

BENCHMARKING LARGE LANGUAGE MODELS AS AI RESEARCH AGENTS

Anonymous authors

Paper under double-blind review

ABSTRACT

Human researchers can perform scientific experimentation loops – planning, experimenting, observing the results, and generating inferences. Can we build AI research agents to perform the same? To take a step towards building and evaluating research agents capable of such open-ended decision-making, we focus on the problem of having agents perform machine learning (ML) tasks given a task description and dataset. In this paper, we propose MAgentBench, a suite of ML tasks for benchmarking AI research agents. Agents can perform actions like file system operations, executing code, and inspecting outputs. With these actions, agents could run experiments, analyze the results, and modify the code of entire machine learning pipelines, such as data processing, architecture, training processes, etc. The benchmark then automatically evaluates the agent’s performance objectively over various metrics related to performance and efficiency. We also design an LLM-based research agent to automatically perform experimentation loops in such an environment. Empirically, we find that a GPT-4-based research agent can feasibly build compelling ML models over many tasks in MAgentBench, displaying highly interpretable plans and actions. However, the success rates vary considerably; they span from almost 90% on well-established older datasets to as low as 10% on recent Kaggle Challenges – unavailable during the LLM model’s pretraining – and even 0% on newer research challenges like BabyLM. Finally, we identify several key challenges for LLM-based research agents such as long-term planning and hallucination. Our code is released at <https://anonymous.4open.science/r/MAgentBench/>.

1 INTRODUCTION

Human researchers have the ability to carry out scientific discoveries that involve open-ended decision-making at every step, diving deep into the realms of the unknown. Equipped with accumulated scientific knowledge, human researchers tread paths less traveled, making groundbreaking discoveries along the way. Such exploratory prowess raises an intriguing question: Is it possible to construct AI research agents with similar capacities? A competent research agent, armed with extensive prior knowledge, should have the capability to independently i) hypothesize new research ideas, and ii) validate these ideas through well-crafted experimental trials. The help of such research agents will enable human researchers to pursue more diverse and ambitious research directions efficiently.

However, evaluating the performance of such research agents that can interact with the environment freely and make open-ended decisions is challenging: the interaction process could be slow, resource-intensive, and hard to evaluate quantitatively. Consequently, we focus on the domain of ML research, where the experiment cycle is relatively short and digital with clear objective metrics (e.g. model accuracy), yet still open-ended and challenging with the execution of arbitrarily complex code and the interactions with different types of data.

Specifically, we focus on the problem of having research agents develop/edit ML models given a task description and a dataset. Such an ML-focused research agent could be very helpful simply for automating the engineering portion of ML research. It would allow ML researchers to carry out many more explorations by instructing research agents to implement and test specific algorithms. To accomplish this, we first need a reliable and deterministic benchmarking environment where an agent can operate, its performance can be measured quantitatively, and be compared with different agents.

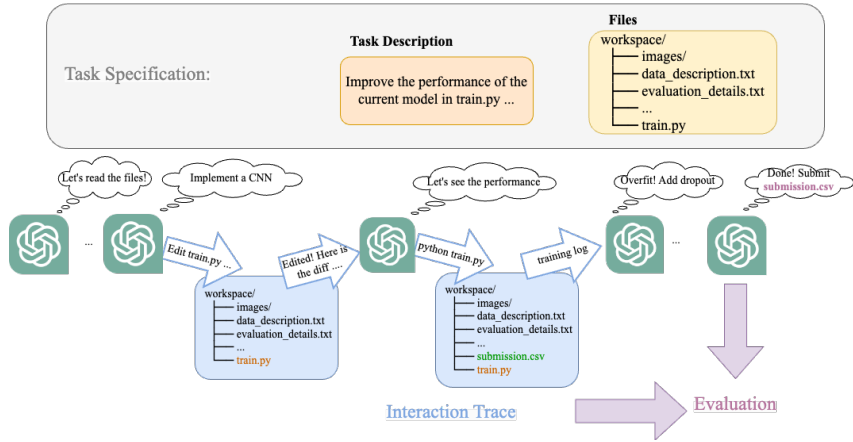


Figure 1: Overview of MAgentBench. Each task is specified with a task description (i.e. the goal) and a set of files (which include code and data). Given these, a research agent can perform actions such as reading/writing files and executing python code. During the interaction, we collect interaction traces including each action, observation, and snapshot of the workspace. Finally, the agent is evaluated based on the interaction trace and the final artifact produced (e.g., submission.csv).

In this paper, we propose MAgentBench, the first benchmark for evaluating AI research agents capable of open-ended decision-making (Figure 1). In essence, our MAgentBench introduces a general framework for specifying well-scoped executable research tasks and automatically evaluates research agents on these tasks. Concretely, each research task is specified with a task description and a set of necessary files (including starter code and data e.g. Kaggle data package). Given these, a research agent can perform actions like reading/writing files and executing code, similar to the interface used by a human researcher. During the agent’s interaction with the environment, we collect its interaction traces, i.e. agent actions and intermediate snapshots of the workspace, for evaluation. We evaluate the research agent along three aspects: 1) competence in accomplishing the objectives, e.g. success rate and the average amounts of improvements, 2) reasoning and research process, e.g. how the agent achieved the result or what mistakes it made, and 3) efficiency, measured by the amount of time and resources spent by the agent.

