

The Explainable AI for Molecules “AiChemist” Project




Newsletter 6 – October 2026

Fifth AiChemist School – Copenhagen Summer School (15-19 June 2026)

Hosted by Prof. Jan H. Jensen and the University of Copenhagen

The 2026 AiChemist Summer School took place in Copenhagen from 15-19 June 2026, bringing together the AiChemist Doctoral Candidates, invited speakers, supervisors and network partners for a week dedicated to explainable artificial intelligence (XAI) and machine learning for chemistry. The School focused on methods for making AI models more transparent, reliable and scientifically meaningful, while highlighting applications across molecular property prediction, drug discovery, catalyst design, toxicology, protein language models and safety assessment. The programme combined lectures, a hackathon, poster presentations, collaboration sessions, governance meetings and networking activities, providing opportunities for both scientific training and interaction among members of the AiChemist network. The School was opened up to the wider scientific community, for both in-person and online attendance, and was announced on the AiChemist website, LinkedIn and Bluesky in the months leading up to the School. The first two days of the School, which were dedicated to XAI, attracted around 80 in-person participants, including Master's and PhD students from the University of Copenhagen. All of the lectures were broadcasted over Zoom for the online attendees – the offering attracted on average 20 online participants per day. Around 70% of the lectures were recorded, with recording being made available on the AiChemist Website and Youtube Channel (<https://aichemist.eu> <https://www.youtube.com/@AiChemist-MSCA-DN>).



The first half of the School explored a broad range of topics spanning decision-tree analyses of chemical datasets, self-explaining graph neural networks, structure-property relationships, chemical language models and generative AI for drug discovery. Subsequent sessions focused on uncertainty quantification, explainability in molecular machine learning, interpretability in protein language models, evaluation of explanations, toxicity prediction and methods for extracting mechanistic insight from black-box models. According to some participants, the School provided a valuable opportunity to exchange ideas and receive feedback from researchers outside of AiChemist working on closely related topics. The School placed a strong emphasis on bridging methodological advances in explainable AI with practical challenges in chemistry and drug discovery. Events were hosted primarily at the University of Copenhagen, with networking activities including a welcome dinner on the day before the school commenced (Sunday, the 14th of June), a welcome reception on the first day, poster sessions and informal collaboration time among the Doctoral Candidates (DCs).

Day 1: Monday, 15 June

The School opened with a strong focus on explainability and machine learning methods for chemistry.

- Jan Jensen (University of Copenhagen) opened the programme with a lecture titled *Do machines dream of atoms?*, introducing decision tree-based analyses of chemical datasets and discussing interpretable machine learning approaches for chemistry.

- Jonas Teufel (KIT) presented approaches for discovering structure-property relationships using self-explaining graph neural networks.
- Inken Fender (University of Bern) discussed how model design choices influence the behaviour and performance of chemical language models.
- Djork-Arné Clevert (Pfizer) concluded the scientific programme with a lecture on generative models for drug discovery.



Jonas Teufel of KIT (left) and Inken Fender from the University of Bern (right) giving their talks on the first day

The day highlighted both the opportunities and challenges associated with developing interpretable AI systems for chemistry, ranging from classical machine learning approaches to modern foundation models and generative methods.

The evening concluded with a welcome reception at the Pavilion Beer Garden, providing an opportunity for the DCs to reconnect and meet the the invited speakers, as well as the students of the UCPH Chemistry Department and other external participants.

Day 2: Tuesday, 16 June

The second day focused on interpretability, uncertainty and trustworthiness in machine learning applications.

- Raquel Rodríguez-Pérez (Novartis) discussed methods for quantifying and explaining uncertainty in molecular machine learning models.
- Parastoo Semnani (TU Berlin) introduced explainability methods for machine-learning-guided catalyst discovery.
- Alex Vicente Sola (CRG) explored the role of interpretability in protein language models.
- Marcel Hiltcher (Sanofi) discussed faithfulness and evaluation of explanations in molecular deep learning.
- AiChemist Doctoral Candidate Dina Khasanova presented her work, *From Toxicophores to Physicochemical Reasoning: Evaluating XAI in Toxicity Prediction*.

- Peter Koo (Cold Spring Harbor Laboratory) presented methods for extracting mechanistic insights from black-box models.
- Michael Wand (IDSIA) closed the day with a lecture on formal reasoning for explainable neural network classification.

