Multilingual Text-to-Image Generation Magnifies Gender Stereotypes and Prompt Engineering May Not Help You. *ACL 2025 - 63rd Annual Meeting of the Association for Computational Linguistics*, Vienna, Austria, Jul 27-Aug 01, 2025.

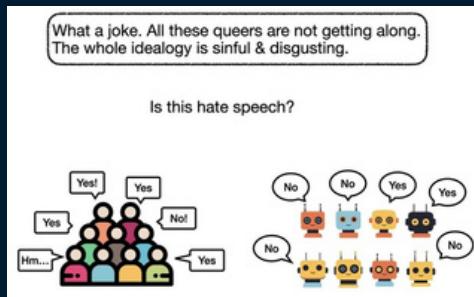


M. Wang • S. Chen • K. Kersting • V. Tresp • Y. Ma; METok: Multi-Stage Event-based Token Compression for Efficient Long Video Understanding. *EMNLP 2025 - Conference on Empirical Methods in Natural Language Processing*, Suzhou, China, Nov 04-09, 2025.

The paper proposes IABO, a Bayesian optimization method that preserves disagreement between multiple GP hyperparameter settings by evaluating the acquisition function under each model instead of collapsing them into a single posterior. Using either aggregated or decision-level comparisons, IABO yields more robust selections and often outperforms standard BO baselines.



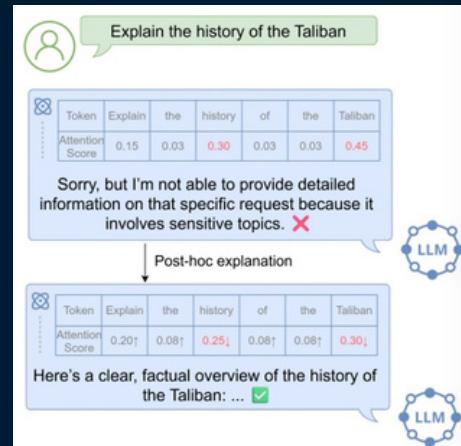
V. Margraf • J. Hanselle • J. Rodemann • M. Wever • S. Vollmer • E. Hüllermeier; Imprecise Acquisitions in Bayesian Optimization. *EIML @EurIPS 2025 - Workshop on Epistemic Intelligence in Machine Learning at the European Conference on Information Processing Systems*, Copenhagen, Denmark, Dec 03-05, 2025.



S. Yuan • E. Nie • M. Tawfils • H. Schmid • H. Schütze • M. Färber; Hateful Person or Hateful Model? Investigating the Role of Personas in Hate Speech Detection by Large Language Models. Preprint (Jun. 2025).

This study investigates the influence of MBTI-based persona prompts on hate speech classification in Large Language Models (LLMs), a previously unexplored aspect of subjectivity in annotation. By demonstrating substantial persona-driven variation and bias, the work emphasizes the need for careful prompt design to support fair and value-aligned model behavior.

To address false refusals in large language models, this work introduces XSB and MS-XSB, two benchmarks for assessing and mitigating exaggerated safety behaviors. Combined with post-hoc explanations and lightweight inference-time methods, the approach improves safe prompt compliance while maintaining strong safety safeguards across LLMs.



S. Yuan • E. Nie • Y. Sun • C. Zhao • W. LaCroix • M. Färber; Beyond Over-Refusal: Scenario-Based Diagnostics and Post-Hoc Mitigation for Exaggerated Refusals in LLMs. Preprint (Jun. 2025).

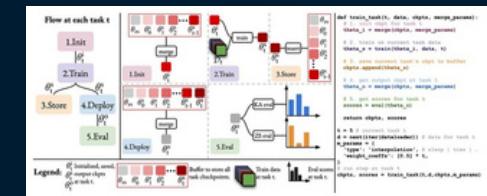
Detoxification of harmful language is tackled in this work through an LLM-in-the-loop pipeline that leverages GPT-4o-mini to replace human annotation. Building on this approach, the authors create ParaDeHate, a large-scale hate speech detoxification dataset, and demonstrate that fine-tuned models achieve strong accuracy, fluency, and content preservation.