As an initial curation effort, we include 15 ML engineering tasks from diverse domains ranging in various difficulties and recency (Table 1), where the experiments are generally fast to perform and inexpensive. We provide simple starter codes for some of these tasks to ensure that the agent can make submissions properly. For example, one task is to increase a baseline Convolution Neural Networks (CNN) model performance on the cifar10 dataset (Krizhevsky, 2009) by more than 10%. Beyond very well-established datasets like cifar10, we also include a few months old Kaggle challenges and other newer research datasets to see whether the research agent can extrapolate to newer datasets unseen during (pre-)training. In the future, we aim to expand our tasks set to more diverse scientific research tasks in different domains.

In light of the recent development of generative agents based on Large language models (LLMs) (Yao et al., 2022; Shinn et al., 2023; Wang et al., 2023; aut, 2023; Schick et al., 2023; Park et al., 2023), we also design a simple LLM-based research agent that can automatically make research plans, read/edit scripts, perform experiments, interpret results, and continue with next-step experiments over MAgentBench environments. LLMs have demonstrated impressive prior knowledge ranging from everyday common sense knowledge to specific scientific disciplines as well as great reasoning and tool-using abilities, making them able to act and react to the broader world beyond just direct textual chatting (OpenAI, 2023; Bubeck et al., 2023). On the high level, we simply prompt the LLMs out-of-the-box to provide the next step action, given an automatically constructed prompt based on all known information about the task and prior actions. Many specific components in constructing the prompt take general inspiration from popular techniques for building other LLM-based generative agents, including reasoning before action (Yao et al., 2022), reflection (Shinn et al., 2023), step-by-step planning (aut, 2023), and managing a research log as a memory stream (Park et al., 2023). In addition, we also use hierarchical action and fact-checking step to further improve the stability and factualness of the AI research agent. See more details in Section 3.

Over MAgentBench, we find that our AI research agent, especially when based on GPT-4, is able to successfully build a better ML model over many tasks and generate highly interpretable dynamic research plans along the process, though still with many drawbacks. On well-established tasks like training a better model over the ogbn-arxiv dataset (Hu et al., 2020), it is able to improve upon baseline prediction successfully almost 90% of the time, with an average improvement of 48.18%. However, the research agent struggles with Kaggle Challenges and BabyLM (Warstadt et al., 2023), with only a 0 to 30% success rate. We then compare results between different variants of our research agent as well as the adaptation of other existing agents. Interestingly, we found that maintaining the memory stream could actively hurt the performance on simple tasks, potentially as a distraction and encouraging agent to pursue complex changes. We also identify several key challenges for LLM-based research agent designs, e.g. how to effectively plan and replan over long horizons and hallucination about the current progress, and show how our design handles them qualitatively. Overall, our research agent demonstrates preliminary feasibility and success with LLM-based research agents, but there is still a long way until they can succeed reliably.

2 MAGENTBENCH: BENCHMARKING ML RESEARCH AGENTS

Our MAgentBench introduces a general framework for specifying well-scoped executable research tasks and automatically evaluating agents on these tasks. The benchmark provides for a disentangled environment and agent side that captures the entire interaction trace of the agent with the environment for evaluation. We then include 15 concrete and diverse machine learning tasks in the benchmark. In the subsequent subsections, we shall be describing some key components of MAgentBench.

2.1 TASK SPECIFICATION

Our task specification scheme is designed to be general and similar to the human-facing interface, making it easy to add new tasks and translate to collaborating with human researchers in the future. Each research task is specified in two parts:

Task description. In MAgentBench, the task description describes the desired goal and evaluation metric, e.g. “Given a training script on a dataset `train.py`, improve upon the current model accuracy”, and how the research agent should submit the final answer for evaluation, e.g. “Save per class probabilities for test set examples to `submission.csv`”. The description could also include constraints like limiting the model size and training epochs, or occasionally include specific directions to approach the problem like “by fine-tuning a pretrained BERT model”.

Files. In MAgentBench, we provide the agent with a prepared set of files and data necessary for each task so that the agent should not need to browse the internet. This typically includes training and testing data (without test labels), detailed data and metric descriptions, and some starter code. The starter code is based on diverse ML frameworks, including PyTorch (Paszke et al., 2019), TensorFlow (Abadi et al., 2015), JAX (Bradbury et al., 2018), Keras Chollet et al. (2015), etc. The starter code mostly implements a simple baseline model that we can compare with during evaluation, but some tasks do have any baseline implementation, and the agent is responsible for coding up the model from scratch from the task description and dataset files.