Together, the talks highlighted the importance of understanding not only what machine learning models predict, but also why they make particular predictions and how confidently those predictions can be trusted.



Parastoo Semnani, PhD student at TU Berlin (left) gave an overview of her work in XAI-guided catalyst discovery and Peter Koo of Cold Spring Harbour Laboratory explained how mechanistic insights can be extracted from black-box models

Day 3: Wednesday, 17 June

Wednesday was dedicated to a full-day hackathon that allowed participants to apply concepts discussed during the first two days.

The challenge tasked participants with reverse-engineering aspects of the training data of machine learning models using only information about the model architecture and learned weights. The exercise demonstrated both the power and limitations of modern machine learning systems and highlighted topics such as model memorisation, data leakage, membership inference attacks and generalisation.

Participants worked collaboratively in teams throughout the day, developing and testing strategies before presenting and discussing their findings.



In the zone: DCs and other participants working on the challenge



At the end of the day, the teams presented their approaches and the winning team (Parastoo, Mateusz and Marcel, left) explained how their careful, and somewhat lucky, selection of datasets to test their approach at the beginning of the hackathon, paved their way to the finish line

Day 4: Thursday, 18 June

The fourth day combined scientific sessions with opportunities for networking and collaboration within the Doctoral Network.

- Jonas Verhellen (University of Copenhagen) discussed computational diversity in the era of explainable AI.
- Joshua Xu (FDA) presented applications of AI for safety assessment.

The afternoon included a poster session, allowing Doctoral Candidates to present their ongoing research and discuss their projects with supervisors, invited lecturers and peers. Dedicated collaboration time enabled participants to further develop ideas sparked by the lectures and hackathon activities earlier in the week. Parallel sessions included meetings of the AiChemist General Assembly and Supervisory Board.



Jonas Verhellen of the University of Copenhagen kicking off the penultimate day (left), which was rounded off with a poster session

Day 5: Friday, 19 June

The final day explored scientific innovation, automation and applications of AI in chemistry and drug discovery.

- Nobel Laureate Morten Meldal delivered a keynote lecture entitled *A Nobel Quest into the Realm of Innovation*. He reflected on his scientific journey, the development of the CuAAC click reaction and perspectives on the future of click chemistry. The talk was highlighted by several participants as one of the memorable moments of the School.
- Miguel Steiner (DTU) discussed understanding chemical reaction mechanisms through automation and machine learning.
- Igor Tetko (HMGU) presented recent developments and applications of the OCHEM platform for model development and interpretation.
- Swapnil Chavan (RISE Sweden) closed the scientific programme with a discussion of challenges in predicting drug metabolism.

The School concluded with some brief closing remarks before participants departed.



Prof. Morten P. Meldal, Professor of Chemistry at UCPH, recounted his journey towards receiving a Nobel Prize for the development of the CuAAC click reaction (left) and Miguel Steiner, Assistant Professor at the Technical University of Denmark (right) discussed his work on ML-assisted elucidation of reaction mechanisms.

The Copenhagen Summer School delivered advanced training in several core areas:

1. Explainable AI and interpretable machine learning

Participants explored methods for understanding, explaining and validating machine learning models used in chemistry, biology and toxicology.

2. Molecular machine learning and generative AI

The programme covered graph neural networks, chemical language models, uncertainty estimation and generative approaches for molecular design and drug discovery.

3. Trustworthy AI for science

Several lectures addressed faithfulness, uncertainty quantification, catalyst discovery, safety assessment and mechanistic understanding of model predictions.

4. Hands-on learning and collaboration

The hackathon provided practical experience with topics such as model memorisation, membership inference attacks and data leakage, while poster sessions and collaboration time encouraged interactions across institutions and disciplines.

5. Innovation and scientific leadership

The keynote lecture from Nobel Laureate Morten Meldal offered valuable perspectives on scientific discovery, perseverance and innovation.

Overall, the Copenhagen Summer School strengthened scientific exchange within the AiChemist network and provided the DCs with both conceptual and practical understanding of explainable AI methods relevant to modern chemistry and drug discovery.

