

Large Language Model Powered Agents in the Web

Yang Deng
National University of Singapore
Singapore
ydeng@nus.edu.sg

An Zhang
National University of Singapore
Singapore
an_zhang@nus.edu.sg

Yankai Lin
Renmin University of China
China
yankailin@ruc.edu.cn

Xu Chen
Renmin University of China
China
xu.chen@ruc.edu.cn

Ji-Rong Wen
Renmin University of China
China
jrwen@ruc.edu.cn

Tat-Seng Chua
National University of Singapore
Singapore
chuats@comp.nus.edu.sg

ABSTRACT

Web applications serve as vital interfaces for users to access information, perform various tasks, and engage with content. Traditional web designs have predominantly focused on user interfaces and static experiences. With the advent of large language models (LLMs), there's a paradigm shift as we integrate LLM-powered agents into these platforms. These agents bring forth crucial human capabilities like memory and planning to make them behave like humans in completing various tasks, effectively enhancing user engagement and offering tailored interactions in web applications. In this tutorial, we delve into the cutting-edge techniques of LLM-powered agents across various web applications, such as web mining, social networks, recommender systems, and conversational systems. We will also explore the prevailing challenges in seamlessly incorporating these agents and hint at prospective research avenues that can revolutionize the way we interact with web platforms.

CCS CONCEPTS

• **Information systems** → *Social networks; Recommender systems; Chat; Language models.*

KEYWORDS

Large Language Model, Social Network, Recommendation, Conversational Agent

ACM Reference Format:

Yang Deng, An Zhang, Yankai Lin, Xu Chen, Ji-Rong Wen, and Tat-Seng Chua. 2024. Large Language Model Powered Agents in the Web. In *Companion Proceedings of the ACM Web Conference 2024 (WWW '24 Companion)*, May 13–17, 2024, Singapore, Singapore. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3589335.3641240>

1 PRESENTERS

Yang Deng is a Postdoctoral Research Fellow at the National University of Singapore. He received his Ph.D. degree from The Chinese University of Hong Kong. His research lies in natural language processing and information retrieval, especially for conversational and interactive systems. He has published over 30 papers on relevant

topics at top venues such as WWW, SIGIR, CIKM, ACL, EMNLP, TKDE, and TOIS. He has rich experience in organizing tutorials at top conferences, including ACL 2023 and SIGIR-AP 2023. Contact via email ydeng@nus.edu.sg.

An Zhang is a Postdoctoral Research Fellow at the NEX++ research center. She earned her Ph.D. from the National University of Singapore. Dr. Zhang's research is primarily focused on large language models (LLMs) and recommender systems, with a specific interest in the deployment of LLM-driven multimodal generative agents for recommendation simulation. She has published more than 20 papers at top-tier conferences and journals, including WWW, KDD, NeurIPS, ICLR, TPAMI, and TOIS. Contact via email an_zhang@nus.edu.sg.

Yankai Lin is an assistant professor at Gaoling School of Artificial Intelligence, Renmin University of China. He received his Ph.D. degree from Tsinghua University. His research interests lie in pre-large language models (LLMs), especially LLM-based tool learning and LLM-based agents. He has published more than 50 papers at top-tier AI and NLP conferences, including ACL, EMNLP, IJCAI, AAAI, and NeurIPS, with over 10,000 Google Scholar citations. He was selected as the most cited Chinese researcher by Elsevier from 2020 to 2023. He has served as area chair for EMNLP and ACL ARR. Contact via email yankailin@ruc.edu.cn.

Xu Chen is a tenure track associate professor at Gaoling School of Artificial Intelligence, Renmin University of China. Before joining Renmin University of China, he was a research fellow at University College London, UK. Xu Chen obtained his PhD degree from Tsinghua University. His research interests lie in large language models, recommender systems, causal inference, and reinforcement learning. He has published more than 70 papers on top-tier conferences/journals like WWW, AIJ, NeurIPS, TKDE, SIGIR, WSDM and TOIS. He has organized many workshops and tutorials on top-tier conferences including SIGIR 2021, SIGIR 2020, SIGIR 2019, WSDM 2021, and WSDM 2018. Contact via email xu.chen@ruc.edu.cn.

