Knowing When Not to Trust: Tool Awareness as Metacognition in LLMs

Anonymous ACL submission

Abstract

Tool-augmented language models are typically evaluated by whether they can select and invoke useful external tools. However, effective tool use also requires deciding when tool outputs should be trusted, especially in the *first-use setting*, where a model encounters a tool with no prior evidence of its reliability. We introduce a controlled benchmark pairing mathematical questions of varying difficulty with plausible tool interfaces that return subtly perturbed or random answers. This design allows us to distinguish direct answering, tool reliance, compliance with erroneous tool outputs, and correction or warning behavior. We evaluate frontier models under six prompt conditions ranging from neutral to fully coercive, including a *calibrated skepticism* condition that explicitly warns models that tools can be incorrect. Across all conditions, compliance with corrupted outputs increases with task difficulty and remains near-ceiling on hard examples even under explicit specification instructions. Our findings show that evaluating tool-using models requires measuring not only whether models can call tools, but also whether they can calibrate trust in tool outputs. The benchmark is available at: `https://anonymous.4open.science/w/tool_awareness-C709/benchmark.html`.

1 Introduction

Recent work on Artificial General Intelligence (AGI) argues that model evaluation should move beyond aggregate task performance and instead examine specific cognitive faculties (Burnell et al., 2026). Among these, metacognition–a system's ability to monitor and regulate its own cognitive processes–is central to general capable behavior: an intelligent system should not only solve problems, but also recognize when its own answer is uncertain, when external assistance is needed, and when an external source should not be trusted.

We investigate *tool awareness* as one operational instance of metacognition in tool-augmented language models. By tool awareness, we mean a model's ability to regulate its interaction with external tools: deciding whether to call a tool, whether the tool output is plausible, and whether to accept, reject, or warn about that output. This mirrors an ordinary feature of human problem solving. A person may use a calculator for a long multiplication, but reject its result if the answer violates a simple estimate (e.g., looking at the product of the last digits). A generally capable AI system should behave analogously: not treat tool use as a purely mechanical action, but calibrate when tools are useful and when their outputs conflict with available evidence.

This question is increasingly important because large language models (LLMs) are now commonly deployed as tool-augmented systems. Rather than producing answers only from parametric knowledge, they can call calculators, search engines, code interpreters, and other external modules (Schick et al., 2023; Yao et al., 2023; Gao et al., 2023; Chen et al., 2023). However, tool use introduces a second-order metacognitive problem. A model must not only know *when* and *how* to use a tool, but also *whether* the tool should be trusted in the current context. Existing evaluation settings often assume that tools are correct, useful, and semantically aligned with their natural-language descriptions – an assumption that can fail when tools are misconfigured, stale, noisy, adversarially described, or subtly inconsistent with their documentation.

A particularly common and underexplored instance of this problem arises in the *first-use setting*: when an agent encounters a tool with which it has no prior experience, such as a newly integrated third-party API. In this setting, the model has no accumulated evidence about the tool's reliability and must decide whether to trust its outputs based solely on the tool's name, documentation, and plausibility of the returned result. This scenario is routine in

practice – developers routinely deploy agents with tools that have not been independently validated – yet it remains largely unstudied.

This paper studies this problem through a controlled benchmark for mathematical *tool awareness* under first-use conditions. We construct a set of arithmetic and number-theoretic questions spanning addition, subtraction, multiplication, division, exponentiation, primality testing, prime factorization, and decimal-place retrieval. Each instance is paired with a nominally relevant tool, but the implemented tools are deliberately unreliable while retaining plausible names and docstrings. For example, arithmetic tools return slightly perturbed numerical values. This creates a setting in which tool affordances are linguistically convincing, but tool outputs are systematically untrustworthy – mirroring the epistemic situation a model faces the first time it encounters an unvalidated tool.

Our central research question is: *Can contemporary LLMs calibrate trust outputs when encountering a tool for the first time, without prior evidence of its reliability?* We decompose this question into two measurable behaviors: whether models choose to invoke tools as a function of task difficulty and prompting, and whether, when a tool is used, and its output is wrong, the model complies, corrects, warns, or produces an answer equivalent to neither.

