NLPCC 2026

THE 15TH CCF INTERNATIONAL CONFERENCE ON NATURAL LANGUAGE PROCESSING AND CHINESE COMPUTING

Handbook

ORGANIZER CHINA COMPUTER FEDERATION

HOST UNIVERSITY OF MACAU

MACAU, CHINA
NOVEMBER 3-5, 2026



中国计算机学会 CHINA COMPUTER FEDERATION

澳門大學 UNIVERSIDADE DE MACAU UNIVERSITY OF MACAU

Primary Sponsors

FinVolution

中业科技 zoneyet

Diamond Sponsors

记忆张量 MemTensor

京东

蚂蚁集团 ANT GROUP

Platinum Sponsors

Tencent 腾讯

荔枝

Golden Sponsors

vivo

Baidu 百度

海天瑞声

復旦大學 FUDAN UNIVERSITY

深绎未来 深译科技 DEEPFUTUREAI DEEPTRANXAI



Contents

About NLPCC 1
Preface by the Program Committee Co-Chairs 2
Organizing Committee 3
Program Committee 4
Venue and Floorplan 6
Program Overview 8
Detailed Program 9
Keynote Talks 12
- Keynote Talk 1 - Advances in Multilingual Large Language Models 12
- Keynote Talk 2 - The Verification Gap: Why Human Participation Persists in an Automated Science 13
- Keynote Talk 3 - Listening the Chinese Way 13
- Keynote Talk 4 - Exploring the Frontier of Scaling Laws for LLMs ... 15

Tutorials 16
Industry Talks 23
Workshops 30
Oral Sessions 43
Poster Sessions 53
University of Macau 64

About NLPCC

The CCF International Conference on Natural Language Processing and Chinese Computing (NLPCC) is the annual meeting of CCF-NLP (Technical Committee of Natural Language Processing, China Computer Federation). NLPCC is a leading international conference specialized in the fields of Natural Language Processing (NLP) and Chinese Computing (CC). NLPCC is in the list of International Academic Conferences and Periodicals Recommended by CCF, and selected for the CCF High-Quality International Academic Conference Promotion Program. It serves as a main forum for researchers and practitioners from academia, industry, and government to share their ideas, research results and experiences, and to promote their research and technical innovations in the fields. Previous NLPCC conferences were successfully held in Beijing (2012), Chongqing (2013), Shenzhen (2014), Nanchang (2015), Kunming (2016), Dalian (2017), Hohhot (2018), Dunhuang (2019), Zhengzhou (2020), Qingdao (2021), Guilin (2022), Foshan (2023), Hangzhou (2024), Urumqi (2025) and Macao (2026).

Preface by the Program Committee Co-Chairs

Welcome to NLPCC 2026, the fifteenth CCF International Conference on Natural Language Processing and Chinese Computing. Following the success of previous conferences held in Beijing (2012), Chongqing (2013), Shenzhen (2014), Nanchang (2015), Kunming (2016), Dalian (2017), Hohhot (2018), Dunhuang (2019), Zhengzhou (2020), Qingdao (2021), Guilin (2022), Foshan (2023), Hangzhou (2024), and Urumqi (2025), this year's NLPCC will be held in Macau, China, from 3 to 5 November 2026. Organized by CCF-NLP (Technical Committee of Natural Language Processing, China Computer Federation), NLPCC is a leading international conference specialized in the fields of Natural Language Processing (NLP) and Chinese Computing (CC). NLPCC is in the list of International Academic Conferences and Periodicals Recommended by CCF, and is selected for the CCF High-Quality International Academic Conference Promotion Program. It serves as a main forum for researchers and practitioners from academia, industry, and government to share their ideas, research results and experiences, and to promote their research and technical innovations in the fields.

The fields of natural language processing (NLP) and Chinese computing (CC) have remained among the most active areas of research and development, driven by advances in large language models and their rapid adoption across academia and industry. Following NLPCC's tradition, NLPCC 2026 welcomes original submissions in a wide range of areas, including but not limited to Large Language Models, LLM Agents, AI for Science, AI for Software Engineering, Computational Social Science and Social Media, Machine Translation and Multilinguality, Information Extraction and Knowledge Acquisition, Question Answering, Sentiment Analysis, Speech and Multimodality, and Resources and Evaluation. This year, we received 790 valid submissions to the main conference. After a thorough reviewing process, 262 papers were finally accepted, 78 of which will be presented as oral papers and 184 as poster papers.

Alongside the main conference oral and poster sessions, NLPCC 2026 presents four keynote talks, by Philipp Koehn (Johns Hopkins University), Iryna Gurevych (Technical University of Darmstadt), Haizhou Li (The Chinese University of Hong Kong, Shenzhen), and Junchi Yan (Shanghai Jiao Tong University); industry talks; and eleven shared tasks with an evaluation track, together with tutorials and workshops.

We are honored to have these distinguished speakers share their findings with us, and we would like to thank all the people who have contributed to NLPCC 2026. First of all, we would like to thank the area chairs for their hard work in recruiting reviewers, monitoring the review process, and carefully rating submissions, and we thank all the reviewers for their time and efforts. We are also grateful for the help and support from our shared task organizers, tutorial and workshop organizers, the general chair, the organizing committee, the publication chairs, and our sponsors. We greatly appreciate all your help!

We are happy to see you at NLPCC 2026 in Macau and hope you enjoy the conference!

Roy Ka-Wei Lee, Singapore University of Technology and Design

Qi Zhang, Fudan University

Meishan Zhang, Harbin Institute of Technology (Shenzhen)

NLPCC 2026, Program Committee Chairs

Organizing Committee

General Chair:

Yue Zhang (Westlake University)

Program Committee Chairs:

Roy Ka-Wei Lee (Singapore University of Technology and Design)

Qi Zhang (Fudan University)

Meishan Zhang (Harbin Institute of Technology, Shenzhen)

Workshop Chairs:

Tieyun Qian (Wuhan University)

Yang Zhao (Institute of Automation, Chinese Academy of Sciences)

Evaluation Chairs:

Jie Yang (Harvard University)

Kunli Zhang (Zhengzhou University)

Tutorial Chairs:

Zhaochun Ren (Leiden University)

Wenbin Jiang (Beijing Normal University)

Xiting Wang (Renmin University of China)

Publicity Chairs:

Buzhou Tang (Harbin Institute of Technology, Shenzhen)

Lizi Liao (Singapore Management University)

Poster Chairs:

Aishan Wumaier (Xinjiang University)

Yile Wang (Shenzhen University)

Yanxia Qin (Singapore University of Technology and Design)

OpenReview Co-chairs:

Junwen Duan (Central South University)

Yulong Chen (University of Cambridge)

Chunhua Liu (The University of Melbourne)

Publication Chairs:

Jianxing Yu (Sun Yat-sen University)

Jin Wang (Yunnan University)

Organization Committee Chairs:

Derek F. Wong (University of Macau)

Haofen Wang (Tongji University)

Sponsorship Chairs:

Yan Song (University of Science and Technology of China)

Longyue Wang (Alibaba International)

Webmaster Chair:

Shudong Liu (University of Macau)

Program Committee

Program Committee Chairs:

Roy Ka-Wei Lee (Singapore University of Technology and Design)

Qi Zhang (Fudan University)

Meishan Zhang (Harbin Institute of Technology, Shenzhen)

Area Chairs:

Xuefeng Bai (Harbin Institute of Technology, Shenzhen)

RUI CAO (University of Cambridge)

Pinzhen Chen (Queen's University Belfast)

Yubo Chen (Minzu University of China)

Yang Deng (Singapore Management University)

Lan Du (Monash University)

Yue Feng (University of Birmingham)

Wei Gao (Singapore Management University)

Chen Gong (Soochow University)

Lin Gui (King's College London, University of London)

Shizhu He (Institute of automation, Chinese Academy of Sciences)

Ming Shan Hee (Mohamed bin Zayed University of Artificial Intelligence)

Qi Huang (Jiangxi Normal University)

Ruihong Huang (Texas A&M University)

Tao Ji (Fudan University)

Yuxiang Jia (Zhengzhou University)

Bobo Li (National University of Singapore)

Dongfang Li (Harbin Institute of Technology, Shenzhen)

Fei Li (Wuhan University)

Juntao Li (Soochow University, China)

Ying Li (Kunmimg University of Science and Technology)

Yongqi Li (The Hong Kong Polytechnic University)

Zhuang Li (Royal Melbourne Institute of Technology)

Lizi Liao (Singapore Management University)

Yankai Lin (Renmin University of China)

Bang Liu (Université de Montréal)

Xuebo Liu (Harbin Institute of Technology, Shenzhen)

Yao Lu (University College London)

Yunshan Ma (Singapore Management University)

Lili Mou (University of Alberta)

Chen Qian (Shanghai Jiaotong University)

Libo Qin (Harbin Institute of Technology, Shenzhen)

Lizhen Qu (Monash University)

Xiangqing Shen (Nanjing University)

Zhiliang Tian (National University of Defense Technology)

Chenglong Wang (Northeastern University)

Fanfan Wang (Jiangnan University)

Wenya Wang (Nanyang Technological University)

Yiwei Wang (University of California, Merced)

Shengqiong Wu (University of Oxford)

Zhen Wu (Nanjing University)

Zhijing Wu (Beijing Institute of Technology)

Qiongkai Xu (Macquarie University)

Dingyi Yang (Tianjin University)

Liang Yang (Dalian University of Technology)

Zheng Yuan (University of Sheffield)

Bo Zhang (Nanjing Normal University)

Li Zhang (Drexel University)

Wenxuan Zhang (Singapore University of Technology and Design)

Yaping Zhang (Institute of automation, Chinese Academy of Sciences)

Zhengchen Zhang (Singapore Institute of Technology)

Zhuosheng Zhang (Shanghai Jiao Tong University)

Ruochen Zhao (Singapore University of Technology and Design)

Yi Zhou (Cardiff University)

Qingfu Zhu (Harbin Institute of Technology)

Venue and Floorplan

Galaxy International Convention Center (GICC)

Address: Galaxy Macau, Cotai, Macau

Galaxy International Convention Center (GICC) is the latest addition to Galaxy Entertainment Group's ever-expanding integrated resort precinct and will usher in a new era for the MICE industry in Macau.



Transport

Galaxy Macau provides complimentary shuttle buses between Galaxy and major transport hubs in Macau, including the Border Gate, Macau International Airport, Taipa Ferry Terminal, Hengqin Port, and the Hong Kong-Zhuhai-Macao Bridge, etc.

Border Gate	To Galaxy Macau	09:00 – 23:30	Taipa Ferry Terminal	To Galaxy Macau	09:00 – 23:00
	From Galaxy Macau	09:00 – 23:30		From Galaxy Macau	09:15 – 23:15
	Average 5 - 15 minutes (Around 25 minutes drive)			Average 15-40 minutes (Around 20 minutes drive)	
	Border Gate → Diamond Lobby# → Border Gate			Taipa Ferry Terminal* → Crystal Lobby# → Taipa Ferry Terminal*	
Macau International Airport	To Galaxy Macau	10:00 - 21:00	Hengqin Port	To Galaxy Macau	09:00 - 22:00
	From Galaxy Macau	10:15 - 20:45		From Galaxy Macau	09:15 - 21:45
	Average 35 minutes (Around 20 minutes drive)			Average 5 – 10 minutes, Non-Peak Hours (17:00-22:00) 30 minutes (Around 10 minutes drive)	
	Macau International Airport → Crystal Lobby# → Macau International Airport			Hengqin Port → Crystal Lobby# → Hengqin Port	
Andaz Macau	To Galaxy Macau	09:00 – 02:00	Hong Kong Zhuhai Macau Bridge To Galaxy Macau	To Galaxy Macau	10:00 – 23:00
	From Galaxy Macau	09:00 – 02:00		From Galaxy Macau	10:00 – 23:00
	Average 5-10 minutes (Around 10 minutes drive)			Average 10 - 25 minutes (Around 25 minutes drive)	
	Andaz Macau → Crystal Lobby → Andaz Macau			Hong Kong Zhuhai Macau Bridge → Diamond Lobby → Hong Kong Zhuhai Macau Bridge	
StarWorld Hotel/Border Gate	To StarWorld Hotel	09:00 - 23:30	StarWorld Hotel / Macau Ferry Terminal	To Galaxy Macau	09:30-22:00
	From StarWorld Hotel	09:00 - 23:30		From Galaxy Macau	09:30-22:00
	Average 10-15 minutes			Average 10-20 minutes (Around 35 minutes drive)	
	Border Gate → StarWorld Hotel → Border Gate			Crystal Lobby → StarWorld Hotel → Macau Ferry Terminal → StarWorld Hotel → Crystal Lobby	

NLPCC 2026 Floorplan

2/F, Galaxy International Convention Center



Program Overview

Date	Time	Content
3 November (Tuesday)	All day	Tutorials
4 November (Wednesday) Main Conference and Workshop	8:45-09:15	Opening Ceremony (Session Chair: Derek Wong)
	9:15-10:15	Keynote I: Philipp Koehn
	10:15-10:30	Industry Talk I :Jun Zhou
	10:30-11:00	Coffee Break
	10:30-11:30	Poster Session I
	11:00-12:00	Keynote II: Haizhou Li
	12:00-12:15	Industry Talk II : Jie Li
	12:15-12:30	Industry Talk III :Nan Guan
	12:30-14:00	Networking Lunch
	14:00-16:00	GBA Workshop, Finvolution Workshop, Oral Session 1-3
	15:30-16:30	Poster Session 2
	16:00-16:15	Coffee Break
	16:15-18:15	GBA Workshop, Oral Session 4-6
	17:00-18:00	Poster Session 3
	18:30-21:00	Banquet
5 November (Thursday) Main Conference and Workshop	09:00-10:00	Keynote III: Iryna Gurevych
	10:00-11:00	NCMMSC Posters
	10:00-10:15	Industry Talk IV:Ye Liu
	10:15-10:30	Coffee Break
	10:30-11:30	Keynote IV : Junchi Yan
	11:30-11:45	Industry Talk V: Zhiyu Li
	11:45-12:00	Industry Talk VI :Zenan Xu
	12:00-13:00	Networking Lunch
	13:00-15:00	Evaluation Workshop 1, Oral Session 7-8, Speech & NLP Workshop
	15:00-15:15	Coffee Break
	15:15-17:15	Evaluation Workshop 2, Oral Session 9-10, Speech & NLP Workshop
	17:15-17:45	Best Paper Session
	17:45-18:15	Closing

Detailed Program

Tutorials — Tuesday, 3 November 2026

Time	Content	Speaker	Venue
08:45-10:15	Tutorial 1- The Neuron Mechanisms of Knowledge Transfer and Low-Resource Evolution in Multilingual Large Models	Shaolin Zhu & Kehai Chen	Meeting Room 18+19 2/F
10:15-10:30	Coffee Break		
10:30-11:00	Tutorial 2- Evaluating LLM Agents: Outcomes, Processes, and Open-Ended Research	Linyi Yang	
11:00-12:30	Tutorial 3- Foundation Models in Action: From Embodied to Social Intelligence	Hao Chen	
12:30-13:30	Lunch (Hall FG@LG1 Floor)		
13:30-15:00	Tutorial 4- Rethinking Information Extraction in the Era of Large Language Models	Xian-Ling Mao & Yi-Fan Lu	
15:00-15:30	Coffee Break		
15:30-17:00	Tutorial 5- Reinforced Generative Information Retrieval	Xin Xin	
17:00-18:30	Tutorial 6- Is Parsing Dead? How Syntactic Analysis Survives in the LLM Era	Yuanhe Tian	

