

GECKO: A SIMULATION ENVIRONMENT TO GROUND AGENT TOOL CALLS WITH STATEFUL FEEDBACK FOR REFINEMENT

Anonymous authors

Paper under double-blind review

ABSTRACT

The ability to use tools is fundamental to large language model (LLM) agents. However, when solving complex tasks, current LLMs are prone to incorrect tool selection and invalid tool-call arguments. Although letting LLMs iteratively refine the tool-call sequence using execution results from real tools can help, repeated testing on real tools can be expensive and lead to unintended side effects. To improve LLM tool calls while addressing the issues caused by using real tools for refinement, we introduce Gecko¹, an environment that simulates tool responses using a combination of rules and LLMs. Specifically, Gecko checks the validity of tool calls including input arguments and tool names, synthesizes reasonable responses that adhere to the output schema, and assesses whether all task objectives have been achieved. Such feedback provided by Gecko allows LLMs to refine their tool calls, forming a simple yet effective test-time scaling method named GATS. In addition, we design an automated API schema converter so that Gecko can quickly integrate and simulate a large number of tools. On BFCL and τ^2 -bench, our test-time scaling method GATS enabled by Gecko consistently improves tool calling performance of existing LLMs including GPT-4o and GPT-5 (Fig. 1) and yields new state of the art. We further discuss working mechanisms of our method and share rosy future possibilities.

1 INTRODUCTION

Building agent systems using LLMs to solve complex tasks has become increasingly popular. In this mission, it is critical to let LLMs be able to use external tools, such as `get_weather` and `fetch_stock_data`. While there exist strong LLMs such as GPT-4o OpenAI et al. (2024), Qwen3 Yang et al. (2025), and xLAM-2 Prabhakar et al. (2025), because of long contexts, high task complexity, and rigid tool definitions, it is still challenging for these LLMs to select suitable tools and give accurate arguments Kate et al. (2025); Huang et al. (2024).

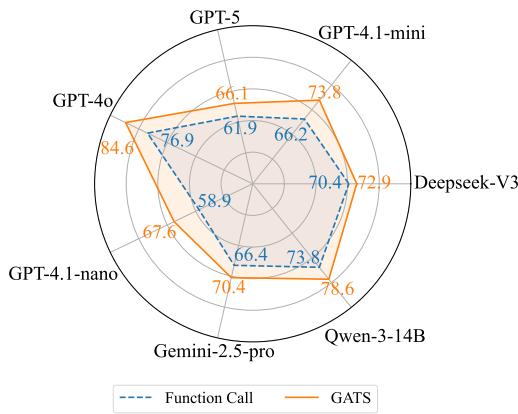


Figure 1: GATS consistently improves the performance of different LLMs on BFCL.

Li & Fung (2025). For example, inappropriate execution of `Tweet_Post` during inference may leak information irrelevant to the task, even if the post is deleted afterward.

¹Gecko comes from keywords **agent** + **feedback** + **environment**.

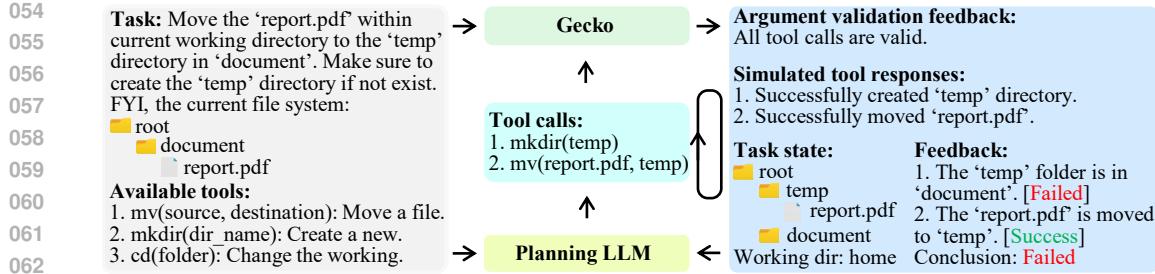


Figure 2: Overview of tool call refinement based on feedback from Gecko. The planning LLM generates tool calls based on a task and available tools. Gecko processes these calls and provides feedback on argument validation, simulated tool responses, and task state. The three types of feedback are used by the planning LLM for iterative refinement.

In this work, we aim to improve the performance of LLM² tool call and avoid using real tools during test-time refinement. To this end, we introduce Gecko, a comprehensive environment hosting a large number of simulated tools that produce semantically reasonable outputs and share the same input and output formats with real tools. As shown in Fig. 2, after receiving a task from user and tool calls from a planning LLM, Gecko will simulate the execution of the tool calls and provide three types of feedback. **First**, Gecko checks whether the input follows predefined formats, *e.g.*, ‘input 2 is invalid because only float numbers are allowed.’ This is implemented by a combination of rules and a helper LLM. **Second**, Gecko simulates tool responses. To ensure the simulated responses are semantically suitable and consistent with prior tool calls, we carefully prompt a helper LLM with the validated tool call, the tool schema, and the current task state. **Third**, Gecko uses a helper LLM to estimate the key states that reflect task progress based on the simulated tool responses. A judge LLM then assesses the inferred task state, determines whether task objectives are met, and provides task feedback on completion status and any outstanding issues.

Further, the above three types of feedback allow us to naturally design a tool-call refinement method named GATS (grounding agent test-time scaling). In this method, feedback from Gecko is sent to the planning LLM to refine the tool calls, which are then fed to Gecko to collect further feedback. This process iterates until task feedback indicates success or until it exceeds the maximum retry times. Through GATS, we observe more correct arguments and a more reasonable selection of tools.

Therefore, by grounding agent tools calls in the Gecko virtual environment, we collect useful feedback for tool-call refinement while avoiding the cost incurred by real tools.

We perform extensive evaluation on the BFCL benchmark Patil et al. (2025) and τ^2 -bench Barres et al. (2025), where Gecko has automatically synthesized 8,578 and 25 tools, respectively. We show that tool calls generated by many LLMs, such as GPT-4o and GPT-5, can be effectively hosted and executed in Gecko. As shown in Fig. 1 and Table 3, improvements are consistent across different agentic LLMs and across both single-turn and multi-turn tasks. For example, the overall performance of GPT-4o is improved from 76.93% to 84.62% on BFCL. We further discuss new tasks and possibilities that can be enabled through the proposed environment. In summary, this paper discusses the following main points.

- We introduce Gecko, a simulation environment which allows virtual tool use and gives informative feedback. Gecko successfully grounds tool calls generated by existing LLMs.
- By providing tool use feedback to the planning LLM, Gecko naturally allows for GATS, a test-time scaling that refines the tool calls during inference.
- We show our method brings consistent improvement to existing LLMs on BFCL and τ^2 -bench.
- We point out exciting insights and future directions made possible by Gecko.

²This paper use ‘planning LLM’ or ‘agentic LLM’ to describe such LLMs.

108 **2 RELATED WORK**

110 **Improving LLMs of their intrinsic tool-calling abilities.** ToolAlpaca Tang et al. (2023) fine-tunes
 111 LLMs on tool-use data generated by strong teacher models like GPT-4. ToolLLM Qin et al. (2023)
 112 collects a large number of real-world APIs and uses an automatic pipeline to construct instruction-
 113 tuning data for tool-use fine-tuning. APIGen Liu et al. (2024b) and ToolACE Liu et al. (2024a)
 114 improve the quality of synthetic tool-use data by format checking and semantic verification to im-
 115 prove fine-tuning. These methods merely focus on training data synthesis, while Gecko, due to its
 116 ability in simulating and grounding the tools, offers much higher flexibility. Gecko naturally sup-
 117 ports test-time scaling while previous methods do not. Gecko also has very good potential in training
 118 data synthesis and reinforcement learning (future work, refer to Section 5).

119 **Test-time scaling for agentic tool use.** Existing methods use feedback loops or self-reflection
 120 grounded in *real* tool execution Shi et al. (2024); Du et al. (2024); Qiao et al. (2024); Singh et al.
 121 (2025); Li et al. (2025); Chen et al. (2025b); Shi et al. (2025); Shinn et al. (2023); Zhou et al. (2025).
 122 For example, ConAgent Shi et al. (2024) iteratively refines tool calls using feedback generated by
 123 an observation LLM from real tool failure messages. TRICE Qiao et al. (2024) combines behavior
 124 cloning with reinforcement learning guided by real tool execution feedback, teaching the model to
 125 refine its tool calls during inference. These methods rely on repeatedly calling real tools, leading to
 126 tool-call costs and potential side effects. In contrast, Gecko removes the need for real tool executions
 127 in test-time scaling. Moreover, while these methods provide feedback on correcting individual failed
 128 tool calls, without maintaining a task state, they are unable to provide task-level feedback.

129 **Simulation environments for agentic tool use.** Existing methods either provide fixed, domain-
 130 specific mock tools Styles et al. (2024); Liu et al. (2023); Chen et al. (2025a) or simply wrap real
 131 APIs Qin et al. (2023), which has limited general-purpose tool simulation. For example, ToolSand-
 132 box Lu et al. (2025) and BFCL Patil et al. (2025) provide a set of stateful tools whose outputs
 133 depend on history tool executions to simulate multi-turn tasks. τ -bench Yao et al. (2024) and τ^2 -
 134 bench Barres et al. (2025) emulate conversations between a user and an agent in airline and retail
 135 scenarios. While these methods provide precise simulation on the designed use cases, they are lim-
 136 ited by human-written tools and datasets and are hard to generalize. Therefore, they could only be
 137 used for agent evaluation rather than to improve LLM performance at test time.

138 **3 THE GECKO SIMULATION ENVIRONMENT**

140 Gecko has five components: (1) an argument validator that checks the syntactic and semantic val-
 141 idity of tool calls (Section 3.1); (2) a response generator that synthesizes realistic outputs for val-
 142 idated tool calls (Section 3.2); (3) a task state estimator that keeps track of the evolving task state
 143 (Section 3.3); (4) a task feedback generator that judges task completion and identifies remaining
 144 objectives (Section 3.4); and (5) an API schema converter that transforms new tools into OpenAPI
 145 3.1 schemas for integration (Section 3.5).

