main 1 branch 0 tags Go to file Code

WooooDyy Merge pull request #28 from cyl628/main ... 95b2153 yesterday 96 commits

assets	update: TOC and Figure 1	last week
.gitignore	init: first commit	last week
README.md	Merge branch 'WooooDyy:main' into main	2 days ago

README.md

The Rise and Potential of Large Language Model Based Agents: A Survey

🔥 **Must-read papers for LLM-based agents.**

🏃 **Coming soon: Add one-sentence intro to each paper.**

🔔 News

- 🥳 [2023/09/20] This project has been listed on GitHub Trendings! It is a great honor!
- 💥 [2023/09/15] Our survey is released! See The Rise and Potential of Large Language Model Based Agents: A Survey for the paper!
- ✨ [2023/09/14] We create this repository to maintain a paper list on LLM-based agents. More papers are coming soon!

🌟 Introduction

For a long time, humanity has pursued artificial intelligence (AI) equivalent to or surpassing human level, with AI agents considered as a promising vehicle of this pursuit. AI agents are artificial entities that sense their environment, make decisions, and take actions.

Due to the versatile and remarkable capabilities they demonstrate, large language models (LLMs) are regarded as potential sparks for Artificial General Intelligence (AGI), offering hope for building general AI agents. Many research efforts have leveraged LLMs as the foundation to build AI agents and have achieved significant progress.

In this repository, we provide a systematic and comprehensive survey on LLM-based agents, and list some must-read papers.

Specifically, we start by the general conceptual framework for LLM-based agents: comprising three main components: brain, perception, and action, and the framework can be tailored to suit different applications. Subsequently, we explore the extensive applications of LLM-based agents in three aspects: single-agent scenarios, multi-agent scenarios, and human-agent cooperation. Following this, we delve into agent societies, exploring the behavior and personality of LLM-based agents, the social phenomena that emerge when they

About

The paper list of the 86-page paper "The Rise and Potential of Large Language Model Based Agents: A Survey" by Zhiheng Xi et al.

arxiv.org/abs/2309.07864

nlp agent survey large-language-models llm

Readme
Activity
2.3k stars
57 watching
132 forks
Report repository

Releases

No releases published

Packages

No packages published

Contributors 16

+ 5 contributors

form societies, and the insights they offer for human society. Finally, we discuss a range of key topics and open problems within the field.

We greatly appreciate any contributions via PRs, issues, emails, or other methods.

Table of Content (ToC)

1. The Birth of An Agent: Construction of LLM-based Agents

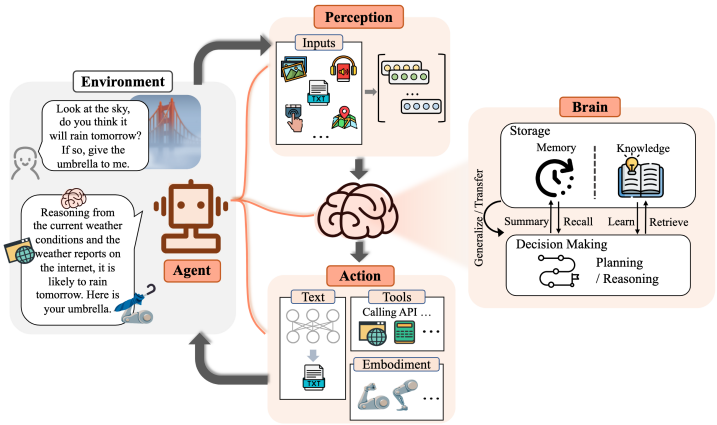

1.1 Brain: Primarily Composed of An LLM

1.1.1 Natural Language Interaction

High-quality generation

- [2023/08] **A Multitask, Multilingual, Multimodal Evaluation of ChatGPT on Reasoning, Hallucination, and Interactivity.** *Yejin Bang et al. arXiv.* [paper]
- [2023/06] **LLM-Eval: Unified Multi-Dimensional Automatic Evaluation for Open-Domain Conversations with Large Language Models.** *Yen-Ting Lin et al. arXiv.* [paper]
- [2023/04] **Is ChatGPT a Highly Fluent Grammatical Error Correction System? A Comprehensive Evaluation.** *Tao Fang et al. arXiv.* [paper]

Deep understanding

- [2023/06] **Clever Hans or Neural Theory of Mind? Stress Testing Social Reasoning in Large Language Models.** *Natalie Shapira et al. arXiv.* [paper]
- [2022/08] **Inferring Rewards from Language in Context.** *Jessy Lin et al. ACL.* [paper]
- [2021/10] **Theory of Mind Based Assistive Communication in Complex Human Robot Cooperation.** *Moritz C. Buehler et al. arXiv.* [paper]

1.1.2 Knowledge

Pretrain model

- [2023/04] **Learning Distributed Representations of Sentences from Unlabelled Data.** *Felix Hill(University of Cambridge) et al. arXiv.* [paper]
- [2020/02] **How Much Knowledge Can You Pack Into the Parameters of a Language Model?** *Adam Roberts(Google) et al. arXiv.* [paper]
- [2020/01] **Scaling Laws for Neural Language Models.** *Jared Kaplan(Johns Hopkins University) et al. arXiv.* [paper]
- [2017/12] **Commonsense Knowledge in Machine Intelligence.** *Niket Tandon(Allen Institute for Artificial Intelligence) et al. SIGMOD.* [paper]
- [2011/03] **Natural Language Processing (almost) from Scratch.** *Ronan Collobert(Princeton) et al. arXiv.* [paper]]

Linguistic knowledge

- [2023/02] **A Multitask, Multilingual, Multimodal Evaluation of ChatGPT on Reasoning, Hallucination, and Interactivity.** *Yejin Bang et al. arXiv.* [paper]
- [2021/06] **Probing Pre-trained Language Models for Semantic Attributes and their Values.** *Meriem Beloucif et al. EMNLP.* [paper]
- [2020/10] **Probing Pretrained Language Models for Lexical Semantics.** *Ivan Vulić et al. arXiv.* [paper]
- [2019/04] **A Structural Probe for Finding Syntax in Word Representations.** *John Hewitt et al. ACL.* [paper]
- [2016/04] **Improved Automatic Keyword Extraction Given More Semantic Knowledge.** *H Leung. Systems for Advanced Applications.* [paper]

Commonsense knowledge

- [2022/10] **Language Models of Code are Few-Shot Commonsense Learners.** *Aman Madaan et al.arXiv.* [paper]
- [2021/04] **Relational World Knowledge Representation in Contextual Language Models: A Review.** *Tara Safavi et al. arXiv.* [paper]
- [2019/11] **How Can We Know What Language Models Know?** *Zhengbao Jiang et al.arXiv.* [paper]