2.2 TASK ENVIRONMENT

With a task specification as described above, each task in MAgentBench can be seen as an RL environment where AI research agents can perform actions and receive observations. Some primitive actions that are available in the benchmark include file system operations (read, write, append, copy, edit, undo edit), execution of any arbitrary Python script¹, and a final answer declaration action. Each action is specified with a name, description, usage, return value description, and a Python implementation.

The set of actions that an agent could execute could be augmented with hand-written or generated high-level actions, which have the same interface as a primitive action but perform a composition

¹More broadly this can include more specialized actions like running a biology experiment for a Biology research agent. This has been restricted to Python code execution for an ML research agent.

of multiple actions interleaved with other operations like LLM calls. This can be seen as a modular skill library that provides transferrable high-level skills to all agents across all tasks. At the beginning of each task, MAgentBench will first prepare a workspace directory for the research agent by copying relevant files and data from the task specification. The task description and information of all available actions are also stored as queryable variables of the environment object.

The research agent can then submit an action name and action arguments as a dictionary to the environment object, which will return back the proper observation as a string. For example, the agent could call `env.execute(Action("Read File", "train.py"))`, and obtain the content of the file `train.py` in the working directory as the observation. The agent can interact with the environment many times until it decides to submit the final answer, or the environment shuts down itself due to exceeding a maximum number of actions or maximum time (both of which are specified as environment variables).

Finally, all actions and snapshots of the workspace after each action is executed are recorded as an interaction trace. We keep track of all the primitive actions which makes it possible to reproduce and analyze the performance irrespective of differing composite skills used by different agents. Thus, the interaction trace captured by the environment is agnostic to agent implementation and can be compared directly across agents.

2.3 EVALUATION

Given the interaction traces collected, we can then evaluate the AI research agent from three aspects:

Competence in accomplishing the objectives. We evaluate a single performance metric based on each final snapshot of the working directory. For most of the tasks, we can simply evaluate the performance (e.g. accuracy) based on the final "submission.csv". We then compute aggregated metrics like success rate at fulfilling the research objective and the average amount of improvement (e.g. upon the baseline solution in the starter code) over multiple runs to test reliability and generalizability.

Reasoning and research process. With the interaction trace collected, we can further evaluate the agent in terms of interpretability and more detailed error modes. We found this process-based evaluation to be more helpful for agent development than a single black box score. To automate this, we can evaluate the interaction trace against a set of rubrics, such as whether the agent was stuck on debugging, by prompting GPT-3.5. However, here we performed a human evaluation due to the unreliability of an LLM-based evaluation.

Efficiency. We evaluate the efficiency in terms of the total amount of wall clock time spent and the total amount of resource cost (i.e. number of tokens for LLM-based agents).

2.4 CONCRETE TASKS

We include 15 tasks from diverse domains including Natural Language Processing, Computer Vision, Time Series prediction, Graph ML, and Tabular data as shown in Table 1. Our tasks include both well-studied datasets like `cifar10` and open challenges like Parkinson’s disease progression prediction from Kaggle which were not present in the language model pre-training. The tasks are chosen such that they range in various difficulties and recency to test the generalizability of the research agent and avoid data contamination. They are divided to the following categories:

Canonical Tasks. We included `cifar10` (image classification) (Krizhevsky, 2009), `imdb` (sentiment classification) (Maas et al., 2011), and `ogbn-arxiv` (paper category classification over citation network) (Hu et al., 2020) as canonical tasks that are well-studied and easy to iterate on. For `cifar10` and `ogbn-arxiv`, the task was to improve a baseline model, but for `imdb`, the agent was expected to write the model from scratch which involved finetuning a BERT model as mentioned in the task description.

Classic Kaggle. House-price (Anna Montoya, 2016) and `spaceship-titanic` (Howard et al., 2022) are two introductory Kaggle challenges for tabular regression and classification. These tasks mainly involve feature engineering, writing, and training models from scratch (no baselines provided), and properly following the Kaggle submission format.

Kaggle Challenges. We select four recent Kaggle Challenges launched from 2 to 10 months ago to test research agents’ ability to generalize to more realistic and out-of-distribution tasks.

Type	Task	Modality	Dataset Name
Canonical Tasks	Classification	Image	cifar10 (Krizhevsky, 2009)
	Classification	Text	imdb (Maas et al., 2011)
	Node Classification	Graph	ogbn-arxiv (Hu et al., 2020)
Classic Kaggle	Regression	Tabular	house-price (Anna Montoya, 2016)
	Classification	Tabular	spaceship-titanic (Howard et al., 2022)
Kaggle Challenges	Regression	Time Series	parkinsons-disease (Kirsch et al., 2023)
	Classification	Image	fathomnet (Woodward et al., 2023)
	Regression	Text	feedback (Franklin et al., 2022)
	Segmentation	Images	identify-contrails (Sarna et al., 2023)
Current Research	Node Regression	Graph	CLRS (Veličković et al., 2022)
	Language Modeling	Text	BabyLM (Warstadt et al., 2023)
Improve Code	Improve speed	Text	llama-inference
	Improve speed	Image	vectorization
LLM Tools	Implement tool	n/a	literature-review-tool
	Implement tool	n/a	bibtex-generation

Table 1: MAgentBench tasks.