146 **3.1 ARGUMENT VALIDATOR**

147 **Checking argument syntactic validity by manually defined rules.** Our rules verify the pres-
 148 ence of all the required parameters and reject unsupported ones based on the argument definitions.
 149 Moreover, our rules check the input data types, *e.g.*, integer, string, or boolean. Besides, we ensure
 150 that input parameters are within the predefined range. We also have some other rules, listed in the
 151 Appendix A.3. Violations of these rules result in *error feedback*. See examples in Fig. 3(a).

152 **Checking argument semantic correctness by a helper LLM.** ‘Semantic’ means descriptions and
 153 common-sense knowledge about arguments. For example, the helper LLM rejects ‘Seattle’ if the
 154 input requires ‘country’; the helper LLM identifies date formats (*e.g.*, yyyy/mm/dd) from the context
 155 and rejects incorrect date formats (*e.g.*, mm/dd/yy) generated by the planning LLM; it also rejects
 156 unreasonable input values implied by context, *e.g.*, a negative value in the ‘age’ field. When such
 157 semantic inconsistencies occur, *error feedback* will be generated. See examples in Fig. 3(b).

158 Table 1 presents the accuracy of argument validation on BFCL-Live-Simple. We use three metrics:
 159 true positive detection rate (correct arguments detected as correct), syntactic error detection rate,
 160 and semantic error detection rate. All the three metrics are computed based on rules. Details are

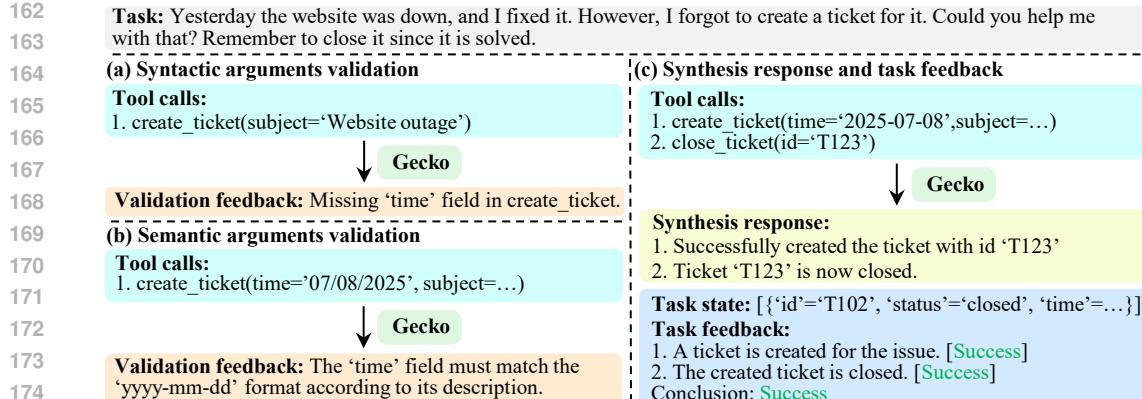


Figure 3: Examples of feedback provided by Gecko through (a) syntactic argument validation, (b) semantic argument validation, and (c) synthetic responses, task state, and task feedback. The validation feedback (orange), synthetic response (yellow), and task state and feedback (blue) will then be fed to the planning LLM during test-time scaling.

Detection rate	GPT-4o	GPT-4.1 nano	Qwen2.5 7B
True positive	100%	100%	100%
Syntactic errors	100%	100%	100%
Semantic errors	71%	69%	63%

	True positive	True negative
Real tools, LLM judge	96.7%	87.5%
Sim tools, LLM judge	91.3%	87.5%

Table 1: Accuracy of argument validation on BFCL-Live-Simple. ‘True positive’ is the rate where correct arguments are determined as correct. We also show percentage of syntactic errors and semantic errors being detected. As helper LLMs, GPT-4o, GPT-4.1-nano and Qwen-2.5-7B are evaluated.

Table 2: Accuracy of task success/failure judgment on BFCL-Multi-Turn-Base. We compare the use of real tools and simulated tools. We use an LLM as judge. Both methods have good accuracy, while using simulated tools has 5.4% lower true positive rate.

provided in the Appendix A.3. Gecko has 100% accuracy in finding correct arguments and detecting syntactic errors for all the three helper LLMs. When there are semantic errors in the tool arguments, Gecko can detect 60% ~ 70% of them. As shown in Fig. 4(a), accuracy drops by 1.5% if argument validation is removed. Results clearly demonstrate the usefulness of this component.

3.2 LLM-BASED TOOL RESPONSE GENERATOR

Beyond using validated arguments and the tool schema (tool description, input schema, and output schema) as inputs, the response generator models the functionality of the tool, producing semantically realistic and schema-compliant outputs. To support multi-turn tool response synthesis, we additionally condition the generator on the current task state (Section 3.3), a compact state that grounds prior tool calls to prevent factual conflicts with earlier tool responses and to preserve cross-call consistency. Responses are produced by a helper LLM; prompts are provided in Appendix A.4. Examples are shown in Fig. 2 and Fig. 3(c). This is done through a helper LLM. Prompts are provided in the Appendix A.4. The responses are part of the input to task state estimator and are an important source of Gecko feedback. Examples of tool responses are shown in Fig. 2 and Fig. 3(c).

It is non-trivial to directly measure the effectiveness of response generation because there is no ground truths. To do so, we designed an indirect experiment where we obtain ground truths of task success/fail using real tools and rule-based task state matching on BFCL-Multi-Turn-Base. We compare the use of real tools and simulated tools in task success/failure estimation. Results are shown in Table 2. Compared with using real tools, tool simulation results in 5.4% accuracy decrease when deciding successful tasks as successful. Interestingly, regardless of using real or simulated tools, we observe the same accuracy of 87.5% in deciding task failure as failure. This experiment indicates that simulated tools and their responses are slightly more erroneous than real tools, but the overall accuracy of task success/failure estimation is acceptable.

216 3.3 LLM-BASED TASK STATE ESTIMATOR
217

218 Task state records the progress of a task. It summarizes the cumulative effects of past tool calls and
219 serves as Gecko’s central reference for grounding tool calls. Given the previous task state and the
220 newest tool call and its response, a helper LLM updates the task state to reflect the effect of that call.
221 For example, given the previous task state ‘apples in cart: 3’ and a tool response ‘successfully delete
222 one apple from cart’, the updated state would be ‘apples in cart: 2’. The most recent task state is
223 used as input to the response generator (Section 3.2). For more examples, see Fig. 3(c) and Fig. 2.

224 Similar to task response generation, there is no direct measurement of task state estimation perfor-
225 mance. In Table 2, the performance of task success/failure estimation is generally very good. This
226 indirectly supports the effectiveness of task state estimation.
227

228 3.4 LLM-BASED TASK FEEDBACK GENERATOR
229

230 We use a judge LLM to create the feedback to be sent to the planning LLM. Two steps are involved.
231 First, we use the task description as input and let the judge LLM generate a checklist specifying what
232 aspects are important for indicating the completion of this task, such as ‘The temp folder is created
233 in document’ and ‘The report.pdf is moved to temp’ for the task in Fig. 2. Second, we let the judge
234 LLM decide whether the Gecko execution results fully satisfy the checklist: if yes, then the task
235 feedback indicates success; if not fully satisfied, the judge LLM identifies remaining objectives, and
236 another round of LLM planning and Gecko simulation will be executed. Specifically, the input to the
237 judge LLM includes the task, tool calls from the planning LLM, the simulated responses (Section
238 3.2), and the task state (Section 3.3). An example of task feedback is shown in Fig. 3(c).
239

240 In Table 2, we present the accuracy of task success/failure judgement on the BFCL-Multi-Turn-Base.
241 When using real tools, the LLM judge achieves a true positive rate of 96.7% and a true negative rate
242 of 87.5%. This indicates that our task feedback generator works well.
243

244 The components described from Section 3.1 to Section 3.4 allow us to finally ground the agent tool
245 calls. That is, after receiving the tool calls from the planning LLM, Gecko checks argument validity,
246 simulates tool responses, estimate task state, and then give task feedback.
247

248 3.5 LLM-BASED API SCHEMA CONVERTER FOR VARIOUS TOOLS
249

250 Tools that have a standard OpenAPI schema can be directly used in Gecko. For Python functions and
251 other non-OpenAPI tool definitions, Gecko uses an LLM-based schema converter to produce Open-
252 API schemas. Given a description of a tool that details its purpose, input parameters, and expected
253 output, an LLM generates an OpenAPI 3.1.0 specification in the JSON format. Our text prompt is
254 provided in the Appendix A.7. This automated conversion allows Gecko to quickly integrate and
255 simulate tools and thus support more tools than StableToolBench. The generated schemas are used
256 by the argument validator (Section 3.1) and response generator (Section 3.2).
257

258 4 GROUNDING AGENT TEST-TIME SCALING (GATS)
259

260 Given a task and tool calls generated by the planning LLM, Gecko provides three types of feedback:
261 argument validation, tool responses, and task feedback. In implementation, argument validation
262 happens before response generation and returns validation feedback to the planning LLM imme-
263 diately. If a tool call is valid, the simulated tool response is also returned immediately. From the
264 planning LLM’s perspective, Gecko mirrors real tools: each call either yields an error message from
265 validation or a tool response. After each valid call, Gecko updates the task state aligned with the
266 user-desired task. Task feedback is produced after the planning LLM finishes its output for the task.
267 The LLM judge (Section 3.4) evaluates the latest task state and the sequence of tool calls against the
268 task, determining success or identifying remaining objectives. This completes one attempt in GATS.
269 If the attempt fails, the tool call sequence and the task feedback will be sent to the planning LLM
in the next retry. Thanks to Gecko’s task state recording and session-based isolation mechanism
(Appendix A.2), retries can restart from exactly the same state snapshot, without interference from
tool calls generated in previous attempts. A diagram of this iterative scaling method is drawn in
Fig. 2. The pseudocode is provided in the appendix (Algorithm 1).

270

5 DISCUSSIONS

271
 272 **Can Gecko be implemented only by prompting?** Technically yes, if we can find a *perfect* prompt
 273 to let the LLM output all the feedback and responses. However, the perfect prompt is almost im-
 274 possible to create, because 1) our system is a combination of rule and LLM use, and 2) even if the
 275 prompt is successfully written, it will be too complex and hard for a LLM to understand.