Voices from the Copenhagen School

"I was happy to give a talk on 'From Toxicophores to Physicochemical Reasoning: Evaluating XAI in Toxicity Prediction'. It's always nice to get feedback in any form, from small corrections to completely new research directions. We then proceeded with a full-day hackathon, where, knowing only the model architecture and weights, we tried to reconstruct the training data using membership inference attacks, which was a really interesting challenge. On the last day, we heard an inspiring talk from Nobel Prize winner Morten Meldal about his path in science, the development of the CuAAC click reaction, and the future of click chemistry. The take-home was: **Never give up and do things that keep you excited, not everything works on the first attempt!**"



Dina Khasanova (DC2), Molecular Networks

"I have found over the years that machine learning people and experimentalists don't talk to each other enough, but the AiChemist XAI Summer School in Copenhagen was a full week of the opposite, and it was a good reminder of what that's worth. I was invited to give a talk on "Class-Aware Explainability for Machine Learning-Guided Catalyst Discovery". My argument was fairly simple: in experimental catalysis, where the interesting catalysts are rare and the data is unbalanced by nature, explainability isn't a box you tick at the end. It's how you find out whether your model actually learned chemistry, or just learned the quirks of how the data was collected. The best part, though, was the mix of people it pulled together. Generative models for drug discovery, self-explaining graph neural networks, faithfulness in molecular deep learning, uncertainty in property prediction, protein language models, toxicity prediction, catalysis. All in one room, and all willing to argue."



Then came the hackathon: if a model memorised its training data, can you make it confess? Framed as a privacy attack, we had to recover the hidden training set of two molecular classifiers from black-box access alone. Softmax outputs are nearly useless here because they saturate for basically every drug-like input, so we worked in logit and gradient space instead. The key idea is that a deliberately overfit model sits at a loss minimum on its own training points, so the gradient of the cross-entropy with respect to the input fingerprint vanishes at exactly those molecules. Rank every candidate by ascending input-gradient norm and the memorised set floats to the top. That beat both the naive baseline and the generative pipeline (REINVENT RL) we'd reached for first. I had a lot of fun with this one, and I'll admit our team winning didn't hurt. Thanks to my teammates Mateusz and Marcel, and to Fabian and for putting the hackathon together."

Parastoo Semnani (invited speaker), TU Berlin

Upcoming Events and Conferences

- AiChemist will be at the [ML4Molecules](#) Workshop at NeurIPS in Paris on the 13th of December! AiChemist PI [Ola Engkvist](#) will co-organise this year's edition, with fellows PIs [Francesca Grisoni](#) and [Philippe Schwaller](#) giving invited talks. Many of the AiChemist DCs will be there too.
- The final AiChemist School will take place at the AstraZeneca site in Mölndal, Sweden, from the 25th to the 29th of January 2027. The school is being co-organised with the [Konrad Zuse School of Excellence in Learning and Intelligent Systems](#) (ELIZA) and will be open to the students of ELIZA, AiChemist and Chalmers University.

Publications

- (1) [Tetko, I.](#) Tox24 Challenge. *Chem. Res. Tox* **2024**. <https://doi.org/10.1021/acs.chemrestox.4c00192>
- (2) Stocco, F.; Artigues-Lleixà, M.; [Hunklinger, A.](#); Widatalla, T.; Güell, M.; Ferruz, N. Guiding Generative Protein Language Models with Reinforcement Learning. *arXiv* **2025**. <https://doi.org/10.48550/arXiv.2412.12979>.
- (3) [Krüger, F. P.](#); Östman, J.; Mervin, L.; Tetko, I. V.; Engkvist, O. Publishing Neural Networks in Drug Discovery Might Compromise Training Data Privacy. *J. Cheminformatics* **2025**. <https://doi.org/10.1186/s13321-025-00982-w>.
- (4) Goldstein, D.; [Alcaide, E.](#); Lu, J.; Cheah, E. RADLADS: Rapid Attention Distillation to Linear Attention Decoders at Scale. *arXiv* **2025**. <https://doi.org/10.48550/arXiv.2505.03005>.
- (5) [Hunklinger, A.](#); Ferruz, N. Toward the Explainability of Protein Language Models for Sequence Design. *Nature Machine Intelligence* **2026**. <https://doi.org/10.1038/s42256-026-01232-w> arXiv version: <https://doi.org/10.48550/arXiv.2506.19532>

