

EMNLP 2026 papers with UIUC Authors

Main Conference Papers

1. **MultiSessionCollab: Learning User Preferences Through Interaction for Long-Term Collaboration.** Shuhaib Mehri, Priyanka Kargupta, Tal August, Dilek Hakkani-Tür.
2. **Measuring, Localizing, and Reducing the AI-Like Writing Style of Post-Trained LLMs.** Aniket Anand, Janvijay Singh, Zhewei Sun, Dilek Hakkani-Tür, Nick Feamster.
3. **The Alignment Veto: How Safety Training Suppresses Cultural Knowledge in LLMs (Oral)** Pardis Sadat Zahraei, Gokhan Tur, Dilek Hakkani-Tur, Ehsaneddin Asgari
4. **Countdown-Code: A Testbed for Studying the Emergence and Generalization of Reward Hacking in RLVR.** Muhammad Khalifa, Zohaib Khan, Omer Tafveez, Hao Peng, Lu Wang.
5. **Gaming the Judge: Unfaithful Chain-of-Thought Can Undermine Agent Evaluation.** Muhammad Khalifa, Lajanugen Logeswaran, Jaekyeom Kim, Sungryull Sohn, Yunxiang Zhang, Moontae Lee, Hao Peng, Lu Wang, Honglak Lee.
6. **ReasoningFlow: Discourse Structures for Understanding LLM Reasoning Traces (Oral).** Jinu Lee, Shivam Agarwal, Amruta Parulekar, Siddarth Madala, Dilek Hakkani-Tür, Julia Hockenmaier
7. **Features in Motion: Periodicity, Homeostasis, and Adaptivity in LLM Few-Shot Prompting** Ikhyun Cho, Sewoong Lee, Julia Hockenmaier
8. **ObSchema: From Local Observations to Global Entity Attribute Schemas in Scientific Literature.** Xueqiang Xu, Wonbin Kweon, Jimeng Shi, Runchu Tian, Yunfan Kang, Wei Hu, Shaowen Wang, Jiawei Han
9. **OLIVIA: Online Learning via Inference-time Action Adaptation.** Sheldon Yu, Junda Wu, Xintong Li, Nikki Lijing Kuang, Sizhe Zhou, Tong Yu, Jiawei Han, Jingbo Shang, Julian McAuley
10. **Predict, Don't Iterate: Efficient Adaptive-Length Infilling for Diffusion Language Models** Haobo Xu, Sirui Chen, Yuanchen Bei, Lingjie Chen, Yuchen Yan, Dongqi Fu, Jingrui He, Hanghang Tong
11. **Subspace Alignment for Vision-Language Model Test-time Adaptation.** Zhichen Zeng, Wenxuan Bao, Xiao Lin, Ruizhong Qiu, Tianxin Wei, Xuying Ning, Yuchen Yan, Chen Luo, Monica Xiao Cheng, Jingrui He, Hanghang Tong
12. **Can Generalist Agents Automate Data Curation?** Feiyang Kang, Hanze Li, Adam Nguyen, Mahavir Dabas, Jiaqi Ma, Frederic Sala, Dawn Song, Ruoxi Jia.
13. **ALERT: Zero-shot LLM Jailbreak Detection via Internal Discrepancy Amplification.** Xiao Lin, Philip Li, Zhichen Zeng, Tingwei Li, Tianxin Wei, Xuying Ning, Gaotang Li, Yuzhong Chen, Hanghang Tong.
14. **PlanBench-XL: Evaluating Long-Horizon Planning of LLM Tool-Use Agents in Large-Scale Tool Ecosystems.** Jiayu Liu, Qihan Lin, Cheng Qian, Rui Wang, Emre Can Acikgoz, Xiaocheng Yang, Jiateng Liu, Zhenhailong Wang, Xiusi Chen, Heng Ji, Dilek Hakkani-Tür.
15. **Evaluating the Hidden Costs of Personalization in Large Language Models.** Yumeng Wang, Yuchen Wu, Cheng Qian, Zhiyuan Fan, Hyeonjeong Ha, Shujin Wu, Jiayu Liu, Heng Ji, Ge Wang.
16. **Geometric-Disentanglement Language Model Unlearning.** Duo Zhou, Yuji Zhang, Tianxin Wei, Ruizhong Qiu, Ke Yang, Xiao Lin, Cheng Qian, Jingrui He, Hanghang Tong, ChengXiang Zhai, Heng Ji, Huan Zhang.