Actionable knowledge

- [2023/07] **Large language models in medicine.** *Arun James Thirunavukarasu et al. nature.* [paper]
- [2023/06] **DS-1000: A Natural and Reliable Benchmark for Data Science Code Generation.** *Yuhang Lai et al. ICML.* [paper]
- [2022/10] **Language Models of Code are Few-Shot Commonsense Learners.** *Aman Madaan et al. arXiv.* [paper]
- [2022/02] **A Systematic Evaluation of Large Language Models of Code.** *Frank F. Xu et al.arXiv.* [paper]
- [2021/10] **Training Verifiers to Solve Math Word Problems.** *Karl Cobbe et al. arXiv.* [paper]

Potential issues of knowledge

- [2023/05] **Editing Large Language Models: Problems, Methods, and Opportunities.** *Yunzhi Yao et al. arXiv.* [paper]
- [2023/05] **Self-Checker: Plug-and-Play Modules for Fact-Checking with Large Language Models.** *Miaoran Li et al. arXiv.* [paper]
- [2023/05] **CRITIC: Large Language Models Can Self-Correct with Tool-Interactive Critiquing.** *Zhibin Gou et al. arXiv.* [paper]
- [2023/04] **Tool Learning with Foundation Models.** *Yujia Qin et al. arXiv.* [paper]
- [2023/03] **SelfCheckGPT: Zero-Resource Black-Box Hallucination Detection for Generative Large Language Models.** *Potsawee Manakul et al. arXiv.* [paper]
- [2022/06] **Memory-Based Model Editing at Scale.** *Eric Mitchell et al. arXiv.* [paper]
- [2022/04] **A Review on Language Models as Knowledge Bases.** *Badr AlKhamissi et al.arXiv.* [paper]
- [2021/04] **Editing Factual Knowledge in Language Models.** *Nicola De Cao et al.arXiv.* [paper]

- [2017/08] **Measuring Catastrophic Forgetting in Neural Networks.** *Ronald Kemker et al.arXiv.* [paper]

1.1.3 Memory

Memory capability

Raising the length limit of Transformers

- [2023/05] **Randomized Positional Encodings Boost Length Generalization of Transformers.** *Anian Ruoss (DeepMind) et al. arXiv.* [paper] [code]
- [2023-03] **CoLT5: Faster Long-Range Transformers with Conditional Computation.** *Joshua Ainslie (Google Research) et al. arXiv.* [paper]
- [2022/03] **Efficient Classification of Long Documents Using Transformers.** *Hyunji Hayley Park (Illinois University) et al. arXiv.* [paper] [code]
- [2021/12] **LongT5: Efficient Text-To-Text Transformer for Long Sequences.** *Mandy Guo (Google Research) et al. arXiv.* [paper] [code]
- [2019/10] **BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension.** *Michael Lewis(Facebook AI) et al. arXiv.* [paper] [code]

Summarizing memory

- [2023/09] **Empowering Private Tutoring by Chaining Large Language Models** *Yulin Chen (Tsinghua University) et al. arXiv.* [paper]
- [2023/08] **ExpeL: LLM Agents Are Experiential Learners.** *Andrew Zhao (Tsinghua University) et al. arXiv.* [paper] [code]
- [2023/08] **ChatEval: Towards Better LLM-based Evaluators through Multi-Agent Debate.** *Chi-Min Chan (Tsinghua University) et al. arXiv.* [paper] [code]
- [2023/05] **MemoryBank: Enhancing Large Language Models with Long-Term Memory.** *Wanjun Zhong (Harbin Institute of Technology) et al. arXiv.* [paper] [code]
- [2023/04] **Generative Agents: Interactive Simulacra of Human Behavior.** *Joon Sung Park (Stanford University) et al. arXiv.* [paper] [code]
- [2023/04] **Unleashing Infinite-Length Input Capacity for Large-scale Language Models with Self-Controlled Memory System.** *Xinnian Liang(Beihang University) et al. arXiv.* [paper] [code]
- [2023/03] **Reflexion: Language Agents with Verbal Reinforcement Learning.** *Noah Shinn (Northeastern University) et al. arXiv.* [paper] [code]
- [2023/05] **RecurrentGPT: Interactive Generation of (Arbitrarily) Long Text.** Wangchunshu Zhou (AIWaves) et al. arXiv.* [paper] [code]

Compressing memories with vectors or data structures

- [2023/07] **Communicative Agents for Software Development.** *Chen Qian (Tsinghua University) et al. arXiv.* [paper] [code]
- [2023/06] **ChatDB: Augmenting LLMs with Databases as Their Symbolic Memory.** *Chenxu Hu(Tsinghua University) et al. arXiv.* [paper] [code]
- [2023/05] **Ghost in the Minecraft: Generally Capable Agents for Open-World Environments via Large Language Models with Text-based Knowledge and Memory.** *Xizhou Zhu (Tsinghua University) et al. arXiv.* [paper] [code]
- [2023/05] **RET-LLM: Towards a General Read-Write Memory for Large Language Models.** *Ali Modarressi (LMU Munich) et al. arXiv.* [paper] [code]
- [2023/05] **RecurrentGPT: Interactive Generation of (Arbitrarily) Long Text.** Wangchunshu Zhou (AIWaves) et al. arXiv.* [paper] [code]

Memory retrieval

- [2023/08] **Memory Sandbox: Transparent and Interactive Memory Management for Conversational Agents.** *Ziheng Huang(University of California—San Diego) et al. arXiv.* [paper]
- [2023/08] **AgentSims: An Open-Source Sandbox for Large Language Model Evaluation.** *Jiaju Lin (PTA Studio) et al. arXiv.* [paper] [project page] [code]
- [2023/06] **ChatDB: Augmenting LLMs with Databases as Their Symbolic Memory.** *Chenxu Hu(Tsinghua University) et al. arXiv.* [paper] [code]
- [2023/05] **MemoryBank: Enhancing Large Language Models with Long-Term Memory.** *Wanjun Zhong (Harbin Institute of Technology) et al. arXiv.* [paper] [code]
- [2023/04] **Generative Agents: Interactive Simulacra of Human Behavior.** *Joon Sung Park (Stanford) et al. arXiv.* [paper] [code]
- [2023/05] **RecurrentGPT: Interactive Generation of (Arbitrarily) Long Text.** Wangchunshu Zhou (AIWaves) et al. arXiv.* [paper] [code]

1.1.4 Reasoning & Planning

Reasoning

- [2023/05] **Self-Polish: Enhance Reasoning in Large Language Models via Problem Refinement.** *Zhiheng Xi (Fudan University) et al. arXiv.* [paper] [code]