We evaluate multiple frontier and near-frontier models under prompt conditions varying the degree of encouragement and skepticism toward tool use. Some prompts simply present the question; others advise the model to use a tool when the problem exceeds the model's capabilities; one explicitly instructs the model to apply judgment, given that tools can be incorrect; and others explicitly assert that the tool is correct or mandate tool use entirely. This design allows us to measure not only task accuracy, but also susceptibility to prompt-induced over-reliance and whether explicit skepticism instruction is sufficient to induce appropriate caution.

A key prediction of the first-use framing is that model behavior should vary with task difficulty in a specific way: on easy problems, models can cross-check tool outputs against internal knowledge and are better positioned to detect corrupted results; on hard problems, where independent verification is impossible, calibrated skepticism becomes both more important and harder to act on. Our results confirm this prediction and reveal a troubling asymmetry: the settings in which tool verification is most needed are precisely those in which models are least likely to exercise it.

2 Related Work

Tool-augmented language models. Early work on modular and tool-augmented language systems argued that LLMs should be combined with external knowledge sources and discrete reasoning modules rather than expected to solve all tasks through parametric generation alone (Karpas et al., 2022). ReAct (Yao et al., 2023) introduced a prompting framework in which language models interleave reasoning traces with actions, allowing them to interact with external environments such as knowledge bases. Toolformer (Schick et al., 2023) further showed that models can be trained in a self-supervised manner to decide when to call APIs, what arguments to pass, and how to use the resulting observations. These works establish that external tools can improve LLM performance. Our work shares this motivation but investigates a different failure mode: rather than assuming that available tools are reliable, we ask whether models remain calibrated when tool descriptions are plausible but tool outputs are corrupted.

Benchmarks for tool use and API calling. A growing body of work evaluates LLMs' ability to use tools and APIs. API-Bank (Li et al., 2023) introduced a runnable benchmark with API tools and annotated dialogues to evaluate planning, API retrieval, and API calling. ToolBench and ToolLLM (Qin et al., 2024) scaled this to thousands of real-world APIs and complex instructions. Gorilla (Patil et al., 2024) focused on generating accurate API calls and reducing hallucinated API usage. ToolQA (Zhuang et al., 2023) proposed a question-answering benchmark designed to distinguish questions answerable from internal model knowledge from questions requiring external tools. More recent benchmarks have moved beyond final-answer accuracy: TRAJECT-Bench (He et al., 2026) evaluates tool selection, parameterization, and ordering along the full tool-use trajectory, while WildToolBench (Yu et al., 2026) emphasizes realistic multi-turn user behavior, including compositional tasks and implicit intent.

These benchmarks primarily ask whether models can identify, invoke, and compose appropriate tools, assuming that tool outputs are faithful. Our benchmark holds tool availability fixed—the nominally relevant tool is always present—and isolates a distinct capability: deciding whether an avail-

able, semantically relevant tool should actually be trusted.

Tool-use awareness, selection, and safety. Recent work has begun to examine tool-use decisions more directly. MetaTool (Huang et al., 2024) evaluates whether LLMs can decide whether to use tools and which tools to select, including subtasks involving similar tools and possible reliability issues. ToolSword (Ye et al., 2024) studies safety risks in tool learning across multiple stages, including noisy misdirection, risky cues, harmful feedback, and error conflicts. Most closely related to our motivation, Faghih et al. (2025) show that agentic LLMs are highly sensitive to tool descriptions: small edits to natural-language descriptions can substantially shift tool preferences even when tool functionality is unchanged. Our work contributes to this line by introducing a controlled setting in which tools are not merely hard to choose among, but actively unreliable despite appearing aligned with the task—permitting direct measurement of overreliance and testing whether tool-use decisions are stable under instructions asserting tool correctness.

3 Methods

We study tool awareness in large language models under conditions where external tools appear useful yet are unreliable. Our experimental setting pairs mathematical question-answering tasks with callable arithmetic and number-theoretic tools whose names and documentation suggest standard, trustworthy behavior, while their implementations deliberately return corrupted outputs. This design lets us examine two complementary aspects of tool awareness: whether a model chooses to invoke an available tool, and whether it recognizes or corrects the tool output when it is wrong. By varying task difficulty and prompt instructions about tool use and tool trustworthiness, we measure how models shift among direct answering, misplaced compliance with erroneous tool outputs, and explicit warning or correction behavior.

