

StableToolBench: Towards Stable Large-Scale Benchmarking on Tool Learning of Large Language Models

Zhicheng Guo^{1,2}, Sijie Cheng^{1,2,3}, Hao Wang⁴, Shihao Liang⁵, Yujia Qin¹,
Peng Li², Zhiyuan Liu¹, Maosong Sun¹, Yang Liu^{1,2,6}

¹Dept. of Comp. Sci. & Tech., Institute for AI, Tsinghua University, Beijing, China

²Institute for AI Industry Research (AIR), Tsinghua University, Beijing, China

³01.AI ⁴Google ⁵The University of Hong Kong

⁶Jiangsu Collaborative Innovation Center for Language Competence, Jiangsu, China
{guo-zc21, csj23}@mails.tsinghua.edu.cn

Abstract

Large Language Models (LLMs) have witnessed remarkable advancements in recent years, prompting the exploration of tool learning, which integrates LLMs with external tools to address diverse real-world challenges. Assessing the capability of LLMs to utilise tools necessitates large-scale and stable benchmarks. However, previous works relied on either hand-crafted online tools with limited scale, or large-scale real online APIs suffering from instability of API status. To address this problem, we introduce StableToolBench, a benchmark evolving from ToolBench, proposing a virtual API server and stable evaluation system. The virtual API server contains a caching system and API simulators which are complementary to alleviate the change in API status. Meanwhile, the stable evaluation system designs solvable pass and win rates using GPT-4 as the automatic evaluator to eliminate the randomness during evaluation. Experimental results demonstrate the stability of StableToolBench, and further discuss the effectiveness of API simulators, the caching system, and the evaluator system.

1 Introduction

With the rapid developments of Large Language Models (LLMs; Brown et al., 2020; Gemini Team, 2023; OpenAI, 2023; Touvron et al., 2023), tool learning which leverage LLMs to schedule a variety of external tools has attracted enormous attention (Nakano et al., 2022; Yao et al., 2022b; Lu et al., 2023). Previous studies (Hao et al., 2023; Hsieh et al., 2023; Schick et al., 2023; Tang et al., 2023) aim to augment LLMs with tools to enhance performance on conventional natural language processing (NLP) downstream tasks, while recent work (Qin et al., 2023a; Yao et al., 2022a; Cai et al., 2024) primarily focus on solving real-world scenarios that require the use of tools. In

general, tool learning complements the capabilities of vanilla LLMs and bridges the gap to real-world applications.

To assess the capability of LLMs to use tools, a series of tool learning benchmarks have been introduced. Several pioneering studies have heavily relied on human-crafted offline tools (Yang et al., 2023b; Xu et al., 2023) or hand-selected online tools (Li et al., 2023b,a; Chen et al., 2023b). While these tools are high-quality, their scale remains relatively small, thereby limiting their ability to accurately reflect real-world scenarios. To address this limitation, subsequent studies (Tang et al., 2023; Ye et al., 2024; Qin et al., 2023b) have advocated for leveraging extensive collections of online tools that span across various domains. Owing to the increased scale, the automatic evaluation of tool learning has moved closer to real-world scenarios. However, concerns have been raised regarding the stability of these online tools, which has implications for the reproducibility and comparability of benchmark performance over time¹. For instance, the well-recognised ToolBench² (Qin et al., 2023c) has shown performance discrepancies that cannot be reproduced months after its release, as analysed in Section 2.1. This is even more important when faced with a complex environment, where APIs and tools keep changing while the evaluation should maintain its consistency across time.

Existing large-scale benchmarks may struggle to provide stable evaluations for various reasons. We propose several hypotheses for this issue. Firstly, the complexity of tasks involving tool usage makes it challenging for the common automatic evaluator, gpt-3.5, to function effectively as a discriminator. As discussed in Section 2.2, the evaluator cannot reliably determine whether a task is solvable or

¹According to its definition, benchmarks should remain stable, and the model performance assessed on them must be comparable over time.

²We use ToolEval2 in ToolBench as the benchmark.

Project: [zhichengg.github.io/stb.github.io/](https://github.com/stb-guozc21/stb-guozc21)
GitHub: [THUNLP-MT/StableToolBench](https://github.com/THUNLP-MT/StableToolBench)

unsolvable, leading to variability in model performance due to this capability limitation. Secondly, the stability of API status for a significant portion of online tools (55.6% in ToolBench) is inconsistent. Users may be required to authorise the use of these tools or APIs, and tools provided by developers may be accessible during the initial construction of the benchmark but become unavailable later. This fluctuation further undermines the reliability and reproducibility of model performance assessments over time. This situation results in a problem where the constructed queries in the benchmarks may no longer be completed with their originally referenced tools. Consequently, it is crucial to strike a balance between enhancing the stability of these benchmarks and maintaining their diversity and scope.

To address these issues, we propose a new benchmark named StableToolBench, which incorporates a virtual API system and a stable evaluation system. We first build a virtual API system to replace the real one. As a start, we build a caching system to store the outputs of API calls. This approach ensures the stability and reproducibility of API behaviours. Given the limited number of benchmark questions, our caching system can cover a significant number of API call scenarios. However, relying solely on a cache is insufficient because many APIs remain unavailable. To resolve this problem, we use large language models (LLMs) to simulate the behaviours of these APIs. Specifically, we feed the documentation and few-shot real API calls if available in the cache to LLMs and ask LLMs to mock the behaviour of the APIs given a request. As a result, users can always get responses from APIs in an indistinguishable way as long as the LLMs are accessible. On the whole, our system first tries to find a hit in the cache. Unless there is a cache miss and a real API call is not received, the simulated server will be used.

We then improve the evaluation system to make it more stable. We design two metrics (i.e., SoPR and SoWR) after judging solvable tasks and replace all the automatic evaluators with GPT-4 to mitigate the randomness and indistinguishability during evaluation. Experiments demonstrate that our virtual API system, when combined with the improved evaluation system, can provide stable evaluation against API modifications. Furthermore, our system exhibits significant reliability in terms of realism, diversity, and documentation following accuracy.

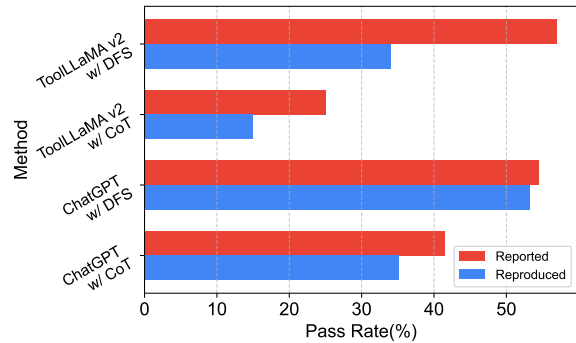


Figure 1: Comparison of performance (Pass Rate) reported in the paper and reproduced by us on the I1-Instruction group of ToolBench.