- [2023-03] **Large Language Models are Zero-Shot Reasoners.** *Takeshi Kojima (The University of Tokyo) et al. arXiv.* [paper][code]
- [2023/03] **Self-Refine: Iterative Refinement with Self-Feedback.** *Aman Madaan (Carnegie Mellon University) et al. arXiv.* [paper] [code]
- [2022/05] **Selection-Inference: Exploiting Large Language Models for Interpretable Logical Reasoning.** *Antonia Creswell (DeepMind) et al. arXiv.* [paper]
- [2022/03] **Self-Consistency Improves Chain of Thought Reasoning in Language Models.** *Xuezhi Wang(Google Research) et al. arXiv.* [paper] [code]
- [2022/01] **Chain-of-Thought Prompting Elicits Reasoning in Large Language Models.** *Jason Wei (Google Research,) et al. arXiv.* [paper]

Planning

Plan formulation

- [2023/05] **Tree of Thoughts: Deliberate Problem Solving with Large Language Models.** *Shunyu Yao (Princeton University) et al. arXiv.* [paper] [code]
- [2023/05] **Plan, Eliminate, and Track -- Language Models are Good Teachers for Embodied Agents.** *Yue Wu(Carnegie Mellon University) et al. arXiv.* [paper]
- [2023/05] **Reasoning with Language Model is Planning with World Model.** *Shibo Hao (UC San Diego) et al. arXiv.* [paper] [code]
- [2023/05] **SwiftSage: A Generative Agent with Fast and Slow Thinking for Complex Interactive Tasks.** *Bill Yuchen Lin (Allen Institute for Artificial Intelligence) et al. arXiv.* [paper] [code]
- [2023/04] **LLM+P: Empowering Large Language Models with Optimal Planning Proficiency.** *Bo Liu (University of Texas at Austin) et al. arXiv.* [paper] [code]
- [2023/03] **HuggingGPT: Solving AI Tasks with ChatGPT and its Friends in Hugging Face.** *Yongliang Shen (Microsoft Research Asia) et al. arXiv.* [paper] [code]
- [2023/02] **Describe, Explain, Plan and Select: Interactive Planning with Large Language Models Enables Open-World Multi-Task Agents.** *ZiHao Wang (Peking University) et al. arXiv.* [paper] [code]
- [2022/05] **Least-to-Most Prompting Enables Complex Reasoning in Large Language Models.** *Denny Zhou (Google Research) et al. arXiv.* [paper]
- [2022/05] **MRKL Systems: A modular, neuro-symbolic architecture that combines large language models, external knowledge sources and discrete reasoning.** *Ehud Karpas (AI21 Labs) et al. arXiv.* [paper]
- [2022/04] **Do As I Can, Not As I Say: Grounding Language in Robotic Affordances.** *Michael Ahn (Robotics at Google) et al. arXiv.* [paper]
- [2023/05] **Agents: An Open-source Framework for Autonomous Language Agents.** Wangchunshu Zhou (AIWaves) et al. arXiv.* [paper] [code]

Plan reflection

- [2023/08] **SelfCheck: Using LLMs to Zero-Shot Check Their Own Step-by-Step Reasoning.** *Ning Miao (University of Oxford) et al. arXiv.* [paper] [code]
- [2023/05] **ChatCoT: Tool-Augmented Chain-of-Thought Reasoning on Chat-based Large Language Models.** *Zhipeng Chen (Renmin University of China) et al. arXiv.* [paper] [code]
- [2023/05] **Voyager: An Open-Ended Embodied Agent with Large Language Models.** *Guanzhi Wang (NVIDA) et al. arXiv.* [paper] [code]
- [2023/03] **Chat with the Environment: Interactive Multimodal Perception Using Large Language Models.** *Xufeng Zhao (University Hamburg) et al. arXiv.* [paper] [code]
- [2022/12] **LLM-Planner: Few-Shot Grounded Planning for Embodied Agents with Large Language Models.** *Chan Hee Song (The Ohio State University) et al. arXiv.* [paper] [code]
- [2022/10] **ReAct: Synergizing Reasoning and Acting in Language Models.** *Shunyu Yao (Princeton University) et al. arXiv.* [paper] [code]
- [2022/07] **Inner Monologue: Embodied Reasoning through Planning with Language Models.** *Wenlong Huang (Robotics at Google) et al. arXiv.* [paper] [code]
- [2021/10] **AI Chains: Transparent and Controllable Human-AI Interaction by Chaining Large Language Model Prompts.** *Tongshuang Wu (University of Washington) et al. arXiv.* [paper]

1.1.5 Transferability and Generalization

Unseen task generalization

- [2023/05] **Training language models to follow instructions with human feedback.** *Long Ouyang et al. NeurIPS.* [paper]
- [2023/01] **Multitask Prompted Training Enables Zero-Shot Task Generalization.** *Victor Sanh et al. ICLR.* [paper]
- [2022/10] **Scaling Instruction-Finetuned Language Models.** *Hyung Won Chung et al. arXiv.* [paper]
- [2022/08] **Finetuned Language Models are Zero-Shot Learners.** *Jason Wei et al. ICLR.* [paper]

In-context learning

- [2023/08] **Images Speak in Images: A Generalist Painter for In-Context Visual Learning.** *Xinlong Wang et al. IEEE.* [paper]
- [2023/08] **Neural Codec Language Models are Zero-Shot Text to Speech Synthesizers.** *Chengyi Wang et al. arXiv.* [paper]
- [2023/07] **A Survey for In-context Learning.** *Qingxiu Dong et al. arXiv.* [paper]
- [2023/05] **Language Models are Few-Shot Learners.** *Tom B. Brown (OpenAI) et al. NeurIPS.* [paper]

Continual learning

- [2023/07] **Progressive Prompts: Continual Learning for Language Models.** *Razdaibiedina et al. arXiv.* [paper]
- [2023/07] **Voyager: An Open-Ended Embodied Agent with Large Language Models.** *Guanzhi Wang et al. arXiv.* [paper]
- [2023/01] **A Comprehensive Survey of Continual Learning: Theory, Method and Application.** *Liyuan Wang et al. arXiv.* [paper]
- [2022/11] **Continual Learning of Natural Language Processing Tasks: A Survey.** *Zixuan Ke et al. arXiv.* [paper]