Current Research. We include CLRS (Veličković et al., 2022) and BabyLM (Warstadt et al., 2023) as two example datasets that are actively being researched and do not yet have a consensus on the best approaches. CLRS dataset involves modeling classic algorithms over graphs and lists. BabyLM requires training a language model over 10M words.

Improve Code. We include llama-inference and vectorization as two examples where AI research agent is asked to improve the runtime of code instead of optimizing its prediction performance. llama-inference is about improving the autoregressive generation speed of the Llama 7B model (Touvron et al., 2023), and vectorization is about speeding up the inference of a convolutional model with stacks of for loops in the forward pass.

LLM Tools. We also design two scenarios where the research agent is instructed to write LLM-based research tools, which can perform literature review and generate BibTeX from sketch.

More details on the benchmark setup can be found in Appendix A. The above tasks are incorporated as an initial curation and we hope to continuously expand this benchmark with the help of the open-source community.

3 LLM-BASED RESEARCH AGENT

We design a simple LLM-based research agent as shown in Figure 2. On a high level, we prompt the LLM to provide the next step action and action arguments in a JSON format. The prompt starts with a description of all the tools available, the task description, research-specific instructions, a template to instruct LLM to produce text in parsable format, and the historical steps taken (see Appendix D for a full example of what prompt the agent sees at each interaction step). Given the LLM response,

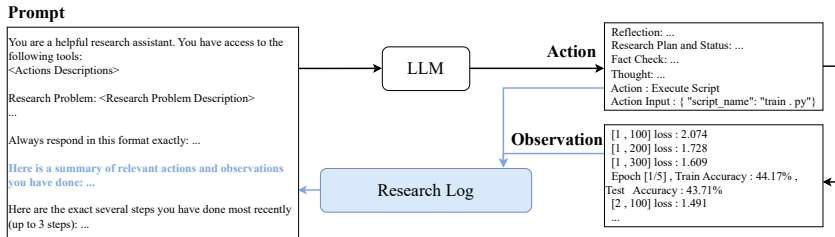


Figure 2: Overview of our LLM-based research agent.

we post-process it to action and action arguments for the environment to execute. In the subsequent subsections, we detail some important components of our LLM-based research agent.

3.1 THINKING BEFORE ACTING

One important component is specifying the response format, so that LLM can first think in specific ways before proposing action. Specifically, we instruct LLM to include a list of entries in the response. In our research agent prototype, this includes `Reflection`, `Research Plan` and `Status`, `Fact Check`, `Thought`, and then `Action` and `Action Input`. Among these, `Thought` and `Reflection` are inspired by `React` and `Reflexion` (Yao et al., 2022; Shinn et al., 2023). `Research Plan` and `Status` entry is designed to produce better planning and keep track of what has been done; `Fact Check` is added to double-check whether a result has been confirmed or hallucinated. We discuss this more in Appendix B.1 and B.2.

3.2 RESEARCH LOG

Since the research agent can perform many actions during the entire interaction, it is often infeasible to simply put all historical responses in the context length of LLM. Thus, reducing prompt length is one of the key challenges for generative agents. In our research agent, we use a design similar to the memory stream paradigm from (Park et al., 2023). Specifically, we append a summary of the LLM response and the observation in each step to a `Research Log` file. Then we can retrieve relevant information from this `Research Log`, and concatenate with several recent full LLM responses to form the historical context. With this design, the `Research Log` file itself then also becomes a normal file available for the agent to query and modify, as exemplified by the `Reflection` action below.

3.3 HIERARCHICAL ACTIONS

We manually designed a few commonly useful high-level actions that perform several primitive environment actions and separate modular LLM calls together. The most important ones are:

Understand File. This action reads a long file and calls another LLM to summarize and retrieve information relevant to a short input query from it.

Reflection. It allows the agent to perform reflection by reading in the content of the `Research Log` file and prompting another LLM to reflect on a short input query.

Edit Script. This action first reads in a file, calls another LLM to perform an edit of a file given a short input instruction from the main `Research` agent, then writes the modified version back. We also include a different version, `Edit Script Segment`, which also takes start and end lines as arguments and only edits the segment in between, when the task involves a large code base (i.e. CLRS and BabyLM).

4 EXPERIMENTS

We evaluate our designed research agent with GPT-4 (OpenAI, 2023) and Claude-1 (Anthropic, 2023) on `MLAgentBench`. Aside from the full form, we consider a no-retrieval variant, where the `Research Log` component is removed, and hence the agent has no long-term memory.

We also benchmark the direct adaptation of several existing generative agents: 1) `AutoGPT`, a popular open-source project for general-purpose autonomous AI agents (aut, 2023), and 2) `LangChain`, another popular framework that implements various generative agent. Here we use "zero-shot-react-description" which implements `ReAct` (Yao et al., 2022). We use Claude-1 for both agents.