The main contributions of our work are summarised as follows:

- A tool-learning benchmark featured a large number of cached stable simulated APIs, well-balancing stability and reality of the APIs and much more stable evaluation metrics.
- Extensive experiments show that our benchmark provides much more stable model performance, robust to various types of API failures.
- Besides enhanced stability, our virtual API system exhibits reality, diversity and reliability comparable to that of the real API system.

2 Stability Analysis on ToolBench

In this section, we initiate a comprehensive analysis to reveal the stability of established tool benchmarks, using Toolbench (Qin et al., 2023b) as a case study. We examine the stability of ToolBench by investigating three key dimensions: performance, evaluation, and API status.

2.1 Stability of Performance

Benchmarks are designed to consistently evaluate the performance of various models over time. To test this consistency, we reproduce the model performances and record any variations. Our study employs Chain-of-Thought (CoT; Wei et al., 2023) and Depth First Search (DFS) strategies, leveraging ChatGPT and ToolLLaMA for comparative analysis. We adhere strictly to the configurations detailed in ToolBench, utilising the ChatGPT version gpt-3.5-turbo-0613 and ToolLLaMA-v2. As depicted in Figure 1, we compare the original Pass Rates for the I1-Instruction group reported by ToolBench with the Pass Rates we reproduced for four conventional methods. Our findings indicate a no-

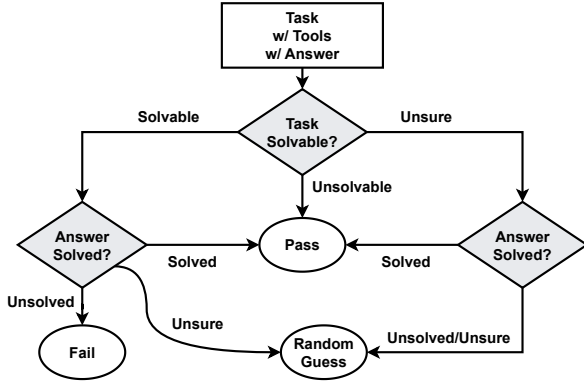


Figure 2: Pass Rate evaluation in ToolBench paper.

Method	Task			Answer			Pass	Win
	S	US	UE	S	US	UE		
CoT	168	23	9	19	170	11	33.0	50.0
	165	29	6	16	174	10	31.5	46.5
	151	40	9	20	167	13	37.5	53.0
DFS	116	68	16	17	167	16	50.5	54.0
	122	59	19	20	162	18	46.5	48.0
	132	54	14	22	157	21	55.0	56.0

Table 1: Experiments use GPT-3.5-Turbo-0613 with CoT and DFS. S, US, and UE indicate solvable (solved), unsolvable (unsolved), and unsure. Pass and Win denote pass rate and win rate, respectively. Win rates are evaluated against the first run of CoT. This experiment is run on 4 Feb 2024.

table decline in the performance of all methods over time, which raises concerns about the stability of ToolBench as a benchmark.

2.2 Stability of Evaluation

In ToolBench, there are two types of metrics, including Pass Rate (PR) and Win Rate (WR). PR is calculated based on using gpt-3.5-turbo-16k to determine if a task is solvable and whether the generated answer can solve the corresponding task. Figure 2 details the computation process of PR. Specifically, a solvable task results in a pass if the answer is solved, a failure if unsolved, and is randomly determined if unsure. For tasks deemed unsure, a pass is assigned only if the answer is solved; otherwise, a random outcome is chosen. If a task is unsolvable, the result defaults to a pass regardless of the answer status. Moreover, WR is derived from the comparative PR of paired candidates. A candidate’s WR increases by one each time it passes while the other fails. In all other situations, WR relies on gpt-3.5-turbo-16k to automatically evaluate the paired candidates.

To assess the stability of evaluation, we perform both CoT and DFS using gpt-3.5-turbo-0613

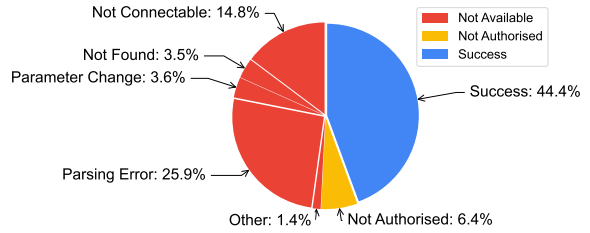


Figure 3: Statistics of API changes. Parsing errors are caused by post-processing errors of local API documentation, which have been solved in our benchmark.

on the I1-Instruction group dataset. These analyses are conducted using the provided tools and repeat over three iterations each. The resulting PR and WR are presented in Table 1, with detailed task and answer items. Despite PRs of DFS being generally higher than CoT, the contribution of the task is larger than the answer. However, it is worth noting that the tasks are the same in both CoT and DFS, where their results are expected to be consistent. On the contrary, the discrimination of answers between CoT and DFS is weak, where a considerable proportion are unsolved. Moreover, WR does not reflect the same trend as PR, where the second run of WR in DFS (48.0) is even lower than the first run of CoT. Therefore, all the phenomena reflect that gpt-3.5-turbo-16k can not assume the role of the automatic evaluator in tool learning, which will be discussed in Section 4.6.

2.3 Stability of API Status

We investigate the change of API status in ToolBench. In detail, we scan the original APIs downloaded from ToolBench, and use gpt-4-turbo to automatically write calls via the prompts as shown in Appendix H. According to the keyword in API feedback, we classify these APIs into three categories: success, not availability, and not authorisation³. The API status and the detailed errors of not availability are presented in Figure 3. As can be seen, only 44.4% of API calls are successful, while other API calls are mostly not available with various errors and some are not authorised.

Furthermore, to validate the impact of API call failures on the stability of model performance, we manually make some success tools⁴ down by returning a special failure call. Specifically, we ran-

³Note that we use the toolbench-key provided in ToolBench to simulate the real running process in the benchmark.

⁴In ToolBench, a tool is composed of several APIs. For example, a database tool can have two APIs: a writing API and a reading API.

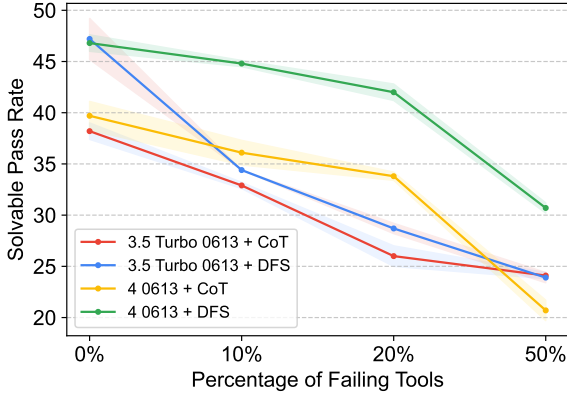


Figure 4: Solvable Pass Rate (SoPR) change when manually making APIs down on the I1 Instruction group.

domly sample a proportion of success tools containing success APIs found in Figure 3. At testing time, when sampled tools are called, a special response will be thrown: {"error": "", "response": "This API did not return any useful information..."} to simulate the API call failures. We conduct baseline models with different proportions (i.e., 0%, 10%, 20%, and 50%) of sampled APIs on the I1 Instruction set. Due to the issues in evaluation, we use our stable evaluation system proposed in Section 3.2 as the same as our main experiments. For each experiment, we evaluate three times and report the average scores as shown in Figure 4. It can be seen that the performance degrades a lot when the proportion of successful APIs is down, thus the impact of API status on stability is considerable.