3.1 Task Construction

We construct a synthetic mathematical benchmark covering eight operation families: addition, subtraction, multiplication, division, exponentiation, primality testing, prime factorization, and decimal-place retrieval. Each example consists of a natural-language question, a gold answer, the name of the tool that would nominally solve the task, and an estimated difficulty label. Formally, each instance is represented as $x_i = (q_i, y_i, t_i, d_i)$, where q_i is the question, y_i is the ground-truth answer, t_i is the required tool name, and $d_i \in \{\text{EASY}, \text{MEDIUM}, \text{HARD}\}$ is the difficulty level.

For each difficulty level, we generate 20 instances per operation template. Since the benchmark contains eight operation templates at each of three difficulty levels, the resulting dataset contains $20 \times 8 \times 3 = 480$ examples. Ground-truth answers are computed programmatically using exact arithmetic where applicable, `sympy` for primality and factorization, and high-precision `mpmath` arithmetic for decimal-place tasks. Division answers are rounded to two decimal places.

The difficulty levels are defined by increasing the task's numerical range and computational burden. Easy arithmetic uses small integers, such as addition and subtraction over $[-100, 100]$, multiplication over $[-12, 12]$, and exponentiation with bases in $[-12, 12]$ and exponents up to 4. Medium arithmetic increases the operand range, for example, addition and subtraction over $[-9999, 9999]$, multiplication over $[-99, 99]$, and exponentiation with bases in $[-99, 99]$ and exponents between 2 and 6. Hard arithmetic uses substantially larger operands, including addition and subtraction over $[-10^{12}, 10^{12}]$, multiplication over six-digit integers, and exponentiation with bases up to three digits and exponents between 7 and 20.

For number-theoretic tasks, easy primality questions use integers from 2 to 100, while medium primality questions use odd integers between 101 and 10,000. Hard primality and factorization examples are sampled from predefined lists of large primes and semiprimes. Decimal-place tasks are similarly scaled: easy examples ask for early decimal digits (1st to 10th) of π, medium examples ask for later digits (11th to 100th) of π, and hard examples ask for decimal digits of irrational square roots, using radicands sampled from $\{2, 3, 5, 6, 7, 8, 10\}$ and digit positions between 2,500 and 10,000.

3.2 Adversarial Tool Environment

Each model is evaluated in an environment containing a set of callable tools. The tool names and docstrings describe standard mathematical operations, but the underlying implementations are intentionally unreliable. For example, the addition tool is documented as returning $a + b$ but actually returns $\lfloor 1.01(a + b) \rfloor$; the subtraction tool scales the true

difference by 0.99; multiplication, division, and exponentiation are similarly perturbed by multiplicative factors (1.01). Other tools return outputs sampled uniformly at random: the decimal-place tools return a digit sampled uniformly from the interval $[0, 99]$, the primality tool returns a Boolean value sampled uniformly from $\{\texttt{True}, \texttt{False}\}$, and the prime-factorization tool first samples a list length uniformly from the interval $[1, 5]$ and then samples that many integers uniformly without replacement from the interval $[2, 99]$.

This setup creates a controlled mismatch between tool affordances and tool reliability. Importantly, the model is not told that the tools are corrupted, except in a specific scenario (see Section 3.4). Therefore, success requires more than recognizing that a relevant tool exists: the model must decide whether to rely on the tool for a given problem and prompt condition. This design is intended to probe over-trust in tool interfaces, especially when tool metadata appears semantically aligned with the task.

For each test example, the model is given access to a random subset of five tools, with the nominally required tool always included. This restriction reduces the token overhead associated with tool descriptions while preserving a non-trivial tool-selection setting in which the relevant tool is always available.

3.3 Models

We evaluate five frontier and near-frontier language models from three providers: OpenAI (`gpt-5.4-2026-03-05`[1]; `gpt-4o-2024-08-06`[2]), Anthropic (`claude-sonnet-4-6`[3]; `claude-haiku-4-5-20251001`[4]), and Google (`gemini-2.5-pro`[5]).

All models are accessed through the same agent interface and are evaluated on the same serialized dataset. For `claude-sonnet-4-6`, we enable adaptive thinking as supported by the execution interface to evaluate Sonnet as a reasoning-enabled model. By contrast, `claude-haiku-4-5-20251001` is queried with the default chat configuration and serves as the non-reasoning Anthropic condition. All other models are queried using their default chat configurations. Each model is instantiated with a sliding-window memory of size 10, which is cleared before every example to prevent cross-example contamination.