Ji-Rong Wen is a full professor, the dean of School of Information, and the executive dean of Gaoling School of Artificial Intelligence at Renmin University of China. He has been working in the big data and AI areas for many years, and publishing extensively in prestigious international conferences and journals. He serves as the Program Chair of SIGIR 2020 and the Associate Editor of TOIS and TKDE. He has previously served as a senior researcher at Microsoft Research Asia and the group manager of the Web Search and Mining



This work is licensed under a Creative Commons Attribution International 4.0 License.

Group. He was elected as a National Distinguished Professor in 2013 and Beijing's Distinguished Young Scientist in 2018. He is a Chief Scientist at the Beijing Academy of Artificial Intelligence. Contact via email jrwen@ruc.edu.cn.

Tat-Seng Chua is the KITHCT Chair Professor with the School of Computing, National University of Singapore, where he was the Founding Dean of the School. His main research interests include multimedia information retrieval and social media analytics. He is the 2015 winner of the ACM SIGMM Technical Achievement Award and has received over 10 best papers awards or nominations in top conferences (WWW, SIGIR, MM, etc). He serves as the general co-chair of ACM MM 2005, ACM SIGIR 2008, WSDM 2023, WWW 2024, etc, and the editor of multiple journals (TOIS, TMM, etc). He has given invited keynote talks at varying international conferences, including the recent ones on the topics of Multimodal Conversational Search and Recommendation and Generative Search and Recommendation. Contact via email chuats@comp.nus.edu.sg.

2 OVERVIEW OF TUTORIAL TOPICS

According to the representative set of papers listed in the selected bibliography, this tutorial will contain about 25% of work that involves at least one of the six presenters. The rest of the tutorial will present a comprehensive overview of the tutorial topic by discussing the related work as much as possible from other researchers.

2.1 Background of LLM-powered Agents

Autonomous AI agents have long been regarded as stepping stones towards artificial general intelligence (AGI), with capabilities for self-guided task execution. Traditional approaches employed heuristic policy functions, which often lacked human-level adeptness in open-domain scenarios, largely due to heuristic limitations and constrained training data. Recently, LLMs have shown impressive strides towards human-like intelligence [32]. This advancement has spurred a growing trend in integrating LLMs as central components in developing autonomous AI agents [34, 35, 37, 52].

- **LLM-based Agent's Architecture.** The architectures of existing LLM-based AI agents can be distilled into a consolidated framework, extensively covered in recent survey literature on AI agents [39]. This unified structure comprises four primary modules: profiling, memory, planning, and action. The profiling module determines the agent's role, while the memory and planning modules immerse the agent in a dynamic environment, facilitating recall and strategizing of future action. The action module then converts decisions into concrete outputs. Notably, the profiling module influences both memory and planning modules, which in turn guide the action module.
- **LLM-based Tool Learning.** LLM-based tool learning seeks to meld the prowess of specialized tools and LLMs, enabling LLM-based agents to use external tools, and bringing in better autonomous problem-solving. Recent studies highlight foundation models' adeptness in tool utilization, such as web search automation [31], online shopping [45], neural model integration [37], computer task execution [23], and embodied robotic learning [2, 20].

2.2 LLM-powered Agents in Social Network

The social network connects different people by allowing them to share opinions and exchange information. Recent years have witnessed many AI techniques to solve social network problems like user connection prediction [47] and social information propagation [4], where the key challenge lies in understanding human intrinsic cognitive processes and behavior patterns. Recently, by learning huge amounts of web knowledge, LLMs have obtained remarkable success in achieving human-level intelligence. This sheds new light on solving social network problems, and several attempts have been made to incorporate LLM-based agents into this field.