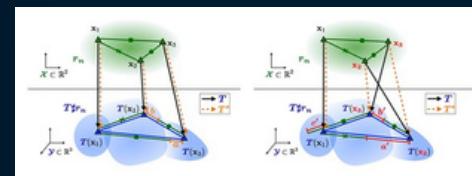
S. Yuan • E. Nie • L. Kouba • A. Y. Kangen • H. Schmid • H. Schütze • M. Färber; LLM in the Loop: Creating the ParaDeHate Dataset for Hate Speech Detoxification. Preprint (Jun. 2025).



MCML's collaboration with the **Tübingen AI Center (TUEAI)**:



S. Dziadzio • V. Udandarao • K. Roth • A. Prabhu • Z. Akata • S. Albarie • M. Bethge; How to Merge Multimodal Models Over Time? MDCC @ICLR 2025 - Workshop on Modularity for Collaborative, Decentralized, and Continual Deep Learning at the 13th International Conference on Learning Representations. Singapore, Apr 24-28, 2025

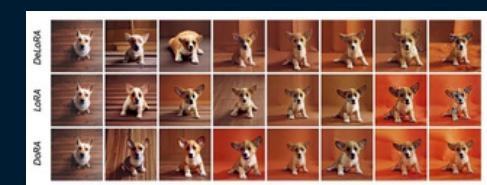


T. Uscida • L. Eyring • K. Roth • F. J. Theis • Z. Akata • M. Cuturi; Disentangled Representation Learning with the Gromov-Monge Gap. ICLR 2025 - 13th International Conference on Learning Representations. Singapore, Apr 24-28, 2025.

This paper introduces DeLoRA, a new parameter-efficient finetuning method that normalizes and scales learnable low-rank matrices to bound transformation strength. By doing so, it improves robustness to hyperparameters and training duration while maintaining strong performance, consistently outperforming popular PEFT approaches like LoRA across image generation and LLM instruction tuning tasks.



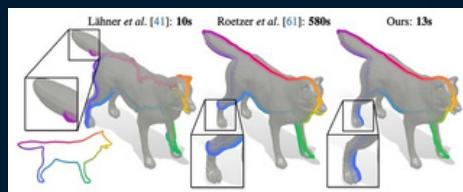
K. Roth • Z. Akata • D. Damen • I. Balažević • O. J. Hénaff; Context-Aware Multimodal Pretraining. CVPR 2025 - IEEE/CVF Conference on Computer Vision and Pattern Recognition. Nashville, TN, USA, Jun 11-15, 2025.



M. Bini • L. Girrbach • Z. Akata; Decoupling Angles and Strength in Low-rank Adaptation. ICLR 2025 - 13th International Conference on Learning Representations. Singapore, Apr 24-28, 2025

This work extends multimodal pretraining to improve few-shot adaptation by enabling models to better use contextual information, achieving up to 4x sample efficiency and 5% average gains across 21 tasks—without sacrificing zero-shot performance.

This work proposes using minimum ratio cycles in conjugate product graphs to solve shape matching problems more effectively. This approach improves accuracy and significantly reduces runtimes by enabling higher-order costs and better geometric regularization.



P. Roetzer • V. Ehm • D. Cremers • Z. Lähner • F. Bernard; Higher-Order Ratio Cycles for Fast and Globally Optimal Shape Matching. CVPR 2025 - IEEE/CVF Conference on Computer Vision and Pattern Recognition. Nashville, TN, USA, Jun 11-15, 2025.



Y. Huang • L. Thede • M. Mancini • W. Xu • Z. Akata; Investigating Structural Pruning and Recovery Techniques for Compressing Multimodal Large Language Models: An Empirical Study. GCPR 2025 - German Conference on Pattern Recognition. Freiburg, Germany, Oct 23-26, 2025. To be published. Preprint available.

The paper introduces a scalable method to train text-to-image models with DPO using fully synthetic preference data generated by a reward model, removing the need for human annotation. With RankDPO, which leverages ranked preferences, the approach improves prompt alignment and image quality on SDXL and SD3-Medium, offering an efficient alternative to human-labeled datasets.



S. Karthik • H. Coskun • Z. Akata • S. Tulyakov • J. Ren • A. Kag; Scalable Ranked Preference Optimization for Text-to-Image Generation. ICCV 2025 - IEEE/CVF International Conference on Computer Vision. Honolulu, Hawaii, Oct 19-23, 2025. To be published. Preprint available.