276 **Why not use responses from real tools in Section 3.2?** While real tool responses are accurate,
 277 using them during test-time scaling has a few drawbacks. **First**, real tool execution may incur
 278 substantial cost, including computational overhead and API usage fees (e.g., RapidAPI charges
 279 per request). **Second**, performing iterative refinement directly on real tools increases the risk of
 280 unintended side effects, such as sending wrong emails.

281 **New research possibilities enabled by Gecko.** **First**, Gecko is complementary to existing tool-
 282 call data synthesis pipelines Liu et al. (2024b;a) as a verifier to improve dataset quality. A typical
 283 tool-call data point contains a task, tool definitions, and a tool-call sequence. These could be fed
 284 into Gecko, which would simulate tool responses, estimate task state, and return task feedback
 285 indicating whether the tool-call sequence solves the task. This feedback can be used to filter out
 286 or correct erroneous data points. **Second**, Gecko can turn existing SFT tool-call datasets into RL
 287 environments. Given tool definitions, Gecko can form an action space by converting each tool
 288 definition to a callable tool. After each action, Gecko returns an observation that simulates the
 289 tool execution result. Rewards are produced by an LLM judge via checklist-state comparison,
 290 allowing multiple valid action sequences and yielding fine-grained, stepwise signals for reward-
 291 function design. The resulting trajectories (states, actions, observations, rewards) can be used for
 292 offline RL, and the same interface supports online exploration in Gecko.

293 **Comparison with StableToolBench (STB) Guo et al. (2025).** While both Gecko and STB can
 294 simulate API responses, Gecko has a few key advantages. **First**, to simulate an API, STB needs to
 295 collect real responses from this API, which can be costly and less flexible. In comparison, Gecko
 296 directly supports new APIs using only API descriptions. **Second**, API-call data in STB has quality
 297 issues: it contains many erroneous responses due to invalid API calls, timeout errors, and server-
 298 side unavailability (30-40% error rates in our preliminary investigation). Gecko does not have these
 299 issues because the responses are simulated based on our converted OpenAPI schemas. **Third**, STB
 300 lacks API argument validations and generates responses for all API calls, including invalid ones,
 301 whereas real-world API servers reject such invalid calls. In contrast, Gecko has an argument val-
 302 idity checker that rejects invalid API calls and returns meaningful error messages, thus precisely
 303 simulating real-world API server behavior. **Finally**, STB focuses on individual API calls without
 304 considering multi-turn conversation history, while Gecko considers history from both task states and
 305 conversation, especially in multi-turn scenarios, which ensures logical coherence.

306 **Limitations.** Gecko currently only supports text-out tools, such as `get_tempreature` and does
 307 not yet support tools that produce non-text outputs, whose outputs are non-text media, such as
 308 `download_video`. In addition, for tools that rely on complex, dynamic external databases, such
 309 as airline reservation systems, simulation outputs (e.g., available flights or a user’s booked tickets)
 310 may diverge from the real-world state. For real-world deployments of GATS, a possible mitigation
 311 is a hybrid execution mode: simulate state-changing (write) tools within Gecko while directly host-
 312 ing read-only/query tools without simulation. This reduces simulation-reality drift for information
 313 retrieval while retaining Gecko’s sandbox benefits for test-time scaling.

314

6 EXPERIMENTS

315

6.1 EXPERIMENTAL SETUP

316 **Benchmark.** We evaluate Gecko and the grounding agent test-time scaling (GATS) method on the
 317 Berkeley Function Call Leaderboard (BFCL) and the τ^2 -bench. **BFCL** evaluates LLM tool-use
 318 ability in three categories: non-live single-turn, live single-turn, and multi-turn. Single-turn means
 319 that a task must be completed in one user-assistant round; non-live indicates that tasks and tools are
 320 designed by experts, while live indicates that tasks and tools are sourced from real-world scenarios.
 321 Multi-turn requires the model to plan and generate tool calls across several rounds based on tool-
 322 execution results and the user feedback. Within single-turn, there are four task types: simple (one
 323

324 Table 3: Method comparison on BFCLv3. We select eight most important metrics from the BFCL
 325 website. Overall accuracy is computed as the average of average ‘Non-live single turn’, average
 326 ‘Live single turn’, and ‘Multi-turn’ categories. GATS consistently improves various planning LLMs.
 327

328 Overall Acc	329 Model	330 Non-live single turn				331 Live single turn			332 Multi-turn 333 base
		334 simple	335 parallel	336 multiple	337 irrelevance	338 simple	339 multiple	340 irrelevance	
334 State-of-the-art reference models									
331 73.12	332 ToolACE-2-8B	333 88.00	334 92.50	335 92.50	336 95.41	337 70.93	338 79.01	339 84.80	340 49.00
331 79.27	332 watt-tool-70B	333 98.25	334 85.50	335 94.00	336 84.16	337 86.04	338 83.47	339 68.48	340 68.00
331 80.96	332 xLAM-2-70b	333 94.75	334 92.00	335 94.50	336 83.33	337 77.13	338 71.13	339 74.48	340 77.50
341 Baseline models and our proposed method									
341 66.20	342 GPT-4.1-mini	343 91.50	344 84.50	345 88.00	346 78.33	347 79.45	348 70.94	349 68.70	350 40.00
341 73.84	342 +GATS	343 96.25	344 88.00	345 95.50	346 84.58	347 84.49	348 74.54	349 80.83	350 50.50
341 58.85	342 GPT-4.1-nano	343 82.25	344 78.50	345 75.00	346 80.83	347 65.11	348 58.97	349 72.22	350 32.00
341 67.59	342 +GATS	343 93.25	344 88.50	345 95.00	346 81.25	347 77.13	348 69.80	349 80.38	350 37.50
341 76.93	342 GPT-4o	343 92.75	344 92.50	345 92.50	346 84.16	347 81.00	348 78.53	349 78.45	350 61.00
341 84.62	342 +GATS	343 96.50	344 95.00	345 95.50	346 95.83	347 84.10	348 81.01	349 93.42	350 72.00
341 73.78	342 Qwen-3-14B	343 95.50	344 92.50	345 95.00	346 84.58	347 86.04	348 80.81	349 77.44	350 48.00
341 78.60	342 +GATS	343 96.75	344 93.50	345 95.00	346 92.50	347 87.59	348 83.00	349 91.50	350 54.00
341 66.44	342 Gemini-2.5-pro	343 86.25	344 69.00	345 86.00	346 91.66	347 77.90	348 62.20	349 89.68	350 39.50
341 70.44	342 +GATS	343 92.25	344 75.00	345 89.00	346 92.50	347 80.62	348 67.99	349 91.83	350 44.00
341 70.40	342 Deepseek-V3	343 97.00	344 92.00	345 94.00	346 80.41	347 86.04	348 79.48	349 72.56	350 41.00
341 72.90	342 +GATS	343 97.25	344 92.00	345 95.50	346 83.75	347 88.75	348 81.76	349 78.79	350 43.50
341 61.94	342 GPT-5-thinking	343 78.00	344 84.00	345 76.00	346 92.91	347 61.62	348 57.45	349 89.70	350 33.50
341 66.08	342 +GATS	343 85.00	344 90.50	345 83.00	346 93.75	347 67.44	348 63.24	349 90.38	350 36.50

351
 352 tool call is executed to answer a user query), multiple (multiple tool calls are executed sequentially
 353 to answer a user query), parallel (multiple tool calls are executed in parallel to answer a user query),
 354 and irrelevance (none of the provided tools is appropriate, so the correct behavior is to avoid tool
 355 use). In total, BFCL contains 3,633 tasks involving 8,578 tools. τ^2 -bench is specifically designed
 356 to assess agent abilities in the real world in the retail scenario (τ^2 -retail) and airline scenario (τ^2 -
 357 airline). In τ^2 -bench, the agent must communicate with an LLM-simulated user, call domain APIs
 358 and follow domain policy rules (e.g., refund and booking rules) to complete tasks. τ^2 -retail has 13
 359 APIs and 114 tasks; τ^2 -airline has 12 APIs and 50 tasks.

360 **Evaluation metric.** For BFCL, we report *accuracy*, defined as the percentage of tasks completed
 361 correctly. For **single-turn** tasks, a prediction is counted as correct only if the tool calls produced
 362 by the model exactly match the reference solution. For **multi-turn** tasks, correctness is judged by
 363 comparing the task outcome after each turn, such as tools results and updated file contents, with the
 364 ground truth. A multi-turn task is considered successful only if the task outcomes match the ground
 365 truth at every turn. For τ^2 -bench, we use pass@1 averaged over 5 independent runs per task. Each
 366 task has an annotated goal database state, and a run is successful only if the agent responses provide
 367 all required information and the final database state matches the annotated goal.

368 **Implementation details.** We use GPT-4o-nano as helper LLM for argument validation because
 369 it is a relatively easy task. We use GPT-4o as helper/judge LLM response generation, task state
 370 estimation and task feedback generation, because these tasks are more complex. For BFCL, we
 371 execute GATS with a maximum of 3 times of retry in Gecko and directly use the resulting tool call
 372 sequence as the final answer for each task. On τ^2 -bench, because the internal database contents
 373 are not exposed to the agent and must be discovered via the native tools in τ^2 -bench (e.g., prior
 374 reservations or available flights), simulated tool calls of Gecko may differ in details. To bridge this
 375 gap, for each user message, we first run GATS to generate up to 3 rollouts, then pass all attempts
 376 (including failures) and their task feedback as in-context examples to the τ^2 -bench agent, so the
 377 agent can learn to avoid potential errors in the τ^2 -bench environment. In our runs, the user simulator
 378 is configured to GPT-4.1. For result stability, we repeated the pass@1 evaluation five times and
 379 report the mean. The LLM temperature is fixed to 0 for all requests if applicable.

378 6.2 MAIN EVALUATION
379

380 **GATS consistently improves tool call capabilities of existing LLMs.** On BFCL and τ^2 -bench, we
381 use various existing LLMs as the planning LLM, such as GPT-4.1-mini OpenAI (2025a), Deepseek
382 V3 DeepSeek-AI et al. (2025), watt-tool-70B watt-ai (2025), Qwen3-14B Yang et al. (2025), Kimi-
383 K2-Instruct Team et al. (2025), Claude Opus 4 Anthropic (2025) and GPT-5 OpenAI (2025b). We
384 apply the proposed GATS on top of these planning LLMs and demonstrate the performance gain in
385 Table 3 and Table 4. We have three observations.