- **Social Network Simulation with LLM-based Agents.** Social network simulation is a fundamental problem. If one can accurately simulate a social network, then its underlying mechanism and running rules can be easily understood and utilized. However, due to the intrinsic nature of human minds, it is quite hard to predict how people may behave in social networks. Recently, there have been several attempts [16, 26, 33] to leverage LLM-based agents to solve this problem. The key to these papers is leveraging LLMs as user brains, and designing profile, memory, and planning modules to make LLMs act like humans.
- **Social Network Problem Solving with LLM-based Agents.** Another research line on combining LLM-based agents with social networks is solving specific problems. People have leveraged agents to discover social system dynamics [17], analyze social principles between different agents [3], and so on. This direction is still rapidly growing, and we foresee that there will be much more promising work in the future.

2.3 LLM-powered Agents in Recommendation

Recommender systems play a pivotal role in contemporary information dissemination, actively shaping individual preferences [25]. With the recent advancements in LLMs, LLM-powered agents demonstrate remarkable achievements in autonomous interaction and user preference understanding [29]. This impressive capability can, on one hand, be harnessed to simulate authentic human behavior within recommender systems at both individual and population levels by scaling their deployment. On the other hand, it opens the potential for leveraging LLM-powered agents in the construction of a new paradigm of personalized recommenders [44].

- **User Behavior Simulation with LLM-powered Agents.** Simulating user behavior in recommender systems is a complex endeavor that requires a deep understanding of human preference and behavior patterns [5, 40, 49]. Bridging this gap necessitates not only the incorporation of agent modules that are tailored for recommendation contexts but also accommodation of the multimodal nature of such environments [27, 48]. Hence, agents driven by LLMs must be equipped with and further fine-tuned for multimodal comprehension to approximate the fidelity of real-world user interactions.
- **Recommender Agents.** While contemporary recommender systems are proficient in predicting domain-specific recommendations leveraging user behavioral data, they typically lack capabilities for explaining their recommendations, engaging in user conversations, and integrating rich user data [21]. To create a dynamic and interactive recommender system, LLMs serve as

the ‘brain’, with the recommender model acting as a tool [30, 42]. This research direction is dedicated to developing user-oriented recommender agents for the recommendation ecosystem [38].

2.4 LLM-powered Conversational Agents

LLM-powered conversational agents [13] not only redefine user interaction but also introduce innovative functionalities that push the boundaries of traditional web engagements.

- **LLM-powered Conversational Agents for User Simulation.** Building user simulators [36, 51] has emerged as an effective and efficient technique for evaluating conversational systems, thereby mitigating the high cost of interacting with real users. Inspired by the recent success of leveraging LLMs for role-play scenarios, researchers design LLM-powered conversational agents, which can be flexibly adapted to different dialogue evaluations, including open-domain dialogues [24], task-oriented dialogues [18], and conversational recommendation [41].
- **LLM-powered Proactive Conversational Agents.** Despite the exceptional proficiency in context understanding and response generation in various dialogue problems, LLM-based conversational agents typically prioritize accommodating users’ intentions as LLMs are trained to passively follow users’ instructions. Therefore, LLM-powered conversational agents often face challenges in handling proactive dialogue problems that require the conversational agent to strategically take the initiative to steer the conversation towards an anticipated goal [10]. To this end, recent works investigate prompt-based policy planning methods that prompt an actor LLM to either conduct self-thinking of strategy planning for each turn [11, 50] or generate AI feedback given the whole dialogue history to iteratively improve the dialogue policy planning for proactive dialogues [12, 14, 46].

2.5 Open Challenges and Beyond

In the last part, we will discuss the main open challenges in developing LLM-powered agents in web applications and several potential research directions for future studies.