- (6) Cirino, T.; Pinto, L.; Iwan, M. *et al.* Consensus Modeling Strategies for Predicting Transthyretin Binding Affinity from Tox24 Challenge Data. *Chem. Res. Toxicol.* **2025**. <https://doi.org/10.1021/acs.chemrestox.5c00018>.
- (7) Eytcheson, S. A.; Tetko, I. V. Which Modern AI Methods Provide Accurate Predictions of Toxicological Endpoints? Analysis of Tox24 Challenge Results. *Chem. Res. Tox.* **2025**. <https://doi.org/10.1021/acs.chemrestox.5c00273>.
- (8) Tetko, I.V.; Clevert, D.A. Advanced Machine Learning for Innovative Drug Discovery. *J. Cheminform.* **2025**. <https://doi.org/10.1186/s13321-025-01061-w>
- (9) Ball, M.; Horvath, D.; Kogej, T.; Kabeshov, M.; Varnek, A. Predicting Reaction Conditions: A Data-Driven Perspective. *Chem. Sci.* **2025**. <https://doi.org/10.1039/D5SC03045E>.
- (10) Krüger, F. P.; Österbacka, N.; Kabeshov, M.; Engkvist, O.; Tetko, I. MolEncoder: Towards Optimal Masked Language Modeling for Molecules. *ChemRxiv.* **2025**. <https://doi.org/10.26434/chemrxiv-2025-h4w9d>.
- (11) Singh, I.; Onogi, Y.; Menezes, F.; Khasanova, D.; Kang, L.; Wang, C.; Ruiz-Trave, J.; Sharma, S.; Khalil, A.; Reichenbach, V. K.; Shi, Y.; Flatley, A.; Yan, X.; Israel, A.; Dragano, N. R. V.; Aguilar-Pimentel, J. A.; Hoffmann, A.; Ghosh, A.; Noé, F.; Wolfrum, C.; Cucuruz, S.; König, A.-C.; Bartscher, I.; Hauck, S. M.; Lickert, H.; Hofmann, S. M.; Feederle, R.; Schriever, S. C.; Hernandez-Bautista, R.; Sancar, G.; Cebrian-Serrano, A.; Tetko, I.; Fuchs, H.; Gailus-Durner, V.; Blüher, M.; Hrabě de Angelis, M.; Ussar, S. NRAC Controls CD36-Mediated Fatty Acid Uptake in Adipocytes and Lipid Clearance in Vivo. *EMBO J.* **2025**. <https://doi.org/10.1038/s44318-025-00520-2>.
- (12) Krüger, F. P.; Österbacka, M.; Kabeshov, M.; Engkvist, O.; Tetko, I. MolEncoder: Improved Masked Language Modeling for Molecules. *International Conference on Artificial Neural Networks* **2025**. https://link.springer.com/chapter/10.1007/978-3-032-04552-2_6
- (13) Krüger, F. P.; Österbacka, M.; Kabeshov, M.; Engkvist, O.; Tetko, I. MolEncoder: Improved Masked Language Modeling for Molecules. *Digital Discovery* **2025**. <https://pubs.rsc.org/en/content/articlelanding/2025/dd/d5dd00369e>
- (14) Park, J.; Rashid, S.; Copsey, H.; Tran, L.; Zabeo, A.; Hristozov, D.; Gakis, G. P.; Charitidis, C.; Yoon, S.; Shin H.K. NanoToxRadar: A Multitarget Nano-QSAR Model for Predicting the Cytotoxicity Values of Multicomponent Nanoparticles. *ACS Nanoscience Au.* **2025**. <https://doi.org/10.1021/acsnanoscienceau.5c00035>
- (15) Viganò, E. L.; Iwan, M.; Colombo, E.; Ballabio, D.; Roncaglioni, A. Mixture of experts for multitask learning in cardiotoxicity assessment. *J. Cheminformatics* **2025**. <https://doi.org/10.1186/s13321-025-01072-7>
- (16) Roncaglioni, A.; Kovarich, S.; Mansorui, K.; Tetko, I.V. Advancing Human and Environmental Safety Science Using *In Silico* Methods. *Chem. Res. Tox* **2025**.
- (17) Alcaide, E.; Gao, Z.; Ke, G.; Li, Y.; Zhang, L.; Zheng, H.; Zhou, G. Uni-Mol Docking V2: Towards Realistic and Accurate Binding Pose Prediction. *International Conference on Artificial Neural Networks* **2025**. https://link.springer.com/chapter/10.1007/978-3-032-04552-2_5