Main Conference — Wednesday, 4 November 2026

Time	Content	Venue
08:45-09:15	**Opening Ceremony** Host: Derek Wong	Auditorium, 2/F
09:15-10:15	**Keynote Speech 1** Philipp Koehn (Johns Hopkins University) *Advances in Multilingual Large Language Models*	
10:15-10:30	**Industry Talk 1** Jun Zhou (Ant Group) *The Intelligence-Efficiency Ratio: Open Intelligence for the Real World*	
10:30-11:00	Group Photo & Tea Break	
10:30-11:30	Poster Session 1	
11:00-12:00	**Keynote Speech 2** Haizhou Li (The Chinese University of Hong Kong, Shenzhen) *Listening the Chinese Way*	
12:00-12:15	**Industry Talk 2** Jie Li (Zhengzhou Zoneyet Technology) *Next-Generation Agent Networks*	
12:15-12:30	**Industry Talk 3** Jiaqi Wang (JD Explore Academy) *The Evolution of AI: From Understanding and Generation to Interaction and Action*	
12:30-14:00	Lunch (Hall FG @LG1 Floor)	
14:00-16:00	Greater Bay Area Workshop @Auditorium	Auditorium Meeting Room 1 Meeting Room 4+5 Meeting Room 8 Meeting Room 9 2/F
14:00-16:00	FinVolution Global Data Science Competition Workshop @Meeting Room 1	
14:00-16:00	Oral Session 1 @Meeting Room 4+5	
14:00-16:00	Oral Session 2 @Meeting Room 8	
14:00-16:00	Oral Session 3 @Meeting Room 9	
16:00-16:15	Coffee Break	
15:30-16:30	Poster Session 2	
16:15-18:15	Greater Bay Area Workshop @Auditorium	
16:15-18:15	Oral Session 4 @Meeting Room 4+5	
16:15-18:15	Oral Session 5 @Meeting Room 8	
16:15-18:15	Oral Session 6 @Meeting Room 9	
17:00-18:00	Poster Session 3	
18:30-21:00	Conference Banquet	

Main Conference — Thursday, 5 November 2026

Time	Content	Venue
09:00-10:00	**Keynote Speech 3** Iryna Gurevych (Technical University of Darmstadt) *The Verification Gap: Why Human Participation Persists in an Automated Science*	Auditorium, 2/F
10:00-10:15	**Industry Talk 4** Ye Liu, (Lizhi Inc., part of Sound Group Inc.) *Technical Exploration of Sound AI Composite Models in Audio Internet Scenarios*	
10:00-11:00	NCMMSC Poster Session	
10:15-10:30	Coffee Break	
10:30-11:30	**Keynote Speech 4** Junchi Yan (Shanghai Jiao Tong University) *Exploring the Frontier of Scaling Laws for LLMs*	
11:30-11:45	**Industry Talk 5** Zhiyu Li (MemTensor) *From Memory Operating Systems to Memory-Native Models: MemOS and MemT 1.0*	
11:45-12:00	**Industry Talk 6** Zenan Xu (Tencent) *Agents That Remember, Anticipate, and Collaborate*	
12:00-13:00	Lunch (Hall FG@LG1 Floor)	
13:00-15:00	Evaluation workshop I @Auditorium	Auditorium Meeting Room 1 Meeting Room 8 Meeting Room 9 2/F
14:00-15:00	Speech and NLP Workshop @Meeting Room 1	
13:00-15:00	Oral Session 7 @Meeting Room 8	
13:00-15:00	Oral Session 8 @Meeting Room 9	
15:00-15:15	Coffee Break	
15:15-17:15	Evaluation workshop II @Auditorium	
15:15-17:20	Speech and NLP Workshop @Meeting Room 1	
15:15-17:15	Oral Session 9 @Meeting Room 8	
15:15-17:15	Oral Session 10 @Meeting Room 9	
17:15-17:45	Best Paper Session @Auditorium	
17:45-18:15	Closing @Auditorium	

Keynote Talks

Keynote Talk 1 - Advances in Multilingual Large Language Models

Wednesday, 4 November 2026 · 09:15-10:15 @Auditorium



Philipp Koehn

Professor (Johns Hopkins University)

Abstract: Large language models have taken over the field of natural language processing. However, they are not equally accessible across all languages. Frontier models tend to focus on a couple of dozen commercially relevant languages. The talk will report on research showing how well knowledge is transferred across languages in these models, how words in different languages are represented, and how to make them more multilingual.

Biography: Philipp Koehn is a professor in computer science at Johns Hopkins University. He is a leading researcher in the field of machine translation, published two textbooks in the field and over 200 research papers as well as numerous datasets and open source tools. His research extends also to mulitlinguality in large language models, speech translation, and neuro-symbolic processing. Prof. Koehn received the Award of Honor from the International Association for Machine Translation in 2015 and was named ACL Fellow in 2024. He was the 2025 ACL General Chair.

Keynote Talk 2 - Listening the Chinese Way

Wednesday, 4 November 2026 • 11:00-12:00 @Auditorium



Haizhou Li

Professor (The Chinese University of Hong Kong, Shenzhen)

Abstract: Humans have a remarkable ability to pay their auditory attention only to a sound source of interest, that we call selective auditory attention, in a multi-talker environment or a Cocktail Party. As discovered in neuroscience and psychoacoustics, the auditory attention is achieved by a modulation of top-down and bottom-up attention. However, signal processing approach to target speaker extraction in a multi-talker environment remains a challenge. Human listening is more of an art than a science. The Chinese character for the word 'listening' carries the ancient wisdom about important aspects of what people believe it means to really 'listen' to someone, i.e. hearing, seeing, and understanding. In this talk, we study the deep learning solutions to target speaker extraction that enable selective auditory attention. We review the findings from human audio-visual speech perception to motivate the design of speech perception algorithms. We will also discuss the computational auditory models, technical challenges and the recent advances in the field.

Biography: Haizhou Li is a Presidential Chair Professor and Dean at the School of Artificial Intelligence, CUHKSZ. He is also an adjunct professor at NUS. He earned his B.Sc., M.Sc., and Ph.D. from South China University of Technology. A veteran in speech and language technology since 1988, Dr. Li has held global academic posts and significant research/industry leadership roles, including at Apple-ISS Research Centre, Lernout & Hauspie, and the Institute for Infocomm Research (I2R), Singapore. He co-founded the Baidu-I2R Research Centre. His technical contributions include award-winning products like Apple's Chinese Dictation Kits and major deployments like Baidu Music Search. His research interests are ASR (Automatic Speech Recognition), NLP (Natural Language Processing), and IR (Information Retrieval). He has held numerous prestigious service positions, including Editor-in-Chief of IEEE/ACM Transactions on Audio, Speech and Language Processing, and President of International Speech Communication Association (ISCA) and Asia Pacific Signal and Information Processing Association (APSIPA). He has chaired major conferences like ACL (Association for Computational Linguistics) 2012, INTERSPEECH 2014, and ICASSP (International Conference on Acoustics, Speech, and Signal Processing) 2022. Dr. Li is an IEEE Fellow, an ISCA Fellow, a Member of Academy of Engineering Singapore, and a recipient of Singapore's President's Technology Award (2013).

Keynote Talk 3 - The Verification Gap: Why Human Participation Persists in an Automated Science

Thursday, 5 November 2026 · 09:00-10:00 @Auditorium



Iryna Gurevych
Professor (Technical University of Darmstadt)

Abstract: The dominant narrative around "AI scientists" treats scale as the main remaining obstacle to automated science. This talk argues that scale alone will not close the gap: the binding constraint is science verification, understood as the capacity for scrutiny. As AI produces scientific claims and outputs at growing speed and volume, our ability to check what it has produced — to detect a contradiction between introduction and results, to catch a paper–code discrepancy, to ground a review comment in the underlying evidence — becomes the real bottleneck. And it is a bottleneck from which the human cannot be removed without loss. Drawing on recent work from our lab, I characterize this science verification gap; for instance, even frontier LLMs detect fewer than half of real paper–code discrepancies. I then present a set of complementary tools designed to support science verification: ProReviewer (proactive AI review), ReGround (evidence-based grounding of review comments in multimodal content), LazyReviewPlus (reviewer augmentation), and REVAS (deployed at ACL). The talk closes with a theoretical framing: human participation in science verification is constitutive, because the very targets of verification emerge through that participation.

Biography: Iryna Gurevych is a full professor of computer science at the Technical University of Darmstadt, Germany, where she founded and heads the Ubiquitous Knowledge Processing (UKP) Lab. She is a pioneer of computational argumentation and cross-document text modeling, with landmark applications to fact verification and misinformation detection, and her research extends to large language models and interdisciplinary NLP for the social sciences and humanities. Prof. Gurevych served as President of the Association for Computational Linguistics (ACL) in 2023 and is an ACL Fellow, a member of the German National Academy of Sciences Leopoldina and Academia Europaea, and a recipient of the 2025 Royal Society Milner Award.

Keynote Talk 4 - Exploring the Frontier of Scaling Laws for LLMs

Thursday, 5 November 2026 • 10:30-11:30 @Auditorium



Junchi Yan

Professor (Shanghai Jiao Tong University)

Abstract: This talk shows our recent progress on improved LLM architectures, namely JTok and xHC that push the frontier of scaling laws. Both stem from an insight: model capacity can be substantially boosted with marginal compute overhead, beyond conventional width-depth scaling. JTok augments the token dimension: each token carries a learnable per-layer memory vector injected via lookup-and-light-modulation side branches, leaving the backbone intact. Evaluated at 61B total parameters (17B base + 44B memory parameters) with 570B training tokens, it achieves ~35% training-FLOPs savings for equal performance, at merely 6.8% throughput loss. On the other hand, xHC advances along the residual-flow dimension. To break the long-standing (N=4) limit of Hyper-Connections variants, we tackle two bottlenecks — insufficient write-back information and cubic-scaling residual-mapping cost — via temporal feature enhancement and sparse-residual-flow (dense-read, sparse-write) design, pushing expansion factor to (N=16). On 18B / 28B MoE models, it yields +4.0 / +3.1 average gains over DeepSeek-mHC across twelve downstream benchmarks with only 4.1% / 3.0% extra training FLOPs, delivering ~33% compute savings versus vanilla Transformers under matching-performance scaling-law conditions. Finally, we compare JTok and xHC dimension-wise and discuss their orthogonal composition potential.

Biography: Junchi Yan is a tenured professor at the School of Artificial Intelligence, Shanghai Jiao Tong University, where he directs the SJTU-ReThinkLab. His research centers on machine learning, in particular learning-based combinatorial optimization, graph matching, and AI for science, with applications in computer vision and autonomous driving. Prof. Yan is an IAPR Fellow, a recipient of the IEEE CIS Outstanding Early Career Award and the IEEE Computer Society AI's 10 to Watch, and in 2025 was elected the first ICML board member from mainland China. He serves on the editorial boards of IEEE TPAMI and JMLR.

Tutorials

Tuesday, 3 November 2026 • 08:45-18:30 @Meeting Room 18+19

Time	Content	Speaker
08:45-10:15	Tutorial 1: The Neuron Mechanisms of Knowledge Transfer and Low-Resource Evolution in Multilingual Large Models	Shaolin Zhu & Kehai Chen
10:30-11:00	Tutorial 2: Evaluating LLM Agents: Outcomes, Processes, and Open-Ended Research	Linyi Yang
11:00-12:30	Tutorial 3: Foundation Models in Action: From Embodied to Social Intelligence	Hao Chen
13:30-15:00	Tutorial 4: Rethinking Information Extraction in the Era of Large Language Models	Xian-Ling Mao & Yi-Fan Lu
15:30-17:00	Tutorial 5: Reinforced Generative Information Retrieval	Xin Xin
17:00-18:30	Tutorial 6: Is Parsing Dead? How Syntactic Analysis Survives in the LLM Era	Yuanhe Tian

Tutorial Chairs



Biography: Zhaochun Ren is an Associate Professor at Leiden University, the Netherlands. He is interested in information retrieval and natural language processing, with an emphasis on generative information retrieval and recommender systems. He aims to develop intelligent agents that can address complex user requests and solve core challenges in NLP and IR towards that goal. His research has been recognized with multiple awards at RecSys, SIGIR, WSDM, EMNLP, and CIKM. Prior to joining Leiden, he was a Professor at Shandong University and a Research Scientist at JD.com.



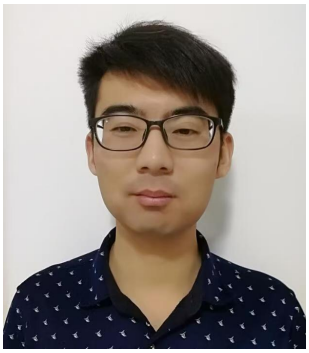
Biography: Wenbin Jiang is a Professor at the School of Artificial Intelligence, Beijing Normal University. His research focuses on large language models and general reasoning, as well as knowledge graphs and natural language processing. He has presided over and participated in multiple research projects including those supported by the National Natural Science Foundation of China and the 863 Program of the Ministry of Science and Technology. He has published more than 50 academic papers and holds over 20 authorized patents. Some of his research outcomes have been deployed in many institutions and well-known enterprises.



Biography: Xiting Wang is an Associate Professor at the Gaoling School of Artificial Intelligence, Renmin University of China. She previously served as a Principal Researcher at Microsoft Research Asia. She received the Wu Wenjun Artificial Intelligence Award for Young Scientists and the Second Prize of the CCF Natural Science Award. Her research focuses on the interpretability and alignment of large language models. Her research has been deployed in Alipay's Bailing large language model, Microsoft Bing, Outlook, and MSN. She serves as an Area Chair for ICML, NeurIPS, ICLR, and AAAI.

Tutorial Speakers

Tutorial 1: The Neuron Mechanisms of Knowledge Transfer and Low-Resource Evolution in Multilingual Large Models



Shaolin Zhu
Tianjin University

Abstract: The current mainstream paradigm of big language models highly relies on the brute force stacking of computing power and massive corpora, but in multilingual, especially low resource language scenarios, this method encounters data bottlenecks and performance ceilings. Due to a lack of understanding of the internal micro mechanisms of the model, existing cross language transfer and fine-tuning often lead to catastrophic forgetting or concept drift. The billions of parameters inside the large model are not chaotic black boxes. How do feedforward neural networks, as key value memory units of knowledge, encode concepts with single neurons or even neural circuits? How is factual knowledge in one language projected onto another language at the microscale? Which neurons are language independent abstract semantic centers, and which are language specific grammatical surfaces? We aim to break away from the homogenized training optimization and agent pipeline, and delve deeper into the micro neuron level of large models. We will analyze the operating mechanism of multilingual models at the scale of single neurons and subnetworks based on the latest advances in mechanism interpretability, and use this as a scientific basis to speculate optimization solutions for multilingual and low resource languages.

Biography: Shaolin Zhu is currently an Associate Research Professor and Ph.D. Supervisor at Tianjin University. His research interests include artificial intelligence and natural language processing. He has published more than 50 papers in leading international conferences and journals, including ACL, AAAI, KDD, Science China information science, and TASLP. He has served as an Area Chair or Senior Area Chair for several major international conferences, including ACL and COLING, and received the Outstanding Paper Award at COLING 2024.

Co-speaker: Kehai Chen, Harbin Institute of Technology

Tutorial 2: Evaluating LLM Agents: Outcomes, Processes, and Open-Ended Research



Linyi Yang
Southern University of Science and Technology

Abstract: As large language models evolve into agents that plan, use tools, and interact with dynamic environments, evaluating their capabilities requires more than measuring answer accuracy. This tutorial introduces three complementary perspectives on LLM agent evaluation: outcome-based evaluation, process-level evaluation, and evaluation of open-ended tasks. Case studies of Memento and Research Claw Arena illustrate how these perspectives apply to agents that learn from experience and support scientific research, including idea generation and experimental design. Throughout the tutorial, we highlight the limitations of individual metrics and explain how complementary evaluation methods can reveal failures hidden by aggregate success rates. Participants will gain a practical framework for designing agent evaluations that connect measurable performance with reliable behavior and meaningful user value.

