

DENG Yang

School of Computing and Information Systems
Singapore Management University (SMU)
80 Stamford Road
Singapore 178902

Email: [ydeng@smu.edu.sg](mailto:y deng@smu.edu.sg)

**Education**

PhD, Chinese University of Hong Kong, China, 2023

Master of Science, Peking University, China, 2019

Bachelor of Engineering, Beijing University of Posts & Telecommunications, China, 2016

Academic Appointments

Assistant Professor of Computer Science, School of Computing and Information Systems, SMU, Jul 2024 - Present

Research Fellow, National University of Singapore, Singapore, Apr 2023 - Jun 2024

Awards and Honors

New Faculty Highlights, AAAI 2026, 2025

Outstanding Area Chairs, EMNLP 2024, 2024

Google South Asia & Southeast Asia Research Awards, Google, 2024

RESEARCH

PublicationsJournal Articles [Refereed]

EXPLAINHM++: Explainable harmful meme detection with retrieval-augmented debate between large multimodal models, by LIN, Hongzhan; GAO, Wei; MA, Jing; DENG, Yang; LUO, Ziyang; WANG, Bo; YANG, Ruichao; CHUA, Tat-Seng. (2025). *IEEE Transactions on Knowledge and Data Engineering*, 1-14. <https://doi.org/10.1109/TKDE.2025.3637552> (Advance Online)

Proactive conversational AI: A comprehensive survey of advancements and opportunities, by DENG, Yang; LIAO, Lizi; LEI, Wenqiang; YANG, Grace Hui; LAM, Wai; CHUA, Tat-seng. (2025). *ACM Transactions on Information Systems*, 43 (3), 1-45. <https://doi.org/10.1145/3715097> (Published)

A comprehensive survey on relation extraction: Recent advances and new frontiers, by ZHAO, Xiaoyan; DENG, Yang; YANG, Min; WANG, Lingzhi; ZHANG, Rui; CHENG, Hong; LAM, Wai; SHEN, Ying; XU, Ruifeng. (2024). *ACM Computing Surveys*, 56 (11), 1-39. <https://doi.org/10.1145/3674501> (Published)

Nonfactoid question answering as query-focused summarization with graph-enhanced multihop inference, by DENG, Yang; ZHANG, Wenxuan; XU, Weiwen; SHEN, Ying; LAM, Wai. (2024). *IEEE Transactions on Neural Networks and Learning Systems*, 35 (8), 11231-11245. <https://doi.org/10.1109/TNNLS.2023.3258413> (Published)

A unified framework for contextual and factoid question generation, by DONG, Chenhe; SHEN, Ying; LIN, Shiyang; LIN, Zhenzhou; DENG, Yang. (2024). *IEEE Transactions on Knowledge and Data Engineering*, 36 (1), 21-34. <https://doi.org/10.1109/TKDE.2023.3280182> (Published)

A survey on aspect-based sentiment analysis: Tasks, methods, and challenges, by ZHANG, Wenxuan; LI, Xin; DENG, Yang; BING, Lidong; LAM, Wai. (2023). *IEEE Transactions on Knowledge and Data Engineering*, 35 (11), 11019-11038. <https://doi.org/10.1109/TKDE.2022.3230975> (Published)

Leveraging long short-term user preference in conversational recommendation via multi-agent reinforcement learning, by DENG, Yang; LI, Yaliang; DING, Bolin; LAM, Wai. (2023). *IEEE Transactions on Knowledge and Data Engineering*, 35 (11), 11541-11555. <https://doi.org/10.1109/TKDE.2022.3225109> (Published)

A unified multi-task learning framework for multi-goal conversational recommender systems, by DENG, Yang; ZHANG, Wenxuan; XU, Weiwen; LEI, Wenqiang; CHUA, Tat-Seng; LAM, Wai. (2023). *ACM Transactions on Information Systems*, 41 (3), 21-34. <https://doi.org/10.1145/3570640> (Published)

Toward personalized answer generation in e-commerce via multi-perspective preference modeling, by DENG, Yang; LI, Yaliang; ZHANG, Wenxuan; DING, Bolin; LAM, Wai. (2022). *ACM Transactions on Information Systems*, 40 (4), 1-28. <https://doi.org/10.1145/3507782> (Published)

Contextualized knowledge-aware attentive neural network: Enhancing answer selection with knowledge, by DENG, Yang; XIE, Yuexiang; LI, Yaliang; YANG, Min; LAM, Wai; SHEN, Ying. (2022). *ACM Transactions on Information Systems*, 40 (1), 1-33. <https://doi.org/10.1145/3457533> (Published)