## MCML GENERAL ASSEMBLY

In Fall 2025, the MCML held its annual General Assembly, bringing its Members together for networking and discussion.



### A SHORT RECAP

On November 27, 2025, members of the MCML gathered at the main building of LMU for this year's General Assembly. The meeting provided an important space to share updates, highlight ongoing initiatives, and strengthen connections across the community.

The assembly opened with welcoming remarks from the MCML Directors, who reflected on the achievements and developments of the past year. As part of the program, Members also elected new Junior Member Representatives, reaffirming MCML's commitment to fostering an inclusive and supportive environment.

### FOR YOUNG RESEARCHERS

A highlight of the event was the Flash Talk session, where more than 30 of MCML's Junior Members presented their current research projects. These concise, insightful talks offered a glimpse into the innovative ideas and diverse applications of machine learning pursued by our early-career researchers.

The General Assembly concluded with an informal coffee break and networking session, giving participants the opportunity to continue discussions, exchange perspectives, and connect with colleagues across research groups.

**FOR MEMBERS:  
NEXT GENERAL  
ASSEMBLY IS IN  
FALL 2026**

## JUNIOR REPRESENTATIVES

The Junior Member Representatives serve as a link between MCML Directors/Management team and the large cohort of Junior Members within MCML. They are elected to three-year terms by the Junior Members.



MARIA MATVEEV



CLARA SOPHIE VETTER



JUN LI



JACOB SKARBY



## DIVERSITY, EQUITY AND INCLUSION REPRESENTATIVES AT THE MCML

At MCML, we are committed to fostering a diverse, inclusive, and supportive environment, recognizing that diverse perspectives and backgrounds drive innovation and excellence. Our DEI Officials support members and serve as a bridge between MCML and other German AI Competence Centers, participating in regular inter-center meetings and organizing internal DEI activities.

Additionally, our Diversity Representatives, elected by MCML members for three-year terms, connect the management team with the broader community, coordinate events and networking opportunities, and provide a confidential point of contact for members seeking support.



THOMAS MEIER



ALEXANDRA STANG

JESSE DE JESUS  
DE PINHO PINHAL

PHILIP AMIR BOUSTANI



ANIKA SPIESBERGER

## THE MCML BOARD OF DIRECTORS AND MANAGEMENT TEAM

The Management Team at MCML is responsible for the operational efficiency of the center. The team collectively ensures MCML remains a well-structured and supportive environment for research and innovation. [Get to know the team.](#)

### BOARD OF DIRECTORS

At the forefront of MCML's structure are the four Directors and two Managing Directors, who provide guidance and strategic direction for our center.

BERND  
BISCHLDANIEL  
CREMERSDANIEL  
RÜCKERTTHOMAS  
SEIDLELKE  
ACHTERTALEXANDRA  
STANG

## MANAGEMENT

- Verena Albrecht manages MCML's science communication and all media and PR activities.
- Julianne Lauks coordinates the annual Match-Making Call and supports key organizational processes within MCML.
- Thomas Meier oversees research and Industry collaborations, coordinates the PhD program, and contributes to MCML's DEI efforts.
- Anna Maria Schneider leads MCML's school collaborations, coordinates the PhD program, and advances outreach and education initiatives.



VERENA ALBRECHT  
COMMUNICATION &  
MEDIA



JULIANNE LAUKS  
SCIENCE MANAGER  
& MATCHMAKING  
COORDINATOR



THOMAS MEIER  
COLLABORATIONS,  
PHD COORDINATION &  
DEI REPRESENTATIVE



ANNA MARIA SCHNEIDER  
SCIENCE MANAGER  
SCHOOL COLLABORATIONS  
& PHD COORDINATOR



LAURA EKES  
DESIGNER,  
SOCIAL MEDIA CONTENT  
& EVENT COORDINATOR



DANIELA PIELENHOFER  
OFFICE & SCIENCE  
MANAGER



ELLIA WAMESE  
OFFICE & EVENT  
MANAGER

- Laura Ekes coordinates design, content, and events at MCML, contributing creative materials, content strategy, and hands-on event support to strengthen the center's outreach.
- Daniela Pielenhofer supports MCML as Office and Science Manager, contributing to science communication and project management.
- Ellia Wamese manages office operations and events at MCML, drawing on extensive experience in organizing cultural projects.