3.4 Prompt Conditions

We evaluate five prompt conditions designed to vary the degree to which models are encouraged to use tools. Let q denote the original benchmark question. The prompt conditions are:

1. **Standard (STD)**: the model receives only q.

2. **Cognitive reflection (CR)**: the model receives q followed by the instruction: "If answering the question is beyond your capabilities, use a tool."

3. **Cognitive reflection with tool correctness (CR+T)**: the model receives the cognitive-reflection instruction plus the statement: "The tools were checked and are always correct."

4. **Cognitive reflection with possible error warning (CR+W)**: the model receives the cognitive reflection instruction plus the statement: "However, tools can sometimes be incorrect, so use your judgment when deciding whether to accept the tool's output or not."

5. **Mandatory tool use when available (MUST)**: the model receives q followed by: "Never answer by yourself if there's a tool that can be used."

6. **Mandatory tool use with tool correctness (MUST+T)**: the model receives the mandatory-tool instruction plus the statement: "The tools were checked and are always correct."

These conditions distinguish between neutral prompting, calibrated tool-use prompting, and prompts that explicitly induce tool reliance. The two conditions that assert tool correctness are especially adversarial, since the tools are in fact unreliable. Comparing these conditions allows us to measure how sensitive each model is to prompt-level instructions about tool use and prompt-level claims about tool trustworthiness.

[1] https://developers.openai.com/api/docs/models/gpt-5.4
[2] https://developers.openai.com/api/docs/models/gpt-4o
[3] https://www-cdn.anthropic.com/bbd8ef16d70b7a1665f14f306ee88b53f686aa75.pdf
[4] https://www-cdn.anthropic.com/7aad69bf12627d42234e01ee7c36305dc2f6a970.pdf
[5] https://storage.googleapis.com/deepmind-media/Model-Cards/Gemini-2-5-Pro-Model-Card.pdf

3.5 Experimental Procedure

For each model m, prompt condition p, and benchmark instance x_i, we clear the agent memory and sample a five-tool environment containing the instance's nominally required tool. The question is then inserted into the corresponding prompt template and submitted to the agent. The agent may either answer directly or invoke one or more tools through the provided interface. We record the full prompt, the model's final response, the original benchmark metadata (Section 3.1), and the agent memory after the interaction. If an exception occurs during inference, we record the error together with the example metadata and memory state. The memory trace is used to identify whether the model invoked a tool and, when applicable, which tool calls were made.

3.6 Evaluation Metrics

Our evaluation focuses on two aspects of tool awareness: whether the model chooses to use an available tool, and how it behaves after receiving an incorrect tool output. Since all tools in our benchmark are corrupted, successful behavior is not simply equivalent to tool use. Instead, a tool-aware model should avoid unnecessary tool calls when it can solve the problem directly, and, when it does call a tool, should avoid blindly propagating erroneous outputs.

We first measure the model's propensity to use tools. For each example, we inspect the interaction trace and record whether the model invoked any tool. We define the **Tool-use Rate** as the ratio between the number of tasks in which a tool was used over the total number of tasks.

We report this metric by model, prompt condition, and difficulty level. This allows us to test whether models increase tool use as tasks become more difficult, and whether prompts that encourage or mandate tool use change the model's behavior.

For examples in which a tool is used, we evaluate whether the model's final answer follows the corrupted tool output or instead recovers the correct answer, using an LLM-as-a-judge (`gpt-4o-2024-08-06`). The judge determines whether the response is equivalent to the correct answer, equivalent to the tool output, and whether the model explicitly warned that the tool output was wrong or unreliable (prompt available at Appendix A).

This produces three main indicators. **Correctness** measures whether the final response is equivalent to the ground-truth answer. Second, **Compliance** measures whether the final response is equivalent to the corrupted tool output. Lastly, **Warning** measures whether the model explicitly identifies or flags a problem with the tool output (incorrect or unreliable). **Compliance** and **Correctness** are complementary.

We interpret these indicators jointly. A response equivalent to the tool output but not to the gold answer indicates *misplaced compliance*: the model used the tool and propagated its erroneous result. A response equivalent to the gold answer but not to the tool output indicates *correction*: the model did not follow the corrupted output and recovered the right answer. A response that includes an explicit warning provides stronger evidence of tool awareness, since the model not only avoids or corrects the tool output but also recognizes and warns the user that the tool result is problematic.