Conference Proceedings

The rise of parameter specialization for knowledge storage in large language models, by HONG, Yihuai; ZHAO, Yiran; TANG, Wei; DENG, Yang; RONG, Yu; ZHANG, Wenxuan. (2025.0). *Proceedings of the Thirty-Ninth Annual Conference on Neural Information Processing Systems, San Diego, California, 2025 December 2-7*, (pp. 1-19) USA: <https://doi.org/10.48550/arXiv.2505.17260> (Published)

When less language is more: Language-reasoning disentanglement makes LLMs better multilingual reasoners, by ZHAO, Weixiang; GUO, Jiahe; DENG, Yang; WU, Tongtong; ZHANG, Wenxuan; HU, Yulin; SUI, Xingyu; ZHAO, Yanyan; CHE, Wanxiang; QIN, Bing; CHUA, Tat-Seng; LIU, Ting. (2025.0). *Proceedings of the Thirty-Ninth Advances in Neural Information Processing Systems (NeurIPS 2025), San Diego, CA, USA, November 30 - December 5*, (pp. 1-35) USA: <https://openreview.net/pdf?id=fleQIZ2VTx> (Published)

AdaSteer: Your aligned LLM is inherently an adaptive jailbreak defender, by ZHAO, Weixiang; GUO, Jiahe; HU, Yulin; DENG, Yang; ZHANG, An; SUI, Xingyu; HAN, Xinyang; ZHAO, Yanyan; QIN, Bing; CHUA, Tat-Seng; LIU, Ting. (2025.0). *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing, Suzhou, China, November 4-9*, (pp. 24570-24588) USA: ACL. <https://aclanthology.org/2025.emnlp-main.1248/> (Published)

From personas to talks: Revisiting the impact of personas on LLM-synthesized emotional support conversations, by WU, Shenghan; ZHU, Yimo; HSU, Wynne; LEE, Mong-Li; DENG, Yang. (2025.0). *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing, Suzhou, China, November 4-9*, (pp. 5439-5453) USA: ACL. <https://aclanthology.org/2025.emnlp-main.277/> (Published)

Chain of strategy optimization makes large language models better emotional supporter, by ZHAO, Weixiang; SUI, Xingyu; HAN, Xinyang; DENG, Yang; HU, Yulin; GUO, Jiahe; QIN, Libo; DU, Qianyun; WANG, Shijin; ZHAO, Yanyan; QIN, Bing; LIU, Ting. (2025.0). *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing, Suzhou, China, November 4-9*, (pp. 15361-15381) USA: ACL. <https://doi.org/10.18653/v1/2025.findings-emnlp.831> (Published)

Why stop at one error? Benchmarking LLMs as data science code debuggers for multi-hop and multi-bug errors, by YANG, Zhiyu; WANG, Shuo; YAN, Yukun; DENG, Yang. (2025.0). *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing, Suzhou, China, November 4-9*, (pp. 17377-17392) USA: ACL. <https://aclanthology.org/2025.emnlp-main.878/> (Published)

MPO: Multilingual safety alignment via reward gap optimization, by ZHAO, Weixiang; HU, Yulin; DENG, Yang; WU, Tongtong; ZHANG, Wenxuan; GUO, Jiahe; ZHANG, An; ZHAO, Yanyan; QIN, Bing; CHUA, Tat-Seng; LIU, Ting. (2025.0). *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics, Vienna, Austria, 2025 July 27 - August 1*, (pp. 23564-23587) USA: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2025.acl-long.1149> (Published)

FACT-AUDIT: An adaptive multi-agent framework for dynamic fact-checking evaluation of large language models, by LIN, Hongzhan; DENG, Yang; GU, Yuxuan; ZHANG, Wenxuan; MA, Jing; NG, See-Kiong; CHUA, Tat-Seng. (2025.0). *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics, Vienna, Austria, 2025 July 27 - August 1*, (pp. 360-381) USA: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2025.acl-long.17> (Published)

How to enable effective cooperation between humans and NLP models: A survey of principles, formalizations, and beyond, by HUANG, Chen; DENG, Yang; LEI, Wenqiang; LV, Jiancheng; CHUA, Tat-Seng; HUANG, Jimmy. (2025.0). *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics, Vienna, Austria, 2025 July 27 - August 1*, (pp. 466-488) USA: Association for Computational Linguistics. (Published)

EvoWiki: Evaluating LLMs on evolving knowledge, by TANG, Wei; CAO, Yixin; DENG, Yang; YING, Jiahao; WANG, Bo; YANG, Yizhe; ZHAO, Yuyue; ZHANG, Qi; HUANG, Xuanjing; JIANG, Yu-Gang; LIAO, Yong. (2025.0). *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics, Vienna, Austria, 2025 July 27 - August 1*, (pp. 948-964) USA: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2025.acl-long.47> (Published)

Knowledge boundary of large language models: A survey, by LI, Moxin; ZHAO, Yong; ZHANG, Wenxuan; LI, Shuaiyi; XIE, Wenya; NG, See-Kiong; CHUA, Tat-Seng; DENG, Yang. (2025.0). *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics, Vienna, Austria, 2025 July 27 - August 1*, (pp. 5131-5157) USA: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2025.acl-long.256> (Published)