3 StableToolBench

Considering that stability is a crucial feature of benchmarking, in this paper, we specifically design a virtual API server and stable evaluation system to improve the stability based on ToolBench, and propose a new benchmark, named StableToolBench.

3.1 Virtual API Server

With real APIs, many of the failures encountered when reproducing its experiments are caused by expired APIs, network issues, or failed authentication. To address this problem, we specifically propose a virtual API server with two components as illustrated in Figure 5, including a caching system and API simulator. Moreover, we design API calling rules to combine these two components to ensure the virtual API server is stable.

Caching System. We first implement a caching system that stores responses of all API callings to ensure consistency. The caching system uses keys composed of their category, tool, API name, and arguments. As a start, we populate the initial cache using the API call history from the training data and the reproduced test data released in ToolBench⁵. To ensure the quality of cached APIs, only valid records following the rule in Appendix D will be saved. It is worth noting that we will also reserve some APIs with exceptions to keep the reality. In this way, most API responses will be readily available, allowing the benchmark focus on probing the tool usage ability of designed methods with minimal impact on tool availability. Furthermore, the API call in new experiments will also be continuously updated in the cache to ensure scalability. The statistics of cache are shown in Table 2. Additionally, as an extra benefit, this approach re-

Source	Train Set	Test Set	New Exp	Total
Before	58,105	5,921	255,828	352,630
After	25,995	2,393	136,592	164,980

Table 2: Cache components and their sizes before and after filtration. The cache of new experiments is updated until 12 Feb 2024.

duces the latency introduced by interacting with real APIs, and also saves the costs for the API simulator discussed below.

API Simulator. Due to the limited coverage of the caching system, we propose to use LLMs to simulate API responses that are not in the cache and unavailable. Specifically, we ask gpt-4-turbo to simulate the API behaviour based on the original documentation in ToolBench. The API documentation includes the descriptions of the functions and their corresponding parameters. To mitigate the difference between simulated and real APIs, we use real API calls in the caching system as few-shot examples (Brown et al., 2020) for the LLM to better mock the behaviours. We keep the maximum number of examples at five. When less than five examples exist in the cache, all of them will be used. Detailed prompts can be found in Appendix F.

API Calling Rules. Based on the caching system and the API simulator, we create API calling rules to ensure the stability of the virtual API server.

⁵<https://drive.google.com/drive/folders/1yBUQ732mPu-Kc1JnuQEhTkaKdXFc3J>

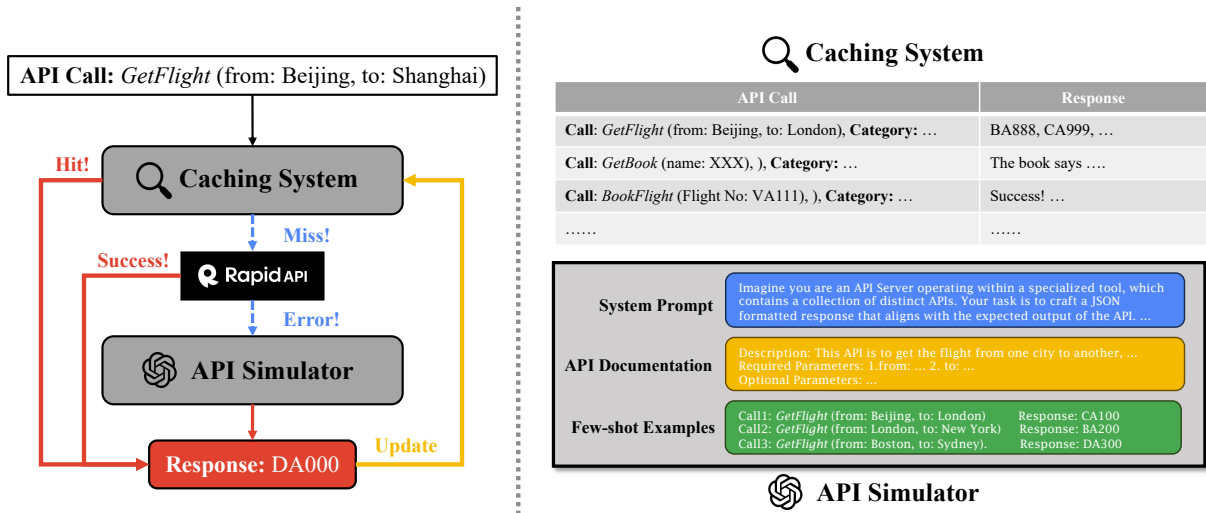


Figure 5: The process of calling APIs in our proposed virtual API server.

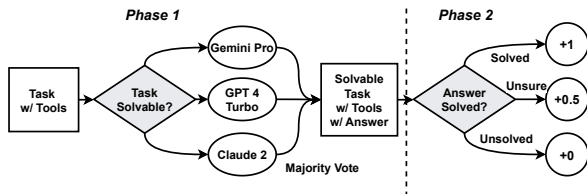


Figure 6: The process of our SoPR evaluation.

	I1-I	I1-C	I1-T	I2-I	I2-C	I3-I	Total
Full	200	200	200	200	200	100	1,100
Solvable	163	153	158	106	124	61	765

Table 3: Statistics of original full and solvable tasks before and after judging. C, I, T stands for the Category, Instruction and Tool subgroup of the test set. Experiments below follow the denotation.

When a call request ($e, args$), where e is the API endpoint and $args$ is the arguments for that endpoint, is received, the system will first search the caching system for ($e, args$) pair. If a cache hit exists, the cached response will be directly returned. When there is a cache miss, then the system will try to call the real API for a response to maintain the reality of the whole system. If the real API calling is successful, the response will be returned. However, when the caching system does not contain and the real API is not available, the system will finally call the simulated API. The final response whether from the real API or the simulated API will be saved to update in the caching system.

3.2 Stable Evaluation System

In this section, we propose a two-phase evaluation process, including judging solvable tasks and evaluating with two metrics, as shown in Figure 6. Moreover, we replace all the automatic evaluators with gpt-4-turbo.

Judging Solvable Tasks. Considering that both unsolvable and unsure tasks would introduce enormous randomness while the results of tasks fluctuate wildly, we try to obtain a fixed collection of solvable tasks to eliminate the problems. To

achieve this, we use three state-of-the-art LLMs, i.e., gpt-4-turbo, gemini-pro, and claude-2, to determine whether a task is solvable or unsolvable. The prompt is shown in Appendix G. The task will be judged as solvable when more than two models evaluate it as solvable. The statistics of solvable tasks across all test datasets in ToolBench are shown in Table 3.