386 **First**, GATS and consistently improves tool-call performance of these LLMs. For example, on
387 BFCL, the overall accuracy of GPT-4o and Qwen-3-14B is improved from 76.93% and 73.78%
388 to 84.62% and 78.60%, respectively. On τ^2 -bench, our method improves GPT-4o from 54.3% to
389 56.7%, and GPT-5-thinking from 71.0% to 72.9%. We also note that there is less improvement
390 on τ^2 -bench. The reason is that τ^2 -bench tasks provide the agent with much less contexts, such
391 as available flights, than BFCL, making it much more challenging for Gecko to simulate accurate
392 responses (see Limitations in Section 5).**Second**, our method is effective for both single-turn and
393 multi-turn tasks. For example, the improvement of GPT-4o on ‘Live single turn’ is +3.10% and
394 +2.48% for ‘simple’ and ‘multiple’, respectively, while its improvement on ‘Multi-turn’ tasks is
395 +8%. **Third** and interestingly, while some planning LLMs have different performance on single-
396 turn tasks, GATS may bring them to similar levels. For example, on ‘Multiple’ under ‘Non-live
397 single turn’, the performance of GPT-4.1-mini, GPT-4.1-nano, Deepseek-V3, and GPT-4o becomes
398 ~95% from 88%, 75%, 94%, and 92.5%, respectively. It suggests there exists some upper limit of
399 Gecko or the benchmark itself (e.g., annotation errors). We leave its explanation to future work.

400 Table 4: Method comparison on τ^2 -bench. We
401 report success rate (%) under τ -retail and τ -airline
402 subsets and their average accuracy (Overall).

403 Model	404 τ^2 -retail	405 τ^2 -airline	406 Overall
407 State-of-the-art reference models			
Claude Opus 4	81.8%	60.0%	70.9%
Claude Sonnet 4	75.0%	55.5%	65.25%
Kimi-K2-Instruct	70.6%	56.5%	63.55%
408 Baseline models and our proposed method			
GPT-4o-mini	46.1%	28.4%	37.3%
+GATS	48.4%	30.8%	39.6%
GPT-4o	63.7%	44.8%	54.3%
+GATS	65.8%	47.6%	56.7%
GPT-5-thinking	81.2%	60.8%	71.0%
+GATS	82.6%	63.2%	72.9%

410 erator. Among the four, response generation cannot be removed³, so our ablation studies are for the
411 rest three. We experiment on the BFCL-Multi-Turn-Base with GPT-4o. Results are shown in Fig. 4
412 (a). *w/o arg validation* removes both rule-based and LLM-based argument validation in Gecko.
413 *w/o task state est.* does not estimate the task states. *w/o task feedback* replaces the judge LLM
414 with a naive gating mechanism: if tool calls are generated, we give a success feedback; if no tool
415 calls are generated, then failure feedback. Results show that removing argument validation slightly
416 decreases accuracy (from 72.0% to 70.5%), while removing task state estimation has a greater im-
417 pact (68.0%). Eliminating task feedback causes the most significant drop (61.5%), indicating that
418 iterative feedback is most important when solving multi-turn tasks.

419 **Comparing different LLMs used in different components in Gecko.** To investigate the impact
420 of using different helper/judge LLMs, we replace them with a different LLM while keeping the
421 other components unchanged. Results on BFCL-Multi-Turn-Base are shown in Fig. 4(b), where the
422 performance of the default setting is 72.0%. Replacing GPT-4.1-nano in argument validation with
423 GPT-4o results in a very minor drop (71.5%). Because argument validation is relatively simple, it

424 **Comparison with the state of the art.** We apply GATS to GPT-4o and report new state of
425 the art on BFCL: **overall accuracy = 84.62%**, which is +3.66% higher than xLAM-2-70B.
426 Moreover, on various subsets, GPT-4o+GATS also reports very competitive performance, e.g.,
427 96.50% on simple single turn, 95.00% on parallel single turn and 95.50% on multiple sin-
428 gle turn. The multi-turn-base performance, 72.00%, is the second best among all the meth-
429 ods. For τ^2 -bench, GPT-5-thinking+GATS achieves an **overall accuracy of 72.9%**, which
430 indicates state-of-the-art performance.

431 6.3 FURTHER ANALYSIS

Ablation studies. Gecko contains four key
432 components: argument validation, task state es-
433 timator, response generator, and feedback gen-
434 erator.

435 **Gecko** contains four key components: argument validation, task state es-
436 timator, response generator, and feedback gen-
437 erator. Among the four, response generation cannot be removed³, so our ablation studies are for the
438 rest three. We experiment on the BFCL-Multi-Turn-Base with GPT-4o. Results are shown in Fig. 4
439 (a). *w/o arg validation* removes both rule-based and LLM-based argument validation in Gecko.
440 *w/o task state est.* does not estimate the task states. *w/o task feedback* replaces the judge LLM
441 with a naive gating mechanism: if tool calls are generated, we give a success feedback; if no tool
442 calls are generated, then failure feedback. Results show that removing argument validation slightly
443 decreases accuracy (from 72.0% to 70.5%), while removing task state estimation has a greater im-
444 pact (68.0%). Eliminating task feedback causes the most significant drop (61.5%), indicating that
445 iterative feedback is most important when solving multi-turn tasks.

446 ³Nonetheless, Table 2 presented a variant of response generator by replacing simulated tools with real tools.
447 We observe that simulated tools have a reasonably lower but acceptable true positive rate.

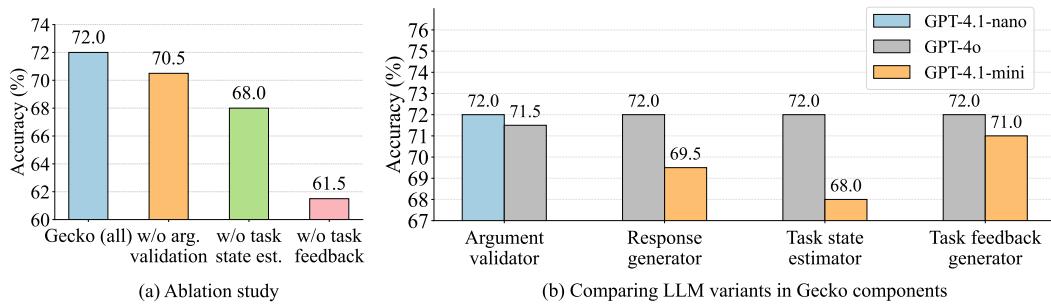


Figure 4: (a) Ablation study of Gecko on BFCL-Multi-Turn-Base. The full system (72.0%) is compared against variants with one component removed: argument validation (70.5%), task state estimation (68.0%), and task feedback (61.5%); (b) LLM replacement study on Gecko evaluated on BFCL ‘Multi-turn base’. For each component, the bar on the left is the original performance 72.0%. Under LLM replacement, e.g., replacing GPT-4.1-nano with GPT-4o for argument validation.

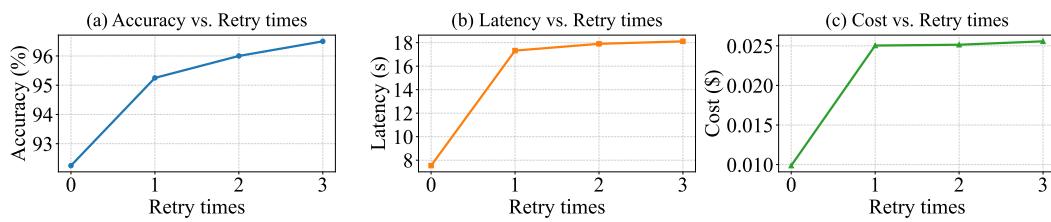


Figure 5: Test-time scaling behaviours. We evaluate GATS on the BFCL-Non-Live-Simple, using GPT-4o as the planning LLM. *Retry times* is the maximum number of feedback-based refinement steps allowed in GATS (Section 4). We report (a) accuracy (%), (b) average latency (s) and (c) average cost (\$) per user task versus maximum retry times.

does not require strong LLMs like GPT-4o. If we replace GPT-4o with GPT-4.1-mini in response generation and task state estimation, performance drops from 72.0% to 69.5% and 68.0%, respectively. For the task feedback components, replacing GPT-4o with GPT-4.1-mini leads to a small decrease in performance (-1.0%). It shows the robustness of this component to weaker LLMs.

Scaling behavior of GATS. GATS allows a maximum of times of retry, where each retry includes Gecko feedback and then tool call refinement. We examine how the retry budget affects performance on the BFCL-Non-Live-Simple, where maximum retry times vary from 0 to 3. As shown in Fig. 5, increasing the max retry times improves accuracy, from 92.25% with no refinement to 96.50% with three times of refinement. Most accuracy gain comes from the first retry, while further retries add less improvement. Accordingly, latency and cost increase with retry times: runtime increases from 7.54 s to 18.11 s, and cost from \$0.00987 to \$0.02557. Both also demonstrate decreasing margin: most user tasks are resolved in the first retry, so they will not use up the maximum retry times. These results clearly demonstrate a trade-off between accuracy gain and cost.

7 CONCLUSION

This paper introduces Gecko, a comprehensive simulation environment that takes tool calls from planning LLMs as input and outputs a variety of feedback. Early feedback is the validity of tool calls, while task-level feedback considers simulated responses and an estimate of the task state. Building on Gecko, we propose a test-time method that iteratively refines tool calls via feedback, named GATS. Our method is shown to consistently improve the performance of various LLMs on agent tool call benchmarks. In the future, Gecko can serve as foundational infrastructure for agentic tool use, enabling the community to (1) improve agents’ tool use at test time via tool simulation and stateful, task-aware feedback; (2) synthesize higher-quality tool-call data for supervised fine-tuning (SFT) by using Gecko as a verifier or within the data-synthesis loop; and (3) use Gecko to turn SFT tool-call datasets into reinforcement learning environments.

486 REFERENCES
487

488 Anthropic. Claude opus 4 / opus 4.1. Anthropic model page / announcement, 2025. URL <https://www.anthropic.com/clause/opus>.