Biography: Linyi Yang is an Assistant Professor and Ph.D. supervisor at Southern University of Science and Technology (SUSTech), where he leads the GenAI Lab. He previously conducted research at University College London (UCL) and Westlake University. His research focuses on AI scientists and intelligent agents. His work has received over 14,000 citations on Google Scholar, and he serves as an Area Chair for leading conferences, including NeurIPS, ICLR, ICML, and ACL.

Tutorial 3: Foundation Models in Action: From Embodied to Social Intelligence



Hao Chen

City University of Macau

Abstract: Foundation models are opening new possibilities for AI systems that perceive their surroundings, act in physical environments, and interact with people and other agents. These developments connect Embodied AI with social intelligence, bringing new opportunities for building agents that can adapt to changing environments and participate in increasingly complex interactions.

This tutorial examines foundation-model applications along this progression. We begin with Embodied AI, exploring how multimodal understanding and reasoning support navigation, visual tracking, and action planning. We then extend the discussion to social intelligence, examining how agents interpret intentions, model behavior, and communicate and cooperate with others. Across these settings, we discuss the roles of memory, environmental feedback, and learning from interaction, together with challenges in generalization, long-horizon decision-making, and evaluation.

To connect these ideas with practice, the tutorial draws on representative applications and research platforms, including UnrealZoo, which provides photorealistic virtual environments for studying embodied perception and multi-agent interaction. These examples illustrate how simulation can support the development and evaluation of interactive agents and help investigate the relationship between physical competence and social behavior. The tutorial aims to give researchers and practitioners a practical perspective on foundation models in action and identify research opportunities at the intersection of Embodied AI and social intelligence.

Biography: Hao Chen is an assistant professor at the Faculty of Data Science, City University of Macau. Before that, he worked as a postdoctoral research fellow at The Hong Kong Polytechnic University. He received his Ph.D degree from Beihang University, China in 2022. He has published over 40 papers, such as TKDE, NeruIPS, ACL, SIGKDD, WWW, and SIGIR, etc. He received the Best Paper Award Honorable Mention at SIGIR 2023. His research interests include Embodied AI, Robotics, and Billion-scale Recommendation Systems.

Tutorial 4: Rethinking Information Extraction in the Era of Large Language Models



Xian-Ling Mao

Beijing Institute of Technology

Abstract: Large language models (LLMs) are reshaping how information extraction (IE) is performed and prompting a reassessment of its role. As general-purpose models can perform a wide range of extraction tasks and process documents directly, a central question is where dedicated IE research can continue to create value. This tutorial examines IE in the LLM era through three questions. First, how do LLMs change established assumptions about tasks, schemas, supervision, and benchmarks? We review representative evidence on the capabilities and limitations of general-purpose extraction. Second, when is explicit structure worth constructing and maintaining? Drawing on studies of structured retrieval, memory, and domain-specific applications, we examine when extracted representations provide benefits over using text directly, and what trade-offs accompany these benefits. Third, what research opportunities emerge from these changes? We discuss how these developments reshape the research priorities of IE and open new directions for its application. The tutorial aims to provide researchers and practitioners with a broader perspective on the changing role of IE, helping them identify meaningful research questions and opportunities for applying IE in the LLM era.

Biography: Xian-Ling Mao is a professor at Beijing Institute of Technology, specializing in natural language processing and big data mining. He has published over 80 papers in top-tier conferences and journals, including SIGIR, WWW, SIGMOD, ACL, TOIS, and TKDE, and has led the development of scientific intelligent systems such as Hammer Scholar and Hammer PDF. He has received 2 Best Paper Awards and 5 Provincial/Ministerial-Level Science and Technology Progress Award, and has led or participated in over 10 national research projects, including National Key R&D Program of China and the Key Program of the National Natural Science Foundation of China.

Co-speaker: Yi-Fan Lu, Beijing Institute of Technology

Tutorial 5: Reinforced Generative Information Retrieval



Xin Xin

Shandong University

Abstract: Generative retrieval (GR) and generative recommendation (GenRec) are reshaping information access by replacing the traditional index–retrieve–then–rank pipeline with a single sequence-to-sequence model that directly generates document or item identifiers from a query or user context. In GR, the corpus is consolidated into model parameters, and a query is mapped to its relevant docids autoregressively. In GenRec, the same paradigm is applied to recommendation: items are encoded as semantic IDs—often via residual quantization—and the model generates the next item or a slate of items directly from a user's history, eliminating cascading errors and enabling end-to-end optimization. However, both paradigms face a fundamental optimization gap: the autoregressive generation is typically trained with token-level likelihood or exact-match supervision, which does not directly align with the retrieval or recommendation objective. Reinforcement learning (RL) has emerged as a principled remedy. In GR, RL is used to optimize docid generation toward retrieval metrics and downstream utility, rather than mere token accuracy. In GenRec, RL addresses the multi-answer nature of recommendation by shaping serendipity rewards, mitigating reward sparsity and aligning generation with potential new user preference. This tutorial provides a unified treatment of reinforced generative retrieval, covering the foundations of GR and GenRec, the limitations of likelihood-based training, and the design of RL objectives—reward shaping, policy optimization, and stability—that make generative retrieval and recommendation systems more effective, robust, and aligned with their end goals.

Biography: Xin Xin is now a tenured Associate Professor in the School of Computer Science and Technology of Shandong University, as a member of the General AI Lab (GAIL). He also serves as the Dean Assistant for the School of Artificial Intelligence, Shandong University. Before that, he got his PhD degree in computing science from University of Glasgow, under the supervision of Prof. Joemon Jose. Formerly, he got his master degree from School of Software Engineering in SJTU and his bachelor degree from XJTU. His research interests span recommender systems, reinforcement learning, information retrieval, NLP, etc. He has published more than 50 papers in top-tier ML & IR conferences including SIGIR, WSDM, CIKM, IJCAI, ACL, KDD, etc. His research is supported by the National Science Foundation of China. He also conducted collaborative research with companies, such as Tecent and Kuaishou.

Tutorial 6: Is Parsing Dead? How Syntactic Analysis Survives in the LLM Era



Yuanhe Tian

Zhongguancun Academy

Abstract: Syntactic parsing makes sentence structure explicit and has long been central to computational language understanding. With large language models (LLMs) now handling many language tasks without explicit parsing pipelines, the role of syntax deserves renewed attention. This tutorial traces the evolution of parsing from rule-based grammars through statistical and neural approaches to the LLM era. Drawing on foundational parsing research and recent developments, it introduces the core ideas behind constituency and dependency parsing, then examines how LLMs perform syntactic analysis, support resource construction, and broaden parsing across languages and domains. It also explains why fluent generation does not guarantee structurally correct analysis. This tutorial provides a historical perspective on parsing, an overview of current methods and limitations, and an understanding of how explicit syntax support the evaluation, interpretation, and control of LLMs.

Biography: Yuanhe Tian is currently an Assistant Professor at Zhongguancun Academy. He received his Ph.D. in Linguistics from the University of Washington in 2025. His research interests include the integration of linguistic knowledge into natural language processing, large language models, multimodal understanding and generation, and AI agents, with particular applications to social media and healthcare. He has published more than 70 papers in leading international conferences and journals, including ACL, EMNLP, NAACL, IEEE TNNLS, and IEEE TMI. His publications have received over 2,500 citations on Google Scholar, with an h-index of 27. He has served as an Area Chair for ACL, EMNLP, NAACL, and EACL, and as a reviewer for journals including IEEE TPAMI, TMM, TNNLS, and TMI.

Industry Talks

Venue: Auditorium

Time	Talk	Speaker and Affiliation
Wednesday, 4 Nov 10:15-10:30	Industry Talk 1: The Intelligence-Efficiency Ratio: Open Intelligence for the Real World	Jun Zhou, Vice President of Ant Group
Wednesday, 4 Nov 12:00-12:15	Industry Talk 2: Next-Generation Agent Networks	Jie Li, General Manager of Zhengzhou Zoneyet Technology Co., Ltd.
Wednesday, 4 Nov 12:15-12:30	Industry Talk 3: The Evolution of AI: From Understanding and Generation to Interaction and Action	Jiaqi Wang (JD Explore Academy)
Thursday, 5 Nov 10:00-10:15	Industry Talk 4: Technical Exploration of Sound AI Composite Models in Audio Internet Scenarios	Ye Liu, Vice President of Technology at Lizhi Inc., part of Sound Group Inc.
Thursday, 5 Nov 11:30-11:45	Industry Talk 5: From Memory Operating Systems to Memory-Native Models: MemOS and MemT 1.0	Zhiyu Li, Co-Founder & CTO of MemTensor
Thursday, 5 Nov 11:45-12:00	Industry Talk 6: Agents That Remember, Anticipate, and Collaborate	Zenan Xu, Senior Researcher of Tencent

Industry Talk 1: The Intelligence-Efficiency Ratio: Open Intelligence for the Real World



Jun Zhou

Vice President of Ant Group

Abstract: Large models are moving from "can they get smarter" to "can stronger intelligence enter the real world at lower cost and higher reliability." This talk proposes the Intelligence-Efficiency Ratio: the real, reliable, professional, and sustainable intelligence value delivered per unit of total resources invested. Drawing on practice with the Ant Ling model families, it shows how three mutually reinforcing paths - efficient architectures, real-task evaluation, and an open ecosystem - form a flywheel that brings high-level intelligence to more people, organizations, and industries at scale.

Biography: Dr. Jun Zhou is Vice President of Ant Group. His research focuses on foundation models, machine learning, and other AI technologies, as well as their industrial applications. He has published over 100 papers at top-tier AI conferences, with an H-index of 60, and has received three international best paper awards and 46 granted U.S. patents.

As an early member of Alibaba Cloud's technical team, he contributed to the development and large-scale commercial deployment of Apsara, Alibaba's independently developed cloud computing operating system. As one of the founding members of AI at Alibaba and Ant Group, he has been deeply involved in and led major projects including personalized recommendation for the Double 11 shopping festival, the Platform for AI (PAI), the Ant Graph Learning system (AGL), AI-based risk management, and the intelligent evolution of Alipay. Under his leadership, his teams have ranked first in four international algorithm competitions. His honors include the First Prize of the Ministry of Education Award for Outstanding Scientific Research Achievements, the CCF Science and Technology Progress Award (First Prize), and the Zhejiang Province Science and Technology Progress Award (First Prize), among others.

Industry Talk 2: Next-Generation Agent Networks



Jie Li

General Manager of Zhengzhou Zoneyet Technology Co., Ltd.

Abstract: AI Agents are evolving toward autonomous collaborative systems. This lecture investigates multi-agent collaboration and inter-agent communication mechanisms, with Mindline serving as a human-AI shared thought container. The agent Echous is deployed on the EchousGrid network, enabling inter-agent coordination via structured communication protocols. We further elaborate the prospective paradigm in which agents interface with human cognitive space and the physical world. Open challenges regarding cognitive-physical integration are analyzed, and a forward-looking framework is proposed for an interconnected intelligent ecosystem integrating multi-agent coordination, human-AI shared cognition and physical-world execution.

Biography: Jie Li, formerly with Shaanxi Normal University. General Manager, Zhengzhou Zhongye Technology Co., Ltd., Standing Committee Member, CCF Technical Committee on Natural Language Processing, Director, Institute of Artificial Intelligence, National Supercomputing Zhengzhou Center, Co-Director, Zhengzhou University–Zhongye NLPLab, Committee Member, Henan Technical Committee for Standardization of Artificial Intelligence Medical Devices, Committee Member, Henan Technical Committee for Standardization of Robotics. Principal drafter of White Paper on Artificial Intelligence Standardization (2021), Knowledge Graph Selection and Implementation Guide (2021), Guidelines for the Standardization of Artificial Intelligence Ethics Governance (2023).

Industry Talk 3: The Evolution of AI: From Understanding and Generation to Interaction and Action



Jiaqi Wang
Head of NLP Lab at JD Explore Academy

Abstract: In this talk, we will share our perspectives on the evolving landscape of AI, tracing its progression from understanding and generation toward interactive intelligence and real-world action. We will introduce our recent advances in JoyAI-VL-Interaction, JoyAI-Video-Edit, and JoyAI-Echo, showcasing our explorations in multimodal intelligence, interactive content creation, and human-AI interaction. We will also highlight JD's efforts in embodied AI data collection, emphasizing the importance of high-quality real-world data in enabling AI systems to learn, interact with, and act in physical environments. Together, these developments reflect our vision for the next generation of AI systems: moving beyond understanding and generating content toward more natural interaction and embodied intelligence.

Biography: Jiaqi Wang heads the Natural Language Processing Lab at JD Explore Academy, leading the development of large language models (LLMs) and vision-language models (VLMs). He has published over 70 papers in top-tier international journals and conferences, with more than 30,000 citations on Google Scholar. Several of these papers have been recognized among the year’s most influential papers at top-tier conferences. He has also served as an Area Chair for NeurIPS, ICLR, ICML, and CVPR.

Industry Talk 4: Technical Exploration of Sound AI Composite Models in Audio Internet Scenarios



Ye Liu

Vice President of Technology at Lizhi Inc., part of Sound Group Inc.

Abstract: In recent years, the technology of AI large models has been advancing rapidly, and it has been widely applied in the fields of text, images, and videos. However, its practical application in the audio domain still has a considerable room for development. Due to the higher difficulty in obtaining high-quality audio data compared to other fields, traditional methods usually involve training multiple models in the audio field and conducting complex system integration to complete specific sound scene tasks. Clearly, such methods are costly and have low reusability in practical applications. Therefore, building sound AI composite models based on the existing AI large models and applying the models to audio scenarios has become an significant direction worthy of research and exploration.

Biography: Dr. Ye Liu currently serves as the Vice President of Technology of Lizhi Inc. under Sound Group Inc., a Nasdaq-listed company. He directs the R&D teams for Artificial Intelligence(AI), Big Data, Business Technology, Fundamental Research, and Audio Systems. With extensive experience in the industrial deployment of artificial intelligence and big data, he possesses deep expertise in applying large-scale AI models to audio-domain internet scenarios. He has authored over 20 papers in international conferences and journals, holds more than 20 granted patents, and demonstrates profound proficiency in sound AI technology research and industry scenario applications.

Industry Talk 5: From Memory Operating Systems to Memory-Native Models: MemOS and MemT 1.0



Zhiyu Li

Co-Founder & CTO of MemTensor

Abstract: As large language models move toward sustained interaction and long-running tasks, memory is becoming essential infrastructure for personalized, continuous, and adaptive AI agents. This talk presents two complementary efforts at MemTensor: MemOS and MemT 1.0. MemOS unifies plaintext, activation, and parametric memory as system resources with coordinated management, scheduling, and governance, covering their lifecycle from writing and retrieval to updating, migration, and forgetting. MemT 1.0 embeds memory into foundation models through persistent, dynamically evolving native memory states. It maintains and accesses historical memory through forward computation alone, with frozen weights and no backpropagation. The talk will examine selective memory writing, compression and capacity limits, knowledge conflicts, the stability–plasticity trade-off, long-horizon evaluation, training and inference efficiency, and privacy, security, and auditable governance. It will conclude with an outlook on the co-evolution of memory operating systems and memory-native models.

Biography: Zhiyu Li, Ph.D., is Co-founder and CTO of MemTensor and Technical Lead and Researcher at the Large Model Center of the Institute for Advanced Algorithms Research, Shanghai. He is a Shanghai Oriental Talent Program (Top Talent) honoree and a 36Kr "36Under36" Science and Technology Innovator. His research focuses on LLM pretraining, memory augmentation, efficient evaluation, and applications. He previously led core algorithm teams at Alibaba and Xiaohongshu. With his team, he introduced Memory^3, a model with hierarchical memory, in 2024; MemOS, a memory operating system for LLMs, in 2025; and MemT 1.0, a memory-native foundation model, in 2026. MemOS has attracted over 11,000 GitHub stars and more than 20,000 developers. Related technologies have been deployed at Bank of China, China Merchants Securities, China Telecom, and Xinhua News Agency. Li has published over 70 papers in venues including Patterns, NeurIPS, ICLR, ACL, and TKDE, and holds more than 10 granted patents. He also authored the Chinese-language monograph Memory Engineering: Building Agent Memory Systems in Practice. His research has been featured in a cover story in MIT Technology Review.