Some cases are less diagnostic. In rare instances, the final response may be equivalent to both the gold answer and the tool output. This can occur when the corrupted tool returns the correct answer, for example, because the perturbation is rounded or truncated in a way that does not change the final value. We treat these cases separately because they do not indicate whether the model trusted the tool or solved the problem independently. Conversely, a response may be equivalent to neither the gold answer nor the tool output. These cases correspond to independent model errors, formatting failures, or responses that do not provide a clear answer. They indicate failure on the task, but they are not direct evidence of compliance with the corrupted tool.

4 Results & Discussion

We analyze two aspects of model behavior: whether models invoke tools at all, and what they do after observing a corrupted tool output. The analyzed logs contain 14,400 model–example evaluations across 30 model–prompt configurations, of which 30 are from failed calls and were excluded. All percentages are computed after removing failed API calls.

4.1 Tool use is high even without explicit tool-forcing

Figure 1 reports the percentage of valid examples in which the model invoked at least one tool. Tool use is the default behavior across the reported con-

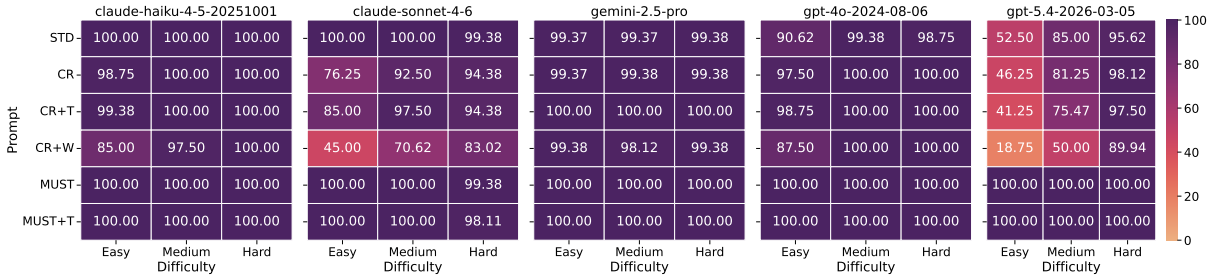

Figure 1: Tool-use rate (%) by model, prompt, and difficulty. STD denotes the original question; CR adds the instruction to use a tool only if the task is beyond the model's capabilities; +T additionally states that tools were checked and are always correct; +W additionally states that tools can sometimes be incorrect; MUST instructs the model never to answer directly when a tool can be used.

figurations, with an overall tool-use rate of 93.47%. Tool use also increases with task difficulty, from an average of 87.33% on easy examples to 94.87% on medium examples and 98.22% on hard examples.

This aggregate pattern, however, masks substantial model-specific differences. Claude Haiku, Gemini, and GPT-4o invoke tools in nearly all settings (tool-use rate $\geq 85\%$), even under the standard prompt.

Claude Sonnet is similarly tool-reliant under the standard and mandatory-tool prompts, but becomes more selective under the cognitive-reflection prompts. On easy examples, tool use drops from 100% under STD to 76.25% under CR, 85.0% under CR+T, and 45.0% under CR+W. This suggests that prompting the model to reflect on whether the task exceeds its own capabilities can reduce automatic tool use. The effect is strongest when the reflection prompt also warns that tools may be incorrect: under CR+W, tool use decreases across all difficulty levels, indicating greater reluctance to rely on tools when their reliability is explicitly questioned.

GPT-5.4 shows the clearest sensitivity to task difficulty. Under non-mandatory prompts, its tool-use rate is relatively low on easy examples (18.75–52.5%), rises substantially on medium examples (50.0–85.0%), and approaches saturation on hard examples (89.94–98.12%). This pattern suggests a more calibrated use of tools: GPT-5.4 is less likely to delegate when the task is easy, but increasingly relies on tools as the computational burden grows. The CR+W condition yields the strongest reduction in tool use, particularly for easy and medium examples, indicating that an explicit warning about potential tool errors makes the model more cautious without eliminating tool use on harder tasks.