489

490 Victor Barres, Honghua Dong, Soham Ray, Xujie Si, and Karthik Narasimhan. τ^2 -bench: Evaluating
491 conversational agents in a dual-control environment, 2025. URL <https://arxiv.org/abs/2506.07982>.

492

493

494 Chen Chen, Xinlong Hao, Weiwen Liu, Xu Huang, Xingshan Zeng, Shuai Yu, Dexun Li, Shuai
495 Wang, Weinan Gan, Yuefeng Huang, Wulong Liu, Xinzhi Wang, Defu Lian, Baoqun Yin, Yasheng
496 Wang, and Wu Liu. Acebench: Who wins the match point in tool usage?, 2025a. URL <https://arxiv.org/abs/2501.12851>.

497

498 Guoxin Chen, Zhong Zhang, Xin Cong, Fangda Guo, Yesai Wu, Yankai Lin, Wenzheng Feng, and
499 Yasheng Wang. Learning evolving tools for large language models, 2025b. URL <https://arxiv.org/abs/2410.06617>.

500

501

502 DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Cheng-
503 gang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang,
504 Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting
505 Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Haowei Zhang, Honghui
506 Ding, Huajian Xin, Huazuo Gao, Hui Li, and Others. Deepseek-v3 technical report, 2025. URL
507 <https://arxiv.org/abs/2412.19437>.

508

509 Yu Du, Fangyun Wei, and Hongyang Zhang. Anytool: Self-reflective, hierarchical agents for large-
510 scale api calls, 2024. URL <https://arxiv.org/abs/2402.04253>.

511

512 Zhicheng Guo, Sijie Cheng, Hao Wang, Shihao Liang, Yujia Qin, Peng Li, Zhiyuan Liu, Maosong
513 Sun, and Yang Liu. Stabletoolbench: Towards stable large-scale benchmarking on tool learning
514 of large language models, 2025. URL <https://arxiv.org/abs/2403.07714>.

515

516 Yue Huang, Jiawen Shi, Yuan Li, Chenrui Fan, Siyuan Wu, Qihui Zhang, Yixin Liu, Pan Zhou, Yao
517 Wan, Neil Zhenqiang Gong, and Lichao Sun. Metatool benchmark for large language models:
518 Deciding whether to use tools and which to use, 2024. URL <https://arxiv.org/abs/2310.03128>.

519

520 Minki Kang, Jongwon Jeong, and Jaewoong Cho. T1: Tool-integrated self-verification for test-time
521 compute scaling in small language models, 2025. URL <https://arxiv.org/abs/2504.04718>.

522

523 Kiran Kate, Tejaswini Pedapati, Kinjal Basu, Yara Rizk, Vijil Chenthamarakshan, Subhajit Chaud-
524 hury, Mayank Agarwal, and Ibrahim Abdelaziz. LongfuncEval: Measuring the effectiveness of
525 long context models for function calling, 2025. URL <https://arxiv.org/abs/2505.10570>.

526

527 Chengpeng Li, Mingfeng Xue, Zhenru Zhang, Jiaxi Yang, Beichen Zhang, Xiang Wang, Bowen Yu,
528 Binyuan Hui, Junyang Lin, and Dayiheng Liu. Start: Self-taught reasoner with tools, 2025. URL
529 <https://arxiv.org/abs/2503.04625>.

530

531 Guohao Li, Hasan Abed Al Kader Hammoud, Hani Itani, Dmitrii Khizbulin, and Bernard Ghanem.
532 Camel: Communicative agents for "mind" exploration of large language model society, 2023.
533 URL <https://arxiv.org/abs/2303.17760>.

534

535 Miles Q. Li and Benjamin C. M. Fung. Security concerns for large language models: A survey,
536 2025. URL <https://arxiv.org/abs/2505.18889>.

537

538 Weiwen Liu, Xu Huang, Xingshan Zeng, Xinlong Hao, Shuai Yu, Dexun Li, Shuai Wang, Weinan
539 Gan, Zhengying Liu, Yuanqing Yu, Zezhong Wang, Yuxian Wang, Wu Ning, Yutai Hou, Bin
540 Wang, Chuhan Wu, Xinzhi Wang, Yong Liu, Yasheng Wang, Duyu Tang, Dandan Tu, Lifeng
541 Shang, Xin Jiang, Ruiming Tang, Defu Lian, Qun Liu, and Enhong Chen. Toolace: Winning the
542 points of llm function calling, 2024a. URL <https://arxiv.org/abs/2409.00920>.

540 Xiao Liu, Hao Yu, Hanchen Zhang, Yifan Xu, Xuanyu Lei, Hanyu Lai, Yu Gu, Hangliang Ding,
 541 Kaiwen Men, Kejuan Yang, Shudan Zhang, Xiang Deng, Aohan Zeng, Zhengxiao Du, Chenhui
 542 Zhang, Sheng Shen, Tianjun Zhang, Yu Su, Huan Sun, Minlie Huang, Yuxiao Dong, and Jie
 543 Tang. Agentbench: Evaluating llms as agents, 2023. URL <https://arxiv.org/abs/2308.03688>.

544

545 Zuxin Liu, Thai Hoang, Jianguo Zhang, Ming Zhu, Tian Lan, Shirley Kokane, Juntao Tan, Weiran
 546 Yao, Zhiwei Liu, Yihao Feng, Rithesh Murthy, Liangwei Yang, Silvio Savarese, Juan Carlos
 547 Niebles, Huan Wang, Shelby Heinecke, and Caiming Xiong. Apigen: Automated pipeline for
 548 generating verifiable and diverse function-calling datasets, 2024b. URL <https://arxiv.org/abs/2406.18518>.

549

550 Jiarui Lu, Thomas Holleis, Yizhe Zhang, Bernhard Aumayer, Feng Nan, Felix Bai, Shuang Ma,
 551 Shen Ma, Mengyu Li, Guoli Yin, Zirui Wang, and Ruoming Pang. Toolsandbox: A state-
 552 ful, conversational, interactive evaluation benchmark for llm tool use capabilities, 2025. URL
 553 <https://arxiv.org/abs/2408.04682>.

554

555 OpenAI. Introducing gpt-4.1 in the api. OpenAI blog / API docs, 2025a. URL <https://openai.com/index/gpt-4-1/>.

556

557 OpenAI. Gpt-5 system card. Technical report, OpenAI, 2025b. URL <https://cdn.openai.com/gpt-5-system-card.pdf>.

558

559 OpenAI, Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan
 560 Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Mądry, Alex Baker-
 561 Whitcomb, Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, Alex Nichol,
 562 Alex Paino, Alex Renzin, Alex Tachard Passos, Alexander Kirillov, Alexi Christakis, Alexis Con-
 563 neau, Ali Kamali, Allan Jabri, Allison Moyer, Allison Tam, Amadou Crookes, Amin Tootoochian,
 564 Amin Tootoonchian, Ananya Kumar, Andrea Vallone, Andrej Karpathy, et al. Gpt-4o system card,
 565 2024. URL <https://arxiv.org/abs/2410.21276>.

566

567 Shishir G Patil, Huanzhi Mao, Fanjia Yan, Charlie Cheng-Jie Ji, Vishnu Suresh, Ion Stoica, and
 568 Joseph E. Gonzalez. The berkeley function calling leaderboard (BFCL): From tool use to agen-
 569 tic evaluation of large language models. In *Forty-second International Conference on Machine
 570 Learning*, 2025. URL <https://openreview.net/forum?id=2GmDdhBdDk>.

571

572 Akshara Prabhakar, Zuxin Liu, Ming Zhu, Jianguo Zhang, Tulika Awalgaonkar, Shiyu Wang, Zhi-
 573 wei Liu, Haolin Chen, Thai Hoang, Juan Carlos Niebles, Shelby Heinecke, Weiran Yao, Huan
 574 Wang, Silvio Savarese, and Caiming Xiong. Apigen-mt: Agentic pipeline for multi-turn data
 575 generation via simulated agent-human interplay, 2025. URL <https://arxiv.org/abs/2504.03601>.

576

577 Shuofei Qiao, Honghao Gui, Chengfei Lv, Qianghuai Jia, Huajun Chen, and Ningyu Zhang. Making
 578 language models better tool learners with execution feedback, 2024. URL <https://arxiv.org/abs/2305.13068>.

579

580 Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru
 581 Tang, Bill Qian, Sihan Zhao, Lauren Hong, Runchu Tian, Ruobing Xie, Jie Zhou, Mark Gerstein,
 582 Dahai Li, Zhiyuan Liu, and Maosong Sun. Toolllm: Facilitating large language models to master
 583 16000+ real-world apis, 2023. URL <https://arxiv.org/abs/2307.16789>.

584

585 Zhengliang Shi, Shen Gao, Xiuyi Chen, Yue Feng, Lingyong Yan, Haibo Shi, Dawei Yin, Pengjie
 586 Ren, Suzan Verberne, and Zhaochun Ren. Learning to use tools via cooperative and interactive
 587 agents, 2024. URL <https://arxiv.org/abs/2403.03031>.

588

589 Zhengliang Shi, Shen Gao, Lingyong Yan, Yue Feng, Xiuyi Chen, Zhumin Chen, Dawei Yin, Suzan
 590 Verberne, and Zhaochun Ren. Tool learning in the wild: Empowering language models as auto-
 591 matic tool agents. In *Proceedings of the ACM on Web Conference 2025*, pp. 2222–2237, 2025.

592

593 Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao.
 594 Reflexion: language agents with verbal reinforcement learning. In A. Oh, T. Nau-
 595 mann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neu-
 596 ral Information Processing Systems*, volume 36, pp. 8634–8652. Curran Associates, Inc.,

594 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/1b44b878bb782e6954cd888628510e90-Paper-Conference.pdf.

595

596

597 Joykirat Singh, Raghav Magazine, Yash Pandya, and Akshay Nambi. Agentic reasoning and tool

598 integration for llms via reinforcement learning, 2025. URL <https://arxiv.org/abs/2505.01441>.

599

600 Olly Styles, Sam Miller, Patricio Cerda-Mardini, Tanaya Guha, Victor Sanchez, and Bertie Vidgen.

601 Workbench: a benchmark dataset for agents in a realistic workplace setting, 2024. URL <https://arxiv.org/abs/2405.00823>.