Industry Talk 6: Agents That Remember, Anticipate, and Collaborate



Zenan Xu

Senior Researcher of Tencent

Abstract: Imagine a personal agent that does more than recap your day. It remembers your commitments, tracks unresolved dependencies, and notices that a delayed confirmation could put next week's delivery at risk. Without waiting for you to ask, it flags the risk. It then joins the project group chat, with your permission, to check on the dependency, help the team agree on next steps, and follow up until the issue is resolved.

We believe memory, proactive assistance, and lifelong learning are essential to building such agents. Beyond personal assistance, we'll explore how agents learn the context of ongoing group discussions and use that understanding to support collaboration. Join us as we explore what it takes to build agents that remember, anticipate, and collaborate.

Biography: Zenan Xu, Senior Researcher in the Foundation Model Department at Tencent Hunyuan, works on large language models and personal agents. His current work focuses on building agents that understand users over time and provide proactive support, with an emphasis on memory and lifelong learning. He joined Tencent in 2023, and his work has included LLM post-training, using synthetic data generation and reinforcement learning to improve instruction following, long-context understanding, and complex multi-turn interactions. He received his Ph.D. in Computer Science from Sun Yat-sen University in 2023. Previously, he was a research intern at Microsoft Research Asia and Huawei Noah's Ark Lab, where he worked on natural language processing and pre-trained models.

Workshops

Workshop	Time	Room	Organisers
Greater Bay Area Workshop	Wednesday 4 Nov, 14:00-17:00	Auditorium	Prof. Kehai Chen and Prof. Tong Zhang
FinVolution Global Data Science Competition Workshop	Wednesday 4 Nov, 14:00-16:00	Meeting Room 1	Stella (Jiali) Tu
Speech and NLP Workshop	Thursday 5 Nov, 14:00-17:20	Meeting Room 1	Li Liu (The Hong Kong University of Science and Technology, Guangzhou)
Evaluation workshop I	Thursday 5 Nov, 13:30-15:30	Auditorium	Dr. Jiangpeng Yan (Frontier Technology Researcher, E Fund Management Co., Ltd.)
Evaluation workshop II	Thursday 5 Nov, 15:30-17:30	Auditorium	Dr. Runzhe Zhan (Nanyang Technological University, Singapore)

Greater Bay Area Workshop

Wednesday, 4 November 2026 • 14:00 – 17:00 @Auditorium

Time	Content	Speaker
14:00-14:40	AlayaDB.AI: The Data Foundation for Large Model Era	Bo Tang
14:40-15:20	Token-Efficient Reasoning in LLMs: Technical Landscape and Future Directions	Xuefeng Bai
15:20-16:00	Beyond Agents: From Execution to RSI	Chao Huang
16:00-16:15	Tea break	
16:15-17:00	Towards Reliable Language and Multimodal Models: Knowing, Retrieving, and Grounding	Hou Pong (Ken) Chan

Chairs



Biography: **Kehai Chen** is a Professor and doctoral supervisor at the School of Computer Science and the Institute of Computing and Intelligence, Harbin Institute of Technology (Shenzhen). He received the 2022 National NSFC Excellent Young Scientists Fund for Overseas Scholars at 2022. His research interests include natural language processing, large language model-based agents, and trustworthy multimodal reasoning. He has published over 100 papers in leading international conferences and journals, including ACL, ICML, NeurIPS, and ICLR. He has contributed to the ITU international standard “Cluster Systems for Foundation Model Training and Inference” and has received several awards, including the Special Prize of the 2025 Wu Wenjun AI Science and Technology Progress Award from the Chinese Association for Artificial Intelligence, the First Prize of the 2024 Qian Weichang Chinese Information Processing Science and Technology Progress Award, the Second Prize of the 2022 Beijing Science and Technology Progress Award, the 2020 CIPS Outstanding Doctoral Dissertation Award, and the Third Prize of the 2025 Guangdong Computer Society Higher Education Teaching Achievement Award. He has served as an Area Chair or Senior Area Chair at major international conferences in natural language processing and as a Senior Program Committee Member at international artificial intelligence conferences. He also serves on the editorial board of Neural Networks.

Chairs



Biography: **Dr. Tong Zhang** is a Professor and Doctoral Supervisor, as well as Vice Dean of the School of Computer Science and Engineering, South China University of Technology. He is an IEEE Senior Member and a recipient of the Excellent Young Scientists Fund from the NSFC and the Distinguished Young Scholars Fund of Guangdong Provincial Natural Science Foundation. His primary research interests lie in fundamental artificial intelligence algorithms and their industrial and interdisciplinary applications, with a particular focus on affective computing and large-scale foundation models. Dr. Zhang has led over 30 research projects as the principal investigator, including key programs, Excellent Young Scientists Fund Programs, and General Programs funded by NSFC, as well as key projects supported by the Guangdong Provincial Natural Science Foundation. He has authored and co-authored more than 200 peer-reviewed papers and holds over 50 authorized invention patents. He has been consecutively listed in Elsevier's Top 2% World's Most Influential Scientists for three consecutive years (2024–2026). Dr. Zhang currently serves as an Associate Editor for four JCR Q1 journals, namely IEEE Transactions on Affective Computing, IEEE Transactions on Computational Social Systems, Brain Informatics, and Journal of Intelligent Manufacturing. He also holds several prominent academic community positions, including Vice Chair of the IEEE Technical Committee on Computational Society and Social Intelligence, and Deputy Secretary-General of the Youth Working Committee of the Chinese Association of Automation. Additionally, he serves as Deputy Director of the Engineering Research Center of Intelligent Health Perception and Parallel Digital Human (Ministry of Education) and the Guangdong Provincial Key Laboratory of Computational Intelligence and Cyberspace Information. For his outstanding academic achievements, Dr. Zhang has received numerous prestigious honors and awards. His research accolades include the First Prize of Guangdong Natural Science Award (second recipient), the First Prize of Guangdong Science and Technology Progress Award, and the First Prize of Natural Science Award issued by the Guangdong Artificial Intelligence Industry Association. He is also the winner of the 10th Wu Wenjun Outstanding Youth Award in Artificial Intelligence and the ACM Guangzhou Rising Star Award. Furthermore, he has earned two flagship paper awards from the IEEE Systems, Man, and Cybernetics (SMC) Society: the 2026 Andrew P. Sage Best Transactions Paper Award and the Franklin V. Taylor Best Paper Award.

Speakers



Bo Tang

Talk Title: AlayaDB.AI: The Data Foundation for Large Model Era

Abstract: In the large model era, vectors—serving as metadata for large models—are being applied across various aspects of these models, such as Retrieval-Augmented Generation (RAG), sparse attention, and KV caching. Our team has taken vectors as the entry point to develop an ultra-efficient LLM inference infrastructure. In this talk, I will introduce two research work: AlayaLite and AlayaJet, and briefly discuss the key challenges in building AlayaDB.AI, the data foundation for the large model era.

Biography: Dr. Bo Tang is a tenured associate professor at Southern University of Science and Technology. His research interests are big data/large model systems. He always published at top-tier conference and journals (e.g., SIGMOD, PVLDB, TODS). His research outputs have been widely used in leading IT companies (e.g., Microsoft, Tencent, and Huawei). He won Huawei Spark Award (3 times). He is awarded SIGMOD China Rising Star at 2021 and sponsored by NSFC excellent young researcher grant at 2024.



Xuefeng Bai

Talk Title: Token-Efficient Reasoning in LLMs: Technical Landscape and Future Directions

Abstract: Reasoning-oriented large language models (LLMs) have recently demonstrated striking capabilities on complex tasks, yet their inference suffers from a fundamental bottleneck: reasoning is inefficient and expensive. Efficient reasoning has therefore become the decisive step that takes reasoning LLMs from “workable” to “practical”. This talk surveys the emerging technical landscape of efficient reasoning, covering chain-of-thought compression, token-efficient RL, and implicit reasoning, and concludes with an outlook on promising research directions.

Biography: Xuefeng Bai is an Associate Professor at the School of Computer Science and Technology, Harbin Institute of Technology, Shenzhen, and a member of CCF-NLP and CIPS-LMG. His research focuses on natural language processing (NLP) and reasoning with LLMs. He has published more than 50 papers at leading international conferences and journals in machine learning and artificial intelligence, including ACL, ICML, NeurIPS, ICLR, etc (5 oral papers, 1 ICML Spotlight paper). He serves as a Workshop Co-Chair of AACL, as an Area Chair for ACL, EMNLP, NAACL, NLPCC, and CCL, and as an editorial board member of the journal Large Language Model (Scilight Press).

Speakers



Chao Huang

Talk Title: Beyond Agents: From Execution to RSI

Abstract: Today's AI agents can converse, call tools, and complete complex tasks—but can they learn from success and failure and truly evolve? This talk explores the journey from instruction-following agents to recursive self-improvement (RSI), focusing on how agents propose strategies, execute and verify them, learn from feedback, and turn each task into fuel for the next stage of evolution. We will examine the possibilities of agents that continuously reflect, adapt, and improve themselves.

Biography: Chao Huang is a PhD supervisor at The University of Hong Kong, working on AI agents, large language models, and graph machine learning. His research has received over 21,000 Google Scholar citations. He leads HKUDS, an open-source AI organization behind projects such as nanobot, CLI-Anything, LightRAG, DeepTutor, RAG-Anything, DeepCode, AutoAgent, and AI-Researcher. HKUDS has surpassed 370,000 GitHub stars, ranked among the world's Top 23 open-source organizations, and appeared on GitHub Trending over 180 times. Chao has been recognized as a 2025 AI100 Young Pioneer and a 2025–2026 AI 2000 Most Influential Scholar. He has also received the 2024 WAIC Yunfan Award and the 2024 ICBS Frontier Science Award. His work has received multiple Most Influential Paper recognitions and Best Paper nominations at leading international conferences.

Speakers



Hou Pong (Ken) Chan

Talk Title: Towards Reliable Language and Multimodal Models: Knowing, Retrieving, and Grounding

Abstract: As large language models (LLMs) and multimodal large language models (MLLMs) become increasingly integrated into real-world applications, their reliability depends on recognizing what they do not know, acquiring relevant evidence, and reasoning faithfully with that evidence. This talk brings together three studies addressing these complementary challenges. First, I will examine how LLMs encode their knowledge boundaries across languages. By probing internal representations, we uncover cross-lingual structure in knowledge-boundary recognition and develop methods to transfer this ability across languages, helping identify hallucination risk. Second, I will discuss how LLMs can learn to formulate effective retrieval queries through reinforcement learning. Our findings show that different retrievers favor distinct query strategies, highlighting the importance of adapting query formulation to the retrieval environment for effective retrieval-augmented generation (RAG). Third, I will turn to multimodal reasoning and examine the reference gap: the failure to maintain stable bindings between language-based reasoning and the relevant visual entities, regions, or locations. Through controlled evaluations, we show that task-appropriate visual references, such as numbering, shading, and grids, can narrow this gap. These findings motivate a training-free visual skill that constructs suitable references to support multimodal reasoning. The talk will conclude with future research opportunities in trustworthy MLLMs.

Biography: Hou Pong (Ken) Chan is an Assistant Professor in the Department of Artificial Intelligence at the University of Macau. Previously, he was a Research Scientist at Alibaba DAMO Academy, Singapore. He also served as a Macao Fellow at the University of Macau and as a Visiting Postdoctoral Researcher at the University of Illinois Urbana-Champaign (UIUC), where he worked with Prof. Heng Ji. He received his Ph.D. from The Chinese University of Hong Kong in 2020 under the supervision of Prof. Irwin King. His research focuses on improving the trustworthiness of large language models (LLMs) and multimodal large language models (MLLMs).

FinVolution Global Data Science Competition Workshop

Wednesday, 4 November 2026 • 14:00-16:00 @Meeting Room 1

Host: Stella (Jiali) Tu

The FinVolution Global Data Science Competition Workshop is an official partner event of NLPCC 2026, offering a backstage pass to this year's 11th FinVolution Global Data Science Competition. Under the theme “LISTEN BEYOND 循声知意”, this year's challenge zeroed in on turn-taking modeling for multi-dialect Chinese conversations — the heart of natural human-machine dialogue: knowing when to speak, and how to respond. The competition drew 776 participants across 553 teams, culminating in an onsite finals in Shanghai this July, where finalists from Nanyang Technological University, Shanghai Jiao Tong University, University of Science and Technology of China and Jilin University, alongside teams from Ant International and Megvii, competed on stage.

During the workshop, FinVolution's data scientists will offer an in-depth look at this year's challenge and share notable cases from the forefront of fintech AI, while our award-winning teams present the solutions behind their success. We look forward to welcoming you to the FinVolution Workshop on November 4!

Time	Agenda	Description	Speaker	Title
14:00-14:20	Opening Remarks	About FinVolution and the Journey of Its Global Data Science Competition	Dr. Chunping WANG	Chief Scientist & Head of FinV Lab, FinVolution Group
14:20-14:35	2026 Challenge Statement	2026 Challenge Statement: Turn-taking Modeling in Conversation	Dr. Boyi NI	Senior Algorithm Scientist, FinVolution Group
14:35-15:45	Excellent Proposals Sharing ×4 (2 online & 2 onsite)	Excellent Proposals Sharing - 1 (15 mins presentation, online)	Chunhan Li	Gold Medalist
		Excellent Proposals Sharing - 2 (15 mins presentation)	Bo WANG	Silver Medalist
		Excellent Proposals Sharing - 3 (15 mins presentation, online)	Junjun Wang	4th Place Winner
		Excellent Proposals Sharing - 4 (15 mins presentation)	Bingkun Wei	6th Place Winner

Time	Agenda	Description	Speaker	Title
		10 mins QA	All Attendees	
15:45-15:55	Summary	Summary and Analysis of Winning Proposals	Dr. Feng GAO	Senior Algorithm Expert, FinVolution Group
15:55-16:00	Closing Remarks			

Speaker Information		
Chunping WANG	Chief Scientist & Head of FinV Lab, FinVolution Group	Ph.D. from Duke University; Deputy Director of the ZJU-FinVolution Joint AI Research Center. Work centers on applying AI to finance — user operations, risk control & anti-fraud, text-based customer service and voice outbound bots, intelligent quality inspection, multimodal identity verification and visual KYC, spanning structured data, computer vision, speech, NLP and graph algorithms. 20+ papers in top AI venues; honors include Shanghai "Smart Craftsman" (2022) and "Outstanding Women in Tech" (2023).
Boyi NI	Senior Algorithm Scientist, FinVolution Group	B.S. & Ph.D. in Automation from Tsinghua University; postdoctoral research in Germany and France. Leads NLP research and engineering for financial applications — natural language understanding, voice assistants, information extraction and knowledge extraction. Previously held research and senior algorithm roles at SAP Labs Shanghai, working on data analytics & modeling, database time-series engines and chatbot applications.
Feng GAO	Senior Algorithm Expert, FinVolution Group	Ph.D. from the University of Chinese Academy of Sciences; Adjunct Master's Supervisor at Shanghai University of Finance and Economics. Over a decade in natural language processing, with several papers and patents; leads R&D of multi-scenario voice outbound-calling and text-based dialogue bots for domestic and international markets.
Chunhan LI	Gold Medalist	Kaggle Grandmaster | 6 Gold, 3 Silver · 2025 Digital China Innovation Contest | Second Prize · 2024 Smart Education of China Platform Innovation Challenge (Track 2 & 5) | 1st Place · KDD Cup 2023 Recommendation Task | 4th Place · Zhejiang Energy Science Innovation Competition | First Prize

Time	Agenda	Description	Speaker	Title
Bo WANG	Silver Medalist	Algorithm Researcher at Megvii Technology, holds a Master's degree from Xidian University, with expertise in research of large model technologies and chip acceleration.		
Junjun WANG	4th Place Winner	Ph.D. student at the University of Science and Technology of China, majoring in Computer Science and Technology. His research interests include computer vision, AI-generated content detection, as well as image, audio and video tampering detection.		
Bingkun WEI	6th Place Winner	CS Master's student at Huazhong University of Science and Technology, with research interests in MLLMs, edge computing and intelligent agents.		