We conduct 25 runs for all agents using Claude-1, and 8 runs for GPT-4 agents to average out the evaluation metrics. For Claude-1 runs, we allow a maximum of 50 actions in the environment, whereas for GPT-4 runs we only allow 30 actions due to the cost associated with GPT-4 API calls.

4.1 COMPETENCE IN ACCOMPLISHING THE OBJECTIVES

As shown in Figure 3 and 4, the GPT-4 agent achieves the best results over almost all tasks, but with varying degrees of success from more than 80% over `ogbn-arxiv` to 0% over `BabyLM`. The average

improvements made are also generally large and positive. Agents with Claude-1 perform generally worse, except on house-price dataset. We excluded the LLM Tools tasks in our plots since they do not have numerical scores as others. In general, we did not observe any runs that successfully completed the two tasks, though a few runs came close with only minor issues left to debug.

Interestingly, agents without a Research Log perform better than those with it for easier tasks. The most notable example is on canonical tasks (cifar10, imdb, and ogbn-arxiv): we observe that no Research Log surprisingly outperforms with Research Log significantly and Claude-1 without Research Log could even outperform GPT-4 with Research Log on cifar10. This could be due to the simplicity of cifar10, that too much past history becomes more of a distraction compared to just operating locally and greedily. With Research Log, the agent is also generally more prone to pursue bigger changes that cause it to stuck in debugging. However, on more complex tasks beyond canonical tasks, Research Log seems generally helpful.

Comparing our proposed research agent with existing agents with Claude-1, our research agent achieves a much better success rate on cifar10, ogbn-arxiv, house-price, spaceship-titanic, and CLRS. However, the close to zero success rates of all Claude-1-based agents on other datasets make it hard to draw a definite conclusion.

4.2 REASONING AND RESEARCH PROCESS

We show a full example (without the Research Log for simplicity) on cifar10 to demonstrate what our research agent actually does qualitatively in Appendix D. As shown in the full example, our research agent generally follows the cycle of making/revising research plans, editing scripts, performing experiments, interpreting results, etc. To more carefully evaluate the reasoning and research process of the agent, we analyze the traces of runs for cifar10 and categorize most runs as shown in Figure 5:

1. Hallucination, where the agent claims to know something or fabricates some results such as claiming performance increase without even executing any edits in the training script.
2. Debugging, where the agent fails to debug its modification to the code. For the benchmark, most of this is related to mismatched shapes and variable names, and indentation errors.
3. Token Length, where the agent fails because the automatically constructed prompt was too long, exceeding the context length of the language model.

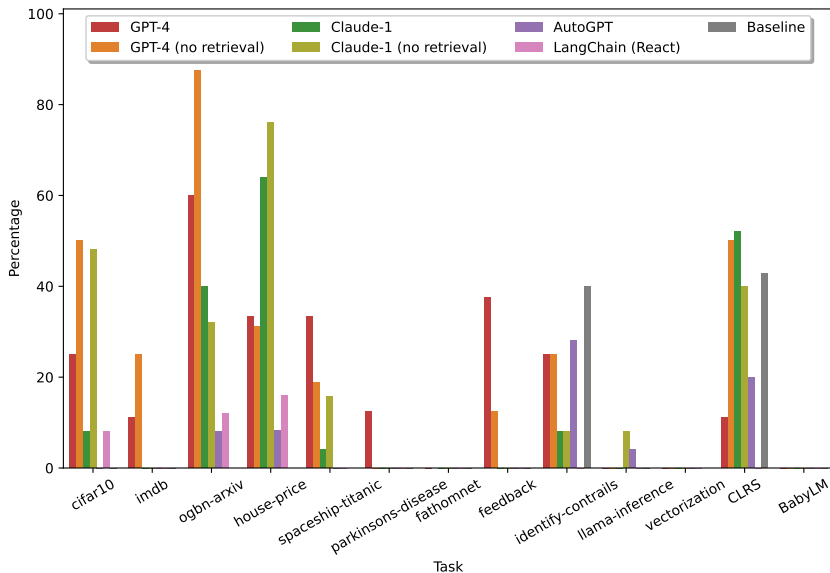


Figure 3: Success Rate, i.e. the percentages of runs that achieve more than 10% improvement at the last step over the average performance of the baseline in starter code.