1.2 Perception: Multimodal Inputs for LLM-based Agents

1.2.1 Visual

- [2023/05] **Language Is Not All You Need: Aligning Perception with Language Models.** *Shaohan Huang et al. arXiv.* [paper]]
- [2023/05] **InstructBLIP: Towards General-purpose Vision-Language Models with Instruction Tuning.** *Wenliang Dai et al. arXiv.* [paper]
- [2023/05] **MultiModal-GPT: A Vision and Language Model for Dialogue with Humans.** *Tao Gong et al. arXiv.* [paper]
- [2023/05] **PandaGPT: One Model To Instruction-Follow Them All.** *Yixuan Su et al. arXiv.* [paper]
- [2023/04] **Visual Instruction Tuning.** *Haotian Liu et al. arXiv.* [paper]
- [2023/04] **MiniGPT-4: Enhancing Vision-Language Understanding with Advanced Large Language Models.** *Deyao Zhu. arXiv.* [paper]
- [2023/01] **BLIP-2: Bootstrapping Language-Image Pre-training with Frozen Image Encoders and Large Language Models.** *Junnan Li et al. arXiv.* [paper]
- [2022/04] **Flamingo: a Visual Language Model for Few-Shot Learning.** *Jean-Baptiste Alayrac et al. arXiv.* [paper]
- [2021/10] **MobileViT: Light-weight, General-purpose, and Mobile-friendly Vision Transformer.** *Sachin Mehta et al.arXiv.* [paper]
- [2021/05] **MLP-Mixer: An all-MLP Architecture for Vision.** *Ilya Tolstikhin et al.arXiv.* [paper]
- [2020/10] **An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale.** *Alexey Dosovitskiy et al. arXiv.* [paper]
- [2017/11] **Neural Discrete Representation Learning.** *Aaron van den Oord et al. arXiv.* [paper]

1.2.2 Audio

- [2023/06] **Video-LLaMA: An Instruction-tuned Audio-Visual Language Model for Video Understanding.** *Hang Zhang et al. arXiv.* [paper]
- [2023/05] **X-LLM: Bootstrapping Advanced Large Language Models by Treating Multi-Modalities as Foreign Languages.** *Feilong Chen et al. arXiv.* [paper]
- [2023/05] **InternGPT: Solving Vision-Centric Tasks by Interacting with ChatGPT Beyond Language.** *Zhaoyang Liu et al. arXiv.* [paper]
- [2023/04] **AudioGPT: Understanding and Generating Speech, Music, Sound, and Talking Head.** *Rongjie Huang et al. arXiv.* [paper]
- [2023/03] **HuggingGPT: Solving AI Tasks with ChatGPT and its Friends in Hugging Face.** *Yongliang Shen et al. arXiv.* [paper]
- [2021/06] **HuBERT: Self-Supervised Speech Representation Learning by Masked Prediction of Hidden Units.** *Wei-Ning Hsu et al. arXiv.* [paper]
- [2021/04] **AST: Audio Spectrogram Transformer.** *Yuan Gong et al. arXiv.* [paper]

1.3 Action: Expand Action Space of LLM-based Agents

1.3.1 Tool Using

- [2023/07] **ToolLLM: Facilitating Large Language Models to Master 16000+ Real-world APIs.** *Yujia Qin et al. arXiv.* [paper] [code] [dataset]
- [2023/05] **Large Language Models as Tool Makers.** *Tianle Cai et al. arXiv.* [paper] [code]
- [2023/05] **CREATOR: Disentangling Abstract and Concrete Reasonings of Large Language Models through Tool Creation.** *Cheng Qian et al. arXiv.* [paper]
- [2023/04] **Tool Learning with Foundation Models.** *Yujia Qin et al. arXiv.* [paper] [code]
- [2023/04] **ChemCrow: Augmenting large-language models with chemistry tools.** *Andres M Bran (Laboratory of Artificial Chemical Intelligence, ISIC, EPFL) et al. arXiv.* [paper] [code]

- [2023/04] **GeneGPT: Augmenting Large Language Models with Domain Tools for Improved Access to Biomedical Information.** *Qiao Jin, Yifan Yang, Qingyu Chen, Zhiyong Lu. arXiv.* [paper] [code]
- [2023/04] **OpenAGI: When LLM Meets Domain Experts.** *Yingqiang Ge et al. arXiv.* [paper] [code]
- [2023/03] **HuggingGPT: Solving AI Tasks with ChatGPT and its Friends in Hugging Face.** *Yongliang Shen et al. arXiv.* [paper] [code]
- [2023/03] **Visual ChatGPT: Talking, Drawing and Editing with Visual Foundation Models.** *Chenfei Wu et al. arXiv.* [paper] [code]
- [2023/02] **Augmented Language Models: a Survey.** *Grégoire Mialon et al. arXiv.* [paper]
- [2023/02] **Toolformer: Language Models Can Teach Themselves to Use Tools.** *Timo Schick et al. arXiv.* [paper]
- [2022/05] **TALM: Tool Augmented Language Models.** *Aaron Parisi et al. arXiv.* [paper]
- [2022/05] **MRKL Systems: A modular, neuro-symbolic architecture that combines large language models, external knowledge sources and discrete reasoning.** *Ehud Karpas et al. arXiv.* [paper]
- [2022/04] **Do As I Can, Not As I Say: Grounding Language in Robotic Affordances.** *Michael Ahn et al. arXiv.* [paper]
- [2021/12] **WebGPT: Browser-assisted question-answering with human feedback.** *Reiichiro Nakano et al. arXiv.* [paper]
- [2021/07] **Evaluating Large Language Models Trained on Code.** *Mark Chen et al. arXiv.* [paper] [code]