Speech and NLP Workshop

Thursday, 5 November 2026 • 14:00-17:20 @Meeting Room 1

Host: Li Liu (The Hong Kong University of Science and Technology, Guangzhou)

Time	Content	Speaker
14:00-14:30	Toward Context-Aware and Interactive Speech Language Models	Zhizheng Wu, The Chinese University of Hong Kong, Shenzhen
14:30-15:00	Bridging Simulation, Understanding, and Generation in AI-Driven Audio and Music	Qiuqiang Kong, The Chinese University of Hong Kong
15:00-15:30	Full-Duplex Speech Large Model with Empathetic Interaction Capability	Yang Feng, Institute of Computing Technology, Chinese Academy of Sciences
15:30-16:00	Learning and Accelerating Generative Trajectories for Image and Audio Generation	Tong Xiao, Northeastern University
16:00-16:10	Tea and coffee break	
16:10-17:20	Roundtable discussion, moderated by the host	Li Liu (moderator), with Zhizheng Wu, Qiuqiang Kong, Yang Feng and Tong Xiao

Evaluation Workshop I

Thursday, 5 November 2026 • 13:30-15:30 @Auditorium

Host: Jiangpeng Yan (Frontier Technology Researcher, E Fund Management Co., Ltd.)

Time	Content
13:31-13:50	**Session 1 — Shared Task 1: Multilingual and Multimodal Medical Instructional Video Understanding Evaluation (DA-MIVQA)**
	Overview of the NLPCC 2026 Shared Task 1: Difficulty-Aware Multilingual and Multimodal Medical Instructional Video Understanding Evaluation *Shenxi Liu, Kan Li, Mingyang Zhao, Yuhang Tian, Bin Li*
	Difficulty-Aware Instructional Step Reasoning for Temporal Answer Grounding in Single Video *Meimei Zhang, Xuetao Zhang, Hong Lei, Ziyang Xu, Xuan Qi, Sa Liu*
	SVP-Rerank: Difficulty-Conditioned Speech-Visual-Procedural Evidence Reranking for Medical Instructional Video Retrieval *Peng Zhang, Hongyou Zhou, and Jia Huang*
	DAEP: Difficulty-Aware Evidence Planning for Medical Video Corpus Temporal Answer Grounding *Tianjian He, Yujie Liu, Zhiping Huang, and Changbo Xu*
13:56-14:10	**Session 2 — Shared Task 2: Schwartz's Basic Human Values' Detection and Alignment with LLMs**
	Overview of NLPCC 2026 Shared Task 2: Schwartz's Basic Human Value Detection and Alignment with LLMs *Changtong Ge, Zhipeng Li, Xinglin Lyu, Lingling Mu, Hongfei Xu*
	Fine-Grained Human Value Detection and Value-Aligned Response Generation Based on MoCo-SCL and GRPO *Dejin Yang, Jin Wang, Xiaobing Zhou, Xuejie Zhang*
	ValueRank and TwoStageDPO for Fine-Grained Human Value Detection and Value-Aligned Response Generation *Xin Lan, You Zhang, Kangjian He, Hao Zhang, Dan Xu*
14:16-14:35	**Session 3 — Shared Task 4: LLM-based Investment Advisor Agents for Asset Allocation in the Chinese Market**
	Benchmarking LLM-Based Investment Advisor Agents for Asset Allocation in the Chinese ETF Market *Liyuan Chen, Shilong Li, Jiangpeng Yan, Shuoling Liu, Xi Cheng, Haoran Zhang, Jianqiu Deng, Xiu Li, Jian Li, Wu Zhu, Yingzi Zhu, Li Yuan, Qianqian Xie, Nan Tang, Yang Veronica Liu, Qiang Yang*
	NTU-QRS AlphaAllocator: Bounded Transformer News Overlays for Chinese Macro-Asset Allocation *Yuhe Sui*

Time	Content
	FinArbi: A Multi-Agent Investment Advisor with Signal Arbitration for ETF Sector Rotation *Yitian Liu, Lu Li, Yi Liu, Leyi Su, Zhou Huang*
	FinanceAgent: Bounded LLM Delegation for Risk-Controlled Asset Allocation in the Chinese Market *Bin Huang, Yi Luo， Ziliang Yin, Jiayang Luo, Kunli Zhang*
14:41-15:00	**Session 4 — Shared Task 5: Chinese BabyLM: Data-Efficient and Developmentally Plausible Language Models for Chinese**
	Overview of the NLPCC 2026 Chinese BabyLM shared task: data-efficient and developmentally plausible language models for Chinese *Siyuan Song, Zhiheng Qian, Yunhao Zhang, Linyang He, Xiaozhe Ji, Yingxin Lin, Hongao Zhu, Chongtian Shao, Chuhan Lang, Luan Li, Rui Wang, Renfen Hu, Shaonan Wang, Hai Hu*
	Beyond Natural Text: Teaching a BabyLM Chinese Character Form and Sound *Jiajie Cai, Xingkai Deng, Yijing Chen, Kunyu Xu*
	RepDuet at NLPCC 2026 Chinese BabyLM: Task-Specific Hybrid Representations for the Cognitive Track *Zhenhao Liu, Qiujie Xie, Minjun Zhu, Zhiyuan Ning, Yixuan Weng, Pinzhang Li, Jiajun Wang, and Yue Zhang*
	Compressed Pinyin-Based Representation for Mandarin Language Modeling *Timo Robrecht, Peng Xie, Maotong Sun, Johannes Gebert, Lukas Edman, Zhen Zhang, Yanliang Huang, Wenyuan Wu, and Zhuocheng Zhang*
15:06-15:25	**Session 5 — Shared Task 6: The 2nd Shared Task on LLM-Generated Text Detection**
	Overview of NLPCC 2026 Shared Task 6: Ternary Chinese LLM-Generated Text Detection in the Era of Human-Machine Collaboration *Junchao Wu, Zeyu Wu, Runzhe Zhan, Zhiwen Xie, Yulin Yuan, Yichao Du, Longyue Wang, Zhao Xu, Weihua Luo, Binkai Ou, Derek F. Wong*
	A Multi-Source LoRA and Pairwise-Ranking Fusion Framework for Chinese LLM-Generated Text Detection *Jiaming Hou, Xiangheng Li, Bin Huang, Guoyu Xu, Xiaoli Li, Kunli Zhang, Hongying Zan*
	Data Augmentation, Balanced Sampling and Sliding Window Majority Voting Ensemble for LLM-generated Chinese Text Detection *Xu Zhao, Hongfei Xu*
	CEDAR: Creator-Editor Disentanglement and Recalibration for LLM-Generated Chinese Text Detection *Bangyuan Wang, Boyuan Song, Jiacheng Ye*

Evaluation Workshop II

Thursday, 5 November 2026 • 15:30-17:30 @Auditorium

Host: Runzhe Zhan (Nanyang Technological University, Singapore

Time	Content
15:30-15:45	**Session 6 — Shared Task 3: Multi-dimensional Evaluation of Traditional Chinese Medicine LLMs**
	Overview of the NLPCC 2026 Shared Task 3: Multi-dimensional Evaluation of Traditional Chinese Medicine LLMs: Capability, Reliability and Robustness *Chuang Liu, Niepengfei, XiaoRuizheng, Zhe Chen, Zhaochen Ji, Meng Xiangfei*
	Textbook-Guided Multi-View Data Curation and Post-Training for TCM LLM Evaluation *Guanghui Zhao, Bin Huang, Xiangheng Li, Jiaming Hou, Kaixuan Yuan, Hongying Zan, Kunli Zhang*
15:46-16:00	**Session 7 — Shared Task 7: ELSSTBenchmark — Implicit Concept Retrieval and Knowledge Discovery for LLMs**
	Overview of the NLPCC 2026 Shared Task 7: Implicit Sociological Concept Retrieval and Knowledge Discovery *Zeqiang Wang*
	SAGE: Semantic Alignment and Guided Extraction for ELSST Concept Indexing *Bo Li, You Zhang, Kangjian He, and Dan Xu*
	Retrieval-Aligned Hard-Negative Mining and Consensus-Aware Fusion for Implicit Social-Science Concept Retrieval *Kejun Wu, Guoyu Xu, Yunhao Li, Bohan Yu, Yu Song, Hongying Zan, and Kunli Zhang*
16:06-16:20	**Session 8 — Shared Task 8: Factivity Inference Inconsistency Attack (FIIA)**
	Overview of the NLPCC 2026 Shared Task 8: Factivity Inference Inconsistency Attack *Daohuan Liu, Jiaoyang Su, Guanliang Cong, Junchao Wu, Yuer Wang, Yulin Yuan, Bin Li, Xuri Tang*
	Localized Commitment Perturbations for Factivity Inference Inconsistency Attacks *Yi Luo, Zhongtian Hua, Zhenkuan Bai, Chi Zhang, Yuxiang Jia, Yingjie Han*
	Breaking Self-Consistency in Chinese Factivity Inference: An Agent Dual-Stage Attack System for NLPCC 2026 Task 8 (FIIA) *Jiakun Liu, Zheng Li, Bin Huang, Xuedan Feng, Kunli Zhang, Yu Song*

16:26-16:40	**Session 9 — Shared Task 9: AISB — AI Scientist Benchmark**
	Overview of NLPCC 2026 Task 9: Auditing AI Scientist Systems through Executable Research Artifacts *QiYao Sun, Yixuan Weng, Minjun Zhu, Qiujie Xie, Zhen Lin, Yue Zhang*
	ADMETWeave: Multi-View ADMET Prediction with Similarity-Conditioned Cascades *Yi Luo, Bin Huang, Ziliang Yin, Chi Zhang, Kunli Zhang, Yingjie Han,*
	Lab2Paper: An Agent-Agnostic Workflow for Evidence-Grounded Scientific Writing from Experiments *Bin Huang, Yi Luo, Jiaming Hou, Ziliang Yin, Xiangheng Li, Yingjie Han, Hongying Zan, Kunli Zhan*
16:46-17:05	**Session 10 — Shared Task 10: Reliability of AI-Assisted Scientific Reporting**
	Overview of the NLPCC 2026 Shared Task 10: Reliability of AI-Assisted Scientific Reporting *Runzhe Zhan, Yutong Yao, Junchao Wu, Jingkun Ma, Yanming Sun, Fengying Ye, Lidia S. Chao, Derek F. Wong*
	ScienceJudge: Pattern Perturbation Augmentation for Sentence-Level Scientific Claim Verification *Chuhai Cai, Xin Chen, Min Yang, Feng Jiang*
	DualFaith at NLPCC-2026 Task 10: Metric-Aware and Domain-Robust Evidence-Grounded Scientific Report Verification *Wenbiao Zhao, Longkun Shi, Zicheng Zuo, Turdi Tohti*
	Internal Training Optimization and External Evidence Augmentation for Scientific Reporting Reliability Verification *Zhongtian Hua, Chenxin Ru, Yi Luo, Kunli Zhang, Meijia Yu, Yingjie Han*
17:11-17:25	**Session 11 — Shared Task 11: Agent-Based Experiment Reproduction from Scientific Papers**
	Overview of the NLPCC 2026 Shared Task 11: Agent-Based Experiment Reproduction from Scientific Papers *Hanhua Hong, Yizhi Li, Luu Gia Huy, Jian Yang, Ming Zhou, Chenghua Lin*
	A Breadth-First and Rubric-Grounded Multi-Agent Pipeline for Scientific Paper Reproduction *Junhao Fu, Jin Wang, Xiaobing Zhou, Xuejie Zhang*
	Stage-Aware LLM Routing and Execution-Guided Repair for Automated Experiment Reproduction *Kejun Wu, Guoyu Xu, Yunhao Li, Zhongtian Hua, Kunli Zhang, Hongying Zan, Yingjie Han, Yu Song*

Oral Sessions

Session 1

Wednesday, 4 November 2026 • 14:00-16:00 @Meeting Room 4+5

No.	Paper ID	Paper Title	Authors
1	126	N:M-Aware Sparsification of Large Language Models via Shared-Axis Rotation	Chenhui Zhu; Hongxu Jiang; Yaochen Han; Runhua Zhang
2	272	Refining Large Language Model Quantization via Asymmetrically Scaled Affine Transformation	xuan xie; Bin Ji; Xiaodong Liu; Jinzhu Kong; Fan Xiaofeng; Linbo Li; Shasha Li; Ma Jun; Jie Yu
3	386	Outlier-Aware Mixed Sparsity Pruning for Encoder-Decoder Speech Models	Zhaolai Zhou; Zhe Wang; Chuanbo Liu; Lin Li
4	635	NASPEFT: Neural Architecture Search for Parameter-Efficient Fine-Tuning of Pre-trained Large Language Models	Martin Homweyetsuro; LALIT; Miao Zhang
5	194	Single-Pass Ensemble: Efficient Text Embeddings via KV-Shared Postfix Prompt Tuning	Weixin Wang; Yeqin Zhang; Yuzi Yang; Cam-Tu Nguyen
6	307	BaScan: Boundary-Aware Sequence Automaton for Retrieval-Based Speculative Decoding	Ang Li; Jiao Wu; Yuliang Lei; Kuangyu Zheng; Bingzhen Xie; Pan Lu; An Huang; Weihang Xu
7	557	LiveThinking: Enabling Real-Time Efficient Reasoning for AI-Powered Livestreaming via Reinforcement Learning	Junfeng Ma; Yuhan Sun; Wanqing Cui; Zhiwei Huang; Shaopan Xiong; Meiguang Jin; Xiaohu Yang
8	135	TFDT-LLM: Time-Frequency Dual-Tower Autoregressive Forecaster	Zhuoran Duan; Yongbin Liu; Chunping Ouyang; Ying Yu

Session 2

Wednesday, 4 November 2026 • 14:00-16:00 @Meeting Room 8

No.	Paper ID	Paper Title	Authors
1	748	Retrieval-Feedback Aligned Distillation for Deployable LLM Query Expansion	Minghan Li; Guodong Zhou
2	330	Do List-wise Rerankers Ignore Document Order? Position Bias Persists	NanHang Luo; Dongdong Xie; Jingcheng Yuan; Donghong Ji
3	654	EHHR-RAG: Evidence-Aware Hierarchical Hypergraph Retrieval Framework for Multi-Hop Adaptive Retrieval-Augmented Generation	Yunfei Chen; Yu Liu; Yongbiao Han; Siwei Wan; Teng Huang; Chen Qiu; Jinguang Gu
4	830	HyBridge-RAG: Robust Graph Retrieval via Hybrid Initialization and Entity Bridging Inference	Jinjie Hao; Guixuan Zhang; Ying Huang; Jie Liu; Shuwu Zhang
5	354	SAGE: Scalable Knowledge Integration for Large Language Models via Distributed Semantic Graphs	Yi-Chun Chang; Zixi Wang; Xinyu Wang
6	377	DisentRAG: Explicitly Disentangling Parametric and Contextual Knowledge in Retrieval-Augmented Generation	Yuning Wan; Jiaang Li; Quan Wang
7	406	S-CAD: Source-Calibrated Adaptive Decoding for Trustworthy Retrieval-Augmented Generation	Yalong Guo; Yangsen Zhang; Jing Kang; 王思维
8	770	TASC: A Temporal Alignment and Snapshot-Consistency Framework for Time-Sensitive Question Answering	Jingyu Zhang; Meiqi Huang; luqing; Xiling Yang