Metrics. We then report Solvable Pass Rate (SoPR) and Solvable Win Rate (SoWR) based on our obtained solvable tasks. Due to the limitation of gpt-3.5-turbo-16k in tool learning, we uniformly adopt gpt-4-turbo as the automatic evaluator. SoPR is in essence PR with all tasks solvable and only assesses the answers using the same prompt in ToolBench. The evaluator assigns outcomes of answers categorised as Solved, Unsolved, or Unsure, which respectively contribute scores of 1, 0.5, and 0 to the overall SoPR calculation. As for SoWR, when one is solved and the other is unsolved, the solved one wins. Under other circumstances, gpt-4-turbo will be used to make a win-lose decision.

Method	I1 Instruction	I1 Category	I1 Tool	I2 Category	I2 Instruction	I3 Instruction	Average
3.5 0613 (C)	52.2±1.1	47.3±0.6	53.6±1.3	42.5±2.1	35.8±2.0	48.1±0.8	46.6±1.3
3.5 0613 (D)	60.3±1.3	66.2±1.2	67.1±0.0	59.1±0.4	51.3±1.2	73.8±2.3	63.0±1.1
4 0613 (C)	45.5±0.4	57.4±0.3	48.8±0.7	43.0±0.7	46.5±0.9	48.1±1.5	48.2±0.8
4 0613 (D)	57.3±0.6	57.3±0.3	60.9±1.0	57.9±1.0	51.3±0.8	66.4±2.4	58.5±1.0
T-LLaMA (C)	32.3±1.0	40.3±0.8	36.7±0.5	34.7±0.7	25.2±0.4	33.9±1.5	33.9±0.8
T-LLaMA (D)	44.5±0.9	49.6±1.3	48.9±2.7	50.8±1.1	31.9±1.9	53.6±2.0	46.6±1.7
3.5 1106 (C)	50.4±0.5	45.1±1.4	50.8±0.3	48.7±0.8	42.1±0.4	55.7±0.0	48.8±0.6
3.5 1106 (D)	62.8±0.3	63.9±1.2	65.6±0.3	56.5±0.7	56.9±1.2	67.2±1.3	62.2±0.8
4 Turbo (C)	52.8±1.3	56.6±0.9	51.9±0.5	51.9±1.0	52.8±0.8	52.5±0.0	53.1±0.8
4 Turbo (D)	59.2±0.5	61.7±0.7	65.7±1.0	55.6±0.6	55.2±0.4	66.1±4.3	60.6±1.3

Table 4: Solvable pass rate scores. We run all models once, evaluate three times and take the average results. “3.5 0613”, “4 0613”, “3.5 1106”, “4 Turbo”, “T-LLaMA” stands for gpt-3.5-turbo-0613, gpt-4-0613, gpt-3.5-turbo-1106, gpt-4-turbo-preview, ToolLLaMA v2 respectively. C and D stand for CoT and DFS respectively. The experiments below follow the denotation. We use gpt-4-turbo-2024-04-09 as the evaluator. Evaluation done on May 2024.

Method	I1-I	I1-C	I1-T	I2-I	I2-C	I3-I	Avg
3.5 0613 (D)	60.7	67.3	59.5	63.2	62.1	75.4	64.7
4 0613 (C)	54.6	58.8	58.2	75.5	60.5	62.3	61.7
4 0613 (D)	62.6	62.7	58.2	74.5	62.9	67.2	64.7
T-LLaMA (C)	31.3	28.1	33.5	35.8	33.9	24.6	31.2
T-LLaMA (D)	44.8	45.8	44.3	59.4	41.1	50.8	47.7
3.5 1106 (C)	47.2	47.7	44.9	50.9	54.0	62.3	51.2
3.5 1106 (D)	55.8	53.6	51.9	68.9	59.7	68.9	59.8
4 Turbo (C)	71.2	77.1	61.4	79.2	71.8	67.2	71.3
4 Turbo (D)	73.0	75.2	68.4	77.4	66.9	60.7	70.2

Table 5: Solvable Win Rate scores. We run all models once against GPT-3.5-Turbo-0613 + CoT and evaluate three times. We follow the ToolBench implementation to take the most frequent result for each query during evaluation. The experiments below follow the denotation. We use gpt-4-turbo-2024-04-09 as the evaluator. Evaluation done on May 2024.

4 Experiment

4.1 Performance

Following ToolBench, we run gpt-3.5-0613, gpt-4-0613, ToolLLaMA-v2 with CoT and DFS, replenishing with latest models gpt-3.5-1106 and gpt-4-turbo. The results of SoPR and SoWR are shown in Tables 4 and 5. Generally, GPT-4 series models outperform GPT-3.5 models, while ToolLLaMA performs worst with the same inference algorithm. Also, DFS significantly outperforms CoT whichever LLMs are used. These phenomena are consistent with ToolBench. Furthermore, probably thanks to the improved function calling capabilities, newer GPT models performed better.

4.2 Stability of Virtual API Server

Following the same setups as in Section 2.3, we manually make the same success tools not available during the running time. In our design, when a

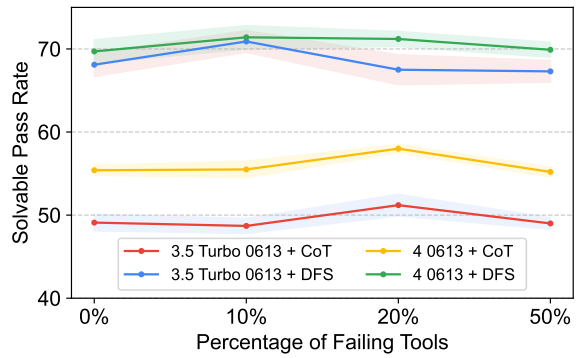


Figure 7: Performance change when manually making APIs down with our virtual online API system. The results are averaged over all six groups. Solving rates are reported. We run each experiment one time and evaluate it three times and take the average score. Unless otherwise stated, gpt-4-turbo-preview at the time of testing is used in this section. This experiment was done in Feb 2024.

call is on an unavailable tool, it will be directed to the simulated API immediately. Compared to Section 2.3, the results as shown in Figure 7 are much more stable with our virtual API server. Even when 50% of APIs are not available, changes in performance are still not significant, which is explainable within the range of variance.

Considering we use gpt-4-turbo as the backbone of the API simulator which may change even with the same version number, we ablate different versions and different temperatures of gpt-4-turbo. The results are shown in Table 6. Under different settings of the backbone LLMs, the performance change is still acceptable within the variance of LLM evaluation, indicating the robustness of our API simulators.