1.3.2 Embodied Action

- [2023/07] **Interactive language: Talking to robots in real time.** *Corey Lynch et al. IEEE(RAL)* [paper]
- [2023/05] **Voyager: An open-ended embodied agent with large language models.** *Guanzhi Wang et al. Arxiv.* [paper]
- [2023/05] **AVLEN: Audio-Visual-Language Embodied Navigation in 3D Environments.** *Sudipta Paul et al. NeurIPS.* [paper]
- [2023/05] **EmbodiedGPT: Vision-Language Pre-Training via Embodied Chain of Thought.** *Yao Mu et al. Arxiv* [paper] [code]
- [2023/05] **NavGPT: Explicit Reasoning in Vision-and-Language Navigation with Large Language Models.** *Gengze Zhou et al. Arxiv* [paper]
- [2023/05] **AlphaBlock: Embodied Finetuning for Vision-Language Reasoning in Robot Manipulation.** *Chuhao Jin et al. Arxiv* [paper]
- [2023/03] **PaLM-E: An Embodied Multimodal Language Model.** *Danny Driess et al. Arxiv.* [paper]
- [2023/03] **Reflexion: Language Agents with Verbal Reinforcement Learning.** *Noah Shinn et al. Arxiv* [paper] [code]
- [2023/02] **Collaborating with language models for embodied reasoning.** *Ishita Dasgupta et al. Arxiv.* [paper]
- [2023/02] **Code as Policies: Language Model Programs for Embodied Control.** *Jacky Liang et al. IEEE(ICRA).* [paper]
- [2022/10] **ReAct: Synergizing Reasoning and Acting in Language Models.** *Shunyu Yao et al. Arxiv* [paper] [code]
- [2022/10] **Instruction-Following Agents with Multimodal Transformer.** *Hao Liu et al. CVPR* [paper] [code]
- [2022/07] **Inner Monologue: Embodied Reasoning through Planning with Language Models.** *Wenlong Huang et al. Arxiv.* [paper]
- [2022/07] **LM-Nav: Robotic Navigation with Large Pre-Trained Models of Language, Vision, and Action.** *Dhruv Shahet al. CoRL* [paper] [code]
- [2022/04] **Do As I Can, Not As I Say: Grounding Language in Robotic Affordances.** *Michael Ahn et al. Arxiv.* [paper]
- [2022/01] **A Survey of Embodied AI: From Simulators to Research Tasks.** *Jiafei Duan et al. IEEE(TETCI).* [paper]
- [2022/01] **Language Models as Zero-Shot Planners: Extracting Actionable Knowledge for Embodied Agents.** *Wenlong Huang et al. Arxiv.* [paper] [code]
- [2020/04] **Experience Grounds Language.** *Yonatan Bisk et al. EMNLP* [paper]
- [2019/03] **Review of Deep Reinforcement Learning for Robot Manipulation.** *Hai Nguyen et al. IEEE(IRC).* [paper]
- [2005/01] **The Development of Embodied Cognition: Six Lessons from Babies.** *Linda Smith et al. Artificial Life.* [paper]

2. Agents in Practice: Applications of LLM-based Agents

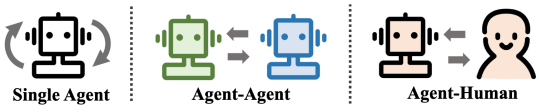

2.1 General Ability of Single Agent

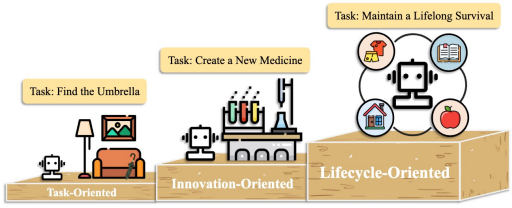

2.1.1 Task-oriented Deployment

In web scenarios

- [2023/07] **WebArena: A Realistic Web Environment for Building Autonomous Agents.** *Shuyan Zhou (CMU) et al. arXiv.* [paper] [code]
- [2023/07] **A Real-World WebAgent with Planning, Long Context Understanding, and Program Synthesis.** *Izzeddin Gur (DeepMind) et al. arXiv.* [paper]
- [2023/06] **SYNAPSE: Leveraging Few-Shot Exemplars for Human-Level Computer Control.** *Longtao Zheng (Nanyang Technological University) et al. arXiv.* [paper] [code]
- [2023/06] **Mind2Web: Towards a Generalist Agent for the Web.** *Xiang Deng (The Ohio State University) et al. arXiv.* [paper] [code]
- [2023/05] **Multimodal Web Navigation with Instruction-Finetuned Foundation Models.** *Hiroki Furuta (The University of Tokyo) et al. arXiv.* [paper]
- [2023/03] **Language Models can Solve Computer Tasks.** *Geunwoo Kim (University of California) et al. arXiv.* [paper] [code]
- [2022/07] **WebShop: Towards Scalable Real-World Web Interaction with Grounded Language Agents.** *Shunyu Yao (Princeton University) et al. arXiv.* [paper] [code]
- [2021/12] **WebGPT: Browser-assisted question-answering with human feedback.** *Reiichiro Nakano (OpenAI) et al. arXiv.* [paper]
- [2023/05] **Agents: An Open-source Framework for Autonomous Language Agents.** Wangchunshu Zhou (AIWaves) et al. arXiv.* [paper] [code]

In life scenarios

- [2023/08] **InterAct: Exploring the Potentials of ChatGPT as a Cooperative Agent.** *Po-Lin Chen et al. arXiv.* [paper]
- [2023/05] **Plan, Eliminate, and Track -- Language Models are Good Teachers for Embodied Agents.** *Yue Wu (CMU) et al. arXiv.* [paper]
- [2023/05] **Augmenting Autotelic Agents with Large Language Models.** *Cédric Colas (MIT) et al. arXiv.* [paper]
- [2023/03] **Planning with Large Language Models via Corrective Re-prompting.** *Shreyas Sundara Raman (Brown University) et al. arXiv.* [paper]
- [2022/10] **Generating Executable Action Plans with Environmentally-Aware Language Models.** *Maitrey Gramopadhye (University of North Carolina at Chapel Hill) et al. arXiv.* [paper] [code]
- [2022/01] **Language Models as Zero-Shot Planners: Extracting Actionable Knowledge for Embodied Agents.** *Wenlong Huang (UC Berkeley) et al. arXiv.* [paper] [code]

2.1.2 Innovation-oriented Deployment

- [2023/08] **The Hitchhiker's Guide to Program Analysis: A Journey with Large Language Models.** *Haonan Li (UC Riverside) et al. arXiv.* [paper]
- [2023/08] **ChatMOF: An Autonomous AI System for Predicting and Generating Metal-Organic Frameworks.** *Yeonghun Kang (Korea Advanced Institute of Science and Technology) et al. arXiv.* [paper]
- [2023/07] **Math Agents: Computational Infrastructure, Mathematical Embedding, and Genomics.** *Melanie Swan (University College London) et al. arXiv.* [paper]
- [2023/06] **Towards Autonomous Testing Agents via Conversational Large Language Models.** *Robert Feldt (Chalmers University of Technology) et al. arXiv.* [paper]
- [2023/04] **Emergent autonomous scientific research capabilities of large language models.** *Daniil A. Boiko (CMU) et al. arXiv.* [paper]
- [2023/04] **ChemCrow: Augmenting large-language models with chemistry tools.** *Andres M Bran (Laboratory of Artificial Chemical Intelligence, ISIC, EPFL) et al. arXiv.* [paper] [code]
- [2022/03] **ScienceWorld: Is your Agent Smarter than a 5th Grader?** *Ruoyao Wang (University of Arizona) et al. arXiv.* [paper] [code]

2.1.3 Lifecycle-oriented Deployment

- [2023/05] **Voyager: An Open-Ended Embodied Agent with Large Language Models.** *Guanzhi Wang (NVIDA) et al. arXiv.* [paper] [code]
- [2023/05] **Ghost in the Minecraft: Generally Capable Agents for Open-World Environments via Large Language Models with Text-based Knowledge and Memory.** *Xizhou Zhu (Tsinghua University) et al. arXiv.* [paper] [code]
- [2023/03] **Plan4MC: Skill Reinforcement Learning and Planning for Open-World Minecraft Tasks.** *Haoqi Yuan (PKU) et al. arXiv.* [paper] [code]