Beware of your Po! Measuring and mitigating AI safety risks in role-play fine-tuning of LLMs, by ZHAO, Weixiang; HU, Yulin; DENG, Yang; GUO, Jiahe; SUI, Xingyu; HAN, Xinyang; ZHANG, An; ZHAO, Yanyan; QIN, Bing; CHUA, Tat-Seng; LIU, Ting. (2025.0). *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics, Vienna, Austria, 2025 July 27 - August 1*, (pp. 5131-5157) USA: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2025.acl-long.544> (Published)

Browsing like human: A multimodal web agent with experiential fast-and-slow thinking, by LUO, Haohao; KUANG, Jiayi; LIU, Wei; SHEN, Ying; LUAN, Jian; DENG, Yang. (2025.0). *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics, Vienna, Austria, 2025 July 27 - August 1*, (pp. 14232-14251) USA: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2025.acl-long.697> (Published)

Think both ways: Teacher-student bidirectional reasoning enhances MCQ generation and distractor quality, by QIU, Yimiao; DENG, Yang; YAO, Quanming; ZHANG, Zhimeng; DONG, Zhiang; YAO, Chang; CHEN, Jingyuan. (2025.0). *Findings of the Association for Computational Linguistics, ACL 2025, Vienna, Austria, July 27 - August 1*, (pp. 8240-8253) USA: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2025.findings-acl.432> (Published)

Unveiling knowledge boundary of large language models for trustworthy information access, by DENG, Yang; LI, Moxin; PANG, Liang; ZHANG, Wenxuan; LAM, Wai. (2025.0). *SIGIR '25: Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval, Padua Italy, July 13 - 18*, (pp. 4086-4089) New York: ACM. <https://doi.org/10.1145/3726302.3731684> (Published)

Query understanding in LLM-based conversational information seeking, by YUAN, Yifei; ABBASIANAE, Zahra; ALIANNEJADI, Mohammad; DENG, Yang. (2025.0). *SIGIR '25: Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval, Padua Italy, July 13 - 18*, (pp. 4098-4101) New York: ACM. <https://dl.acm.org/doi/10.1145/3726302.3731687> (Published)

Hello again! LLM-powered personalized agent for long-term dialogue, by LI, Hao; YANG, Chenghao; ZHANG, An; DENG, Yang; WANG, Xiang; CHUA, Tat-Seng. (2025.0). *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies, Albuquerque, New Mexico, April 29 - May 4*, (pp. 5259-5276) USA: Association for Computational Linguistics. (Published)

Query understanding in LLM-based conversational information seeking, by YUAN, Yifei; ABBASIANAE,

Zahra; DENG, Yang; ALIANNEJADI, Mohammad. (2025.0). *WWW '25: The ACM Web Conference 2025, Sydney, Australia, April 28 - May 2*, (pp. 73-76) New York: ACM. <https://dl.acm.org/doi/10.1145/3701716.3715869> (Published)

ChatCRS: Incorporating external knowledge and goal guidance for LLM-based conversational recommender systems, by LI, Chuang; DENG, Yang; HU, Hengchang; KAN, Min-Yen; LI, Haizhou. (2025.0). *Findings of the Association for Computational Linguistics, NAACL 2025, Albuquerque, New Mexico, April 29 - May 4*, (pp. 295-312) USA: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2025.findings-naacl.17> (Published)

PEARL: Towards permutation-resilient LLMs, by CHEN, Liang; SHEN, Li; DENG, Yang; ZHAO, Xiaoyan; LIANG, Bin; WONG, Kam-Fai. (2025.0). *Proceedings of the Thirteenth International Conference on Learning Representations, Singapore, 2025 April 24*, (pp. 1-20) USA: (Published)

LightPROF: A lightweight reasoning framework for large language model on knowledge graph, by AO, Tu; YU, Yanhua; WANG, Yuling; DENG, Yang; GUO, Zirui; PANG, Liang; WANG, Pinghui; CHUA, Tat-Seng; ZHANG, Xiao; CAI, Zhen. (2025.0). *AAAI'25/IAAI'25/EAAI'25: Proceedings of the Thirty-Ninth AAAI Conference on Artificial Intelligence and Thirty-Seventh Conference on Innovative Applications of Artificial Intelligence and Fifteenth Symposium on Educational Advances in Artificial Intelligence, Philadelphia, Pennsylvania, February 25 - March 4*, (pp. 23424-23432) New York: ACM. <https://doi.org/10.1609/aaai.v39i22.34510> (Published)