GPT-4 Config	1106-preview		0125	
	T=1	T=0.1	T=1	T=0.1
3.5 0613 (C)	49.1±1.0	49.0±0.8	52.1±0.5	50.2±0.8
3.5 0613 (D)	68.1±1.4	67.4±0.9	69.3±1.0	67.9±1.2
4 0613 (C)	55.4±0.6	56.3±0.7	57.7±0.5	54.5±0.6
4 0613 (D)	69.7±1.4	70.3±1.0	71.4±0.7	70.4±1.3
3.5 1106 (C)	52.1±0.7	51.1±0.5	54.5±0.9	52.7±0.6
3.5 1106 (D)	69.9±0.7	71.2±0.9	70.0±0.9	71.0±0.9
4 Turbo (C)	60.8±0.7	62.4±0.8	63.6±0.4	64.0±1.1
4 Turbo (D)	73.2±1.1	76.2±0.9	75.0±0.7	77.3±0.9

Table 6: Performance of baselines with different settings of the LLM server. Results are averaged over all groups and reported in SoPR. We run each experiment one time and evaluate three times and take the average score.

4.3 Turing Test of API Simulator

To test the effectiveness of API simulators, we design a ‘‘Turing Test’’ (Turing, 2009) between the real APIs and the simulated ones. Note that we believe it is not required for API simulators to exactly output the same answers as those of real APIs, where rationality is more important. For example, when a query asks about the weather today, the API simulator does not need to retrieve the ‘‘real’’ temperature. Instead, the API simulator needs to generate a reasonable temperature number.

To do the test, we first sample 70 available real APIs and their corresponding simulated APIs. Given the API callings and their real and simulated response pairs, we ask three human annotators to determine which response more closely resembles an actual API response overall, based on the given descriptions of the API functions. We ask human annotators to evaluate along three dimensions: Overall, Format Accuracy and Answer Relevance. The annotator first need to answer which response is overall more like a real response. When assessing format accuracy, annotators must determine which response more accurately adheres to the format specifications outlined in the documentation. In evaluating answer relevance, they are tasked with identifying which response more effectively fulfills the instruction in accordance with the documentation’s guidelines. The results are shown in Figure 8. Surprisingly, human annotators cannot distinguish simulated and real APIs very well, where the simulated APIs are judged to act more like real situations. Moreover, the proportion of tie is much larger, indicating that simulated APIs can work very similarly to real APIs.

In addition to the Turing Test mentioned above, we assess the quality of LLM simulations by evaluating the adherence of simulated outputs to their

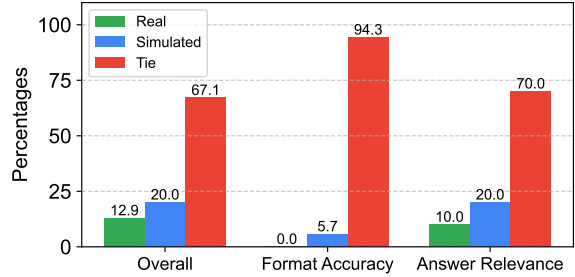


Figure 8: Results of the ‘‘Turing Test’’ for the real and simulated APIs. Results are win-lose-tie percentages.

corresponding documentation. To conduct this evaluation, we randomly sample 50 simulated outputs along with their documentation from the Turing Test dataset. A human evaluator is then tasked with determining whether the LLM simulations reasonably follow the provided documentation. The results indicate that 90% of the simulations are deemed reasonable, 6% are considered unreasonable, and 4% are uncertain. These findings suggest that the LLM is highly capable of generating simulated responses that adhere closely to the provided documentation.

4.4 Diversity of API Simulator

With the LLM simulation, API simulators will not exactly feedback the same as real APIs. Hence, a natural concern is whether the simulated APIs will degrade in diversity in API functionalities. To study the problem, firstly, we explore the distribution of real and simulated API responses. We first use all 246 APIs in the Tool API category from the successful APIs mentioned in Section 2.3. Then, we use the same call arguments to call these real and simulated APIs. All the responses are encoded using S-BERT (Reimers and Gurevych, 2019) and their corresponding embeddings are visualised by UMAP (McInnes et al., 2018). Detailed configuration is shown in Appendix E. The result is shown in Figure 9. As can be seen from the figure, real and simulated APIs occupied similar embedding space, indicating that the diversity of simulated APIs is similar to the real APIs.

Secondly, we try to explore the behaviour when a simulated API is given several calls with different input arguments. We sample 60 APIs from all successful APIs and make 5 different calls to each API, using the same prompt as in Appendix H. We then count the number of APIs that give exactly the same responses in any 2 of the 5 calls. Results show that only 2 of 60 APIs contain such

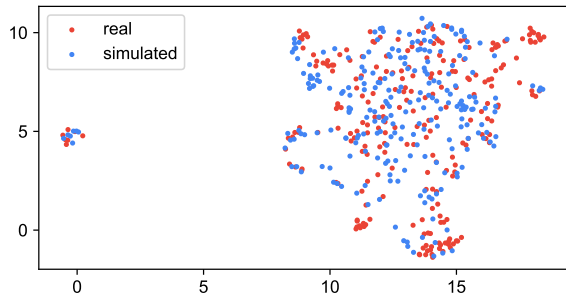


Figure 9: Visualisation of the embeddings of responses from real and simulated APIs.

responses, which supports the sufficient diversity of our simulated APIs.

4.5 Effectiveness of Caching System

To show the effectiveness of our caching system in maintaining the stability of the virtual API server, we run several methods and record the cache hit rates. In detail, we run four methods used in ToolBench, `gpt-3.5-turbo-0613` and `gpt-4-0613` with CoT and DFS. Reproduction data in ToolBench of these methods has been used in the cache. Results are recorded in Table 7, which shows that rerunning these in-domain methods has a very high cache hit rate. This means that most of the call responses are fixed and instability from the API system are much smaller. Nevertheless, models and methods may change over time, and therefore, we further run `gpt-3.5-turbo-1106` and `gpt-4-turbo-preview` with CoT and DFS. As can be seen in the table, although the cache hit rates are smaller with these out-of-domain methods, the scores are still high enough to mitigate the instability significantly.

4.6 Human Evaluation of Evaluator

Considering that GPT-3.5 is limited to evaluating the performance in tool learning, we replace the automatic evaluators with stronger LLMs. In this section, we manually assess the correctness of different automatic evaluators. Here, we sample 100 task-solvable questions, 50 answer-solving questions in the PR / SoPR evaluation, and 50 comparison questions in the WR / SoWR evaluation from the experiments running.

We then collect all the corresponding answers of different LLMs during previous evaluations. These questions are further manually labeled by three human annotators to obtain the ground truth. With the ground truth, we calculate the accuracy scores

Methods	Final	Mid	Start
GPT 3.5 Turbo 0613 + CoT	96.7	36.2	11.7
GPT 3.5 Turbo 0613 + DFS	97.0	34.5	11.6
GPT 4 0613 + CoT	96.5	36.2	11.7
GPT 4 0613 + DFS	97.0	35.0	11.7
GPT 3.5 Turbo 1106 + CoT	91.4	35.1	11.8
GPT 3.5 Turbo 1106 + DFS	75.8	34.5	11.4
GPT 4 Turbo + CoT	88.2	35.0	11.6
GPT 4 Turbo + DFS	77.8	34.5	11.8

Table 7: Cache hit rate (%) with various models and methods. Final, Mid, and Start represent the final version of the cache, the mid-way version containing 151,152 (91.6%) items of the final version, and the starting version containing only the train and test set. Experiments are independent runs of Section 4.1 with fixed cache, run on 13 Feb 2024.

of these models and show the results in Table 8. It can be seen that our used LLMs (i.e., Claude 2, Gemini, and GPT-4) are much better than GPT-3.5 in ToolBench to determine the solvability of tasks, where the Gemini and GPT-4 outperform by a large margin. In both the evaluation of answers and comparison, GPT-4 significantly outperforms GPT-3.5, especially in the comparison to compute WR. It is worth noting that all the accuracies of GPT-3.5 are lower than 70%, indicating that GPT-3.5 cannot assess the performance in tool learning.