- [2023/02] **Describe, Explain, Plan and Select: Interactive Planning with Large Language Models Enables Open-World Multi-Task Agents.** *Zihao Wang (PKU) et al. arXiv.* [paper] [code]
- [2023/01] **Do Embodied Agents Dream of Pixelated Sheep: Embodied Decision Making using Language Guided World Modelling.** *Kolby Nottingham (University of California Irvine, Irvine) et al. arXiv.* [paper] [code]

2.2 Coordinating Potential of Multiple Agents

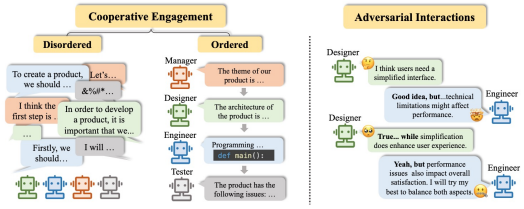

2.2.1 Cooperative Interaction for Complementarity

Disordered cooperation

- [2023/07] **Unleashing Cognitive Synergy in Large Language Models: A Task-Solving Agent through Multi-Persona Self-Collaboration.** *Zhenhailong Wang (University of Illinois Urbana-Champaign) et al. arXiv.* [paper] [code]
- [2023/07] **RoCo: Dialectic Multi-Robot Collaboration with Large Language Models.** *Zhao Mandi, Shreeya Jain, Shuran Song (Columbia University) et al. arXiv.* [paper] [code]
- [2023/04] **ChatLLM Network: More brains, More intelligence.** *Rui Hao (Beijing University of Posts and Telecommunications) et al. arXiv.* [paper]
- [2023/01] **Blind Judgement: Agent-Based Supreme Court Modelling With GPT.** *Sil Hamilton (McGill University). arXiv.* [paper]
- [2023/05] **Agents: An Open-source Framework for Autonomous Language Agents.** Wangchunshu Zhou (AIWaves) et al. arXiv.* [paper] [code]

Ordered cooperation

- [2023/08] **CGMI: Configurable General Multi-Agent Interaction Framework.** *Shi Jinxin (East China Normal University) et al. arXiv.* [paper]
- [2023/08] **ProAgent: Building Proactive Cooperative AI with Large Language Models.** *Ceyao Zhang (The Chinese University of Hong Kong, Shenzhen) et al. arXiv.* [paper] [code]
- [2023/08] **AgentVerse: Facilitating Multi-Agent Collaboration and Exploring Emergent Behaviors in Agents.** *Weize Chen (Tsinghua University) et al. arXiv.* [paper] [code]
- [2023/08] **AutoGen: Enabling Next-Gen LLM Applications via Multi-Agent Conversation Framework.** *Qingyun Wu (Pennsylvania State University) et al. arXiv.* [paper] [code]
- [2023/08] **MetaGPT: Meta Programming for Multi-Agent Collaborative Framework.** *Sirui Hong (DeepWisdom) et al. arXiv.* [paper] [code]
- [2023/07] **Communicative Agents for Software Development.** *Chen Qian (Tsinghua University) et al. arXiv.* [paper] [code]
- [2023/06] **Multi-Agent Collaboration: Harnessing the Power of Intelligent LLM Agents.** *Yashar Talebira (University of Alberta) et al. arXiv.* [paper]
- [2023/05] **Training Socially Aligned Language Models in Simulated Human Society.** *Ruibo Liu (Dartmouth College) et al. arXiv.* [paper] [code]
- [2023/05] **SwiftSage: A Generative Agent with Fast and Slow Thinking for Complex Interactive Tasks.** *Bill Yuchen Lin (Allen Institute for Artificial Intelligence) et al. arXiv.* [paper] [code]
- [2023/05] **ChatGPT as your Personal Data Scientist.** *Md Mahadi Hassan (Auburn University) et al. arXiv.* [paper]
- [2023/03] **CAMEL: Communicative Agents for "Mind" Exploration of Large Scale Language Model Society.** *Guohao Li (King Abdullah University of Science and Technology) et al. arXiv.* [paper] [code]
- [2023/03] **DERA: Enhancing Large Language Model Completions with Dialog-Enabled Resolving Agents.** *Varun Nair (Curai Health) et al. arXiv.* [paper] [code]

2.2.2 Adversarial Interaction for Advancement

- [2023/08] **ChatEval: Towards Better LLM-based Evaluators through Multi-Agent Debate.** *Chi-Min Chan (Tsinghua University) et al. arXiv.* [paper] [code]
- [2023/05] **Improving Factuality and Reasoning in Language Models through Multiagent Debate.** *Yilun Du (MIT CSAIL) et al. arXiv.* [paper] [code]
- [2023/05] **Improving Language Model Negotiation with Self-Play and In-Context Learning from AI Feedback.** *Yao Fu (University of Edinburgh) et al. arXiv.* [paper] [code]
- [2023/05] **Examining the Inter-Consistency of Large Language Models: An In-depth Analysis via Debate.** *Kai Xiong (Harbin Institute of Technology) et al. arXiv.* [paper] [code]
- [2023/05] **Encouraging Divergent Thinking in Large Language Models through Multi-Agent Debate.** *Tian Liang (Tsinghua University) et al. arXiv.* [paper] [code]

2.3 Interactive Engagement between Human and Agent

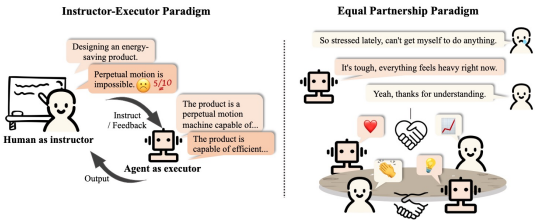

2.3.1 Instructor-Executor Paradigm

Education

- [2023/07] **Math Agents: Computational Infrastructure, Mathematical Embedding, and Genomics.** *Melanie Swan (UCL) et al. arXiv.* [paper]
  - Communicate with humans to help them understand and use mathematics.
- [2023/03] **Hey Dona! Can you help me with student course registration?** *Vishesh Kalvakurthi (MSU) et al. arXiv.* [paper]
  - This is a developed application called Dona that offers virtual voice assistance in student course registration, where humans provide instructions.