Session 3

Wednesday, 4 November 2026 • 14:00-16:00 @Meeting Room 9

No.	Paper ID	Paper Title	Authors
1	214	From Diversity to Retrieval Robustness: Adapting Large Language Models to Complex Table Access	Jia Wang; Mingyu Zheng; Chenwang Zheng; Zheng Lin
2	289	STAR: A Multi-Agent Framework for Table Reasoning via Sub-Table Anchoring and Routing	Shuai Luo; Yanhua Yu; Ruopei Guo; WANG XiDian; Nan Cheng; Xin Gao; Zhiyong Huang; Haoyang Xu
3	711	SchemaRank: Structure-Aware Schema Linking via Graph Propagation for Text-to-SQL	Jingyu Zhang; Xu Zhang; Qing Lu; Shumin Li; Xiling Yang
4	775	Structural and Semantic-Aware Calibration for Semantic Parsing	Robert Ridley; Zhen Wu; Jianbing Zhang; Xinyu Dai
5	516	Multi-View Expansion for Field-Name-Robust Structured Extraction in Small Language Models	Chenzhang Guo; Wenbo Zhao; Xijun Wang; Limei Huang; Honghong Liao; Xiang Chen-SYSU
6	622	Calibrated and Conflict-Aware Rule Injection for Explainable Inductive Knowledge Graph Reasoning	闫志聪; Hui Zhao
7	349	RepoCG: Python Call Graph Construction with Static Analysis and LLM-Assisted Call Resolution	Zhen Zhu; Jiaxin Li; Lin Peng; Lihua Yang; Linjin Wei; Xuwen Wang; Min Xie
8	680	Consistently Wrong: An Operand-Grounding Anatomy of Reliability in LLM Numerical Table Question Answering	Yifan Lu

Session 4

Wednesday, 4 November 2026 • 16:15-18:15 @Meeting Room 4+5

No.	Paper ID	Paper Title	Authors
1	266	Optimizing Diversity and Correctness for Reliable LLM-Based Assertion Generation	ChengFeng Ge; Chunyan Hou; Chen Chen
2	544	LASER: Learning Active Perception via Self-Evolving Reasoning for GUI Grounding	Wanfu Wang; Qipeng Huang; Guangquan Xue; Xiaobo Liang; Qi Lv; XiaoLong Zhang; Zhe Zhao; Zhenxiang Yan; Juntao Li
3	639	Detected Is Not Enough: Risk-State Binding for Reliable Long-Horizon Data Agents	Xueyan Yao; Yuji Dong; Matilda Issac
4	640	FIG-MAC: A Fine-grained Inspiration Graph Empowered Multi-Agent Collaboration for Automated Scientific Hypothesis Generation	SiyuanChen; Yuping Gao; Tai Huang; Yuling Zhou; Dongyu Su; Tongguan Wang; Ying Sha
5	801	MACL-IOMA: A Unified Framework for Multi-LLM Co-Evolution via Cognitive Gap Quantification	Shihua Gou; Xinyu Hou; YIPING SONG
6	827	AutoViewMem: Self-Configuring Orthogonal Views for Conversational Long-Term Memory	Zijie Cao; xijun qu; Zhicheng Gu; Xiaoshu Chen; Duanyang Yuan; Yanning Hou; Sihang Zhou; Jianxing Gong; Jian Huang; YANG Mei
7	515	Evaluating LLM Reliability in External SOP Execution: A Process-Level Study of Context Selection in Task-Oriented Dialogue	Shengyan Cai; Yang He; Renhong Liu; Lujun Yin
8	789	Learning Task Progress in Robot Manipulation from Suboptimal and Failed Policy Rollouts	Huaxi Ai; Li Zeng; Yuhang Guo

Session 5

Wednesday, 4 November 2026 • 16:15-18:15 @Meeting Room 8

No.	Paper ID	Paper Title	Authors
1	218	SimpleREA: A Lightweight Context Selection for Robust Multi-Turn Conversations	Zhaoning Zhang; hongwanyang
2	603	LEAP: LLM-Guided Experience Augmentation with MCTS Planning for Task-Oriented Dialogue Policy Learning	Renchao Xie; Yangyang Zhao
3	591	TRACE: Transition-aware Contrastive LLM-CRF for Emotional Support Strategy Prediction	Yu Zhang; Yulong Li; Jiaoyun Yang; honglin chen; Yan-Yan Chen; Ning An
4	645	WereToM: Benchmarking Theory of Mind for Social Deduction in LLMs	YimingQin; Bobo Li; Chong Teng; Fei Li; Donghong Ji
5	815	Reasoning in Buddhist Hetuvidya Debate: A Benchmark for Adjudication, Source Attribution, and Dynamic Commitment Tracking in LLMs	Ruiyin Qian; JiaxinYou; Xueer Wang; Daifeng Li
6	677	TTC: Term-Tree Coupled Agents for Pedagogical Vocabulary Construction	Ziyi Zhang; Guanlin Li; Guangfei Geng; ZHANG ZHE; Wang Zhimin
7	484	SimLearnerEval: Auditing Whether Simulated Learners Are Valid Measurement Instruments for Tutor-Policy Evaluation	Zijun Mao; Peng Zhang; Liu Lun; Haozhe Yang
8	701	SNEU-Guide: Traceable Evidence-State Evaluation for Chinese Tour-Guide Training	Guanheng Qu; Jingxiang Qiu; Na Li; Jiangming Liu

Session 6

Wednesday, 4 November 2026 • 16:15-18:15 @Meeting Room 9

No.	Paper ID	Paper Title	Authors
1	85	MCEMA: Multi-Granularity Contrastive EM Attention Network for Fake News Video Detection	Yihao Wang; Zhong Qian; Peifeng Li; Qiaoming Zhu
2	351	Explaining to Detect: Improving Multimodal Metaphor Detection with Structured Explanation Supervision	Huiting Gao; Yongtao Wei; Yanquan Zhou
3	407	Multi-Granularity Sentiment Integration for LLM-Based Multimodal Sentiment Analysis	Shanshan Lin; Yuesheng Wu; Chao Chen; Yizhe Yang; Zhihao Chen; Zexian Yang; Xiangwen Liao
4	502	Listen Before You Look: Joint Visual-Audio Mixture of Retrieval Experts for Video Comprehension	Zhiqiao Wang; Xiaomei Chen; Haiwei Xue
5	628	TriCon-MEL: Candidate-Aware Triple Verification for Multimodal Entity Linking	Wenhao Xu; Yafei Jia; Xiaoying Pan; Hongcai Shen; Ruiqiang Lyu; Tingting Du; You Song
6	806	SIEVE: Sifting Hallucinated Visual Arguments in Multimodal Event Extraction via Training-Free Scene-Level Verification	Yuanye He; Shuhao Hu; Ji Xiang; Tianshu Fu; Zeyi Liu; Xin Wang
7	614	Enough to See, Not Enough to Answer: BIND-Probe for Multimodal Scientific Documents	Xuan Wang; Bo Wang; Ye Xiang; Ge Shi; Chong Feng
8	231	CORE-RL: Consistency-Oriented Reasoning Enhancement for Multimodal Mathematical Problem Solving	Xingyun Hong; Yan Shao; Shanshan Yan; Hui Liu

Session 7

Thursday, 5 November 2026 • 13:00-15:00 @Meeting Room 8

No.	Paper ID	Paper Title	Authors
1	42	LiteMedCoT-VL: Parameter-Efficient Adaptation for Medical Visual Question Answering	Runze Ma; Shunbo Jia; Haonan Lyu; Guo Liu; Caizhi Liao
2	374	Didactic knowledge or Clinical Cases? How Data Types Shape Medical Large Language Models	Yuzheng Fan; Haochun Wang; Sendong Zhao; Han Xiao; Ming Ma; Bing Qin
3	316	MedCollab: IBIS-Guided Multi-Agent Collaboration with Hierarchical Disease Relation Chains for Clinical Diagnosis	Yuqi Zhan; Xinyue Wu; Tianyu Lin; Yutong Bao; Xiaoyu Wang; Weihao Cheng; Huangwei Chen; Feiwei Qin; Zhu Zhu
4	721	DiagState: Process-Level Evaluation and Stage-Aware State Tracking for Interactive Medical Diagnosis	Zhuoran Lu; Qihuang Zhong; Juhua Liu; Bo Du; Dehua Peng; Daojun Gong
5	733	ClinCog-Router: Target-Label-Free Answer-Utility Routing over Frozen Medical Adapter Banks	yining huang; Keke Tang; Meilian Chen
6	768	UICD: Uncertainty-Aware Iterative Contrastive Decoding for Radiology Report Generation	Xing Jia; Ailing Shen; Yun Xiong; Haofen Wang; Yichen Liu; Siqi Lin; 李柯曼; Chunlei Tang
7	378	EDSurv-Lite: A Strict Early-Window Benchmark for Multimodal ED Deterioration Prediction	Peilin Shen; Yulong Zheng
8	427	SAGE: Soft-gated Aggregation of Graded Experts for Cognitive Level Prediction in Video Interviews	Yan Luo; Huiyao Chen; Ramen Liu; Yumeng Fu; Kamal Elgallouch; Sanjida Ferdous; Meishan Zhang; Baotian Hu; Min Zhang

Session 8

Thursday, 5 November 2026 • 13:00-15:00 @Meeting Room 9

No.	Paper ID	Paper Title	Authors
1	463	CoReE-Net: Aligning Convolutional and Axial Reasoning for Document-Level Relation Extraction	Rongen Yan; Lixuan Zhou; Jinyi Zhan; Yakun Li; Fu Xu
2	726	CMNIE: An Information Extraction Benchmark for Chinese Military News	Yan Yu; Mengna Zhu; Zhenyu Song; Hao Yang; Haiwen Chen; Mao Wang
3	851	A Feature-Guided Framework for Zero-Shot Event Temporal Relation Extraction with Large Language Models	Qingkai Min; Yue Zhang
4	428	M2DO: A Prior-Guided Multi-View Decision-Optimized Framework for Short Text Classification	Mingle Zhou; Yiyue Liu; Yuan Gao; Min Li; Delong Han
5	731	SAFE: A Semantic Augmentation and Fusion Enhancement Method for Chinese Offensive Language Detection	Jingkai Lin; Ziqi Yin; Xinyi Li; Bo Chen; Jingbo Sun; Yubo Chen; Ge Gao; Xiaobing Zhao
6	821	Threshold-Stable Selective Classification for Freedom of Information Act Privilege Review	Shibo Niu; Ke He; Bolin Hu; Yuejia Li; Yi Zhu
7	771	Beyond Dialect Labels: Validation-Gated Sparse Features in Dialectalized Chinese Writing	Wanqin Zhou; Keyan Jin
8	576	CFEC-Bench: A Chinese Fact Error Correction Benchmark for Large Language Models	Qingcheng Xiao; Tan Yue; Wei Wang; Huishuai Zhang; Dongyan Zhao

Session 9

Thursday, 5 November 2026 • 15:15-17:15 @Meeting Room 8

No.	Paper ID	Paper Title	Authors
1	657	Toward Fine-Grained Value Alignment for Chinese LLMs: A Value-Aware Modular Architecture with Reflection	Shenmeng Chen
2	331	PrefHCL: Human-AI Collaborative Text Detection via Multi-Style Preference Optimization with Hierarchical Contrastive Learning	Wangwang Zhan; HongChen Yu; Wenbo Li; Yangfan Zhang; Qianyu Xu; Wenliang Deng; Wu Guohua
3	360	TOSR: Task-Oriented Subspace Repair for Textual Backdoor Defense	Jiayi Li; Yuanchao Liu
4	437	BlackDecipher: Collaborative Weak Supervision for Robust Black-Box Membership Inference in LLMs	Wei Luo; Ruihan Hu; Yu-Ming Shang; Yong Ma; Sanchuan Guo; Xi Zhang
5	574	TRACE: Real-Time Token-Level Rectification for Safe and Capable LLM Generation	Pingyu Wu; Jiayu Wang; Wenbo Zhou; Weiming Zhang; Nenghai Yu
6	836	Code-ORBench: A Code-Oriented Over-Refusal Benchmark for Large Language Models	Yunxiang Wu; Qian Chen
7	450	Same Exam, Different Scores: A Tokenizer-Grounded Audit of Evaluation-Protocol Sensitivity on Chinese Benchmarks	Kun Zhang; Chunwei Xia
8	303	SCORE: Towards Robust Situational Cognitive Reasoning in LLMs	Yirui Wang; Kaicheng Yang; Yumeng Wang

Session 10

Thursday, 5 November 2026 • 15:15-17:15 @Meeting Room 9

No.	Paper ID	Paper Title	Authors
1	719	Curriculum-Based Incremental Evolution for Chinese ASR Post-Processing	Rui Wang; Ying Li; Yuan Li; Yuxin Huang; Shengxiang Gao; Cunli Mao; Zhengtao Yu
2	526	Text Judges Can't Hear: Prosodic Contrast Sets Expose Systematic Bias in Voice AI Evaluation	Adarsh Agrawal
3	395	BCP: A Lightweight Boundary-Controlled Framework for Online Sign Language Recognition and Translation	Xiao Ma; Lin Li; Shiqing Wu; Wang Gao; Chuanbo Liu; Jianquan Liu
4	49	Adaptive Denoising for Monolingual Domain Adaptation in kNN-MT	Hongyu Liu; yingfeng luo; Abudurexiti Reheman; Tong Xiao; JingBo Zhu
5	107	Lost in Interaction: Systematic Failures of MT and Automated Evaluation in Preserving Interpersonal Discourse Markers	Sirui Chen
6	697	Less Thinking, More Poetry: A Controlled Comparison of Reasoning and Non-Reasoning LLMs for Modern Chinese Poetry Translation	Bailiang Li; Wentao Zhang; Mu You; Jietian Tan; Junbin Zhang; Zehan Tan; Henghua Shen; Derek F. Wong; Tao Fang

Poster Sessions

Poster Session 1

Wednesday, 4 November 2026 · 10:30-11:30 @Poster Foyer

Poster ID	Poster Title
1	Context-Emotion Collaborative Representation Enhancement for Multi-turn Dialogue Value Recognition
2	HAR-Mem: A Hierarchical Memory Framework with Adaptive Retrieval
3	EvoConfig: Collaborative Systems for Efficient Autonomous Environment Configuration
4	DialectS2S: End-to-End Speech Dialogue Modeling for Low-Resource Chinese Dialects
5	MCPO: Metamorphic-Causal Prompt Optimization for LLM-based Multi-Agent System
6	SEAM: A Schema-Guided External Memory System for LLM Agents
7	Knowledge-Boundary-Aware Temporal Character Memory for Role-Playing in Long-Form Narratives
8	MemGraph: A Narrative Memory Graph for Long-Term Role-Playing Dialogue
9	Teaching LLMs to Learn Tool Trialing and Execution through Environment Interaction
10	BVPA: Branch-Verified Preference Augmentation for Long-Horizon LLM Agents
11	PERLU: Pointer-guided Evidence Reranking and Long-term Update for Dialogue Memory
12	Role-Aligned Collaborative Reinforcement Training for Heterogeneous Multi-Role Dialogue
13	EmbodiedCoop: Goal-Driven Multi-Robot Collaboration in Interleaved Task Stream via Progress Memory
14	Temporal-State Agents: Diagnosis-Driven LLM Agents for Time Series Machine Learning Engineering
15	HEKT: Dual-level Heterogeneous Distillation for Knowledge Tracing with Large Language Models
16	AutoEDM: SHAP-Guided Iterative Feature Engineering for Educational Data Mining
17	KAM-ICD: An Abbreviation-Aware and Multi-Source Knowledge-Enhanced Method for Automatic ICD Coding
18	Modeling Multi-Mapping in TCM Prescription Recommendation via Hierarchical Prototype Memory
19	CHARGE: Chinese Healthcare Augmented Retrieval and Generation Framework