Aligning large language models for faithful integrity against opposing argument, by ZHAO, Yong; DENG, Yang; NG, See-Kiong; CHUA, Tat-Seng. (2025.0). *AAAI'25/IAAI'25/EAAI'25: Proceedings of the Thirty-Ninth AAAI Conference on Artificial Intelligence and Thirty-Seventh Conference on Innovative Applications of Artificial Intelligence and Fifteenth Symposium on Educational Advances in Artificial Intelligence, Philadelphia, Pennsylvania, February 25 - March 4*, (pp. 27748-27756) New York: ACM. <https://doi.org/10.1609/aaai.v39i26.34990> (Published)

Unlocking markets: A multilingual benchmark to cross-market question answering, by YUAN, Yifei; DENG, Yang; SOGAARD, Anders; ALLIANNEJADI, Mohammad. (2024.0). *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, Florida, USA, November 12-16*, Texas: ACL. (Published)

Thoughts to target: Enhance planning for target-driven conversation, by ZHENG, Zhonghua; LIAO, Lizi; DENG, Yang; LIM, Ee-Peng; HUANG, Minlie; NIE, Liqiang. (2024.0). *Proceedings of the Conference on Empirical Methods in Natural Language Processing, 19th EMNLP 2024: Miami, FL, November 12-16*, (pp. 21108-21124) Stroudsburg, PA: ACL. <https://doi.org/10.18653/v1/2024.emnlp-main.1175> (Published)

Don't just say "I don't know"! Self-aligning Large Language Models for responding to unknown questions with explanations, by DENG, Yang; ZHAO, Yong; LI, Moxin; NG, See-Kiong; CHUA, Tat-Seng. (2024.0). *Proceedings of the 19th Conference on Empirical Methods in Natural Language Processing (EMNLP 2024) : Miami, Florida, USA, November 12-16*, (pp. 13652-13673) USA: Association for Computational Linguistics. (Published)

Consecutive batch model editing with HooK layers, by LI, Shuaiyi; DENG, Yang; CAI, Deng; LU, Hongyuan; CHEN, Liang; LAM, Wai. (2024.0). *Proceedings of the 19th Conference on Empirical Methods in Natural Language Processing (EMNLP 2024) : Miami, Florida, USA, November 12-16*, (pp. 13817-13833) Miami, USA: Association for Computational Linguistics. (Published)

Experience as source for anticipation and planning : Experiential policy learning for target-driven recommendation dialogues, by DAO, Huy; DENG, Yang; BUI, Khanh-Huyen; LE, Dung D.; LIAO, Lizi. (2024.0). *Proceedings of the 19th Conference on Empirical Methods in Natural Language Processing (EMNLP 2024) : Miami, Florida, USA, November 12-16*, (pp. 14179-14198) USA: Association for Computational Linguistics. (Published)

Ask-before-plan : proactive language agents for real-world planning, by ZHANG, Xuan; DENG, Yang; REN, Zifeng; NG, See-Kiong; CHUA, Tat-Seng. (2024.0). *Proceedings of the Conference on Empirical Methods in Natural Language Processing 19th EMNLP 2024 : Miami, Florida, USA, November 12-16*, (pp. 10836-10863) USA: Association for Computational Linguistics. (Published)

Beyond persuasion : towards conversational recommender system with credible explanations, by QIN, Peixin; HUANG, Chen; DENG, Yang; LEI, Wenqiang; CHUA, Tat-Seng. (2024.0). *Proceedings of the 19th Conference on Empirical Methods in Natural Language Processing (EMNLP 2024) : Miami, Florida, USA, November 12-16*, (pp. 4264-4282) USA: Association for Computational Linguistics. (Published)

Selective annotation via data allocation: These data should be triaged to experts for annotation rather

than the model, by HUANG, Chen; DENG, Yang; LEI, Wenqiang; LV, Jiancheng; DAGAN, Ido. (2024.0). *Proceedings of the 19th Conference on Empirical Methods in Natural Language Processing (EMNLP 2024) : Miami, Florida, USA, November 12-16*, (pp. 301-320) USA: Association for Computational Linguistics. (Published)

Strength lies in differences! improving strategy planning for non-collaborative dialogues via diversified user simulation, by ZHANG, Tong; HUANG, Chen; DENG, Yang; LIANG, Hongru; LIU, Jia; WEN, Zujie; LEI, Wenqiang; CHUA, Tat-Seng. (2024.0). *Proceedings of the Conference on Empirical Methods in Natural Language Processing 19th EMNLP 2024 : Miami, Florida, USA, November 12-16*, (pp. 424-444) USA: Association for Computational Linguistics. (Published)

STYLE: Improving domain transferability of asking clarification questions in large language model powered conversational agents, by CHEN, Yue; HUANG, Chen; DENG, Yang; LEI, Wenqiang; JIN, Dingnan; LIU, Jia; CHUA, Tat-Seng. (2024.0). *Findings of the Association for Computational Linguistics ACL 2024, Bangkok, Thailand, August 11-16*, (pp. 10633-10649) USA: ACL. (Published)