5 Related Work

Tool Learning Benchmarks. Recent studies have shed light on the burgeoning capabilities of LLMs in understanding and mastering tools (Li et al., 2023a; Patil et al., 2023; Yang et al., 2023b; Song et al., 2023; Tang et al., 2023; Ye et al., 2024; Xu et al., 2023). Gaining access to external tools endows LLMs with real-time factual knowledge (Yang et al., 2023a), multimodal functionalities (Gupta and Kembhavi, 2023), and specialised skills in vertical domains (Jin et al., 2024). However, few work has been done to explore the stability of the tool environment in specific benchmarks and how it affects the LLMs’ performance in tool-augmented tasks.

Tool Inference Methods. Recent literature has begun to explore various methodologies for integrating tool functionalities within LLMs. Notably, the robust in-context learning prowess of LLMs, as demonstrated in (Brown et al., 2020), has facilitated the augmentation of LLMs with external tools via in-context tool descriptions and demonstrations (Hsieh et al., 2023; Ruan et al., 2023; Mialon et al.,

Methods	Solvability	Solving	Comparison
Claude 2	71.0	-	-
Gemini Pro	82.0	-	-
GPT 3.5 Turbo	65.0	68.0	56.0
GPT 4 Turbo	80.0	74.0	78.0

Table 8: Human evaluation on task solvability, answer solving (for pass rate) and comparison (for win rate).

2023). An alternative approach involves the explicit training of LLMs (Patil et al., 2023; Tang et al., 2023; Chen et al., 2023a; Qin et al., 2023c; Huang et al., 2023) using datasets enriched with tool interactions, thereby familiarising models with the nuances of tool usage.

Evaluation in Tool Learning. Evaluating the performance of LLMs in tool-augmented tasks presents unique challenges and opportunities. Numerous works have been developed for the assessment of tool utilisation, primarily emphasising response comparison (Zhuang et al., 2023), tool call accuracy (Patil et al., 2023), or a synthesis of these aspects (Li et al., 2023a). Distinguishing itself, (Qin et al., 2023c) introduces an innovative methodology by integrating a large language model (LLM) as a judge to evaluate the comprehensive solution path. Subsequent research (Wang et al., 2023) focuses on the multi-turn interaction capabilities of LLMs with both tools and user feedback. In a departure from the aforementioned approaches, (Chen et al., 2023b) presents itself as the inaugural benchmark specifically tailored for the fine-grained assessment of tool utilisation capabilities. However, there exists a gap in the literature concerning the exploration of evaluation stability when evaluating the tool usage capabilities of LLMs.

6 Conclusion

In this paper, we propose StableToolBench, a benchmark developed to enhance the stability of ToolBench. Our analysis identified instability issues in the evaluation processes of ToolBench and API status, causing variability in model performance assessments. To address this, we implement a caching system for consistent data availability. We also replace the real API server with an LLM-simulated virtual server for reliable API behaviour simulation. Experiments show that StableToolBench significantly improves the stability of model performance evaluations, with the simulated APIs offering realism and the caching system

contributing greatly to the enhanced stability of the benchmark.

Acknowledgement

This work is supported by the National Natural Science Foundation of China (No. 62276152, 61925601). We also extend our gratitude to Jingwen Wu and Yao Li for their assistance with human evaluation and additional suggestions.

Limitations

In this work, we propose StableToolBench, a new tool learning benchmark with increased stability but non-declining reality. However, our work faces the following limitations. Firstly, we used GPT-4 as the automatic evaluator in the evaluation process and as the backbone server, which increase the cost of using our benchmark. Secondly, the GPT-4 backbone server demonstrate strong performance in simulating API behaviours. Nevertheless, the backbone LLM may experience significant upgrades, which may affect the performance. Therefore, we believe that the ultimate solution is to solve the problem with a trained open-source LLM. However, current open-source LLMs are not performant enough to simulate API behaviours well. As a result, closed-source LLMs are the only options. We believe that when open-source LLMs are strong enough to be well suited for this task, In the future, we may turn to open-source LLMs when they are strong enough to be well suited for this task. Thirdly, although the cache hit rates are high with our explored methods, new methods will be developed in the future. Whether this cache will still be effective is unsure. In this regard, we aim to continuously update the cache in the future in a slow pace for both balanced stability and effectiveness.

References

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. *Language Models are Few-Shot Learners*. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Pro-*