Poster ID	Poster Title
20	Rethinking Law: Reflective Experience Augmentation for Legal Judgment Prediction
21	DSCF: Dual-View Semantic Clustering and Cross-Modal Fusion for Radiology Report Generation
22	TCM-AndroLLM: Domain Adaptation and Evaluation of Large Language Models for Traditional Chinese Medicine Andrology
23	MediRec: Enhancing Chinese Medication Recommendation with Explainable Clinical Reasoning
24	IndusSafetyLLM: A Domain-Adapted Large Language Model for Industrial Safety Risk Management
25	Normal-Referenced Global-Local Differential Modeling for Chest X-Ray Report Generation
26	RuleAct-Safe: Evidence Interaction Representation for Auditable Construction Safety Visual Reasoning
27	MathHERMES: Exploring Error Diagnosis and Mathematical Knowledge Remediation with Large Language Models
28	EviFrame: An Evidence-Guided Framework for Medical Text Classification
29	Decoupled Hard-Negative Ranking for Automated Short Answer Grading
30	AdaptSocratic: Adaptive Multi-Agent Socratic Teaching with Dynamic Learner Modeling and Curriculum Grounding
31	JT-RAG: A Jointly Trained Temporal Retrieval-Augmented Generation Framework for Biomedical Hypothesis Generation
32	Beyond Correlation: LLM-Guided Causal Graph Learning for Legal Text Classification
33	PriorStock: Finance-Rule Priors and Return-Fluctuation Learning for News-Driven Stock Movement Prediction
34	Zero-shot cross-domain slot filling via Domain-aware Contrastive prompting with Post-hoc Entity Refining
35	Bidirectional End-to-End Learning of Retriever-Reader Paradigm for Entity Linking
36	Classical Chinese Named Entity Recognition with LLM-Augmented Multi-Source Character Fusion
37	SCMeeting Room: Schema-Conditioned Mention-to-Proxy Routing for Multi-Event Document-Level Event Extraction
38	Structure-Prior-Guided Path Interaction for Uncertain Knowledge Graph Reasoning
39	CoPlan: Dual-Stage Taxonomy Expansion via Fast Location and Slow Reasoning
40	A Code-style Unified Framework for Multimodal Information Extraction with Knowledge Enhancement and Self-Refinement

Poster ID	Poster Title
41	Boundary-Type Decoupled Preservation for Class-Incremental Named Entity Recognition
42	Confusion-Guided Relation Disambiguation for Financial Relation Extraction
43	TAPN: Temporal Adaptive Propagation Network for Interpretable Temporal Knowledge Graph Reasoning
44	SEA: A Structured Evidence Adapter for LLM-Augmented Dialogue Relation Extraction
45	SR-NER: A Self-Evolution Training Framework for Zero-Shot Named Entity Recognition
46	Multi-source Evidence-Guided Zero-Shot Event Temporal Relation Extraction
47	Candidate-to-Event: Social Event Detection with Residual Semantic Anchor Fusion
48	Learning Risk-Calibrated Ranking for Retrieval-Augmented Generation via Three-Tier Evidence Structure
49	DIR-TIR: Synergizing Textual Dialog and Generative Visual Feedback for Interactive Image Retrieval
50	TopoChunker: Topology-Aware Agentic Document Chunking Framework
51	Bridging Semantic Gaps in RAG through Generated Context Knowledge Fusion
52	Learning to Select: Attention-based Residual Selection for Unified Hierarchical Aggregation in Multimodal Recommendation
53	DISD: Domain-Informed Speculative Decoding for Generative Recommendation
54	VeriCite: Type-Aware Verification and Revision for Faithful Citations in Scientific Survey Generation
55	State-Aware Degeneration Detection and Recovery for Repetition Suppression in LLMs
56	GuideRAG: Planning-Guided, Intent-Conditioned Graph Navigation for Retrieval-Augmented Generation
57	ReCQR: Incorporating Conversational Query Rewriting to Improve Multimodal Image Retrieval
58	ABLS-Rec: Aligning Behavioral and Linguistic Spaces for LLM-based Sequential Recommendation
59	Entropy-Gap Guided Evidence Selection for Multi-Turn RAG
60	SLEEV: A Dual-Branch Slot-Level Event--Evidence Verification Framework for Chinese Criminal Judgment Fact Condensation
61	Lost-in-the-Middle in Long-Text Generation: Synthetic Dataset, Evaluation Framework, and Mitigation Paradigm
62	Role-aware Prediction and Structure-aware Preference Alignment for Concept DEF Generation

Poster Session 2

Wednesday, 4 November 2026 · 15:30-16:30 @Poster Foyer

Poster ID	Poster Title
63	M-CREST: Dataset Construction and Two-Stage Reranking for Multilingual Personalized Timeline Summarization
64	More Than Static: Continual Instruction Learning based on Dynamic Orthogonality and Contrastive Learning
65	Teaching VLMs to Think Again: Advancing Mathematical Reasoning via Critique and Correction
66	FOCUS: Decoupling Expert Personas in LLMs to Enhance Domain Expert Capabilities
67	Teacher-Guided On-Policy Distillation for Large Language Models via Interleaved Inference
68	UCO-WRF: Uncertainty-Calibrated Dynamic OptScale and Weighted Reward Frequency Aggregation
69	Causal Evidence for Attention Head Imbalance in Modality Conflict Hallucination
70	Diagnosing Contrastive Decoding Transfer for Multi-Hop Knowledge Editing
71	Distribution-Aware Retraining-Free Visual Token Pruning for Vision-Large Language Models
72	Boosting Post-Training Quantization for Large Language Models via Learnable Grids and Residuals
73	Dynamic Curriculum-Guided Data Scheduling for Multi-Task RLHF
74	ICrossformer A Transformer for Implicit Cross-Channel Dependency Modeling in Multivariate Time Series Forecasting
75	EmoHarness: Diagnosing Candidate-Selection Headroom in Exact-Match Social Reasoning
76	Understanding Before Reasoning: Enhancing Chain-of-Thought with Iterative Summarization Pre-Prompting
77	LG-Sparse: LoRA-Guided Task Protection for Pruning Multi-Dialect LLMs
78	Boosting Post-training Quantization via Salient-aware Pairing and Residual Compensation
79	Chip-Tuning: Build Efficient Classification Models with Probing
80	Understanding the Bottleneck in Multi-Intent RAG: A Detection-Selection Decomposition
81	More Thinking Is Not Always Better: Structural and Semantic Factors of Complex Reasoning Tasks in LLMs
82	CORM: Criteria-Guided Preference Data Synthesis for Omni-Modality Reward Modeling

Poster ID	Poster Title
83	Simplicity Goes Far: Auditing Prompt Optimizers with Demonstration-Conditioned Prompt Search
84	Multi-Head Multi-Token Prediction with Cross-Head Teacher Forcing for Speculative Decoding
85	COPD: Contrastive On-Policy Distillation
86	Tone Negatives and Conditional Glyph Priors for Sample-Efficient Chinese BabyLM Pretraining
87	Local Semantic Enhancement and Knowledge-Guided Annotation Enhancement for Classical Chinese Paragraph Boundary Identification
88	WKFD: Rhetorical Device Recognition via World Knowledge Injection and Feature-Decoupled Reasoning
89	When More Knowledge Is Not Better: Prompt-Level Knowledge Augmentation for NL2SQL in General-Purpose LLM Workflows
90	Evaluating AIGC in Logographic Decoding: A Multimodal Analysis of Visual Attention and Neurological Engagement
91	Identifying Chinese Sentence Modality: What Makes Modality Difficult for NLP Systems?
92	Hidden in Plain Text: A Cue-Visibility Audit for Reader-Perceived Social Comparison
93	DTCD: Dual-Track Constrained Decoding for Constituency Parsing with Large Language Models
94	BMWSD: A Bidirectional Multidimensional Reasoning Framework for LLM-based Word Sense Disambiguation
95	From Measurement Instability to Contextual Adaptation: Investigating the Abstract Malleability of Cognitive Styles in Large Language Models
96	PEEF: Pairwise Event Compatibility Learning and Multi-View Dynamic Evidential Fusion for Script Event Prediction
97	Boosting Chinese Semantic Dependency Graph Parsing via Unified Modeling of Heterogeneous Semantic Parsing Tasks
98	FableGen: Training-Free Story Visualization via Narrative Factorization and Zero-Shot Identity Anchoring
99	Beyond the Encoder: Parameter-Efficient Whisper Encoder-Decoder Joint Representations for Cross-Lingual Audio Deepfake Detection
100	Zero-TTA: Zero-Shot Text-to-Avatar Synthesis
101	A Literal-Paraphrase Dual-Decoder Bidirectional Correction Model for Teochew Dialect Speech Recognition
102	Unified Spoken Language Understanding from Heterogeneous Data: Spoken NER and SA as a Case Study

Poster ID	Poster Title
103	Multimodal Conditioned Diffusive Time Series Forecasting
104	When to Use Context and Vision: Demand-aware Information Selection for Multimodal Dialogue Translation
105	ASR-Guided Preference Optimization for Low-Resource Thai and Lao TTS
106	Acoustic Retrieval-Guided Entity-Aware Speech Translation with GRPO
107	GARE-MoE: Mixture of Experts with Global and Routed Expert Collaboration for Multilingual Speech Translation
108	TriVerify: Sanskrit-Tibetan-Chinese Inter-Translation Methods for Low-Resource Classical Texts
109	Bridging the Modality Gap in Audio-Visual Speech Recognition: Progressively Bidirectional Cross-Modal Transformation
110	Multilingual Speech-to-Instruction Translation via Structured Reasoning
111	Select What to Learn, Rank What to Prefer: Efficient Reinforcement Learning for Literary Translation
112	Unbiased Scene Graph Generation by Type-Aware Message Passing on Heterogeneous and Dual Graphs
113	SGFE: Semantically Guided Feature Enhancement for Unbiased Scene Graph Generation
114	HiPrompt-CV: Scene Graph Generation via Hierarchical Language Prompting and Commonsense Fact Verification
115	Hybrid-MetricGAN: A Cascaded Time-Frequency Wiener Filter and Generative Adversarial Network for Speech Enhancement
116	R2-VideoAgent: Retrieval-Enhanced and RL-Aligned Agent for Long-form Video Question Answering
117	AdaRAG: Adaptive Retrieval-Augmented Generation for Multi-Hop Question Answering
118	MaintRAG: A Layout-Aware Agentic QA Framework for Maintenance Manuals
119	MECo: Multi-source Evidence Fusion and Budget-aware Compilation for Lightweight Conversational Memory Question Answering
120	AdaDRAG: Efficient Adaptive Reasoning for Debate-Augmented RAG
121	MedRoute-RAG: A Routed Text-Graph Retrieval-Augmented Framework for Hallucination Mitigation in Chinese Medical Question Answering
122	CoT-HM: Source-Aware Chain-of-Thought Hallucination Mitigation for Visual Question Answering
123	D2F-ReAG: Dynamic Decomposition and Filtering for Multi-Hop Reasoning-Augmented Generation

Poster Session 3

Wednesday, 4 November 2026 · 17:00-18:00 @Poster Foyer

Poster ID	Poster Title
124	StepGuard: Hierarchical Chain-of-Thought with Local Rollback for Faithful Question Answering
125	TVQA: Table-Text Verification for Reliable Heterogeneous Question Answering
126	Better Decomposition, Free Aggregation: A Synthesizer-Folding Framework for Multilingual Multi-Hop Question Answering
127	Taming the Long Tail: A Simple Head-to-Tail Representation Transfer for Reinforcement Learning in Agricultural VQA
128	When APIs Change: Benchmarking LLM Agent Robustness under Tool Evolution
129	ConfusEval: A Chinese Language Vocabulary Understanding Evaluation Benchmark
130	MobileMem: Evaluating Long-Horizon Memory for Agents in Real-World Mobile Environments
131	Auditing Large Language Models as Educational Assistants and Judges
132	ECWC-Metaphor: An English-Chinese Word-Level Corpus for Cross-Lingual Metaphor Processing
133	Valid Citations, Weak Grounding: Disentangling Evaluation Signals in Chinese RAG
134	Overconfidence in LLM-as-a-Judge: Diagnosis and Confidence-Driven Solution
135	Bridging Modalities and Knowledge: A Diagnostic Dataset for Evaluating Multi-Hop Reasoning in Real-World Scenarios
136	CET-Bench: A Cross-Lingual End-to-End Technical-Intelligence Benchmark for Document-Based Retrieval-Augmented Generation
137	LiveSciBench: A Multi-Agent Dynamic Benchmark for Multimodal Scientific Literature Reasoning
138	SMEC: A Benchmark Dataset for Structured Math Error Chains in Multi-Step Reasoning
139	CCL-DRR: A Chinese Contract Clause-Legal Article Dataset for Dual-Role Risk Review
140	HateScope: Benchmarking Large Language Models for Long-Form Chinese Hate Speech Attribution
141	A Fine-Grained Trustworthiness Assessment of LLM Safety Benchmark Results
142	Large Language Models on Classical Chinese AMeeting Room Annotation: A Multi-Genre Empirical Analysis via Multi-Model Cross-Evaluation
143	LAEP: A Leakage-Audited Evaluation Protocol for Chinese LLM-Refined Text Detection

Poster ID	Poster Title
144	Why Do Metrics Think That? Diagnosing Error Overestimation in LLM-based Machine Translation Evaluation
145	Concept-Centric Evaluation of Grade-Adaptive Educational Science Question Answering
146	MQA-Score: A Fine-grained Dataset for Chinese Multi-Question Single-Answer Scoring
147	Beyond Edit Matching: Benchmarking Chinese Grammatical Error Correction and Explanation for Korean L2 Learners with LLM-based Evaluation
148	TCM-DiagCommit: A TCM-SD-Derived Benchmark for Evidence-Conditioned Syndrome Selection and Commitment
149	ConsisOCR: A Consistency-Driven Framework for Low-Cost Automated Document Annotation
150	Inference-Sequence Guided Chain-of-Thought for Multimodal Desire, Emotion and Sentiment Recognition
151	Exploiting Global Context and Adaptive Selection for Multimodal Sarcasm Detection
152	CCOM: Conflict-Aware Calibration in Ordinal Sentiment Space for Multimodal Sentiment Analysis
153	GDSSD: Gradient-Decoupled Spectral Self-Distillation for Incomplete Multimodal Sentiment Analysis
154	HALT: Hallucination-Augmented Lightweight Triplet Discrimination for Aspect Sentiment Triplet Extraction
155	Grounding-Aware Uncertainty Reasoning for Multimodal Aspect-Based Sentiment Analysis
156	Your Stance Depends on Mine: From Node Encoding to Parent-Child Interaction for Conversational Stance Detection
157	Co-occurring Emotion Prototype Learning for Multi-label Emotion Recognition in Conversation
158	Missing-modality Multi-modal Sentiment Analysis Using Robust Representation Generated by Variational Autoencoders
159	Acoustic Grounding and Conversational Reasoning for Audio-Text Emotion-Cause Triplet Extraction
160	First-principle Driven: Modeling Language Credibility Bias for Multimodal Affective Computing
161	Thought Purity: A Defense Method for Chain-of-Thought Attacks in Local Agent Frameworks
162	RevLeet: Jailbreak Attack on LLMs via Morphological Disguise and Logical Induction