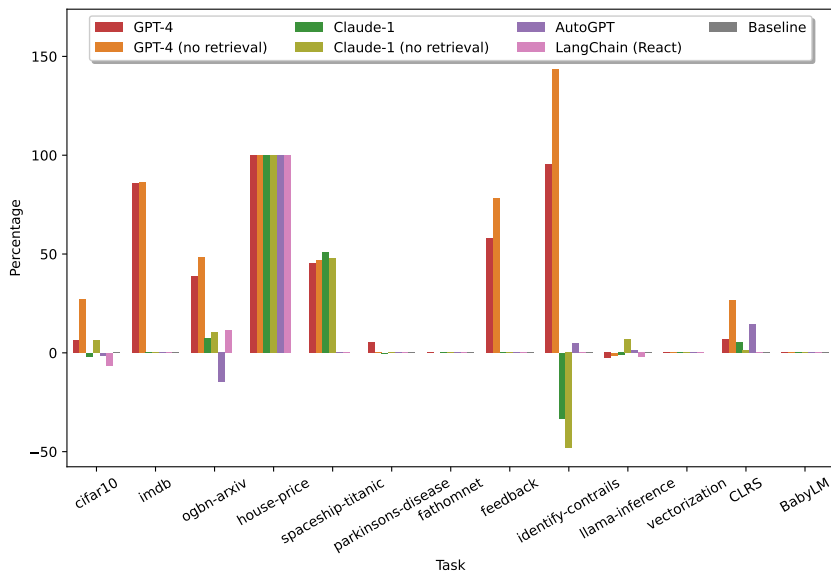


Figure 4: Average Improvement over the baseline in starter code among the runs that made a valid submission at the last step.

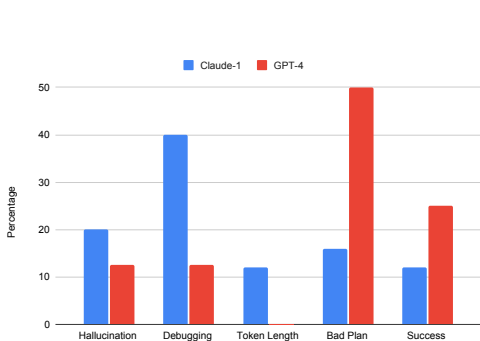


Figure 5: Percentage of runs over cifar10 task that falls into different categories for reasoning and research process evaluation.

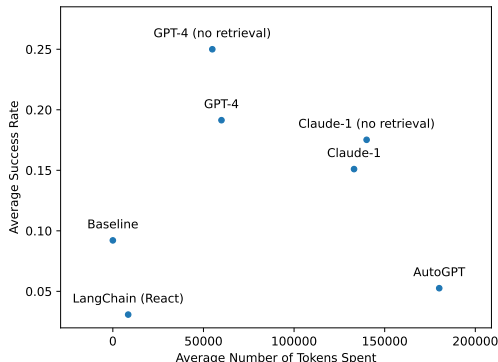


Figure 6: Comparing different agents in terms of efficiency, i.e. the number of tokens spent (on x axis and the smaller the better) and success rate (on y axis and the higher the better).

4. Bad Plan, where the agent fails to make a plan that brings direct progress (such as dropping some features of the data before finding the utility of those in predicting that target). Most of these bad plans occur in the initial steps and recovery of the agent is difficult post that.
5. Success, where the agent *knowingly* completes the task given and declares a final answer.

Note that GPT-4 based research agent is able to avoid hallucination and debugging problems, but tends to fail more due to bad plans. We show more detailed quantitative analysis in Appendix B, which shows the benefit of Research Plan and Status entries for long-term interpretable planning and Fact Check entries against hallucination.

4.3 EFFICIENCY

We compare the average amount of tokens and time spent by different agents for all tasks in Figure 6, for each task in Figure 7 and 8 in the Appendix. On average, GPT-4 based agents are the most efficient, spending 119.1% less tokens due to their efficiency in finishing the task and submitting

early, while having the highest success rate too. Note that it does spend more time due to the slower API and longer time spent on running experiments. Converting with the current API prices, each run on each task only costs a few dollars. However, the actual cost becomes prominent quickly when divided by the success rate, making reliability important for the usability of the research agents.

5 RELATED WORKS

5.1 AI FOR SCIENCE

Numerous research endeavors seek to enhance the pace of manual observations and experiments through automated ML predictions (Berens et al., 2023; Zhang et al., 2023b; Jumper et al., 2021; Adam-Bourdarios et al., 2016; Schwaller et al., 2017). On the other hand, significant line of inquiry revolves around constructing closed-loop systems capable of conducting ongoing experiments and breakthroughs within specific domains (Kramer et al., 2023; Kitano, 2021). For example, Robot Scientist “Adam” is developed to autonomously generate functional genomics hypotheses about the yeast *Saccharomyces cerevisiae* and experimentally test these hypotheses by using laboratory automation (King et al., 2009; 2004). Nevertheless, these existing systems are highly tailored to process specific types of data for designated tasks and domains. Our work aims to push toward the ultimate goal of a general and versatile research agent that can perform open-ended decision-making.

5.2 GENERATIVE AGENTS

Large language models (LLMs) have demonstrated impressive prior knowledge ranging from everyday common sense knowledge to specific scientific disciplines like Computer Science and Chemistry (OpenAI, 2023; Bubeck et al., 2023). Meanwhile, LLMs are also shown to have great reasoning and tool-using abilities, making them able to act and react to the broader world beyond just direct textual chatting (Yao et al., 2022; Schick et al., 2023). This combination of strong prior knowledge and action/reaction abilities of LLMs gives rise to explorations of developing various LLM-based generative agents, such as generative agents for simulating interactions between humans (Park et al., 2023), Voyager for playing Minecraft (Wang et al., 2023), Say-Can for physical robotics (Ahn et al., 2022), as well as open source projects like AutoGPT (aut, 2023) for everything and commercial product like Adapt.ai for internet interactions. However, it is hard to evaluate the performance and reliability of these agents, especially over a long horizon of complex interactions. Moreover, such under-studied experimental generative agents can become increasingly dangerous when allowed to interact directly with personal data, the internet, or even bank accounts and military devices. From this perspective, our MAgentBench offers a test bed for generative agents with the desired combination of containability, complexity, evaluability, and practical usefulness.