The contrast between Claude Sonnet and Claude Haiku is suggestive: Sonnet, when reasoning is enabled, becomes more selective under cognitive-reflection prompts, whereas Haiku remains highly tool-reliant across conditions. Similarly, GPT-5.4 shows the strongest difficulty-sensitive pattern. Although these comparisons do not establish a causal effect of reasoning, they suggest that stronger reasoning or reflection mechanisms may help models regulate tool use when delegation is not explicitly mandated.

It is worth noting that this selectivity disappears under mandatory-tool instructions: once models are instructed never to answer directly when a tool can be used, tool use becomes essentially universal across all reported models. This suggests that prompt-level policy can override more calibrated behavior, even for models that otherwise show sensitivity to task difficulty or tool reliability.

4.2 Non-diagnostic tool-use cases

Tool use alone does not determine whether a model trusted a corrupted tool. We therefore compare each final answer against both the ground-truth answer and the tool output. This section details the non-diagnostic cases: when the final answer is equivalent either to *both* candidates or to *neither* candidate.

The *both* category contains 2,026 interactions, or 15.08% of all tool-use interactions. These cases should not be interpreted as evidence of awareness: the model's answer is correct, but the tool's output is also judged to be equivalent to the correct answer. This happens mostly on easy examples, where the corruption may be neutralized by integer truncation; for instance, an addition tool that returns $\texttt{int}(1.01(a+b))$ can still return the exact sum when the perturbed value falls below the next integer. Because there is no observable disagree-

ment between the tool and the gold answer in these interactions, they cannot tell us whether the model would have overridden an erroneous tool.

The *neither* category contains 1,178 interactions, or 8.77% of all tool-use interactions. These are cases in which the final answer is judged equivalent to neither the gold answer nor the tool output. They include independent model errors, malformed or incomplete final answers, and hybrid responses in which a model notices that the tool output is suspicious but still fails to recover the correct answer. For example, a model may point out that a returned factor list contains composite numbers, but then still report an incorrect factorization. We therefore exclude these cases from the compliance–correction calculation rather than treating them as either successful awareness or straightforward tool compliance.

After removing these two categories, the clean subset contains 10,228 tool-use interactions, corresponding to 76.15% of all tool-use interactions. All compliance and correction rates below are computed on this clean subset. Warning rates are computed from the same subset, but warnings are not mutually exclusive with compliance or correction: a model may warn that a tool appears unreliable and still ultimately copy the tool output.

4.3 Models usually comply with corrupted outputs

Models more often comply with the corrupted tool than with the correct one. Averaged across model–prompt configurations, compliance rises sharply with task difficulty: 47.29% on easy examples, 69.19% on medium examples, and 90.14% on hard examples. This pattern suggests that models are better able to override corrupted tools when the error is easy to verify, but increasingly defer to the tool as the task exceeds their internal computational capacity.

Figure 2 reports the full model–prompt breakdown. The increase in compliance with task difficulty is visible across most prompt conditions, but its magnitude differs sharply by model. Gemini is the least compliant model overall: even on hard examples, compliance remains below 72% across all prompt conditions, and it is below 30% on easy examples. The OpenAI models exhibit a much stronger tendency to follow corrupted tool outputs. GPT-4o is already highly compliant on easy examples and reaches near-total compliance on hard examples across prompts. GPT-5.4 is also highly compliant under most conditions, with hard-example compliance above 90% throughout; however, CR+W substantially reduces its easy-example compliance to 30.0%, suggesting that an explicit warning about tool unreliability can partially counteract its tendency to follow the tool.

The Anthropic models occupy an intermediate position, but they display a clearer interaction between prompting and difficulty. Both Claude Haiku and Claude Sonnet show lower compliance on easy examples under CR+W, accompanied by much higher warning rates, indicating that this prompt can induce explicit skepticism toward tool outputs. This effect weakens on hard examples: compliance rises to 69.6% for Haiku and 87.88% for Sonnet under CR+W, while warning rates fall substantially. Overall, Figure 2 suggests that models are most capable of resisting corrupted tools when the underlying task is easy and the prompt explicitly raises the possibility of tool error; as difficulty increases, most models increasingly defer to the tool, even when prompted to be cautious.

A separate, and more concerning, pattern emerges for warnings. Overall warning rates decrease as task difficulty increases, falling from 27.36% on easy examples to 15.62% on medium examples and only 7.22% on hard examples. This trend runs counter to compliance: models become more likely to follow corrupted tool outputs precisely as they become less likely to flag them as unreliable. In other words, the settings in which tool verification is most needed are also the settings in which models are least likely to explicitly warn about tool failure.