602

603

604 Qiaoyu Tang, Ziliang Deng, Hongyu Lin, Xianpei Han, Qiao Liang, Boxi Cao, and Le Sun. Toolal-

605 pac: Generalized tool learning for language models with 3000 simulated cases, 2023. URL

606 <https://arxiv.org/abs/2306.05301>.

607

608 Kimi Team, Yifan Bai, Yiping Bao, Guanduo Chen, Jiahao Chen, Ningxin Chen, Ruijue Chen,

609 Yanru Chen, Yuankun Chen, Yutian Chen, Zhuofu Chen, Jialei Cui, Hao Ding, Mengnan Dong,

610 Angang Du, Chenzhuang Du, Dikang Du, Yulun Du, Yu Fan, Yichen Feng, Kelin Fu, Bofei Gao,

611 Hongcheng Gao, Peizhong Gao, Tong Gao, Xinran Gu, Longyu Guan, Haiqing Guo, Jianhang

612 Guo, Hao Hu, Xiaoru Hao, Tianhong He, Weiran He, Wenyang He, Chao Hong, Yangyang Hu,

613 and Others. Kimi k2: Open agentic intelligence, 2025. URL <https://arxiv.org/abs/2507.20534>.

614

615 watt-ai. watt-tool-70B model card. Hugging Face model hub, 2025. URL <https://huggingface.co/watt-ai/watt-tool-70B>; model id = "watt-ai/watt-tool-70B";

616 70.6B params; license = Apache-2.0; references arXiv:2406.14868; accessed=2025-09-25.

617

618 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang

619 Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu,

620 Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin

621 Yang, Jiaxi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang,

622 Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui

623 Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang

624 Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger

625 Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan

626 Qiu. Qwen3 technical report, 2025. URL <https://arxiv.org/abs/2505.09388>.

627

628 Shunyu Yao, Noah Shinn, Pedram Razavi, and Karthik Narasimhan. τ -bench: A benchmark for

629 tool-agent-user interaction in real-world domains, 2024. URL <https://arxiv.org/abs/2406.12045>.

630

631 Huichi Zhou, Yihang Chen, Siyuan Guo, Xue Yan, Kin Hei Lee, Zihan Wang, Ka Yiu Lee, Guchun

632 Zhang, Kun Shao, Linyi Yang, and Jun Wang. Memento: Fine-tuning llm agents without fine-

633 tuning llms, 2025. URL <https://arxiv.org/abs/2508.16153>.

634

635

636

637

638

639

640

641

642

643

644

645

646

647

648 **A APPENDIX**
649650 **A.1 THE USE OF LARGE LANGUAGE MODELS (LLMs)**
651652 OpenAI ChatGPT was used for grammar checks and phrasing suggestions. Anthropic Claude
653 (Claude Code) was used to assist with implementing and debugging portions of the code. LLMs
654 did not contribute to research ideation or experimental design.655 **A.2 DESIGN PRINCIPLES OF GECKO**
656658 Gecko is a simulated tool execution environment built on modern web-service principles and accessible
659 over the network. It exposes a RESTful HTTP interface so clients can interact with it using
660 standard JSON payloads.661 Gecko uses **FastAPI**, a high-performance asynchronous web framework, to handle API requests,
662 and **CAMEL** Li et al. (2023) as the LLM agent framework. FastAPI’s native async support and
663 lightweight routing let Gecko handle many concurrent requests with low latency.664 The architecture follows a **middleware-based design pattern** that provides a clear separation of
665 responsibilities. Each incoming request passes through a chain of middleware components that
666 handle different aspects of the processing pipeline. The session middleware manages state isolation
667 between different execution contexts, while the route middleware handles the mapping between API
668 endpoints and their corresponding OpenAPI specifications. This layered approach ensures that each
669 component focuses on a specific responsibility while maintaining loose coupling between different
670 parts of the system.671 To support concurrency and retry mechanisms, we designed a **session-based state management**
672 **system**. Each client interaction runs inside its own session with a unique session ID. The session
673 holds its own configuration, execution history, and task state history so the server can safely retry
674 actions and replay earlier steps without mixing data from different clients. This isolation prevents
675 users from interfering with each other and lets Gecko handle many users at once. Sessions persist
676 across multiple tool calls, enabling multi-turn interactions that remain tied to the same session.677 **A.3 ARGUMENT VALIDATOR IMPLEMENTATION DETAILS**
678679 **A.3.1 SYNTACTIC VALIDATION RULES**
680681 Our syntactic validation enforces comprehensive rule-based checks to ensure argument correctness
682 at the structural level.683 **Presence of required fields.** We verify that all fields marked `required` in the OpenAPI schema
684 are present in the tool call. We also check that the tool call does not include any fields that are not
685 defined in the schema. Tool calls that violate this rule will receive an instant error message.686 **Type checking.** Each provided argument is checked against the type (e.g., `integer`, `float`
687 `number`, `string`, `boolean`, `array`, and `object`) declared in the OpenAPI schema. Type
688 mismatches will result in an error message.689 **Constraint enforcement.** The validator enforces constraints such as numeric bounds
690 (`minimum`/`maximum`, `exclusiveMinimum`/`exclusiveMaximum`), enumeration (`enum`),
691 and length/pattern restrictions for strings (`minLength`/`maxLength`). When constraints are vio-
692 lated, the validator reports the specific constraint failure.694
695
696
697
698
699
700
701

702 A.3.2 SEMANTIC VALIDATION VIA LLM
703

704 The semantic validator takes two primary inputs: the tool call to be verified and the corresponding
705 OpenAPI schema. The validator returns a JSON-parsable message that indicates whether the
706 tool call is semantically acceptable and enumerates any problems found. The main prompt for the
707 semantic validator is provided below.

708

```

709 Please validate the given function call arguments against their
710 parameter schemas.
711
712 **Validation Rules**:
713 1. **Scope**
714   - Only validate arguments defined in the provided schemas.
715   - Ignore arguments not present in the schema (do not treat them as
716     errors).
717   - Type validation has already been handled elsewhere. Just skip
718     type checking.
719
720 2. **Semantic Checks**
721   - Validate according to the parameter description, examples, enums,
722     or format requirements.
723   - If examples are provided (e.g. "full-time, part-time"), treat
724     them as semantic categories. Any value in the same category (e.g.
725     "internship", "contract") is valid.
726   - If the description specifies a format (e.g. 'YYYY-MM-DD'),
727     enforce that exact pattern.
728   - Use common sense to ensure values are within a reasonable range (e.g.
729     interest rate within [0,1]; clock hour within [0,12]).
730   - Detect redundant/overlapping information across arguments (e.g. 'item="large pizza"' and 'size="large"' are overlap).
731   - If uncertain about validity, default to considering the argument
732     valid.
733
734 3. **Error Messages**
735   - Concise, precise, and human-readable.
736   - Do not include or suggest correct values.
737   - Only state which argument is invalid and why.
738
739 **Output Format**:
740   ```
741
742   valid=<true|false> error\_message="<if false, list each invalid
743   argument and reason>"
744
745   ```
746
747 **Example**:
748 - **params_schema**
749   `"[{"location": "The city that you want to go, e.g. 'Beijing, China
750     '}, {"date": "The start date for the booking, format: YYYY-MM-DD
751     "}]`'
752 - **args**
753   `>{"location": "London", "date": "01/01/2024"}`
754 - **Output**
755   `valid=false error_message="location not in required format (should
    include city and country); date not in required format (YYYY-MM-DD
    )"`'
```

756 A.4 RESPONSE GENERATOR IMPLEMENTATION DETAILS
757758 The response generator synthesizes a tool response given three inputs: a tool call, the corresponding
759 OpenAPI schema, and the current estimated task state. The main prompt for the response generator
760 is provided below.

761

762

763 You are an API simulation engine that generates JSON responses
764 strictly following OpenAPI 3.1 schemas.

765

Rules:

1. Schema first. Always match the schema exactly (structure, names, types, formats, required fields).
2. Entity-level consistency. Do not contradict any provided state or any prior successful responses in this session.
3. Open-world reads. For read/query/search operations, if requested entities/data are absent in the provided state, you MUST synthesize realistic, schema-compliant values instead of returning not-found or error responses.
4. Writes remain consistent. For create/modify/delete operations, produce a success result consistent with the schema unless it would contradict previously returned state; do not invent conflicts.
5. No extra rules. Do not invent constraints beyond the tool definition and the provided state.

778

Realism & uniqueness guidelines (domain-agnostic):

- Deterministic diversity: derive identifier-like fields using stable transforms of input arguments (e.g., incorporating parts of arguments or their hashes) so that different arguments yield different values within the session, while the same arguments yield stable values.
- Identifier-like fields (e.g., keys ending with '_id', 'Id', 'code', 'number'): prefer distinct values for distinct entities in the same response unless the schema indicates they refer to the same entity.
- Consistency: when two items share the same identifier-like value in one response, their associated attributes MUST NOT contradict each other within that response.
- Plausible formats: choose values that look realistic when the schema allows free-form text, but always prioritize matching schema types and formats.
- Temporal consistency when both present: end/arrival timestamps should be after start/arrival timestamps; choose plausible intervals without assuming domain-specific constraints.
- Diversity: vary counts and enumerations when optional, within reasonable ranges, while staying schema-compliant.

796

Illustrative synthesis examples:

- Tool name: get_user_details:
Request: {"user_id": "john_doe_001"}
Response (success object):
{
 "user_id": "john_doe_001",
 "name": "John Doe",
 "email": "john.doe@example.com",
 "phone": "+1-222-345-6789",
 "loyalty_status": "silver",
 "miles": 50000,
 "address": "7340 Oak Street, San Francisco, CA 94110"
}

808

809

810 A 5 TASK STATE ESTIMATION IMPLEMENTATION DETAILS

811

813 The task state estimation contains two phases: initialization and progressive updating. First, the
814 task state bootstrapper constructs an initial task state from the task description and the relevant Ope-
815 nAPI schemas. The task state updater then progressively revises this state as the response generator
816 synthesizes tool responses.

817 The main prompt for the task state bootstrapper is as below

618

819

801

822

823

multiple toolkits based on the given background information.