On the multi-turn instruction following for conversational web agents, by DENG, Yang; ZHANG, Xuan; ZHANG, Wenxuan; YUAN, Yifei; NG, See-Kiong; CHUA, Tat-Seng. (2024.0). *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics, Bangkok, Thailand, 2024 August 11-16*, (pp. 8795-8812) USA: ACL. (Published)

WatME: Towards lossless watermarking through lexical redundancy, by CHEN, Liang; BIAN, Yatao; DENG, Yang; CAI, Deng; LI, Shuaiyi; ZHAO, Peilin; WONG, Kam-Fai. (2024.0). *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics, Bangkok, Thailand, 2024 August 11-16*, (pp. 9166-9180) USA: ACL. <https://aclanthology.org/2024.acl-long.496> (Published)

CLAMBER: A benchmark of identifying and clarifying ambiguous information needs in large language models, by ZHANG, Tong; QIN, Peixin; DENG, Yang; HUANG, Chen; LEI, Wenqiang; LIU, Junhong; JIN, Dingnan; LIANG, Hongru; CHUA, Tat-Seng. (2024.0). *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics, Bangkok, Thailand, 2024 August 11-16*, (pp. 10746-10766) USA: ACL. <https://aclanthology.org/2024.acl-long.578> (Published)

Self-chats from large language models make small emotional support chatbot better, by ZHENG, Zhonghua; LIAO, Lizi; DENG, Yang; QIN, Libo; NIE, Liqiang. (2024.0). *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics, Bangkok, Thailand, 2024 August 11-16*, (pp. 11325-11345) USA: ACL. <https://aclanthology.org/2024.acl-long.611> (Published)

Chain-of-exemplar: Enhancing distractor generation for multimodal educational question generation, by LUO, Haohao; DENG, Yang; SHEN, Ying; NG, See-Kiong; CHUA, Tat-Seng. (2024.0). *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics, Bangkok, Thailand, 2024 August 11-16*, (pp. 7978-7993) USA: ACL. <https://aclanthology.org/2024.acl-long.432> (Published)

Broadening the view: Demonstration-augmented prompt learning for conversational recommendation, by DAO, Huy; DENG, Yang; LE, Dung D.; LIAO, Lizi. (2024.0). *SIGIR '24: Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, Washington, DC, July 14-18*, (pp. 785-795) New York: ACM. <https://doi.org/10.1145/3626772.3657755> (Published)

Towards human-centered proactive conversational agents, by DENG, Yang; LIAO, Lizi; ZHENG, Zhonghua; YANG, Grace Hui; CHUA, Tat-Seng. (2024.0). *SIGIR '24: Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, Washington, DC, July 14-18*, (pp. 807-818) New York: ACM. <https://doi.org/10.1145/3626772.3657843> (Published)

Large language model powered agents for information retrieval, by ZHANG, An; DENG, Yang; LIN, Yankai; CHEN, Xu; WEN, Ji-Rong; CHUA, Tat-Seng. (2024.0). *SIGIR '24: Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, Washington, DC, July 14-18*, (pp. 2989-2992) New York: ACM. <https://doi.org/10.1145/3626772.3661375> (Published)

Large language model powered agents in the web, by DENG, Yang; ZHANG, An; LIN, Yankai; CHEN, Xu; WEN, Ji-Rong; CHUA, Tat-Seng. (2024.0). *WWW '24: Companion Proceedings of the ACM on Web Conference 2024, Singapore, May 13-17*, (pp. 1242-1245) New York: ACM. <https://doi.org/10.1145/3589335.3641240> (Published)

Plug-and-play policy planner for large language model powered dialogue agents, by DENG, Yang; ZHANG, Wenxuan; LAM, Wai; NG, See-Kiong; CHUA, Tat-Seng. (2024.0). *Proceedings of The Twelfth International Conference on Learning Representations, Vienna, Austria, 2024 May 7-11*, (pp. 1-33) USA: ICLR. <https://openreview.net/forum?id=MCNqgUFTHI> (Published)

Attack prompt generation for red teaming and defending large language models, by DENG, Boyi; WANG, Wenjie; FENG, Fuli; DENG, Yang; WANG, Qifan; HE, Xiangnan. (2023.0). *Proceeding of the 2023 Findings of the Association for Computational Linguistics, Singapore, December 6-10*, (pp. 2176-2189) USA: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.findings-emnlp.143> (Published)

Prompting and evaluating large language models for proactive dialogues: Clarification, target-guided, and non-collaboration, by DENG, Yang; LIAO, Lizi; CHEN, Liang; WANG, Hongru; LEI, Wenqiang; CHUA, Tat-Seng. (2023.0). *Proceeding of the 2023 Findings of the Association for Computational Linguistics, Singapore, December 6-10*, (pp. 10602-10621) Texas: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.findings-emnlp.711> (Published)