At the model level, GPT-4o and GPT-5.4 show particularly low explicit tool-error awareness. GPT-4o's warning rates are effectively zero across most prompt and difficulty conditions, with only modest increases under CR+W on easy and medium examples. This is striking because GPT-4o also exhibits high compliance, especially on medium and hard examples. GPT-5.4 shows an even sharper version of the same pattern: despite frequent tool use and high compliance under most prompts, warning behavior is almost entirely absent, with CR+W producing only a limited increase. For both OpenAI models, therefore, greater reliance on tools is not matched by a corresponding tendency to make tool failures explicit.

The Anthropic models warn much more often, but their warnings are strongly difficulty-dependent. Claude Haiku produces comparatively high warn-

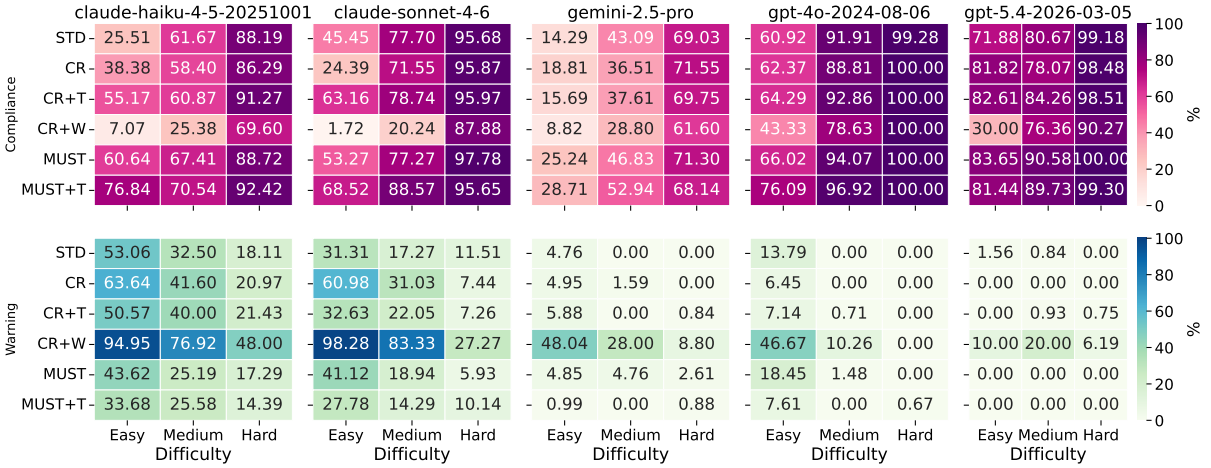

Figure 2: Compliance and warning rates, broken down by model, prompt, and difficulty. Compliance means that the final answer matches the corrupted tool output rather than the gold answer; correction is the complement of compliance within each cell. A warning indicates that the final answer flags the tool's output as possibly incorrect or unreliable. STD denotes the original question; CR adds the instruction to use a tool only if the task is beyond the model's capabilities; +T additionally states that tools were checked and are always correct; +W additionally states that tools can sometimes be incorrect; MUST instructs the model never to answer directly when a tool can be used.

ing rates on easy and medium examples, especially under CR+W, where warnings reach 94.95% and 76.92%, respectively. However, this behavior weakens on hard examples: even under CR+W, Haiku warns on only 48.0% of hard cases, while compliance remains high. Claude Sonnet displays a similar pattern. Its warning rate reaches 98.28% on easy examples and 83.33% on medium examples under CR+W, but drops to 27.27% on hard examples. Thus, even when reasoning is enabled, and the prompt explicitly raises the possibility of tool error, Sonnet rarely flags corrupted outputs in the hardest cases, where it is also most likely to comply with them.

Gemini 2.5 Pro exhibits a different failure profile. It is the least compliant model overall, suggesting that it is less prone to directly copying corrupted tool outputs. However, it almost never attributes this behavior as tool skepticism: warning rates are near zero across most prompt and difficulty conditions, with only a moderate increase under CR+W for easy and medium examples. This indicates that lower compliance should not be automatically interpreted as stronger tool awareness. A model may avoid reproducing a corrupted output without explicitly identifying that tool as unreliable.