- **Trustworthy and Reliable Web Agents.** While LLM-powered web agents are designed to be accurate, hallucination and inconsistency issues [6, 22] can lead to incorrect or inappropriate responses. Ensuring that these agents are both trustworthy (data privacy and ethical considerations [7]) and reliable (consistent and accurate performance) remains a pressing challenge.
- **Multi-agent Collaboration and Competition.** As the web ecosystem grows in complexity, there is a foreseeable future where multiple LLM-powered agents will need to interact with each other, either collaboratively to achieve common goals or competitively. Designing agents that can effectively collaborate requires addressing challenges in communication [43], shared knowledge bases [52], and synchronizing actions in real-time [1]. On the other hand, competitive scenarios [14] necessitate agents that can strategize, negotiate, and adapt to dynamic conditions.

3 RELEVANCE TO THE WEB CONFERENCE

The autonomous AI agent powered by LLMs is a trending topic across various web applications, such as web mining [8, 15, 45], recommendation [21, 40, 48], and conversational systems [12, 14].

This topic receives notably increasing attention from both academia and industry. In academia, the web conference recognizes *Large Language Models* as a new research topic for various research tracks this year due to the revolutionary techniques based on LLMs. In industry, recent years have witnessed many successful web applications that are empowered by the integration of LLMs. For example, Microsoft released a new version of Bing with its integration with ChatGPT. Several tutorials about the integration of LLMs in specific web applications have been given in related top-tier conferences, including but not limited to 1) *Tutorial on Large Language Models for Recommendation* at RecSys 2023 [19], 2) *Proactive Conversational Agents in the Post-ChatGPT World* at SIGIR 2023 [28], and 3) *Goal Awareness for Conversational AI: Proactivity, Non-collaborativity, and Beyond* at ACL 2023 [9].

However, these tutorials mainly introduce advanced designs for building specific web applications with the assistance of LLMs. In our tutorial, we aim to elaborate a comprehensive introduction to cutting-edge research on LLM-powered agents across multiple important web applications.

4 DETAILED SCHEDULE

The following summarizes the detailed schedule of the tutorial:

- (1) Introduction [10 min]
- (2) Background of LLM-powered Agents [35 min]
 - (a) Agent Architecture
 - (b) Tool Learning
- (3) LLM-powered Agents in Social Network [35 min]
 - (a) Social Network Simulation with LLM-based Agents
 - (b) Social Network Problem Solving with LLM-based Agents
- (4) LLM-powered Agents in Recommendation [35 min]
 - (a) User Behavior Simulation with LLM-powered Agents
 - (b) Recommender Agents
- (5) LLM-powered Conversational Agents [35 min]
 - (a) LLM for User Simulation in Conversations
 - (b) Proactive Conversational Agents
- (6) Open Challenges and Beyond [20 min]
 - (a) Trustworthy and Reliable Web Agents
 - (b) Multi-agent Collaboration and Competition
- (7) Summary and Outlook [10 min]

5 STYLE AND INTENDED AUDIENCE

This tutorial is a **lecture-style tutorial**. The target audiences are researchers and practitioners who are interested in web mining, information retrieval, natural language processing, and human-computer interaction. No specific prerequisite knowledge or skill is required. The audience will learn about the state-of-the-art research in web mining, information retrieval, and natural language processing as well as the cutting-edge designs of autonomous AI agents powered by large language models in various web applications.

6 SUPPORTING MATERIALS

(1) **Slides** will be made publicly available; (2) **A survey** [39] is accompanied with this tutorial; (3) **A video teaser**¹ is provided for public promotion.