- (18) Hiltcher, M.; Bianciotto, M.; Grisoni, F. Explaining What Matters: Faithfulness in Molecular Deep Learning. *Chemrxiv* **2026**. <https://chemrxiv.org/doi/full/10.26434/chemrxiv.10001594/v1>
- (19) Krüger, F. P.; Hunklinger, A.; Wolny, A.; Adler, T.J.; Tetko, I.; Villalba, S.D. SEISMO: Increasing Sample Efficiency in Molecular Optimization with a Trajectory-Aware LLM Agent. *arXiv* **2026**. <https://arxiv.org/abs/2602.00663>
- (20) Iwan, M.; Roncaglioni, A.; Grisoni, F. Beyond Molecular Structure: Investigating Demographic Factors in Drug-Induced Cardiotoxicity Prediction Models. *Journal of Chemical Information and Modelling* **2026**. <https://doi.org/10.1021/acs.jcim.6c00418>
- (21) Ball, M.; Faber, F.A.; Pratley, C.; Kogej, T.; Horvath, D.; Varnek, A.; Bergonzini, G.; Jäger, C.M.; Kabeshov, M. Expert-Informed Contrastive Learning of Condition Space for Amide Coupling Reactions. *Chemrxiv* **2026**. <https://chemrxiv.org/doi/full/10.26434/chemrxiv.15001054/v1>
- (22) Xuan-Vu, N.; Armstrong, D.; Wehrbach, M.; Bran, A.M.; Jončev, Z.; Schwaller, P. Synthelite: Chemist-aligned and feasibility-aware synthesis planning with LLMs. *arXiv* **2025**. <https://arxiv.org/abs/2512.16424>
- (23) Khasanova, D.; Tetko, I.V. Benchmarking explainable AI methods for toxicophore detection and toxicity prediction. *Digital Discovery* **2026**. <https://doi.org/10.1039/D5DD00576K>
- (24) Herrmann, V.; Alcaide, E.; Wand, M.; Schmidhuber, J. Multiple Token Divergence: Measuring and Steering In-Context Computation Density. *arXiv* **2025**. <https://arxiv.org/abs/2512.22944>
- (25) Treis, T.; Moshkov, M.; Mqawass, G.; Theis, F. J. CP-BG-1M: A Controlled Multi-View Benchmark for Density and Background Shortcuts in Morphology Profiling. *ICLR 2026 Workshop "MLGenX"*. <https://openreview.net/pdf?id=lrFOBw487r>
- (26) Krüger, F.; Kogej, T.; Kabeshov, M.; Engkvist, O.; Tetko, I.V. SmartReact: An Open-Source Python Package for Scalable Forward Synthesis Enumeration Built on Interpretable SMARTS Templates. *Chemrxiv* **2026**. <https://chemrxiv.org/doi/abs/10.26434/chemrxiv.15006435/v1>
- (27) Xuan-Vu, N.; Armstrong, D. P.; Jončev, Z.; Schwaller, P. TempRe: Template generation for single and direct multi-step retrosynthesis. TempRe: Template generation for single and direct multi-step retrosynthesis. *arXiv* **2026**. <https://doi.org/10.48550/arXiv.2507.21762>
- (28) Armstrong, D.; Xuan-Vu, N.; Susanu, O.; Gibberd, G.; Neukomm, T. A.; Strunden, T.; Forster, D.; Delattre, M.; Teh, S.; Rols, C.; Federice, J.; Leatherwood, H.; Barnes, M. L.; Dobbelaere, M. R.; Wipf, P.; Njardarson, J. T.; Zhu, J.; Schwaller, P. Strategy-first synthesis planning for complex natural products. *arXiv* **2026**. <https://doi.org/10.48550/arXiv.2608.07454>
- (29) Xuan-Vu, N.; Susanu, O.; Armstrong, D.; Schwaller, P. Mechanistic Reaction Prediction via Discrete Flow Matching on Graph-Structured Electron Occupation. *arXiv* **2026**. <https://arxiv.org/abs/2608.27429v2>

(30) Iwan, M.; Grisoni, F.; Garcia de Lomana, M.; Roncaglioni, A. 25 Years of Cardiac Ion Channel QSAR: From hERG Dominance to Multi-Channel Modeling. *Journal of Chemical Information and Modeling* **2026**. <https://doi.org/10.1021/acs.jcim.6c03072>

(31) Sigmund, L. M.; Assante, M.; Ball, M.; Kabeshov, M. BONAFIDE: a Python framework for the calculation of local features for atoms and bonds in molecules. *Journal of Cheminformatics* **2026**. <https://doi.org/10.1186/s13321-026-01283-6>