IMPORTANT: System state should contain ONLY these two types of data:

1. **Domain Data (Databases)**: The actual data that tools operate on
 - * FileSystem toolkit: files, directories structure
 - * Airline toolkit: users, flights, tickets, bookings
 - * Message toolkit: messages, inbox items
 - * These are stored at appropriate top-level or domain-specific keys
2. **Runtime Variables**: Execution context and session state
 - * Store these DIRECTLY under 'runtime_state' (flat structure)
 - * Examples: current_working_directory, current_user, is_logged_in, session_token
 - * IMPORTANT: Read toolkit descriptions carefully for initialization requirements

CRITICAL:

- * NO 'runtime_state.toolkits' structure - keep runtime_state FLAT
- * NO nested toolkit sections within runtime_state
- * NO duplicate concepts (e.g., only ONE current_directory for the whole system)
- * NO static values, validation rules, or schema metadata

Example of CORRECT runtime_state structure:

```
"runtime\_state": {  
  "current\_working\_directory": "/root"  
}
```

Rules:

1. Preserve all existing structures in the background information
2. Add runtime variables DIRECTLY under 'runtime_state' (flat structure)
3. Add domain data at appropriate keys (not in runtime_state)
4. NEVER create 'runtime_state.toolkits' or any similar nesting
5. Avoid duplicating the same concept
6. Output valid JSON only

Background information:

```
{background_information}
```

Toolkits summary:

```
{json.dumps(toolkits\_summary, indent=2)}
```

Return the UPDATED config JSON with necessary domain data and runtime state.

860

861

862

863

The main prompt for the task state updater is as below.

864
865 You are an expert at tracking the execution state of a task.
866 Update the system state based on the tool calls and their effects on
867 the system.
868
869 **IMPORTANT GUIDELINES**
870 1. ****State Tracking Principles****
871 - Update the system state to reflect ALL persistent state changes
872 caused by tool calls
873 - Operations that create, modify, or delete resources MUST update
874 the corresponding structures
875 - {"In synthesis mode Store ALL synthesized data from read
876 operations as ground truth state" if synthesis_mode else "
877 Operations that just query or read data should NOT add their
878 results to the system state"}
879 2. ****System State Organization****
880 - When tool operations modify existing structures, update them
881 directly (e.g., adding a new directory should add it to the
882 directory tree)
883 - For execution context that doesn't fit existing domain structures
884 , use the root-level "runtime_state"
885 - The "runtime_state" section is ONLY for execution context and
886 ephemeral telemetry (e.g., current location/cursor, active
887 selections, session info, temporary counters)
888 - DO NOT store canonical domain data in "runtime_state" (e.g.,
889 files, inbox messages, database rows must live under their domain
890 keys)
891 - If a counter already lives under "runtime_state" (e.g.,
892 runtime_state.toolkits.messageapi.message_count), update it ONLY
893 when the tool call semantics deterministically imply the change;
894 never infer from read-only calls
895 - Never duplicate the same fact both in a domain section and in "
896 runtime_state"
897 3. ****Value Formatting****
898 - When recording locations, positions, or identifiers, use complete
899 , unambiguous values
900 - Avoid partial or relative references that could be misinterpreted
901 - Preserve the format conventions used in the original system state
902 4. ****What Changes to Track****
903 - Resource creation/deletion/modification (files, directories,
904 database records, etc.)
905 - State transitions (status changes, position changes, mode
906 switches)
907 - Context updates (current location, active selections, session
908 data)
909 - DO NOT track query results, search results, temporary
910 computations, or read-only operation outputs
911 5. ****Example Structure with runtime_state****
912 {{
913 "DomainSystem" {{
914 // Domain resources with any modifications from state-changing
915 operations
916 }},
917 "runtime_state" {{
918 // Execution context only (no canonical domain data)
919 "current_context" "...",
920 "current_user" "USR001"
921 }
922 }}
923
924 Output the updated system state in JSON format only.
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
999

918 A.6 TASK FEEDBACK GENERATION IMPLEMENTATION DETAILS
919
920

921 Task feedback generation has two steps. First, given the task description, we generate a detailed
922 checklist that decomposes and clarifies the task intent into clear and verifiable objectives. Second,
923 the checklist is verified item by item by an LLM-based judge, which receives the checklist, current
924 task state, executed tool calls, the response from the planning LLM, and the corresponding OpenAPI
925 schemas. The judge aggregates any checklist objectives that fail verification into a compact task-
926 level feedback message, which is returned to the planning agent to guide subsequent refinements.

927 The main prompt for checklist generation is given below.

```

929
930
931 You produce ATOMIC, STATE-OPERATION-BASED verification checklists for
932 tasks in ANY domain.
933 PREVIOUS TASKS (assumed done; resolve references only, do NOT re-
934 verify): {prev_text}
935 CURRENT TASK: {current_task}

936 CRITICAL MULTI-TURN CONTEXT RULE:
937 When task mentions "values obtained", "results from previous", or
938 specific counts like "three values", these refer to OUTPUT from
939 the LAST task in PREVIOUS TASKS, not data from earlier tasks

940 RULES
941 1) Verify ONLY the current task. Return the MINIMAL set; if one item
942 suffices, return EXACTLY one.
943 2) Each item = one pass/fail assertion about final state (implied by
944 the question, do not guess the answer by yourself) or an executed
945 operation (no procedures).
946 3) IDENTIFY ALL SEMANTIC UNITS: Each complete thought, question (
947 direct or indirect like "I wonder"), or action in the task needs
948 verification
949 4) PRESERVE LOGICAL FLOW: When actions depend on prior information or
950 results, verify each step
951 5) Use explicit identifiers/paths/IDs when inferable; avoid vague
952 pronouns.
953 6) RESOLVE AMBIGUOUS REFERENCES: When the current task contains
954 pronouns like "the file", "it", "that item", etc., resolve them to
955 specific entities based on PREVIOUS TASKS context.
956 7) Do NOT add optional behaviors (saving/exporting/logging/formatting)
957 unless explicitly required.
958 8) Search/lookup/filter. Assert the search was executed with the
959 specified term/criteria; do NOT require matches unless asked.
960 9) Transform/update. Assert the stated post-condition holds; do NOT
961 invent extra artifacts.
962 10) Copy operations. Verify: source file remains intact (copy
963 preserves original); destination file exists with the new name.
964 11) Create/Delete to assert existence/absence as specified.
965 12) Discrete relocation between containers (domain-agnostic). If
966 applicable and implied: destination container exists (if mentioned
967 ); entity absent at source (for move, not copy); entity present at
968 destination.

969 OUTPUT
970 Return ONLY a pure JSON array of objects with a single key "
971 description". No extra text.