Cue-CoT: Chain-of-thought prompting for responding to in-depth dialogue questions with LLMs, by WANG, Hongru; WANG, Rui; MI, Fei; DENG, Yang; WANG, Zezhong; LIANG, Bin; XU, Ruifeng; WONG, Kam-Fai. (2023.0). *Proceeding of the 2023 Findings of the Association for Computational Linguistics, Singapore, December 6-10*, (pp. 12047-12064) USA: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.findings-emnlp.806> (Published)

Large language models as source planner for personalized knowledge-grounded dialogues, by WANG, Hongru; HU, Minda; DENG, Yang; WANG, Rui; MI, Fei; WANG, Weichao; WANG, Yasheng; KWAN, Wai-Chung; KING, Irwin; WONG, Kam-Fai. (2023.0). *Proceeding of the 2023 Findings of the Association for Computational Linguistics, Singapore, December 6-10*, (pp. 9556-9569) USA: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.findings-emnlp.641> (Published)

Unifying text, tables, and images for multimodal question answering, by LUO, Haohao; SHEN, Ying; DENG, Yang. (2023.0). *Proceeding of the 2023 Findings of the Association for Computational Linguistics, Singapore, December 6-10*, (pp. 9355-9367) Texas: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.findings-emnlp.626> (Published)

DepWiGNN: A depth-wise graph neural network for multi-hop spatial reasoning in text, by LI, Shuaiyi; DENG, Yang; LAM, Wai. (2023.0). *Proceeding of the 2023 Findings of the Association for Computational Linguistics, Singapore, December 6-10*, (pp. 6459-6471) USA: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.findings-emnlp.428> (Published)

Beyond factuality: A comprehensive evaluation of large language models as knowledge generators, by CHEN, Liang; DENG, Yang; BIAN, Yatao; QIN, Zeyu; WU, Bingzhe; CHUA, Tat-Seng; WONG, Kam-Fai. (2023.0). *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, Singapore, December 6-10*, (pp. 6325-6341) Texas: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.emnlp-main.390> (Published)

Multimodal fashion knowledge extraction as captioning, by YUAN, Yifei; ZHANG, Wenxuan; DENG, Yang; LAM, Wai. (2023.0). *SIGIR-AP '23: Proceedings of the 11th International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region, Beijing, November 23-28*, (pp. 52 -62) New York: ACM. <https://doi.org/10.1145/3624918.3625315> (Published)

Rethinking conversational agents in the era of LLMs: Proactivity, non-collaborativity, and beyond, by DENG, Yang; LEI, Wenqiang; HUANG, Minlie; CHUA, Tat-Seng. (2023.0). *SIGIR-AP '23: Proceedings of the 11th International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region, Beijing, November 26-28*, (pp. 298-301) New York: ACM. <https://doi.org/10.1145/3624918.3629548> (Published)

A survey on proactive dialogue systems: Problems, methods, and prospects, by DENG, Yang; LEI, Wenqiang; LAM, Wai; CHUA, Tat-Seng. (2023.0). *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, Macao, 2023 August 19-25*, (pp. 6583-6591) California: IJCAI. <https://doi.org/10.24963/ijcai.2023/738> (Published)

Learning to ask clarification questions with spatial reasoning, by DENG, Yang; LI, Shuaiyi; LAM, Wai. (2023.0). *SIGIR '23: Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, Taipei, Taiwan, July 23-27*, (pp. 2113-2117) New York: ACM. <https://doi.org/10.1145/3539618.3592009> (Published)

Towards robust personalized dialogue generation via order-insensitive representation regularization, by CHEN, Liang; WANG, Hongru; DENG, Yang; KWAN, Wai-Chung; WANG, Zezhong; WONG, Kam-Fai. (2023.0). *Proceedings of the 2023 Findings of the Association for Computational Linguistics, Toronto, Canada, July 9-14*, (pp. 7337-7345) USA: Association for Computational Linguistics.

<https://doi.org/10.18653/v1/2023.findings-acl.462> (Published)

PeerDA: Data augmentation via modeling peer relation for span identification tasks, by XU, Weiwen; LI, Xin; DENG, Yang; LAM, Wai; BING, Lidong. (2023.0). *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics, Toronto, Canada, 2023 July 9-14*, (pp. 8681-8699) USA: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.acl-long.484> (Published)

Product question answering in e-commerce: A survey, by DENG, Yang; ZHANG, Wenxuan; YU, Qian; LAM, Wai. (2023.0). *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics, Toronto, Canada, 2023 July 9-14*, (pp. 11951-11964) USA: Association for Computational Linguistics. (Published)