Poster ID	Poster Title
163	TCRD: Trust Modeling and Counterfactual Intervention for Rumor Detection
164	PMAPN and WBGB: A Dual-Stream Multi-Scale Framework and Benchmark for Gender Bias User Profiling
165	MFHP:A Rumor Detection Method based on Multiple Dynamic Features Fusion
166	DualKR: Joint Veracity Detection and Forgery Recognition for Fake News Videos
167	SIDE: Efficient Single-Inference Dual-Sided Jailbreaking for Open-Source LLMs
168	Romanized Harmful-Intent Prompts as Phonetic-Preserving Cross-Script Evaluation for LLM Safety
169	DR.GAP: Mitigating Bias in Large Language Models using Gender-Aware Prompting with Decoupled Reasoning
170	From Detection to Construction: User-Profile-Guided Safe Action Spaces for LLM Alignment
171	BiaGuard: Low-Rank Activation-Space Decoupling for Inference-Time Debiasing of Large Language Models
172	How Little Poisoning Is Enough? Backdoor Vulnerabilities of Low-Resource Large Language Models
173	DP-QVAC: GDP-Guided Query/Value-Wise Adaptive Clipping for Differentially Private LoRA Fine-Tuning
174	Semi-Supervised Chinese Harmful Meme Detection via Feature Decoupling and Robust Pseudo-Labeling
175	ORCA: State-Driven Policy Retrieval for Compliance in Tool-Calling Agents
176	Predictive but Bounded: Quantitative-Linguistic Features in Neural Chinese Social-Media Depression Detection
177	EADD: Entropy-Aware Dynamic Defense Against User-Side Misinformation in LLM Input
178	Quantization-Robust LLM Unlearning via Module-Sensitive Updates
179	BotRHG: Reliability-Guided Hypergraph Learning for Social Bot Detection
180	PropPathSSM: Structure-Aware State Space Modeling over Propagation Paths for Fake News Detection
181	From Alignment to Cultivation: Learning Value Priors before Alignment
182	Balancing User Preferences and Safety Risks in Open-Ended Roleplay Preference Optimization
183	LLM-RaF: Decision-Decoupled Rationale Fusion for Social-Context Fake News Detection
184	NCR-GPO: News-Comment Relation-Guided Preference Optimization for Fake News Detection

NCMMSC 2026 Joint Poster Session (Poster Session 4)

Thursday, 5 November 2026 · 10:00-11:00 @Poster Foyer

Poster ID	Poster Title
NCMMSC 2026 #1	FusionFlow-SVC: A Gated Fusion-Based Framework for Robust and High-Fidelity Singing Voice Conversion
NCMMSC 2026 #2	DeepASMeeting Room-NSpeech: A Fine-Grained Benchmark for Non-Speech ASMeeting Room Generation
NCMMSC 2026 #3	Time-Layer Adaptive Alignment for Speaker Similarity in Flow-Matching Based Zero-Shot TTS
NCMMSC 2026 #4	ST-News: A Shantou News Dataset for Cross-Dialect Speech-to-Subtitle Generation and Non-Verbatim Mapping Analysis
NCMMSC 2026 #5	MDFNet: A Multi-Domain Fusion Network for EEG-Based Cross-Session Auditory Attention Decoding
NCMMSC 2026 #6	Post-fusion Event-aware Time-Frequency Refinement for Audio Moment Retrieval
NCMMSC 2026 #7	SURE-EVAL: A Systematic and Unified Agentic Framework for Reproducible Evaluation
NCMMSC 2026 #8	Meeting RoomCodec: A Lightweight Multiple Bitrate Scalar Quantization Neural Speech Codec for Real-Time Communication
NCMMSC 2026 #9	Band-MixNet-R4: Low-Complexity Real-World Music Restoration with Teacher-Student Distillation and Hierarchical Routing
NCMMSC 2026 #10	Multilingual TASU: Extending Text-Only Alignment to Multilingual Speech Recognition via Simulated CTC Posteriors
NCMMSC 2026 #11	Asynchronous Voice Anonymization by Combining Adversarial Perturbations on Speaker Embedding and Waveform
NCMMSC 2026 #12	VocalMind-Bridge: Acoustic Feature Coordination for Depression Detection from Participant Speech
NCMMSC 2026 #13	CTT-ED: A Unified Morpheme-Tail Framework for Low-Resource Speech Recognition in Morphologically Rich Languages
NCMMSC 2026 #14	From Implicit to Explicit: A Two-Stage Evidence-Driven Framework for LLM-based ASR Error Correction

Evaluation Poster Session (Poster Session 5)

Thursday, 5 November 2026 · 13:30-17:30 @Poster Foyer

Poster ID	Paper Title / Authors
8	EVIL-Detect for NLPCC 2026 Shared Task 6: LLM-Generated Text Detection Hongrui Bao, Hangyu Rong, Zhuoshang Wang, Yubing Ren, and Yanan Cao
10	DAMAR: Difficulty-Aware Adaptive Multi-Agent Retrieval for Medical Instructional Videos Wosheng Deng, Mingyang Wu, Manasa Kumari, and Zhiping Huang
11	AEGIS: An Auditable Evidence-Graph Agent for Drawdown-Shielded Asset Allocation Bohe Chen
12	A Training-Free System for NLPCC 2026 Task 6: Recovering Weak Signals with Multi-pass Voting Lingling Shi, Xingrui Li, Shaozhen Liu, Ruiyu Fang, Shuangyong Song, and Yongxiang Li
28	SchemaTCM: A Schema-Guided Dual-Stream Reasoning Framework for TCM LLMs Tongyu Wang, Yan Wang, Yifeng Xue, Shilin Pan, Zhengtao Feng, Songsheng Wang, and Tongxuan Zhang
29	What a Baby Language Model Generalizes and What It Must Memorize: The Winning System for the NLPCC 2026 Chinese BabyLM Challenge Yiming Zhao, Yue Zhang, Shunda Chen, Yalu Liu, and Ziyi Han
33	NLPCC 2026 Task 10: Citation-Level Faithfulness Verification with DeBERTa Ensembles and Class-Wise Calibration Yanling Li, Zirui Li, and Mingyu Wan
35	FinanceAI: Temporal Macro-Graph Reasoning and Bounded LLM Advice for Daily ETF Allocation Jiawei Yang, Xuesheng Li, Cheng Li, and Xiu Li
37	Curriculum-Enhanced BabyDeBERTa for Natural Language Understanding in Chinese Rui Lu, Jin Wang, Kangjian He, Dan Xu, and Xuejie Zhang
48	FIIA Attack Construction via Narrative Evidence Revision with Qwen3 Scaling Diagnostics Yiyang Peng, Bo Zhang, Qizhao Wang, Yifeng Gou, and Gang Lv
62	Imbalance-Aware Cascades and Label-Preserving Evidence Fusion for Reliable Scientific Reporting Zhijie Liu, Jiacheng Wang, Wenhan Bian, and Ping Jian

University of Macau

澳門大學
UNIVERSIDADE DE MACAU
UNIVERSITY OF MACAU

Macau, China · Est. 1981

A Comprehensive University, Rooted in Macau

Founded in 1981, the University of Macau is a public, research-oriented university with an international faculty, English-medium teaching and a residential college system without parallel in Macao.



The Hengqin campus. The university relocated here in 2014.

An institution of international standing

The University of Macau is a public, comprehensive, research-oriented university. Its teaching units span arts and humanities, business administration, design, education, engineering, information science and computing, law, medicine, science and social sciences. Its research institutes run from advanced studies in the humanities and social sciences to Chinese medical sciences, applied physics, artificial intelligence and brain sciences, microelectronics and smart city technologies.

Education built around the student

Undergraduate education follows a "4-in-1" model that combines discipline-specific study with general education, research and internship experience, and community and peer education. Ten residential colleges, each with its own fellows, tutors and programme of experiential learning. They turn that model into daily practice, and are credited with developing seven competencies, from responsible citizenship to cultural engagement and a healthy lifestyle.

A campus built for scholarship

Since 2014 the university has occupied roughly one square kilometre on Hengqin Island, twenty times the size of its former home in Taipa, with a library and residential colleges laid out along a landscape designed for walking and conversation. A second campus in the Guangdong-Macao In-Depth Cooperation Zone, where ground was broken in December 2024, will make UM the first Macao institution to extend into the zone.

Some 16,900 students, more than 140 programmes

Around 16,900 students are enrolled across more than 140 degree programmes at bachelor's, master's and doctoral level, in fields that run from literature and law to computer science, biomedical sciences and integrated resort management. English is the primary medium of instruction, and four in five academic staff are recruited from outside Macao, an internationalism that shapes both classroom and research culture. Agreements with more than 400 institutions in over 40 countries, and exchange places at more than 200 partner universities, keep that classroom global.

Research, and the reach it has

Research follows a "3+3+3+3" strategy: three state key laboratories; three emerging fields, in precision medicine, advanced materials and regional oceanography; three interdisciplinary fields, among them artificial intelligence and robotics; and three platforms in the humanities and social sciences.

The results are measurable. UM sits in the top 1 per cent of the Essential Science Indicators in sixteen fields. It hosts five Ministry of Education key research institutes and seven joint laboratories with the Chinese Academy of Sciences; its microelectronics group has taken the second prize of the State Scientific and Technological Progress Award, and fifteen papers were accepted at the 2023 International Solid-State Circuits Conference.

RECOGNITION

THE World University Rankings	**154th**	THE Asia University Rankings	**28th**
QS World University Rankings	**267th**	THE Young University Rankings	**14th**

University of Macau · Avenida da Universidade, Taipa, Macau · +853 8822 8833 · info@um.edu.mo

信 息 學 院
Faculdade de Ciência da Informação e Computação
Faculty of Information Science and Computing

University of Macau

Six Disciplines. One Convergence.

Established in 2026 from the former Faculty of Science and Technology, the Faculty of Information Science and Computing brings six departments into a single faculty built for an age of convergence.

A new faculty, on a distinguished foundation

The faculty emerged from a deliberate reorganisation of the former Faculty of Science and Technology. It is, in the Dean's phrase, not a break with the past but a natural evolution, a reimagining of the faculty's strengths and traditions for the challenges of the digital age. Its six departments are at once the core pillars of modern information science and the frontiers where its next breakthroughs are most likely to be found.

Built for convergence

The faculty is designed to work across boundaries: between hardware and software, between theory and application, between human intelligence and machine intelligence. Four research institutes carry that work from laboratory to application and deepen collaboration with industry across the Greater Bay Area.

Vision and mission

The faculty's stated vision is to become a leading centre for information technology education and research in the Greater Bay Area and beyond. Its mission joins the cultivation of innovative talent to advanced work in artificial intelligence, microelectronics, financial technology and computer science, and to the development of Macao as a smart city.

Teaching and programmes

FIC offers three bachelor's, seven master's and three doctoral programmes to more than 1,600 students: around 350 undergraduates, more than 750 taught postgraduates and over 500 doctoral candidates.

Undergraduate programmes are accredited by the Hong Kong Institution of Engineers and recognised under the Washington Accord and the Seoul Accord, giving our graduates standing across a broad group of economies.

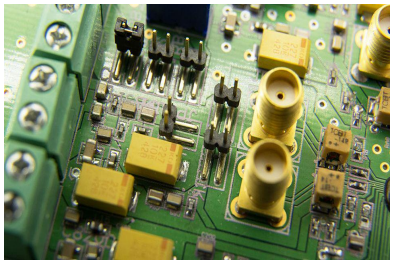
Research and recognition

Seven institutes and centres underpin research, among them the Institute of Artificial Intelligence and Brain Sciences, with centres for artificial general intelligence, brain sciences, data science and robotics, the key lab of Micro-electronics and the the Smart City Technologies. Twenty-two laboratories support the work.

UM ranks in the top 0.1 per cent of the Essential Science Indicators for engineering and computer science. THE places both disciplines within the world's top 101; U.S. News places electronic and computer engineering 36th, artificial intelligence 42nd and computer science 79th. Since 1993, the faculty has produced more than 2,000 alumni.







THE SIX DEPARTMENTS

01	**Artificial Intelligence**	Machine learning, perception and intelligent systems
02	**Computer Science**	Foundations, systems and software engineering
03	**Electronic and Communication Engineering**	Circuits, signals and communication networks
04	**Financial Technology and Business Intelligence**	Data-driven finance and analytics
05	**Microelectronics**	Integrated circuits, analog and mixed-signal design
06	**Robotics**	Autonomous systems and intelligent machines

Faculty of Information Science and Computing · Room 4040, E11 · +853 8822 4963 · fic.enquiry@um.edu.mo

NLPCC 2026 November 3-5 Macau, China | FinVolution

During NLPCC 2026 · Official Partner Workshop

The 2026 FinVolution Global Data Science Competition Workshop

Wed, Nov 4, 2026 · 14:00–16:00 | Meeting Room 1, Galaxy International Convention Center, Macau

About the Workshop

The 2026 FinVolution Global Data Science Competition Workshop is an official partner event of NLPCC 2026, offering a behind-the-scenes look at the 11th FinVolution Global Data Science Competition. Under the theme **“ LISTEN BEYOND · 循声知意 ”**, this year's challenge focused on **turn-taking modeling for multi-dialect Chinese conversations** — the heart of natural human-machine dialogue: knowing when to speak and how to respond. The competition drew 776 participants across 553 teams, with finalists from top universities and tech companies competing at the on-site final in Shanghai this July.

Join FinVolution's data scientists for an in-depth look at the challenge and real-world FinTech AI case studies, and hear the award-winning teams present their winning solutions live. See you on November 4!

Inside the Winning Solutions

Three finalist teams present their full approaches live

Ask the FinTech Experts

On stage with FinVolution's Chief Scientist and algorithm leads.

Agenda at a Glance

1. Welcome: FinVolution & the Competition
2. The 2026 Challenge Unpacked
3. Winning Team Presentations (Sharing + Q&A)
4. What the Winning Solutions Get Right
5. Open Networking with FinTech & AI Experts



Free FinVolution giveaways at the door — first come, first served!

Right by the NLPCC main sessions · Nov 4 · 14:00 · Meeting Room 1 — See you there.

FinVolution

FinVolution Group

Pioneering Credit-Tech Innovation Since 2007

FinVolution Group (NYSE: FINV) is a leading fintech company bridging millions of consumers and small businesses with financial institutions.

Founded in 2007 and listed on the New York Stock Exchange in 2017, we have been a pioneer in the pan-Asian credit technology industry, driving innovation in credit risk assessment, fraud detection, big data, and artificial intelligence.

Our fintech platforms are established in China, Indonesia, the Philippines, Pakistan, and Australia, with our technology footprint extending across Asia, Latin America, Africa, and beyond.

2007
Founded

11
Countries Covered

100+
Financial Institution Partners

45.7 Million
Cumulative Borrowers

198 Billion USD
Cumulative Transaction Volume

Data as of Q2 2026





记忆张量

MemTensor

记忆张量是一家记忆驱动的人工智能模型公司，推动 AI 实现长期智能。

公司由上海算法创新研究院孵化，中科院院士担任公司首席科学顾问，坚持基本原理驱动的系统性创新，以“低成本、低幻觉、高泛化”为核心特色，通过记忆推动 AI 模型范式创新，实现 AI 应用更广泛落地。

公司以 Memory³ 分层记忆研究为技术起点，研发原生记忆模型 MemT、记忆操作系统 MemOS、企业级智能体平台 MemMate 和个人记忆中枢与专属智能体 Memmy，形成从基础研究、模型架构到记忆系统与应用产品的技术体系，帮助企业沉淀和复用专业知识与业务经验，支持个人知识、偏好与工作经验的长期积累，让 AI 更好地理解背景、延续任务、持续协作。

公司核心团队汇聚普林斯顿大学、北京大学、上海交通大学等高校及阿里巴巴、华为、字节跳动等科技企业的人才，兼具基础研究与产业实践经验。公司已获得华为哈勃、荣耀战投、东方富海、深创投等机构投资，服务海尔、联想等 500 余家企业客户，并与招商证券、中国海诚、中国电信等企业开展合作，推动记忆驱动的 AI 在工业、金融、通信等领域落地应用。

智能，因记忆而进化。

Notes

中国计算机学会 CHINA COMPUTER FEDERATION

澳門大學 UNIVERSIDADE DE MACAU UNIVERSITY OF MACAU

Primary Sponsors

FinVolution

中业科技 zoneyet

Diamond Sponsors

记忆张量 MemTensor

京东

蚂蚁集团 ANT GROUP

Platinum Sponsors

Tencent 腾讯

荔枝

Golden Sponsors

vivo

Baidu 百度

海天瑞声

復旦大學 FUDAN UNIVERSITY

深绎未来 深译科技 DEEPFUTUREAI DEEPTRANXAI