¹https://drive.google.com/file/d/1YGIUUMxAW-HcX5pQrx6q7Gf_XZpUUdl/view

REFERENCES

- [1] Saaket Agashe, Yue Fan, and Xin Eric Wang. 2023. Evaluating Multi-Agent Coordination Abilities in Large Language Models. *CoRR* (2023).
- [2] Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, et al. 2022. Do as i can, not as i say: Grounding language in robotic affordances. *ArXiv preprint abs/2204.01691* (2022).
- [3] Jitao Bai, Simiao Zhang, and Zhonghao Chen. 2023. Is There Any Social Principle for LLM-Based Agents? *arXiv preprint arXiv:2308.11136* (2023).
- [4] Meeyoung Cha, Alan Mislove, and Krishna P Gummadi. 2009. A measurement-driven analysis of information propagation in the flickr social network. In *WWW 2009*.
- [5] Jin Chen, Zheng Liu, Xu Huang, Chenwang Wu, Qi Liu, Gangwei Jiang, Yuanhao Pu, Yuxuan Lei, Xiaolong Chen, Xingmei Wang, Defu Lian, and Enhong Chen. 2023. When Large Language Models Meet Personalization: Perspectives of Challenges and Opportunities. *CoRR* (2023).
- [6] Liang Chen, Yang Deng, Yatao Bian, Zeyu Qin, Bingzhe Wu, Tat-Seng Chua, and Kam-Fai Wong. 2023. Beyond Factuality: A Comprehensive Evaluation of Large Language Models as Knowledge Generators. *CoRR* (2023).
- [7] Boyi Deng, Wenjie Wang, Fuli Feng, Yang Deng, Qifan Wang, and Xiangnan He. 2023. Attack Prompt Generation for Red Teaming and Defending Large Language Models. In *Findings of ACL: EMNLP 2023*. 2176–2189.
- [8] Xiang Deng, Yu Gu, Boyuan Zheng, Shijie Chen, Samuel Stevens, Boshi Wang, Huan Sun, and Yu Su. 2023. Mind2Web: Towards a Generalist Agent for the Web. In *NeurIPS 2023*.
- [9] Yang Deng, Wenqiang Lei, Minlie Huang, and Tat-Seng Chua. 2023. Goal Awareness for Conversational AI: Proactivity, Non-collaborativity, and Beyond. In *ACL 2023*. 1–10.
- [10] Yang Deng, Wenqiang Lei, Wai Lam, and Tat-Seng Chua. 2023. A Survey on Proactive Dialogue Systems: Problems, Methods, and Prospects. In *IJCAI 2023*.
- [11] Yang Deng, Lizi Liao, Liang Chen, Hongru Wang, Wenqiang Lei, and Tat-Seng Chua. 2023. Prompting and Evaluating Large Language Models for Proactive Dialogues: Clarification, Target-guided, and Non-collaboration. In *Findings of ACL: EMNLP 2023*. 10602–10621.
- [12] Yang Deng, Wenxuan Zhang, Wai Lam, See-Kiong Ng, and Tat-Seng Chua. 2023. Plug-and-Play Policy Planner for Large Language Model Powered Dialogue Agents.
- [13] Yang Deng, Xuan Zhang, Wenxuan Zhang, Yifei Yuan, See-Kiong Ng, and Tat-Seng Chua. 2024. On the Multi-turn Instruction Following for Conversational Web Agents.
- [14] Yao Fu, Hao Peng, Tushar Khot, and Mirella Lapata. 2023. Improving Language Model Negotiation with Self-Play and In-Context Learning from AI Feedback. *CoRR abs/2305.10142* (2023).
- [15] Hiroki Furuta, Ofir Nachum, Kuang-Huei Lee, Yutaka Matsuo, Shixiang Shane Gu, and Izzeddin Gur. 2023. Multimodal Web Navigation with Instruction-Finetuned Foundation Models. *CoRR* (2023).
- [16] Chen Gao, Xiaochong Lan, Zhihong Lu, Jinzhu Mao, Jinghua Piao, Huandong Wang, Depeng Jin, and Yong Li. 2023. S3: Social-network Simulation System with Large Language Model-Empowered Agents. *arXiv preprint arXiv:2307.14984* (2023).
- [17] Navid Ghaffarzadegan, Aritra Majumdar, Ross Williams, and Niyousha Hoseinichimeh. 2023. Generative Agent-Based Modeling: Unveiling Social System Dynamics through Coupling Mechanistic Models with Generative Artificial Intelligence. *arXiv preprint arXiv:2309.11456* (2023).