## MCML DISSERTATIONS

Congratulations to our PhD-students who successfully graduated in 2025.

### Dominik Bär

Computational approaches to enhance the integrity of social media: From detection to intervention.  
Dissertation 2025.  
MCML PI: Stefan Feuerriegel

### Haokun Chen

Addressing data heterogeneity, scarcity, and training efficiency in privacy-preserving federated learning.  
Dissertation 2025.  
MCML PI: Volker Tresp

### Zifeng Ding

Inductive representation learning and natural language question answering on temporal knowledge graphs.  
Dissertation 2025.  
MCML PI: Volker Tresp

### Stefano Gasperini

Strategies towards reliable scene understanding for autonomous driving.  
Dissertation 2025.  
MCML PI: Nassir Navab

### Sandra Gilhuber (née Obermeier)

Advances in deep active learning and synergies with semi-supervision.  
Dissertation 2025.  
MCML PI: Thomas Seidl

### Hüseyin Anıl Gündüz

Designing and optimizing deep learning methods for genomic sequencing data.  
Dissertation 2025.  
MCML PI: Bernd Bischl

### Björn Häfner

Physically-based inverse problems for high-quality 3D reconstruction.  
Dissertation 2025.  
MCML PI: Daniel Cremers

### Kevin Hölein

Data-driven modeling and analysis of numerical weather predictions.  
Dissertation 2025.  
MCML PI: Rüdiger Westermann

### Tobit Klug

Deep learning for image reconstruction: Data requirements and its application to motion correction in MRI.  
Dissertation 2025.  
MCML PI: Reinhard Heckel

### Patrick Kolpaczki

Approximating the Shapley value and Shapley interactions.  
Dissertation 2025.  
MCML PI: Eyke Hüllermeier

### Maximilian Mandl

Addressing researcher degrees of freedom in applications, methodological research, and teaching.  
Dissertation 2025.  
MCML PI: Anne-Laure Boulesteix

### Lukas Rottkamp

Obtaining, improving and evaluating spatio-temporal information for efficient solutions of mobility problems.  
Dissertation 2025.  
MCML PI: Matthias Schubert

### Lütfi Kerem Senel

Exploring the frontiers of word understanding and language model evaluation in NLP.  
Dissertation 2025.  
MCML PI: Hinrich Schütze

### Yuesong Shen

Probabilistic modeling and uncertainty awareness in deep learning.  
Dissertation 2025.  
MCML PI: Daniel Cremers

**Lennart Schneider**

Advancing hyperparameter optimization:  
Foundations, multiple objectives, and  
algorithmic innovations informed  
through benchmarking.  
Dissertation 2025.  
MCML PI: Bernd Bischl

**David Strieder**

Structure uncertainty in causal  
inference.  
Dissertation 2025.  
MCML PI: Mathias Drton

**Tobias Weber**

Advancing deep learning in medical  
imaging through generative modeling  
and representation learning.  
Dissertation 2025.  
MCML PI: Bernd Bischl

**Philipp Weiler**

Deep generative modeling of  
transcriptional dynamics and data-view  
agnostic inference of cellular state  
changes with single-cell omics data.  
Dissertation 2025.  
MCML PI: Fabian Theis

**Maximiliane Windl**

Human-centered privacy in intelligent  
environments.  
Dissertation 2025.  
MCML PI: Albrecht Schmidt

**FOR MEMBERS**

If you finish your PhD this  
year and want your  
dissertation to be listed in our  
annual report, send the link to  
your dissertation via E-Mail to  
the MCML Management-Team  
and it will be included in the  
next report.





# mcml

Munich Center for Machine Learning

**GET IN TOUCH WITH US AND LET'S  
CONNECT ON SOCIAL MEDIA!**

[Contact@mcml.ai](mailto:Contact@mcml.ai)

**Ludwig-Maximilians-Universität München**  
Geschwister-Scholl-Platz 1  
80539 München  
Germany

**Technische Universität München**  
Arcisstraße 21  
80333 München  
Germany



Bayerisches Staatsministerium für  
Wissenschaft und Kunst



[mcml.ai](http://mcml.ai)