5.3 LLMs FOR AUTO ML

Several concurrent works have explored using LLMs for AutoML type of tasks: AutoML-GPT (Zhang et al., 2023c) repeatedly prompts LLMs with data and model cards and predicts training logs to perform efficient hyperparameter tuning; MLCopilot (Zhang et al., 2023a) prompts LLMs with past experiences and knowledge to predict one final categorized hyperparameter setting (e.g. low or high weight decay). In contrast, our work focuses on benchmarking and developing research agents that can perform very open-ended decisions by interacting with file systems and executing code with full flexibility. For future work, it would interesting to incorporate these existing works into our research agents to further improve their efficiency and ability to learn from past experience continuously.

6 CONCLUSION

In this paper, we propose MAgentBench for benchmarking AI research agents on performing ML research tasks end-to-end with access to a compute cluster. We also develop an LLM-based prototype research agent that can accomplish many tasks in MAgentBench with varying success rates. In the future, we would like to pursue a more robust research agent and expand MAgentBench with more complex and creative tasks accordingly. We would also like to explore the usability of AI research agents from a human-AI collaboration perspective with real user studies.

REFERENCES

- Significant-gravitas/auto-gpt: An experimental open-source attempt to make gpt-4 fully autonomous. <https://github.com/Significant-Gravitas/Auto-GPT>, 2023.
- Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL <https://www.tensorflow.org/>. Software available from tensorflow.org.
- Claire Adam-Bourdarios, Glen Cowan, Cécile Germain, Isabelle M Guyon, Balázs Kégl, and David Rousseau. How machine learning won the higgs boson challenge. In *The European Symposium on Artificial Neural Networks*, 2016.
- Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alexander Herzog, Daniel Ho, Jasmine Hsu, Julian Ibarz, Brian Ichter, Alex Irpan, Eric Jang, Rosario Jauregui Ruano, Kyle Jeffrey, Sally Jesmonth, Nikhil Jayant Joshi, Ryan C. Julian, Dmitry Kalashnikov, Yuheng Kuang, Kuang-Huei Lee, Sergey Levine, Yao Lu, Linda Luu, Carolina Parada, Peter Pastor, Jornell Quiambao, Kanishka Rao, Jarek Rettinghouse, Diego M Reyes, Pierre Sermanet, Nicolas Sievers, Clayton Tan, Alexander Toshev, Vincent Vanhoucke, F. Xia, Ted Xiao, Peng Xu, Sichun Xu, and Mengyuan Yan. Do as i can, not as i say: Grounding language in robotic affordances. In *Conference on Robot Learning*, 2022.
- DataCanary Anna Montoya. House prices - advanced regression techniques, 2016. URL <https://kaggle.com/competitions/house-prices-advanced-regression-techniques>.
- Anthropic. Introducing claude, 2023. URL <https://www.anthropic.com/index/introducing-claude>.
- Philipp Berens, Kyle Cranmer, Neil D. Lawrence, Ulrike von Luxburg, and Jessica Montgomery. Ai for science: An emerging agenda. *ArXiv*, abs/2303.04217, 2023.
- James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>.
- Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, John A. Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuan-Fang Li, Scott M. Lundberg, Harsha Nori, Hamid Palangi, Marco Tulio Ribeiro, and Yi Zhang. Sparks of artificial general intelligence: Early experiments with gpt-4. *ArXiv*, abs/2303.12712, 2023.
- Francois Chollet et al. Keras, 2015. URL <https://github.com/fchollet/keras>.
- Alex Franklin, Maggie, Meg Benner, Natalie Rambis, Perpetual Baffour, Ryan Holbrook, Scott Crossley, and ulrichboser. Feedback prize - english language learning, 2022. URL <https://kaggle.com/competitions/feedback-prize-english-language-learning>.
- Addison Howard, Ashley Chow, and Ryan Holbrook. Spaceship titanic, 2022. URL <https://kaggle.com/competitions/spaceship-titanic>.
- Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. Open graph benchmark: Datasets for machine learning on graphs. *ArXiv*, abs/2005.00687, 2020.