- cessing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual.
- Tianle Cai, Xuezhi Wang, Tengyu Ma, Xinyun Chen, and Denny Zhou. 2024. [Large Language Models as Tool Makers](#). In *Proc. of The Twelfth International Conference on Learning Representations (ICLR 2024)*.
- Baian Chen, Chang Shu, Ehsan Shareghi, Nigel Collier, Karthik Narasimhan, and Shunyu Yao. 2023a. [FireAct: Toward language agent fine-tuning](#). *ArXiv preprint*, abs/2310.05915.
- Zehui Chen, Weihua Du, Wenwei Zhang, Kuikun Liu, Jiangning Liu, Miao Zheng, Jingming Zhuo, Songyang Zhang, Dahua Lin, Kai Chen, et al. 2023b. [T-Eval: Evaluating the Tool Utilization Capability Step by Step](#). *ArXiv preprint*, abs/2312.14033.
- Gemini Team. 2023. [Gemini: A Family of Highly Capable Multimodal Models](#).
- Tanmay Gupta and Aniruddha Kembhavi. 2023. [Visual programming: Compositional visual reasoning without training](#). In *In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR 2023)*, pages 14953–14962.
- Shibo Hao, Tianyang Liu, Zhen Wang, and Zhiting Hu. 2023. [ToolkenGPT: Augmenting Frozen Language Models with Massive Tools via Tool Embeddings](#). *ArXiv preprint*, abs/2305.11554.
- Cheng-Yu Hsieh, Si-An Chen, Chun-Liang Li, Yasuhisa Fujii, Alexander Ratner, Chen-Yu Lee, Ranjay Krishna, and Tomas Pfister. 2023. [Tool documentation enables zero-shot tool-usage with large language models](#). *ArXiv preprint*, abs/2308.00675.
- Yue Huang, Jiawen Shi, Yuan Li, Chenrui Fan, Siyuan Wu, Qihui Zhang, Yixin Liu, Pan Zhou, Yao Wan, Neil Zhenqiang Gong, et al. 2023. [MetaTool benchmark for large language models: Deciding whether to use tools and which to use](#). *ArXiv preprint*, abs/2310.03128.
- Qiao Jin, Yifan Yang, Qingyu Chen, and Zhiyong Lu. 2024. [GeneGPT: augmenting large language models with domain tools for improved access to biomedical information](#). *Bioinformatics*, 40(2):btac075.
- Minghao Li, Feifan Song, Bowen Yu, Haiyang Yu, Zhoujun Li, Fei Huang, and Yongbin Li. 2023a. [API-Bank: A Benchmark for Tool-Augmented LLMs](#).
- Minghao Li, Feifan Song, Bowen Yu, Haiyang Yu, Zhoujun Li, Fei Huang, and Yongbin Li. 2023b. [API-Bank: A Comprehensive Benchmark for Tool-Augmented LLMs](#). *ArXiv preprint*, abs/2304.08244.
- Pan Lu, Baolin Peng, Hao Cheng, Michel Galley, Kai-Wei Chang, Ying Nian Wu, Song-Chun Zhu, and Jianfeng Gao. 2023. [Chameleon: Plug-and-Play Compositional Reasoning with Large Language Models](#). *ArXiv preprint*, abs/2304.09842.
- Leland McInnes, John Healy, Nathaniel Saul, and Lukas Grossberger. 2018. [UMAP: Uniform Manifold Approximation and Projection](#). *The Journal of Open Source Software*, 3(29):861.
- Grégoire Mialon, Roberto Dessì, Maria Lomeli, Christoforos Nalmpantis, Ram Pasunuru, Roberta Raileanu, Baptiste Rozière, Timo Schick, Jane Dwivedi-Yu, Asli Celikyilmaz, et al. 2023. [Augmented Language Models: A Survey](#). *ArXiv preprint*, abs/2302.07842.
- Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, Xu Jiang, Karl Cobbe, Tyna Eloundou, Gretchen Krueger, Kevin Button, Matthew Knight, Benjamin Chess, and John Schulman. 2022. [WebGPT: Browser-assisted question-answering with human feedback](#).
- OpenAI. 2023. [GPT-4 Technical Report](#).
- Shishir G. Patil, Tianjun Zhang, Xin Wang, and Joseph E. Gonzalez. 2023. [Gorilla: Large Language Model Connected with Massive APIs](#). *ArXiv preprint*, abs/2305.15334.
- Yujia Qin, Zihan Cai, Dian Jin, Lan Yan, Shihao Liang, Kunlun Zhu, Yankai Lin, Xu Han, Ning Ding, Huadong Wang, et al. 2023a. [WebCPM: Interactive Web Search for Chinese Long-form Question Answering](#). *ArXiv preprint*, abs/2305.06849.
- Yujia Qin, Shengding Hu, Yankai Lin, Weize Chen, Ning Ding, Ganqu Cui, Zheni Zeng, Yufei Huang, Chaojun Xiao, Chi Han, et al. 2023b. [Tool learning with foundation models](#). *ArXiv preprint*, abs/2304.08354.
- Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, Sihan Zhao, Runchu Tian, Ruobing Xie, Jie Zhou, Mark Gerstein, Dahai Li, Zhiyuan Liu, and Maosong Sun. 2023c. [ToolLLM: Facilitating Large Language Models to Master 16000+ Real-world APIs](#).
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Jingqing Ruan, Yihong Chen, Bin Zhang, Zhiwei Xu, Tianpeng Bao, Guoqing Du, Shiwei Shi, Hangyu Mao, Xingyu Zeng, and Rui Zhao. 2023. [TPTU: Task planning and tool usage of large language model-based AI agents](#). *ArXiv preprint*, abs/2308.03427.
- Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. [Toolformer: Language Models Can Teach Themselves to Use Tools](#). *ArXiv preprint*, abs/2302.04761.

- Yifan Song, Weimin Xiong, Dawei Zhu, Cheng Li, Ke Wang, Ye Tian, and Sujian Li. 2023. [Rest-GPT: Connecting Large Language Models with Real-World Applications via RESTful APIs](#). *ArXiv preprint*, abs/2306.06624.
- Qiaoyu Tang, Ziliang Deng, Hongyu Lin, Xianpei Han, Qiao Liang, and Le Sun. 2023. [ToolAlpaca: Generalized Tool Learning for Language Models with 3000 Simulated Cases](#).
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. [LLaMA: Open and Efficient Foundation Language Models](#).
- Alan M. Turing. 2009. *Computing Machinery and Intelligence*, pages 23–65. Springer Netherlands, Dordrecht.
- Xingyao Wang, Zihan Wang, Jiateng Liu, Yangyi Chen, Lifan Yuan, Hao Peng, and Heng Ji. 2023. [MINT: Evaluating LLMs in multi-turn interaction with tools and language feedback](#). *ArXiv preprint*, abs/2309.10691.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023. [Chain-of-Thought Prompting Elicits Reasoning in Large Language Models](#).
- Qiantong Xu, Fenglu Hong, Bo Li, Changran Hu, Zhengyu Chen, and Jian Zhang. 2023. [On the Tool Manipulation Capability of Open-source Large Language Models](#).
- Linyao Yang, Hongyang Chen, Zhao Li, Xiao Ding, and Xindong Wu. 2023a. [ChatGPT is not Enough: Enhancing Large Language Models with Knowledge Graphs for Fact-aware Language Modeling](#). *ArXiv preprint*, abs/2306.11489.
- Rui Yang, Lin Song, Yanwei Li, Sijie Zhao, Yixiao Ge, Xiu Li, and Ying Shan. 2023b. [GPT4Tools: Teaching Large Language Model to Use Tools via Self-instruction](#).
- Shunyu Yao, Howard Chen, John Yang, and Karthik Narasimhan. 2022a. [WebShop: Towards Scalable Real-World Web Interaction with Grounded Language Agents](#). In *In Proceedings of the Advances in Neural Information Processing Systems (NeurIPS 2022)*, volume 35, pages 20744–20757. Curran Associates, Inc.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2022b. [ReAct: Synergizing reasoning and acting in language models](#). volume abs/2210.03629.
- Junjie Ye, Guanyu Li, Songyang Gao, Caishuang Huang, Yilong Wu, Sixian Li, Xiaoran Fan, Shihan Dou, Qi Zhang, Tao Gui, and Xuanjing Huang. 2024. [ToolEyes: Fine-Grained Evaluation for Tool Learning Capabilities of Large Language Models in Real-world Scenarios](#).
- Yuchen Zhuang, Yue Yu, Kuan Wang, Haotian Sun, and Chao Zhang. 2023. [ToolQA: A Dataset for LLM Question Answering with External Tools](#). *ArXiv preprint*, abs/2306.13304.

A Comparison of Reported and Reproduced Performance

Detailed comparison scores of reported and reproduced performance are shown in Table 9.