```

971 The main prompt for the LLM judge is as follows.

```

972
973 You are an expert in verifying a checklist based on the execution
974 results.
975 You have access to:
976 1. Current system state: The system state after execution
977 2. Tool calls: The list of functions that were called WITH their
978 results
979 3. Agent response: The agent's output/explanation (if available)
980 4. Conversation history: Previous turns showing the context of how
981 data was obtained{history_text}
982
983 IMPORTANT: Some operations (like grep, sort, find, ls) are query
984 operations that don't modify the system state.
985 For these operations, verify their execution by checking if the
986 corresponding tool was called in the tool_calls list OR Looking
987 for evidence in the agent_response (if provided) that the
988 operation was performed and results were obtained
989
990 Guidelines:
991 1. Verify items in the checklist one by one
992 2. For state-modifying operations (mkdir, create file, cd), check the
993 system state for changes
994 3. For query operations (grep, sort, ls, find), check tool_calls and
995 agent_response
996 4. For efficiency-related checklist items, analyze whether multiple
997 tool calls could be merged based on the tool definitions provided
998 5. Status should be one of ["success", "failed", "unknown"]
999 "success": Task completed (evidence in tool_calls/system state/
1000 agent_response)
1001 "failed": Task NOT completed AND agent provided NO explanation
1002 "unknown": Task NOT completed BUT agent explained why (e.g., "missing
1003 information", "need user confirmation", "tool unavailable")
1004 6. Just modify the status and reasoning fields of the checklist items,
1005 do not include any other text outside the json.{tool_defs_text}
1006 7. If a tool call is made as the task required, do not mark it as
1007 failed even if the result is not as you expected
1008 8. Use conversation history to understand data references.
1009 9. When evaluating relevance, consider the full multi-turn context to
1010 understand where numbers/data come from
1011
1012 Evaluation principles (keep these high-level and tool-agnostic):
1013 Only mark "failed" when there is clear evidence that the requirement
1014 was not met and no explanation was provided by the agent; if
1015 evidence is incomplete or you are not sure, use "unknown".
1016 Accept equivalent pipelines that produce the required final outcome,
1017 regardless of operation order or scope.
1018 Do not fail solely because an intermediate step operated on a broader
1019 scope; fail only if the final required subset/result is missing or
1020 incorrect.
1021
1022 The output should be in json format:
1023 [
1024     {"name": "...", "description": "...", "reasoning": "...", "status": "success" or "failed" or "unknown"}}
1025     {"name": "...", "description": "...", "reasoning": "...", "status": "success" or "failed" or "unknown"}}
1026     ...
1027 ]
1028 Do not include any other text outside the json.
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187
1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2100
2101
2102
2103
2104
2105
2106
2107
2108
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159
2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2179
2180
2181
2182
2183
2184
2185
2186
2187
2188
2189
2190
2191
2192
2193
2194
2195
2196
2197
2198
2199
2200
2201
2202
2203
2204
2205
2206
2207
2208
2209
2210
2211
2212
2213
2214
2215
2216
2217
2218
2219
2220
2221
2222
2223
2224
2225
2226
2227
2228
2229
2230
2231
2232
2233
2234
2235
2236
2237
2238
2239
2240
2241
2242
2243
2244
2245
2246
2247
2248
2249
2250
2251
2252
2253
2254
2255
2256
2257
2258
2259
2260
2261
2262
2263
2264
2265
2266
2267
2268
2269
2270
2271
2272
2273
2274
2275
2276
2277
2278
2279
2280
2281
2282
2283
2284
2285
2286
2287
2288
2289
2290
2291
2292
2293
2294
2295
2296
2297
2298
2299
2300
2301
2302
2303
2304
2305
2306
2307
2308
2309
2310
2311
2312
2313
2314
2315
2316
2317
2318
2319
2320
2321
2322
2323
2324
2325
2326
2327
2328
2329
2330
2331
2332
2333
2334
2335
2336
2337
2338
2339
2340
2341
2342
2343
2344
2345
2346
2347
2348
2349
2350
2351
2352
2353
2354
2355
2356
2357
2358
2359
2360
2361
2362
2363
2364
2365
2366
2367
2368
2369
2370
2371
2372
2373
2374
2375
2376
2377
2378
2379
2380
2381
2382
2383
2384
2385
2386
2387
2388
2389
2390
2391
2392
2393
2394
2395
2396
2397
2398
2399
2400
2401
2402
2403
2404
2405
2406
2407
2408
2409
2410
2411
2412
2413
2414
2415
2416
2417
2418
2419
2420
2421
2422
2423
2424
2425
2426
2427
2428
2429
2430
2431
2432
2433
2434
2435
2436
2437
2438
2439
2440
2441
2442
2443
2444
2445
2446
2447
2448
2449
2450
2451
2452
2453
2454
2455
2456
2457
2458
2459
2460
2461
2462
2463
2464
2465
2466
2467
2468
2469
2470
2471
2472
2473
2474
2475
2476
2477
2478
2479
2480
2481
2482
2483
2484
2485
2486
2487
2488
2489
2490
2491
2492
2493
2494
2495
2496
2497
2498
2499
2500
2501
2502
2503
2504
2505
2506
2507
2508
2509
2510
2511
2512
2513
2514
2515
2516
2517
2518
2519
2520
2521
2522
2523
2524
2525
2526
2527
2528
2529
2530
2531
2532
2533
2534
2535
2536
2537
2538
2539
2540
2541
2542
2543
2544
2545
2546
2547
2548
2549
2550
2551
2552
2553
2554
2555
2556
2557
2558
2559
2560
2561
2562
2563
2564
2565
2566
2567
2568
2569
2570
2571
2572
2573
2574
2575
2576
2577
2578
2579
2580
2581
2582
2583
2584
2585
2586
2587
2588
2589
2590
2591
2592
2593
2594
2595
2596
2597
2598
2599
2600
2601
2602
2603
2604
2605
2606
2607
2608
2609
2610
2611
2612
2613
2614
2615
2616
2617
2618
2619
2620
2621
2622
2623
2624
2625
2626
2627
2628
2629
2630
2631
2632
2633
2634
2635
2636
2637
2638
2639
2640
2641
2642
2643
2644
2645
2646
2647
2648
2649
2650
2651
2652
2653
2654
2655
2656
2657
2658
2659
2660
2661
2662
2663
2664
2665
2666
2667
2668
2669
2670
2671
2672
2673
2674
2675
2676
2677
2678
2679
2680
2681
2682
2683
2684
2685
2686
2687
2688
2689
2690
2691
2692
2693
2694
2695
2696
2697
2698
2699
2700
2701
2702
2703
2704
2705
2706
2707
2708
2709
2710
2711
2712
2713
2714
2715
2716
2717
2718
2719
2720
2721
2722
2723
2724
2725
2726
2727
2728
2729
2730
2731
2732
2733
2734
2735
2736
2737
2738
2739
2740
2741
2742
2743
2744
2745
2746
2747
2748
2749
2750
2751
2752
2753
2754
2755
2756
2757
2758
2759
2760
2761
2762
2763
2764
2765
2766
2767
2768
2769
2770
2771
2772
2773
2774
2775
2776
2777
2778
2779
2780
2781
2782
2783
2784
2785
2786
2787
2788
2789
2790
2791
2792
2793
2794
2795
2796
2797
2798
2799
2800
2801
2802
2803
2804
2805
2806
2807
2808
2809
2810
2811
2812
2813
2814
2815
2816
2817
2818
2819
2820
2821
2822
2823
2824
2825
2826
2827
2828
2829
2830
2831
2832
2833
2834
2835
2836
2837
2838
2839
2840
2841
2842
2843
2844
2845
2846
2847
2848
2849
2850
2851
2852
2853
2854
2855
2856
2857
2858
2859
2860
2861
2862
2863
2864
2865
2866
2867
2868
2869
2870
2871
2872
2873
2874
2875
2876
2877
2878
2879
2880
2881
2882
2883
2884
2885
2886
2887
2888
2889
2890
2891
2892
2893
2894
2895
2896
2897
2898
2899
2900
2901
2902
2903
2904
2905
2906
2907
2908
2909
2910
2911
2912
2913
2914
2915
2916
2917
2918
2919
2920
2921
2922
2923
2924
2925
2926
2927
2928
2929
2930
2931
2932
2933
2934
2935
2936
2937
2938
2939
2940
2941
2942
2943
2944
2945
2946
2947
2948
2949
2950
2951
2952
2953
2954
2955
2956
2957
2958
2959
2960
2961
2962
2963
2964
2965
2966
2967
2968
2969
2970
2971
2972
2973
2974
2975
2976
2977
2978
2979
2980
2981
2982
2983
2984
2985
2986
2987
2988
2989
2990
2991
2992
2993
2994
2995
2996
2997
2998
2999
3000
3001
3002
3003
3004
3005
3006
3007
3008
3009
3010
3011
3012
3013
3014
3015
3016
3017
3018
3019
3020
3021
3022
3023
3024
3025
3026
3027
3028
3029
3030
3031
3032
3033
3034
3035
3036
3037
3038
3039
3040
3041
3042
3043
3044
3045
3046
3047
3048
3049
3050
3051
3052
3053
3054
3055
3056
3057
3058
3059
3060
3061
3062
3063
3064
3065
3066
3067
3068
3069
3070
3071
3072
3073
3074
3075
3076
3077
3078
3079
3080
3081
3082
3083
3084
3085
3086
3087
3088
3089
3090
3091
3092
3093
3094
3095
3096
3097
3098
3099
3100
3101
3102
3103
3104
3105
3106
3107
3108
3109
3110
3111
3112
3113
3114
3115
3116
3117
3118
3119
3120
3121
3122
3123
3124
3125
3126
3127
3128
3129
3130
3131
3132
```

1026 A.7 LLM-BASED API SCHEMA CONVERTER IMPLEMENTATION DETAILS
10271028 The main prompt for the API schema converter is given below.
1029

```

1030     Convert this tool description to an OpenAPI 3.1 endpoint specification
1031
1032     Tool description:
1033     {tool_description}
1034
1035     Create an endpoint object with these exact fields:
1036     1. operationId: {tool['name']}
1037     2. summary: one-line description (keep it short)
1038     3. description: brief description of the operation (1 sentence max)
1039     4. requestBody: proper schema based on parameters
1040     5. responses: ONLY USE STATUS 200 with oneOf schema for success/error
1041         - Success response: Based on tool's PURPOSE (not generic "result")
1042         - Error response: Standard error object
1043     6. Analyze the tool's PURPOSE to generate an appropriate response
1044         schema. For example:
1045         If it calculates something (area, factorial, etc.), return the
1046             calculated value
1047         If it fetches data (user info, list of items), return the data
1048             structure
1049         If it performs an action (create, update, delete), return success
1050             confirmation with relevant details
1051         If it searches/filters, return matching results
1052
1053     CRITICAL REQUIREMENTS:
1054     - NEVER use $ref - always inline all schemas
1055     - ONLY use HTTP status 200 for ALL responses
1056     - Use oneOf schema in the 200 response to handle both success and
1057         error cases
1058     - All properties MUST have a "description" field
1059     - The requestBody must have a schema with type "object"
1060     - If no parameters, still include requestBody with empty properties
1061
1062     Return ONLY the endpoint object JSON.

```

1060 A.8 EXPERIMENTS ON THE ACCURACY OF ARGUMENT VALIDATION
10611062 We evaluated the argument validator on the BFCL-Live-Simple subset. GPT-4o was used to generate
1063 tool calls for each example. For tool calls that passed the official BFCL evaluation, we marked them
1064 as correct. For tool calls that failed the BFCL evaluation, human annotators labeled each failure
1065 as either a *syntactic error* (violates the tool schema, e.g., missing required field or wrong field
1066 name/type) or a *semantic error* (schema-conformant but semantically inappropriate, e.g., wrong
1067 granularity or implausible value). These labels formed the ground truth.1068 We then ran our argument validator on the same generated tool calls and compared its predictions
1069 (correct / syntactic error/ semantic error) to the ground truth. Table 1 reports the detection rates.
10701071 A.9 PSEUDOCODE OF GATS
10721073
1074
1075
1076
1077
1078
1079

1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093

Algorithm 1: GATS loop for an individual dialog turn (inline detailed comments).

Input: Task T , tool definitions D , schema converter C , planner LLM π , argument validator V , response generator G , task state estimator E , judge LLM J , initial task state st_{init} , maximum retries R_{max}

Output: Tool-call sequence $trace$

```

 $S \leftarrow C.\text{convert}(D)$  // convert tool definitions  $D$  into schemas  $S$ 
 $checklist \leftarrow J.\text{gen\_checklist}(T)$  // decompose task  $T$  to checklist
 $last\_error \leftarrow \text{null}$  // holds last validation or judge error
for attempt  $\leftarrow 0$  to  $R_{\text{max}}$  do
     $trace \leftarrow []$  // trace records tool call sequence
     $st \leftarrow st_{\text{init}}$  // initialize task state from previous task state
    while true do
         $c \leftarrow \pi.\text{gen\_next}(T, st, trace, last\_error)$  // planner generates tool call
        if  $c == \text{STOP}$  then
             $\quad \text{break}$  // planner finished tool use for this attempt
         $(valid, errors) \leftarrow V.\text{validate}(c, S)$  // arguments validation
        if not valid then
             $\quad last\_error \leftarrow errors$  // return validation feedback
             $\quad \text{continue}$ 
         $r_c \leftarrow G.\text{generate}(c, S, st)$  // simulate tool execution
         $st \leftarrow E.\text{update}(st, c, r_c)$  // update task state
         $\quad \text{append } c \text{ to } trace$  // record executed call
     $feedback \leftarrow J.\text{judge}(checklist, st, trace, S)$  // give task-level feedback
    if  $feedback.\text{success} == \text{true}$  then
         $\quad \text{break}$  // task satisfied
     $last\_error \leftarrow feedback.\text{error}$  // return task-level feedback
return  $trace$ 

```

1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133