Knowledge-enhanced mixed-initiative dialogue system for emotional support conversations, by DENG, Yang; ZHANG, Wenxuan; YUAN, Yifei; LAM, Wai. (2023.0). *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics, Toronto, Canada, 2023 July 9-14*, (pp. 4079-4095) USA: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.acl-long.225> (Published)

Goal awareness for conversational AI: Proactivity, non-collaborativity, and beyond, by DENG, Yang; LEI, Wenqiang; HUANG, Minlie; CHUA, Tat-Seng. (2023.0). *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics, Toronto, Canada, 2023 July 9-14*, (pp. 1-10) USA: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.acl-tutorials.1> (Published)

ConReader: Exploring implicit relations in contracts for contract clause extraction, by XU, Weiwen; DENG, Yang; LEI, Wenqiang; ZHAO, Wenlong; CHUA, Tat-Seng; LAM, Wai. (2022.0). *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, Abu Dhabi, United Arab Emirates, December 7-11*, (pp. 2581-2594) Texas: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.emnlp-main.166> (Published)

PACIFIC: Towards proactive conversational question answering over tabular and textual data in finance, by DENG, Yang; LEI, Wenqiang; ZHANG, Wenxuan; LAM, Wai; CHUA, Tat-Seng. (2022.0). *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, Abu Dhabi, United Arab Emirates, December 7-11*, (pp. 6970-6984) Texas: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.emnlp-main.469> (Published)

User satisfaction estimation with sequential dialogue act modeling in goal-oriented conversational systems, by DENG, Yang; ZHANG, Wenxuan; LAM, Wai; CHENG, Hong; MENG, Helen. (2022.0). *WWW '22: Proceedings of the ACM Web Conference 2022, Virtual Conference, April 25-29*, (pp. 2998-3008) New York: ACM. <https://doi.org/10.1145/3485447.3512020> (Published)

Factual consistency evaluation for text summarization via counterfactual estimation, by XIE, Yuexiang; SUN, Fei; DENG, Yang; LI, Yaliang; DING, Bolin. (2021.0). *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, Virtual Conference, November 7-11*, (pp. 100-110) USA: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.findings-emnlp.10> (Published)

Exploiting reasoning chains for multi-hop science question answering, by XU, Weiwen; DENG, Yang; ZHANG, Huihui; CAI, Deng; LAM, Wai. (2021.0). *Proceedings of the Findings of the Association for Computational Linguistics, EMNLP 2021, Punta Cana, Dominican Republic, November 7-11*, (pp. 1143-1156) Texas: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.findings-emnlp.99> (Published)

Aspect-based sentiment analysis in question answering forums, by ZHANG, Wenxuan; DENG, Yang; LI, Xin; BING, Lidong; LAM, Wai. (2021.0). *Proceedings of the 2021 Findings of the Association for Computational Linguistics, Virtual Conference, November 7-11*, (pp. 4582-4591) USA: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.findings-emnlp.390> (Published)

Aspect sentiment quad prediction as paraphrase generation, by ZHANG, Wenxuan; DENG, Yang; LI, Xin; YUAN, Yifei; BING, Lidong; LAM, Wai. (2021.0). *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, Virtual Conference, November 7-11*, (pp. 9209-9219) USA: Association for Computational Linguistics. (Published)

Towards generative aspect-based sentiment analysis, by ZHANG, Wenxuan; LI, Xin; DENG, Yang; BING, Lidong; LAM, Wai. (2021.0). *Proceedings of the ACL-IJCNLP 2021: 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, Virtual, August 1-6*, (pp. 504-510) Texas: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.acl-short.64> (Published)

