

Xufeng Zhao / Ph.D. Candidate

📍 Hamburg, Germany ✉ xufeng.zhao@uni-hamburg.com 🔗 <https://xf-zhao.github.io>
 in xufeng-zhao-961a82286 🎓 Google Scholar: sLwQ22MAAAAJ 🌐 xf-zhao

SUMMARY

I am an AI scientist and engineer with interests in machine learning, generative AI (AIGC), and embodied agents/robotics. With a solid interdisciplinary background: i.e. BSc in Electronic Information Engineering, MSc in Digital Signal Processing, and then 6+ years of AI research and engineering experience across industry and academia (PhD), my work spans large language models, multimodal fusion, reinforcement learning, and intelligent agents.

I approach research from a systems perspective, grounded in first principles, and seek deep theoretical insights to drive impactful, cross-disciplinary innovation.

EXPERTISE

- **Theoretical Methods** Information theory, Diffusion models, Probability theory, Reinforcement learning, Machine learning, Language models, Signal processing
- **Programming** Python (Numpy, Scipy, Scikit-learn, Pandas, Pytorch, etc), C, Matlab
- **Simulation** Isaac Sim, CoppeliaSim (PyRep), ThreeDWorld (Unity), ROS2

EDUCATION

Ph.D.	University of Hamburg , Artificial Intelligence (AI) Thesis title: “ <i>Environment Exploration and Autonomous Adaptation in Embodied Agents</i> ”	Hamburg, Germany Sept 2021 - Present
M.Sc.	University of Chinese Academy of Sciences , Signal Processing Thesis title: “ <i>Multipath Clutter Suppression and Multi-Frame Signal Processing for Passive Radar</i> ”	Beijing, China Sept 2015 - June 2018
B.E.	Xidian University , Electronic Information Engineering (EIE) Thesis title: “ <i>Design of Digital Filters Using the Least Squares Method</i> ”	Xi'an, China Sept 2010 - June 2014

EXPERIENCE

Research Associate University of Hamburg, Department of Informatics Advisor: Prof. Stefan Wermter	Hamburg, Germany Mar 2021 - Present
<ul style="list-style-type: none"> • Co-lead research on curriculum-RLAIF for post-training of LLMs, exploring LLM alignment with curricula. (under review) • Collaborative research on a diffusion-based framework for joint protein surface and structure design. (under review) • Co-research <i>reasoning reward models</i> for LLM personalization. (under review) • Agentic skill discovery framework for embodied agents to explore skills automatically. (CoRL 2024 Workshop & ICRA@40) • Logic-enhanced Chain-of-Thought (inference-time scaling) sampling method to improve reasoning abilities of LLMs. (COLING 2024) • Multimodal perception and decision-level fusion for embodied agents. (IROS 2023) • Co-lead research on <i>reinforcement learning (RL) with inductive rewards</i>, explored how jointly training of a reward model and an RL agent introduces noise and bias, and how to mitigate it. (ICML 2023) • Collaborative research on <i>RL with deductive rewards</i>, which explores how LLMs can be employed to guide RL agents for efficient learning. (CoRL 2023 Workshop) 	



- Research on multimodal perception and feature-level fusion for intrinsically motivated RL agents to perform active exploration and policy warm-up for downstream tasks. (*IROS 2022*)
- Other Collaborative works on active learning, bimanul robot task planning, XAI, etc. (TMLR, Humanoids, ICANN, etc.)

AI Engineer | JD.COM, Core AI Division (originally AI & Big Data, later JD New Retail)
Leveraging JD.com's extensive e-commerce platform and big data infrastructure, I contributed to multiple projects aimed at driving profit growth, reducing operational costs, and enhancing user experience. Key projects included:

- Sales forecasting and product recommendation for smart vending machines (*time-series forecasting & recommender systems*)
- Medical information recognition for smart medical boxes (*computer vision*)
- Automated customer service system for after-sales orders and inquiries (*natural language processing & machine learning*)
- Harmful content detection in real-time messaging between merchants and customers (*natural language processing*)
- Quality assessment and ranking of consumer reviews with images and text (*computer vision*)

Beijing, China
 July 2018 - June 2020



Graduate Researcher | Institute of Electronics, Chinese Academy of Sciences
 Advisor: Prof. Daojing Li

- Conducted research on multipath suppression, focusing on signal denoising and filtering techniques to improve passive radar signal quality.
- Investigated multi-frame radar target processing methods for enhanced detection and tracking performance.