17. **VISPA: Pluralistic Alignment via Automatic Value Selection and Activation.** Shenyang Zheng, Jiayou Zhong, Anudeex Shetty, Heng Ji, Preslav Nakov, Usman Naseem.
18. **DuplexGen: Adaptive Synthesis of Human-AI Turn-Taking Dialogues** (Oral). Takyoung Kim, Kang-wook Kim, Sang Hoon Woo, Julia Hirschberg, Gunhee Kim, Dilek Hakkani-Tür.
19. **Unsupervised Speech Recognition at the Syllable Level.** Liming Wang, Kai-Wei Chang, Kunio Kashino, David Harwath, Mark A. Hasegawa-Johnson, James R. Glass.
20. **Efficient Iterative Retrieval with Heterogeneous Batching.** Dohyun Park, Hubertus Franke, Daniel G Waddington, Swaminathan Sundararaman, Yongjoo Park.
21. **VGI-Bench: Probing Visual Intelligence in Video Generation Models.** Xuan He, Cong Wei, Yuhao Cheng, Linrui Ma, Yuxuan Zhang, Zuojun Li, Yuhao Wen, Jize Jiang, Zeyi Liu, Yuren Hao, Songcheng Cai, Keming Wu, Penghui Du, Kai Zou, Rui Yang, Chenkai Sun, Ke Yang, Ping Nie, Kelsey R Allen, Chenglong Wang, Michel Galley, Jianfeng Gao, ChengXiang Zhai.
22. **Evaluating Coding Agents on Kernel Exploit Generation** (Oral). Junyoung Jang, Gwanhyun Lee, Hwiwon Lee, Kyuheon Kim, Jongseong Kim, Jinho Jung, Lingming Zhang
23. **AI Psychosis: Does Conversational AI Amplify Delusion-Related Language?** Soorya Ram Shingekar, Vipin Gunda, Jiwon Kim, Violeta J. Rodriguez, Hari Sundaram, Koustuv Saha
24. **The Hidden Toll of Social Media News on Psychosocial Wellbeing.** Olivia Pal, Agam Goyal, Eshwar Chandrasekharan, Koustuv Saha
25. **VASTU: Language Models Struggle to Recognize Online Community Values** (Oral). Agam Goyal, Xianyang Zhan, Charlotte Lambert, Koustuv Saha, Eshwar Chandrasekharan
26. **FOXGLOVE: Understanding Goal-Oriented and Anchored Writing Feedback from Experts and LLMs on Argumentative Essays.** Yijun Liu, Yifan Song, John R. Gallagher, Sarah Sterman, Tal August
27. **LLUMI: Improving LLM Writing Assistance for Mental Health Support with Online Community Feedback** (Oral, *Best Paper Nominee*). Jiwon Kim, Maya Ajit, Sherry Gong, Soorya Ram Shingekar, Dong Whi Yoo, Eshwar Chandrasekharan, and Koustuv Saha
28. **Multi-Agent Evolve: LLM Self-Improve through Co-evolution.** Yixing Chen, Yiding Wang, Siqi Zhu, Weixian Xu, Haoifei Yu, Tao Feng, Muhan Zhang, Mostofa Patwary, Jiaxuan You.
29. **Can Large Language Models Forecast What Researchers Study Next?** Fenghai Li, Zihan Tang, Haoifei Yu, Yining Zhao, Jiaxuan You.
30. **LongSocialBench: Do Long-Context LLMs Understand Online Discussion Threads?** Xinyi Liu, Rinat Khaziev, Dilek Hakkani-Tür, Tarek F. Abdelzaher.