- Unified conversational recommendation policy learning via graph-based reinforcement learning, by DENG, Yang; LI, Yaliang; SUN, Fei; DING, Bolin; LAM, Wai. (2021.0). *SIGIR '21: Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Conference, July 11-15*, (pp. 1431-1441) New York: ACM. <https://doi.org/10.1145/3404835.3462913> (Published)
- Multi-hop inference for question-driven summarization, by DENG, Yang; ZHANG, Wenxuan; LAM, Wai. (2020.0). *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), Virtual Conference, November 16-20*, (pp. 6734-6744) USA: ACL. <https://doi.org/10.18653/v1/2020.emnlp-main.547> (Published)
- AnswerFact: Fact checking in product question answering, by ZHANG, Wenxuan; DENG, Yang; MA, Jing; LAM, Wai. (2020.0). *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, Virtual Conference, November 16-20*, (pp. 2407-2417) USA: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-main.188> (Published)
- Opinion-aware answer generation for review-driven question answering in e-commerce, by DENG, Yang; ZHANG, Wenxuan; LAM, Wai. (2020.0). *CIKM '20: Proceedings of the 29th ACM International Conference on Information & Knowledge Management, Virtual Conference, October 19-23*, (pp. 255-264) New York: ACM. <https://doi.org/10.1145/3340531.3411904> (Published)
- Bridging hierarchical and sequential context modeling for question-driven extractive answer summarization, by DENG, Yang; ZHANG, Wenxuan; LI, Yaliang; YANG, Min; LAM, Wai; SHEN, Ying. (2020.0). *Proceedings of the 43rd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual, Online, 2020 July 25-30*, (pp. 1693-1696) New York: ACM. <https://doi.org/10.1145/3397271.3401208> (Published)
- Answer ranking for product-related questions via multiple semantic relations modeling, by ZHANG, Wenxuan; DENG, Yang; LAM, Wai. (2020.0). *Proceedings of the 43rd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual, Online, 2020 Jul 25-30*, (pp. 569-578) New York: ACM. <https://doi.org/10.1145/3397271.3401166> (Published)
- Review-guided helpful answer identification in e-commerce, by ZHANG, Wenxuan; LAM, Wai; DENG, Yang; MA, Jing. (2020.0). *WWW '20: Proceedings of The Web Conference 2020, Taipei, Taiwan, April 20-24*, (pp. 2620-2626) New York: ACM. <https://doi.org/10.1145/3366423.3380015> (Published)
- Joint learning of answer selection and answer summary generation in community question answering, by DENG, Yang; LAM, Wai; XIE, Yuexiang; CHEN, Daoyuan; LI, Yaliang; YANG, Min; SHEN, Ying. (2020.0). *Proceedings of the 34th AAAI Conference on Artificial Intelligence, New York, 2020 Feb 7-12*, (pp. 7651-7658) Washington: AAAI Press. <https://doi.org/10.1609/aaai.v34i05.6266> (Published)
- MedTruth: A semi-supervised approach to discovering knowledge condition information from multi-source medical data, by DENG, Yang; LI, Yaliang; SHEN, Ying; DU, Nan; FAN, Wei; YANG, Min; LEI, Kai. (2019.0). *CIKM '19: Proceedings of the 28th ACM International Conference on Information and Knowledge Management, Beijing, China, November 3-7*, (pp. 719-728) New York: ACM. <https://doi.org/10.1145/3357384.3357934> (Published)
- Multi-task learning with multi-view attention for answer selection and knowledge base question answering, by DENG, Yang; XIE, Yuexiang; LI, Yaliang; YANG, Min; DU, Nan; FAN, Wei; LEI, Kai; SHEN, Ying. (2019.0). *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence, Honolulu, Hawaii, 2019 January 27 - February 1*, (pp. 6318-6325) USA: AAAI Press. <https://doi.org/10.1609/AAAI.V33I01.33016318> (Published)
- Knowledge as a bridge: Improving cross-domain answer selection with external knowledge, by DENG, Yang; SHEN, Ying; YANG, Min; LI, Yaliang; DU, Nan; FAN, Wei; LEI, Kai. (2018.0). *Proceedings of the 27th International Conference on Computational Linguistics, New Mexico, USA, 2018 August 20-26*, (pp. 3295-3305) USA: Association for Computational Linguistics. <https://aclanthology.org/C18-1279> (Published)
- Knowledge-aware attentive neural network for ranking question answer pairs, by SHEN, Ying; DENG, Yang; YANG, Min; LI, Yaliang; DU, Nan; FAN, Wei; LEI, Kai. (2018.0). *SIGIR '18: The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, Ann Arbor, MI, USA, July 8-12*, (pp. 901-904) New York: ACM. <https://doi.org/10.1145/3209978.3210081> (Published)

Research Grants

Singapore Management University

Proactive Conversational Agents for Real-world Planning, SMU Internal Grant, Ministry of Education (MOE) Tier 1 , PI (Project Level): DENG Yang, 2024, S\$119,980

Robust and Efficient Alignment for Integrity in Conversational Agents, SMU Internal Grant, Ministry of Education (MOE) Tier 1 , PI (Project Level): DENG Yang, 2024, S\$241,000

Other Institutions

Culturally-Aware Proactive Conversational AI for Enhancing Social Resilience in Singapore, AI RESEARCH - GOVERNANCE JOINT GRANT, AISG ZHANG Wenxuan, Co-PI (Project Level): DENG Yang, YOW Wei Quin, 2025, SGD1,019,408

TEACHING

Courses Taught

Singapore Management University

Postgraduate Professional Programmes :

Natural Language Processing for Smart Assistants

Postgraduate Research Programmes :

Empirical Research Project 4

OTHER ACADEMIC AND PROFESSIONAL ACTIVITIES

Other Grants and Funding Received (Non-Research)

Google South Asia & Southeast Asia Research Awards, Google, Nov 2024

EXTERNAL SERVICE – PROFESSIONAL

Program Chair, SSNLP 2024, 2024

Area Chair, NAACL 2025, 2024

Area Chair, COLING 2025, 2024

Area Chair, EMNLP 2024, 2024

Area Chair, ACL 2024, 2024