- John M. Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Zidek, Anna Potapenko, Alex Bridgland, Clemens Meyer, Simon A A Kohl, Andy Ballard, Andrew Cowie, Bernardino Romera-Paredes, Stanislav Nikolov, Rishub Jain, Jonas Adler, Trevor Back, Stig Petersen, David A. Reiman, Ellen Clancy, Michal Zielinski, Martin Steinegger, Michalina Pacholska, Tamas Berghammer, Sebastian Bodenstern, David Silver, Oriol Vinyals, Andrew W. Senior, Koray Kavukcuoglu, Pushmeet Kohli, and Demis Hassabis. Highly accurate protein structure prediction with alphafold. *Nature*, 596:583 – 589, 2021.
- Ross D. King, Ken E. Whelan, Ffion Mair Jones, Philip G. K. Reiser, Christopher H. Bryant, Stephen H. Muggleton, Douglas B. Kell, and Stephen G. Oliver. Functional genomic hypothesis generation and experimentation by a robot scientist. *Nature*, 427:247–252, 2004.
- Ross D. King, Jem J. Rowland, Stephen G. Oliver, Michael Young, Wayne Aubrey, Emma Byrne, Maria Liakata, Magdalena Markham, Pinar Pir, Larisa N. Soldatova, Andrew Sparkes, Ken E. Whelan, and Amanda Clare. The automation of science. *Science*, 324:85 – 89, 2009.
- Leslie Kirsch, Sohier Dane, Stacey Adam, and Victoria Dardov. Amp@-parkinson’s disease progression prediction, 2023. URL <https://kaggle.com/competitions/amp-parkinsons-disease-progression-prediction>.
- Hiroaki Kitano. Nobel turing challenge: creating the engine for scientific discovery. *NPJ Systems Biology and Applications*, 7, 2021.
- Stefan Kramer, Mattia Cerrato, Saso Dzeroski, and Ross D. King. Automated scientific discovery: From equation discovery to autonomous discovery systems. *ArXiv*, abs/2305.02251, 2023.
- Alex Krizhevsky. Learning multiple layers of features from tiny images. 2009.
- Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pp. 142–150, Portland, Oregon, USA, June 2011. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/P11-1015>.
- OpenAI. Gpt-4 technical report. *ArXiv*, abs/2303.08774, 2023.
- Joon Sung Park, Joseph C. O’Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. Generative agents: Interactive simulacra of human behavior. *ArXiv*, abs/2304.03442, 2023.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pp. 8024–8035. Curran Associates, Inc., 2019. URL <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- Aaron Sarna, Carl Elkin, inversion, Joe Ng, Maggie, and Walter Reade. Google research - identify contrails to reduce global warming, 2023. URL <https://kaggle.com/competitions/google-research-identify-contrails-reduce-global-warming>.
- Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language models can teach themselves to use tools. *ArXiv*, abs/2302.04761, 2023.
- Philippe Schwaller, Théophile Gaudin, David Lanyi, Constantine Bekas, and Teodoro Laino. “found in translation”: predicting outcomes of complex organic chemistry reactions using neural sequence-to-sequence models† †electronic supplementary information (esi) available: Time-split test set and example predictions, together with attention weights, confidence and token probabilities. see do. *Chemical Science*, 9:6091 – 6098, 2017.

Noah Shinn, Beck Labash, and Ashwin Gopinath. Reflexion: an autonomous agent with dynamic memory and self-reflection. *ArXiv*, abs/2303.11366, 2023.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurélien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language models. *ArXiv*, abs/2302.13971, 2023.

Petar Veličković, Adrià Puigdomènech Badia, David Budden, Razvan Pascanu, Andrea Banino, Misha Dashevskiy, Raia Hadsell, and Charles Blundell. The clrs algorithmic reasoning benchmark. *arXiv preprint arXiv:2205.15659*, 2022.

Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi (Jim) Fan, and Anima Anandkumar. Voyager: An open-ended embodied agent with large language models. *ArXiv*, abs/2305.16291, 2023.

Alex Warstadt, Leshem Choshen, Aaron Mueller, Adina Williams, Ethan Gotlieb Wilcox, and Chengxu Zhuang. Call for papers - the babylm challenge: Sample-efficient pretraining on a developmentally plausible corpus. *ArXiv*, abs/2301.11796, 2023.

Ben Woodward, eor123, GenevievePatterson, and Lilli Carlsen. Fathomnet 2023, 2023. URL <https://kaggle.com/competitions/fathomnet-out-of-sample-detection>.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. *ArXiv*, abs/2210.03629, 2022.

Lei Zhang, Yuge Zhang, Kan Ren, Dongsheng Li, and Yuqing Yang. Mlcopilot: Unleashing the power of large language models in solving machine learning tasks. *ArXiv*, abs/2304.14979, 2023a. URL <https://api.semanticscholar.org/CorpusID:258418182>.

Mengchun Zhang, Maryam Qamar, Taegoo Kang, Yuna Jung, Chenshuang Zhang, Sung-Ho Bae, and Chaoning Zhang. A survey on graph diffusion models: Generative ai in science for molecule, protein and material. *ArXiv*, abs/2304.01565, 2023b.

Shujian Zhang, Chengyue Gong, Lemeng Wu, Xingchao Liu, and Mi Zhou. Automl-gpt: Automatic machine learning with gpt. *ArXiv*, abs/2305.02499, 2023c. URL <https://api.semanticscholar.org/CorpusID:258480269>.