Health

- [2023/08] **Zhongjing: Enhancing the Chinese Medical Capabilities of Large Language Model through Expert Feedback and Real-world Multi-turn Dialogue.** *Songhua Yang (ZZU) et al. arXiv.* [paper] [code]
- [2023/05] **HuatuoGPT, towards Taming Language Model to Be a Doctor.** *Hongbo Zhang (CUHK-SZ) et al. arXiv.* [paper] [code] [demo]
- [2023/05] **Helping the Helper: Supporting Peer Counselors via AI-Empowered Practice and Feedback.** *Shang-Ling Hsu (Gatech) et al. arXiv.* [paper]
- [2020/10] **A Virtual Conversational Agent for Teens with Autism Spectrum Disorder: Experimental Results and Design Lessons.** *Mohammad Rafayet Ali (U of R) et al. IVA '20.* [paper]

Other Application

- [2023/08] **RecMind: Large Language Model Powered Agent For Recommendation.** *Yancheng Wang (ASU, Amazon) et al. arXiv.* [paper]
- [2023/08] **Multi-Turn Dialogue Agent as Sales' Assistant in Telemarketing.** *Wanting Gao (JNU) et al. IEEE.* [paper]
- [2023/07] **PEER: A Collaborative Language Model.** *Timo Schick (Meta AI) et al. arXiv.* [paper]
- [2023/07] **DIALGEN: Collaborative Human-LM Generated Dialogues for Improved Understanding of Human-Human Conversations.** *Bo-Ru Lu (UW) et al. arXiv.* [paper]
- [2023/06] **AssistGPT: A General Multi-modal Assistant that can Plan, Execute, Inspect, and Learn.** *Difei Gao (NUS) et al. arXiv.* [paper]
- [2023/05] **Agents: An Open-source Framework for Autonomous Language Agents.** Wangchunshu Zhou (AIWaves) et al. arXiv.* [paper] [code]

2.3.2 Equal Partnership Paradigm

Empathetic Communicator

- [2023/08] **SAPIEN: Affective Virtual Agents Powered by Large Language Models.** *Masum Hasan et al. arXiv.* [paper] [code] [project page] [dataset]
- [2023/05] **Helping the Helper: Supporting Peer Counselors via AI-Empowered Practice and Feedback.** *Shang-Ling Hsu (Gatech) et al. arXiv.* [paper]
- [2022/07] **Artificial empathy in marketing interactions: Bridging the human-AI gap in affective and social customer experience.** *Yuping Liu-Thompkins et al.* [paper]

Human-Level Participant

- [2023/08] **Quantifying the Impact of Large Language Models on Collective Opinion Dynamics.** *Chao Li et al. CoRR.* [paper]
- [2023/06] **Mastering the Game of No-Press Diplomacy via Human-Regularized Reinforcement Learning and Planning.** *Anton Bakhtin et al. ICLR.* [paper]
- [2023/06] **Decision-Oriented Dialogue for Human-AI Collaboration.** *Jessy Lin et al. CoRR.* [paper]
- [2022/11] **Human-level play in the game of Diplomacy by combining language models with strategic reasoning.** *FAIR et al. Science.* [paper]

3. Agent Society: From Individuality to Sociality

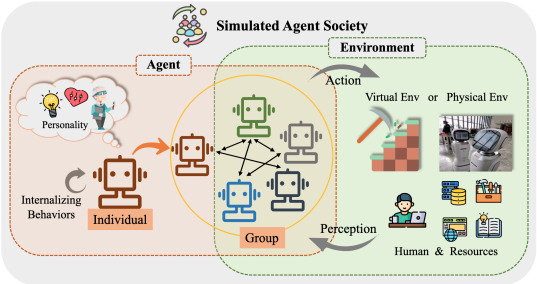

3.1 Behavior and Personality of LLM-based Agents

3.1.1 Social Behavior

Individual behaviors

- [2023/05] **Voyager: An Open-Ended Embodied Agent with Large Language Models.** *Guanzhi Wang (NVIDA) et al. arXiv.* [paper] [code]
- [2023/04] **LLM+P: Empowering Large Language Models with Optimal Planning Proficiency.** *Bo Liu (University of Texas) et al. arXiv.* [paper] [code]
- [2023/03] **Reflexion: Language Agents with Verbal Reinforcement Learning.** *Noah Shinn (Northeastern University) et al. arXiv.* [paper] [code]
- [2023/03] **PaLM-E: An Embodied Multimodal Language Model.** *Danny Driess (Google) et al. ICML.* [paper] [project page]
- [2023/03] **ReAct: Synergizing Reasoning and Acting in Language Models.** *Shunyu Yao (Princeton University) et al. ICLR.* [paper] [project page]
- [2022/01] **Chain-of-thought prompting elicits reasoning in large language models.** *Jason Wei (Google) et al. NeurIPS.* [paper]

Group behaviors

- [2023/09] **Exploring Large Language Models for Communication Games: An Empirical Study on Werewolf.** *Yuzhuang Xu (Tsinghua University) et al. arXiv.* [paper]
- [2023/08] **AgentVerse: Facilitating Multi-Agent Collaboration and Exploring Emergent Behaviors in Agents.** *Weize Chen (Tsinghua University) et al. arXiv.* [paper] [code]
- [2023/08] **AutoGen: Enabling Next-Gen LLM Applications via Multi-Agent Conversation Framework.** *Qingyun Wu (Pennsylvania State University) et al. arXiv.* [paper] [code]
- [2023/08] **ChatEval: Towards Better LLM-based Evaluators through Multi-Agent Debate.** *Chi-Min Chan (Tsinghua University) et al. arXiv.* [paper] [code]
- [2023/07] **Communicative Agents for Software Development.** *Chen Qian (Tsinghua University) et al. arXiv.* [paper] [code]
- [2023/07] **RoCo: Dialectic Multi-Robot Collaboration with Large Language Models.** *Zhao Mandi, Shreeya Jain, Shuran Song (Columbia University) et al. arXiv.* [paper] [code]
- [2023/08] **ProAgent: Building Proactive Cooperative AI with Large Language Models.** *Ceyao Zhang (The Chinese University of Hong Kong, Shenzhen) et al. arXiv.* [paper] [code]

3.1.2 Personality

Cognition

- [2023/03] **Machine Psychology: Investigating Emergent Capabilities and Behavior in Large Language Models Using Psychological Methods.** *Thilo Hagendorff (University of Stuttgart) et al. arXiv.* [paper]
- [2023/03] **Mind meets machine: Unravelling GPT-4's cognitive psychology.** *Sifatkaur Dhingra (Nowrosjee Wadia College) et al. arXiv.* [paper]
- [2022/07] **Language models show human-like content effects on reasoning.** *Ishita Dasgupta (DeepMind) et al. arXiv.* [paper]
- [2022/06] **Using cognitive psychology to understand GPT-3.** *Marcel Binz et al. arXiv.* [paper]