- [18] Zhiyuan Hu, Yue Feng, Anh Tuan Luu, Bryan Hooi, and Aldo Lipani. 2023. Unlocking the Potential of User Feedback: Leveraging Large Language Model as User Simulator to Enhance Dialogue System. In *CIKM 2023*.
- [19] Wenyue Hua, Lei Li, Shuyuan Xu, Li Chen, and Yongfeng Zhang. 2023. Tutorial on Large Language Models for Recommendation. In *RecSys 2023*.
- [20] Wenlong Huang, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. 2022. Language Models as Zero-Shot Planners: Extracting Actionable Knowledge for Embodied Agents. In *ICML 2022*.
- [21] Xu Huang, Jianxun Lian, Yuxuan Lei, Jing Yao, Defu Lian, and Xing Xie. 2023. Recommender AI Agent: Integrating Large Language Models for Interactive Recommendations.
- [22] Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Yejin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of Hallucination in Natural Language Generation. *ACM Comput. Surv.* 55, 12 (2023), 248:1–248:38.
- [23] Geunwoo Kim, Pierre Baldi, and Stephen McAleer. 2023. Language Models can Solve Computer Tasks.
- [24] Chuyi Kong, Yaxin Fan, Xiang Wan, Feng Jiang, and Benyong Wang. 2023. Large Language Model as a User Simulator. *CoRR abs/2308.11534* (2023).
- [25] Yehuda Koren, Steffen Rendle, and Robert M. Bell. 2022. Advances in Collaborative Filtering. In *Recommender Systems Handbook*. Springer US, 91–142.
- [26] Xiaochong Lan, Chen Gao, Depeng Jin, and Yong Li. 2023. Stance Detection with Collaborative Role-Infused LLM-Based Agents. *arXiv preprint arXiv:2310.10467* (2023).
- [27] Yuan Li, Yixuan Zhang, and Lichao Sun. 2023. MetaAgents: Simulating Interactions of Human Behaviors for LLM-based Task-oriented Coordination via Collaborative Generative Agents.
- [28] Lizi Liao, Grace Hui Yang, and Chirag Shah. 2023. Proactive Conversational Agents in the Post-ChatGPT World. In *SIGIR 2023*.
- [29] Jianghao Lin, Xinyi Dai, Yunjia Xi, Weiwen Liu, Bo Chen, Xiangyang Li, Chenxu Zhu, Hui Feng Guo, Yong Yu, Ruiming Tang, and Weinan Zhang. 2023. How Can Recommender Systems Benefit from Large Language Models: A Survey.
- [30] Jianghao Lin, Rong Shan, Chenxu Zhu, Kounianhua Du, Bo Chen, Shigang Quan, Ruiming Tang, Yong Yu, and Weinan Zhang. 2023. ReLLa: Retrieval-enhanced Large Language Models for Lifelong Sequential Behavior Comprehension in Recommendation.
- [31] Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, et al. 2021. WebGPT: Browser-assisted question-answering with human feedback. *ArXiv preprint abs/2112.09332* (2021).
- [32] OpenAI. 2023. GPT-4 Technical Report. *arXiv:2303.08774*
- [33] Joon Sung Park, Joseph C O'Brien, Carrie J Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. 2023. Generative agents: Interactive simulacra of human behavior. *arXiv preprint arXiv:2304.03442* (2023).
- [34] Yujia Qin, Shengding Hu, Yankai Lin, Weize Chen, Ning Ding, Ganqu Cui, Zheni Zeng, Yufei Huang, Chaojun Xiao, Chi Han, et al. 2023. Tool learning with foundation models. *arXiv preprint arXiv:2304.08354* (2023).
- [35] Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, et al. 2023. ToolLLM: Facilitating Large Language Models to Master 16000+ Real-world APIs. (2023).