Beijing, China
 July 2015 - June 2018



Signal Processing Engineer | ExtantFuture.com (Wearable Device for Pregnant Women)

- Developed a signal processing system for real-time monitoring of fetal heart rate, fetal movement, and maternal activity.
- Contributed to the mobile app that provides continuous fetal health insights, supporting comprehensive care for both mother and baby.

Beijing, China
 Jan 2015 - June 2015



PUBLICATIONS

- Mengdi Li*, Jiaye Lin*, **Xufeng Zhao**, Wenhao Lu, Peilin Zhao, Stefan Wermter, and Di Wang. "Curriculum-RLAIF: Curriculum Alignment with Reinforcement Learning from AI Feedback". In: arXiv:2505.20075 (under review)
- **Xufeng Zhao**, Cornelius Weber, and Stefan Wermter. "Agentic Skill Discovery". In: 8th Conference on Robot Learning (CoRL 2024) Workshop on Language and Robot Learning: Language as an Interface (LangRob), Munich, Germany. Aug. 2024. & In: ICRA@40.
- Kun Chu, **Xufeng Zhao**, Cornelius Weber, and Stefan Wermter. "LLM+MAP: Bimanual Robot Task Planning Using Large Language Models and Planning Domain Definition Language". In: arXiv:2503.17309 (under review)
- Wenhao Lu, **Xufeng Zhao**, Josua Spisak, Jae Hee Lee, and Stefan Wermter. "Mental Modelling of Reinforcement Learning Agents by Language Models". In: Transactions on Machine Learning Research (TMLR) (2025), pp. 2835-8856
- Honggen Zhang, **Xufeng Zhao**, Igor Molybog, and June Zhang. "REAL: Response Embedding-Based Alignment for LLMs". In: 34th International Joint Conference on Artificial Intelligence (IJCAI 2025) Workshop on Causal Learning RecSys.
- **Xufeng Zhao**, Mengdi Li, Wenhao Lu, Cornelius Weber, Jae Hee Lee, Kun Chu, and Stefan Wermter. "Enhancing Zero-Shot Chain-of-Thought Reasoning in Large Language Models through Logic". In: Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024). Ed. by Nicoletta Calzolari, Min-Yen Kan, Veronique Hoste, Alessandro Lenci, Sakriani Sakti, and Nianwen Xue.

Torino, Italia, May 2024, pp. 6144-6166

- Xiaowen Sun, **Xufeng Zhao**, Jae Hee Lee, Wenhao Lu, Matthias Kerzel, and Stefan Wermter. “*Details Make a Difference: Object State-Sensitive Neurorobotic Task Planning*”. In: Artificial Neural Networks and Machine Learning (ICANN 2024). Ed. by Michael Wand, Kristína Malinová, Jürgen Schmidhuber, and Igor V. Tetko. Cham: Springer Nature Switzerland, 2024. pp. 261-275
- Kun Chu, **Xufeng Zhao**, Cornelius Weber, Mengdi Li, Wenhao Lu, and Stefan Wermter. “*Large Language Models for Orchestrating Bimanual Robots*”. In: The 2024 IEEE-RAS International Conference on Humanoid Robots (Humanoids 2024). 2024.
- Wenhao Lu, **Xufeng Zhao**, Thilo Fryen, Jae Hee Lee, Mengdi Li, Sven Magg, and Stefan Wermter. “*Causal State Distillation for Explainable Reinforcement Learning*”. In: Proceedings of the Third Conference on Causal Learning and Reasoning (CLeaR 2024). Ed. by Francesco Locatello and Vanessa Didelez. Vol. 236. Proceedings of Machine Learning Research. PMLR, Apr. 2024, pp. 106-142
- Kun Chu, **Xufeng Zhao**, Cornelius Weber, Mengdi Li, and Stefan Wermter. “*Accelerating Reinforcement Learning of Robotic Manipulations via Feedback from Large Language Models*”. In: 7th Conference on Robot Learning (CoRL 2023) Workshop on Bridging the Gap Between Cognitive Science and Robot Learning in the Real World: Progresses and New Directions (CRL WS), Atlanta, Georgia USA. 2024
- **Xufeng Zhao**, Mengdi Li, Cornelius Weber, Muhammad Burhan Hafez, and Stefan Wermter. “*Chat with the Environment: Interactive Multimodal Perception Using Large Language Models*”. In: 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2023). IEEE, 2023, pp. 3590-3596
- **Xufeng Zhao***, Mengdi Li*, Jae Hee Lee, Cornelius Weber, and Stefan Wermter. “*Internally Rewarded Reinforcement Learning*”. In: Proceedings of the 40th International Conference on Machine Learning (ICML 2023). Ed. by Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett. Vol. 202. Proceedings of Machine Learning Research. PMLR, July 2023, pp. 20556-20574
- Wenhao Lu, **Xufeng Zhao**, Sven Magg, Martin Gromniak, Mengdi Li, and Stefan Wermter. “*A Closer Look at Reward Decomposition for High-Level Robotic Explanations*”. In: 2023 IEEE International Conference on Development and Learning (ICDL 2023). 2023.
- **Xufeng Zhao**, Cornelius Weber, Muhammad Burhan Hafez, and Stefan Wermter. “*Impact Makes a Sound and Sound Makes an Impact: Sound Guides Representations and Explorations*”. In: 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2022). IEEE, 2022, pp. 2512-2518
- Tingting Wang, **Xufeng Zhao**, Qiujian Lv, Bo Hu, and Degang Sun. “*Density weighted diversity based query strategy for active learning*.” In 2021 IEEE 24th International Conference on Computer Supported Cooperative Work in Design (CSCWD 2021), pp. 156-161. IEEE, 2021.