Method	Reported	Reproduced
GPT 3.5 Turbo 0613 + CoT	41.5	35.2 <i>-32.5%</i>
GPT 3.5 Turbo 0613 + DFS	54.5	53.2 <i>-2.4%</i>
ToolLLaMA v2 + CoT	25.0	15.0 <i>-40%</i>
ToolLLaMA v2 + DFS	57.0	34.0 <i>-40.4%</i>

Table 9: Comparison of performance (Pass Rate) reported in the paper and reproduced by us of ChatGPT and ToolLLaMA v2 on the I1-Instruction group of ToolBench.

B Statistics of API change information

Detailed statistics of API change categories and information are shown in Table 10 and Table 11.

Status Type	Number	Percentage (%)
Not Available	8095	49.2
Not Authorised	1058	6.4
Success	7311	44.4

Table 10: APIs changed in ToolBench.

Status Type	Number	Percentage (%)
Not Connectable	2426	30.0
Not Found	583	7.2
Parameter Change	591	7.3
Parsing Error	4247	52.6
Other	248	3.1
Total	8095	100

Table 11: Categories of Not Availability in ToolBench.

C Stability Test Scores with Virtual API Systems

Detailed scores of stability tests of various models are shown in Table 12. Note that in addition to GPT 3.5 Turbo 0613 and GPT 4 0613, we report the performance of newer versions, namely GPT 3.5 Turbo 1106 and GPT 4 Turbo Preview.

D Call Error Identification and Cache Filtering Rule

We identify call errors and filter out invalid call to RapidAPI based on keyword occurrences. In detail, we identify the following error:

Method	Real API Failure Rate			
	0%	10%	20%	50%
GPT 3.5 Turbo 0613 + CoT	49.1±1.0	48.7±0.9	51.2±1.3	49.0±0.7
GPT 3.5 Turbo 0613 + DFS	68.1±1.4	70.9±1.3	67.5±1.8	67.3±1.3
GPT 4 0613 + CoT	55.4±0.6	55.5±1.0	58.0±0.5	55.2±0.6
GPT 4 0613 + DFS	69.7±1.4	71.4±1.4	71.2±0.9	69.9±0.9
GPT 3.5 Turbo 1106 + CoT	52.1±0.7	52.4±0.8	53.9±0.6	50.2±0.6
GPT 3.5 Turbo 1106 + DFS	69.9±0.7	71.7±0.7	69.4±0.8	71.6±0.9
GPT 4 Turbo preview + CoT	60.8±0.7	62.8±0.5	64.2±0.7	62.4±0.5
GPT 4 Turbo preview + DFS	73.2±1.1	76.7±1.0	76.0±0.8	74.2±1.3

Table 12: Performance change when manually make APIs down with our virtual online API system. The results are averaged over all six groups. Solving rates are reported. We run each experiment one time and evaluate three times and take the average score.

- Not Connected Error: when error information contains HTTP or the response information contains HTTP error, connection, rate limit, time(d) out;
- Not Found Error: when the error information or response contains not found, not available, API doesn't exist, Service Not Found, internal error or 404 error message;
- Parameter Change: when the error information or response contains parameter, parse, is not defined;
- Parsing Error: when the error information starts with Function executing from;
- Not Authorised: when the error information or response contains authoriz(s), unauthoriz(s), blocked user, unsubscribe, credential, disabled for your subscription, ACCESS_DENIED or 401, 403 error message;
- Other Errors: messages with non-empty error messages;
- Success: Other calls.

We consider all types of errors when identifying errors. However, when filtering the cache, we do not consider the "Other Errors".

E Configurations of API Diversity Analysis

The configurations of diversity analysis are as follows:

- Embedding model: all-mpnet-base-v2;

API Simulation Prompt	
System	<p>Imagine you are an API Server operating within a specialized tool, which contains a collection of distinct APIs. Your role is to deeply understand the function of each API based on their descriptions in the API documentation. As you receive specific inputs for individual API calls within this tool, analyze these inputs to determine their intended purpose. Your task is to craft a JSON formatted response that aligns with the expected output of the API, guided by the provided examples.</p> <p>Your responses must adhere to a specific JSON structure, which is as follows: <code>{ "error": "", "response": "Your_Response" }</code></p> <p>The error field should remain empty, indicating no errors in processing. The response field should contain the content you formulate based on the API's functionality and the input provided. Ensure that your responses are meaningful, directly addressing the API's intended functionality. If the provided examples are mostly error messages or lack substantial content, use your judgment to create relevant and accurate responses. The key is to maintain the JSON format's integrity while ensuring that your response is an accurate reflection of the API's intended output within the tool.</p> <p>Please note that your answer should not contain anything other than a json format object, which should be parsable directly to json.</p> <p>Note that:</p> <ul style="list-style-type: none"> - your response should be around 100 to 200 words, containing rich information given the api input parameters. Keep Your answer short and simple. - your response must be effective and have practical content. - if the api response example is null or ineffective, ignore the example and give your independent response.
User	<p>API Documentation: Documentation JSON file</p> <p>API Examples: Example input 1: Example response 1 Example input 2: Example response 2 Example input 3: Example response 3</p> <p>API input: Argument JSON string, e.g: <code>{"category": "Logistics", "tool_name": "SQUAKE", "api_name": "Checkhealth", "tool_input": "{}", "strip": "filter"}</code></p>

Table 13: Prompt used to simulate APIs.

- UMAP metric (distance metric): correlation;
- Num of neighbours: 15;
- Min distance: 0.5.

H Prompt Used to Make API Calls

The prompt used to construct API calls to scan availables is shown in Table 15.

F Prompts of API simulation

The prompt used to simulate API behaviours is shown in Table 13.

G Prompt to Filter Solvable Task

The prompt used to filter solvable tasks is shown in Table 14.

Solvable Task Filtration Prompt

Please check whether the given task solvable with following rules:

1. If the query provide invalid information (e.g. invalid email address or phone number), return Unsolvable
2. If the query needs more information to solve (e.g. the target restaurant name in a navigation task), return Unsolvable
3. If the current available_tools are enough to solve the query, return Solvable
4. Return only Solvable or Unsolvable

Task:{task}

Now please give your answer (only Solvable or Unsolvable):

Table 14: Prompt used to filter solvable tasks.

API Call Writing Prompt

	<p>Imagine you are an API requester, Your role is to deeply understand the function of each API based on their descriptions in the API documentation. Your task is to craft a JSON formatted input that aligns with the expected input of the API, guided by the provided examples.</p> <p>Your responses must adhere to a specific JSON structure, which is as follows: Please note that your answer should not contain anything other than a json format object, which should be parsable directly to json.</p> <p>Note that:</p> <ul style="list-style-type: none">- your response should be around 100 to 500 words, containing rich information given the api input parameters.- your response must be effective and have practical content.- if the api response example is null or ineffective, ignore the example and give your independent response.
System	
User	<p>API Documentation: Documentation JSON file API Examples (if available): Example input 1: Example response 1 Example input 2: Example response 2 Example input 3: Example response 3 one more API Input example:</p>

Table 15: Prompt used to write API calls.