- [36] Ivan Sekulic, Mohammad Aliannejadi, and Fabio Crestani. 2022. Evaluating Mixed-initiative Conversational Search Systems via User Simulation. In *WSDM 2022*.
- [37] Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu, and Yueting Zhuang. 2023. Hugginggpt: Solving ai tasks with chatgpt and its friends in huggingface. *arXiv preprint arXiv:2303.17580* (2023).
- [38] Yubo Shu, Haonan Zhang, Hansu Gu, Peng Zhang, Tun Lu, Dongsheng Li, and Ning Gu. 2023. RAH! RecSys-Assistant-Human: A Human-Centered Recommendation Framework with LLM Agents.
- [39] Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai Tang, Xu Chen, Yankai Lin, et al. 2023. A survey on large language model based autonomous agents. *arXiv preprint arXiv:2308.11432* (2023).
- [40] Lei Wang, Jingsen Zhang, Xu Chen, Yankai Lin, Ruihua Song, Wayne Xin Zhao, and Ji-Rong Wen. 2023. RecAgent: A Novel Simulation Paradigm for Recommender Systems. *arXiv* (2023).
- [41] Xiaolei Wang, Xinyu Tang, Wayne Xin Zhao, Jingyuan Wang, and Ji-Rong Wen. 2023. Rethinking the Evaluation for Conversational Recommendation in the Era of Large Language Models. In *EMNLP 2023*.
- [42] Yancheng Wang, Ziyang Jiang, Zheng Chen, Fan Yang, Yingxue Zhou, Eunah Cho, Xing Fan, Xiaojiang Huang, Yanbin Lu, and Yingzhen Yang. 2023. RecMind: Large Language Model Powered Agent For Recommendation.
- [43] Zhenhailong Wang, Shaoguang Mao, Wenshan Wu, Tao Ge, Furu Wei, and Heng Ji. 2023. Unleashing Cognitive Synergy in Large Language Models: A Task-Solving Agent through Multi-Persona Self-Collaboration. *CoRR* (2023).
- [44] Zhengyi Yang, Jiancan Wu, Yanchen Luo, Jizhi Zhang, Yancheng Yuan, An Zhang, Xiang Wang, and Xiangnan He. 2023. Large Language Model Can Interpret Latent Space of Sequential Recommender.
- [45] Shunyu Yao, Howard Chen, John Yang, and Karthik Narasimhan. 2022. Webshop: Towards scalable real-world web interaction with grounded language agents. *ArXiv preprint abs/2207.01206* (2022).
- [46] Xiao Yu, Maximilian Chen, and Zhou Yu. 2023. Prompt-Based Monte-Carlo Tree Search for Goal-Oriented Dialogue Policy Planning. In *EMNLP 2023*.
- [47] Herman Yuliansyah, Zulaiha Ali Othman, and Azuraliza Abu Bakar. 2020. Taxonomy of link prediction for social network analysis: a review. *IEEE Access* 8 (2020), 183470–183487.
- [48] An Zhang, Leheng Sheng, Yuxin Chen, Hao Li, Yang Deng, Xiang Wang, and Tat-Seng Chua. 2023. On Generative Agents in Recommendation.
- [49] Junjie Zhang, Yupeng Hou, Ruobing Xie, Wenqi Sun, Julian McAuley, Wayne Xin Zhao, Leyu Lin, and Ji-Rong Wen. 2023. AgentCF: Collaborative Learning with Autonomous Language Agents for Recommender Systems.
- [50] Qiang Zhang, Jason Naradowsky, and Yusuke Miyao. 2023. Ask an Expert: Leveraging Language Models to Improve Strategic Reasoning in Goal-Oriented Dialogue Models. In *Findings of ACL: ACL 2023*.
- [51] Shuo Zhang and Krisztian Balog. 2020. Evaluating Conversational Recommender Systems via User Simulation. In *KDD 2020*.
- [52] Xizhou Zhu, Yuntao Chen, Hao Tian, Chenxin Tao, Weijie Su, Chenyu Yang, Gao Huang, Bin Li, Lewei Lu, Xiaogang Wang, Yu Qiao, Zhaoxiang Zhang, and Jifeng Dai. 2023. Ghost in the Minecraft: Generally Capable Agents for Open-World Environments via Large Language Models with Text-based Knowledge and Memory. *CoRR* (2023).