Google Scholar: <https://scholar.google.com/citations?user=sLwQ22MAAAAJ>

PRESENTATIONS

- “*Agentic Skill Discovery*”, **Poster**, ICRA@40 2024, Rotterdam, Netherlands; CoRL 2024 Workshop, Munich, Germany.
- “*Enhancing Zero-Shot Chain-of-Thought Reasoning in Large Language Models through Logic*”, **Oral + Poster**, LREC-COLING 2024, Turin, Italy
- “*Intelligent Agents: from Reinforcement Learning to Large Language Models*”, **Invited Talk**, a robotic company, 2023, Online.
- “*Chat with the Environment: Interactive Multimodal Perception Using Large Language Models*”, **Oral + Poster**, IROS 2023, Detroit, USA
- “*Internally Rewarded Reinforcement Learning*”, **Poster**, ICML 2023, Hawaii, USA
- “*Impact Makes a Sound and Sound Makes an Impact: Sound Guides Representations and Explorations*”, **Oral**, IROS 2022, Kyoto, Japan

AWARDS & GRANTS

- **ICRA@40 Travel Grants**, IEEE Robotics and Automation Society 2024
- **Researcher Access Program**, OpenAI 2024
- **Academic Scholarship**, University of Chinese Academy of Sciences 2015 - 2018
- **Excellent Student Cadre**, University of Chinese Academy of Sciences 2015
- **Triple-A Student**, University of Chinese Academy of Sciences 2015
- **The Second Prize of China Undergraduate Mathematical Contest in Modeling**, China Society for Industrial and Applied Mathematics 2013
- **The Second Prize of Xinghuo Electronic Competition**, Xidian University 2012
- **The First Prize of Xidian's Mathematical Contest in Modeling**, Xidian University 2012

TEACHING

Lectures

- Lecture on *Large Language Models*, Master Course on Neural Networks, University of Hamburg, 2024. (Co-lecture with Dr. [Jae Hee Lee](#).)

Thesis Supervision

- “*Task-Agnostic Policy Distillation: Continual Deep Reinforcement Learning with Alternating Self-Supervised Prediction*”, Kerim Erekmén, BSc thesis, 2023, University of Hamburg.

Seminar Supervision

- “*Robust RGB-D to 3D mesh Construction for Robotic Simulation*”, Neural Networks Seminar 2024, University of Hamburg
- “*LLM Fine-tuning with News Data*”, Bio-inspired Artificial Intelligence Seminar 2023, University of Hamburg
- “*Survey on Deployable LLMs*”, Neural Networks Seminar 2023, University of Hamburg
- “*Unsupervised Skill Discovery Implementation*”, Bio-inspired Artificial Intelligence Seminar 2022, University of Hamburg
- “*Survey on Transformers in Reinforcement Learning*”, Bio-inspired Artificial Intelligence Seminar 2022, University of Hamburg
- ...

PEER REVIEWER

Robotics

- IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2023 - Present
- International Conference on Ubiquitous Robots (UR) 2025 - Present
- Robotics and Autonomous Systems (RAS) 2025 - Present
- IEEE Robotics and Automation Letters (RA-L) 2025 - Present
- IEEE International Conference on Robotics and Automation (ICRA) 2025 - Present
- IEEE-RAS International Conference on Humanoid Robots (Humanoids) 2025 - Present

Artificial Intelligence and Machine Learning

- International Conference on Learning Representations (ICLR) 2025 - Present
- PeerJ Computer Science 2023

Natural Language Processing

- Empirical Methods in Natural Language Processing (EMNLP) 2025 - Present
- International Conference on Computational Linguistics (COLING) 2024