Emotion

- [2023/07] **Emotional Intelligence of Large Language Models.** *Xuena Wang (Tsinghua University) et al. arXiv.* [paper]
- [2023/05] **ChatGPT outperforms humans in emotional awareness evaluations.** *Zohar Elyoseph et al. Frontiers in Psychology.* [paper]
- [2023/02] **Empathetic AI for Empowering Resilience in Games.** *Reza Habibi (University of California) et al. arXiv.* [paper]
- [2022/12] **Computer says “No”: The Case Against Empathetic Conversational AI.** *Alba Curry (University of Leeds) et al. ACL.* [paper]

Character

- [2023/07] **Do LLMs Possess a Personality? Making the MBTI Test an Amazing Evaluation for Large Language Models.** *Keyu Pan (ByteDance) et al. arXiv.* [paper] [code]
- [2023/07] **Personality Traits in Large Language Models.** *Mustafa Safdari (DeepMind) et al. arXiv.* [paper] [code]
- [2022/12] **Does GPT-3 Demonstrate Psychopathy? Evaluating Large Language Models from a Psychological Perspective.** *Xingxuan Li (Alibaba) et al. arXiv.* [paper]
- [2022/12] **Identifying and Manipulating the Personality Traits of Language Models.** *Graham Caron et al. arXiv.* [paper]

3.2 Environment for Agent Society

3.2.1 Text-based Environment

- [2023/08] **Hoodwinked: Deception and Cooperation in a Text-Based Game for Language Models.** *Aidan O'Gara (University of Southern California) et al. arXiv.* [paper] [code]
- [2023/03] **CAMEL: Communicative Agents for "Mind" Exploration of Large Scale Language Model Society.** *Guohao Li (King Abdullah University of Science and Technology) et al. arXiv.* [paper] [code]
- [2020/12] **Playing Text-Based Games with Common Sense.** *Sahith Dambekodi (Georgia Institute of Technology) et al. arXiv.* [paper]
- [2019/09] **Interactive Fiction Games: A Colossal Adventure.** *Matthew Hausknecht (Microsoft Research) et al. AAAI.* [paper] [code]
- [2019/03] **Learning to Speak and Act in a Fantasy Text Adventure Game.** *Jack Urbanek (Facebook) et al. ACL.* [paper] [code]
- [2018/06] **TextWorld: A Learning Environment for Text-based Games.** *Marc-Alexandre Côté (Microsoft Research) et al. IJCAI.* [paper] [code]

3.2.2 Virtual Sandbox Environment

- [2023/08] **AgentSims: An Open-Source Sandbox for Large Language Model Evaluation.** *Jiaju Lin (PTA Studio) et al. arXiv.* [paper] [project page] [code]
- [2023/05] **Training Socially Aligned Language Models in Simulated Human Society.** *Ruibo Liu (Dartmouth College) et al. arXiv.* [paper] [code]
- [2023/05] **Voyager: An Open-Ended Embodied Agent with Large Language Models.** *Guanzhi Wang (NVIDA) et al. arXiv.* [paper] [code]
- [2023/04] **Generative Agents: Interactive Simulacra of Human Behavior.** *Joon Sung Park (Stanford University) et al. arXiv.* [paper] [code]
- [2023/03] **Plan4MC: Skill Reinforcement Learning and Planning for Open-World Minecraft Tasks.** *Haoqi Yuan (PKU) et al. arXiv.* [paper] [code]
- [2022/06] **MineDojo: Building Open-Ended Embodied Agents with Internet-Scale Knowledge.** *Linxi Fan (NVIDIA) et al. NeurIPS.* [paper] [project page]

3.2.3 Physical Environment

- [2023/09] **RoboAgent: Generalization and Efficiency in Robot Manipulation via Semantic Augmentations and Action Chunking.** *Homanga Bharadhwaj (Carnegie Mellon University) et al. arXiv.* [paper] [project page]
- [2023/05] **AVLEN: Audio-Visual-Language Embodied Navigation in 3D Environments.** *Sudipta Paul et al. NeurIPS.* [paper]
- [2023/03] **PaLM-E: An Embodied Multimodal Language Model.** *Danny Driess (Google) et al. ICML.* [paper] [project page]
- [2022/10] **Interactive Language: Talking to Robots in Real Time.** *Corey Lynch (Google) et al. arXiv.* [paper] [code]

3.3 Society Simulation with LLM-based Agents

- [2023/08] **AgentSims: An Open-Source Sandbox for Large Language Model Evaluation.** *Jiaju Lin (PTA Studio) et al. arXiv.* [paper] [project page] [code]
- [2023/07] **S^3 : Social-network Simulation System with Large Language Model-Empowered Agents.** *Chen Gao (Tsinghua University) et al. arXiv.* [paper]
- [2023/07] **Epidemic Modeling with Generative Agents.** *Ross Williams (Virginia Tech) et al. arXiv.* [paper] [code]
- [2023/06] **RecAgent: A Novel Simulation Paradigm for Recommender Systems.** *Lei Wang (Renmin University of China) et al. arXiv.* [paper]
- [2023/05] **Training Socially Aligned Language Models in Simulated Human Society.** *Ruibo Liu (Dartmouth College) et al. arXiv.* [paper] [code]
- [2023/04] **Generative Agents: Interactive Simulacra of Human Behavior.** *Joon Sung Park (Stanford University) et al. arXiv.* [paper] [code]
- [2022/08] **Social Simulacra: Creating Populated Prototypes for Social Computing Systems.** *Joon Sung Park (Stanford University) et al. UIST.* [paper]

Citation

If you find this repository useful, please cite our paper:

```
@misc{xi2023rise,
      title={The Rise and Potential of Large Language Model Based Agents: A Survey},
      author={Zhiheng Xi and Wenxiang Chen and Xin Guo and Wei He and Yiwen Ding and Boyang Hong and
      year={2023},
      eprint={2309.07864},
      archivePrefix={arXiv},
      primaryClass={cs.AI}
}
```

Project Maintainers & Contributors

- Zhiheng Xi （奚志恒, @WooooDyy）
- Wenxiang Chen （陈文翔, @chenwxOggai）
- Xin Guo （郭昕, @XinGuo2002）
- Wei He（何为, @hewei2001）
- Yiwen Ding （丁怡文, @Yiwen-Ding）
- Boyang Hong（洪博杨, @HongBoYang）
- Ming Zhang （张明, @KongLongGeFDU）
- Junzhe Wang（王浚哲, @zsxmwjz）
- Senjie Jin（金森杰, @Leonnnnnn929）

Contact

- Zhiheng Xi: zhxi22@m.fudan.edu.cn

Star History