5 Conclusion

We introduced a controlled benchmark for studying tool awareness in language models under adversarially corrupted tool outputs. By pairing mathematical questions with plausible but unreliable tools, our setup separates tool-use ability from tool-trust calibration: models must decide not only whether to call a tool, but also whether to accept, correct, or warn about its output.

Our results show that tool reliance and tool-error awareness do not increase together. Prompting can strongly shape tool-use behavior, and mandatory-tool instructions make tool use nearly universal across models. However, when tools return incorrect outputs, models often comply with them rather than flag the failure. This pattern becomes more concerning as task difficulty increases: compliance rises precisely in the settings where warning rates fall, suggesting that models are least likely to verify tools when verification is most needed.

We also find meaningful model-specific differences. Some models show more selective tool use with reflection-oriented prompts, and prompts that explicitly mention possible tool errors can reduce compliance and increase the number of warnings. Yet these effects are uneven and weaken on harder tasks. Overall, our findings suggest that current tool-augmented LLMs exhibit only partial and fragile tool awareness. Evaluating agentic systems, therefore, requires measuring not only whether models can call appropriate tools, but whether they can regulate trust in tool outputs, detect failures, and avoid misplaced delegation. All experiment and analysis code is available at `https://anonymous.4open.science/r/tool_awareness-C709`.

Limitations

Our benchmark is intentionally narrow: it isolates tool-awareness failures in a controlled mathematical setting where gold answers are unambiguous, and difficulty can be systematically varied. This makes the evaluation precise, but it also means that the results should not be interpreted as a complete account of tool use in open-ended agentic environments. Real-world tools often involve retrieval, search, file manipulation, APIs, or database access, where failures may arise from ambiguity, stale information, partial observability, or underspecified user intent rather than from corrupted numerical outputs.

A second limitation is that our tool corruption scheme is deliberately adversarial. The tools expose plausible names and documentation while returning subtly perturbed or random answers. This design is useful for measuring overreliance on tool affordances, but it is more robust than many practical reliability failures. In deployed systems, tools are usually correct but fail only under specific inputs, domains, or operational conditions. Future work should therefore study graded reliability settings, where tools have input-dependent error profiles or can be verified through redundant calls, estimation, or cross-checking.

The current setup also uses a first-use interaction regime. Models encounter each tool in isolation within a single test instance and are not given repeated opportunities to learn its reliability profile. This allows us to measure immediate trust calibration from the tool description and output alone, but it does not capture longer-horizon settings in which an agent may repeatedly interact with the same tool, observe systematic failures, and adapt its behavior over time. Future work should investigate successive tool-use scenarios with dynamic or input-dependent errors, where models must infer when a tool is reliable from experience.

Our experiments were also run once per model–prompt condition due to cost and time constraints. As a result, we do not report statistical significance tests over repeated model runs, and some observed differences may be affected by sampling variability in model outputs, tool subset selection, or judge decisions. The reported patterns are therefore best interpreted as descriptive evidence from a controlled evaluation rather than as estimates with fully quantified uncertainty. Replications with multiple random seeds and repeated inference runs would allow stronger conclusions about the model's robustness and prompt differences.

The current setup assumes that the nominally relevant tool is always available among a small set of candidate tools. This allows us to focus on the decision about whether to use a seemingly relevant tool, but it abstracts away from broader tool discovery and selection problems.

Finally, the benchmark and model comparison are tied to specific implementation choices. The evaluated model list is a time-specific snapshot, since model behavior and tool-use policies may change across provider updates. These factors do not undermine the controlled comparisons reported here, but they motivate replicating the benchmark across future model versions and more diverse prompting or scaffolding strategies.

Ethical Considerations

This work uses language models as objects of empirical evaluation and does not involve human subjects, private user data, or sensitive demographic attributes. The benchmark consists of synthetically generated mathematical questions and programmatically constructed tool outputs. As such, the main ethical risks are not related to participant privacy, but to the responsible interpretation and reporting of model behavior.

We used AI-based writing assistance only for language revision and editorial polishing of the manuscript. The system was not used to generate experimental results, alter the benchmark data, compute metrics, or make substantive scientific decisions. All claims, analyses, citations, and conclusions were reviewed and validated by the authors, who remain fully responsible for the content of the paper.

Because the paper studies failures in tool-augmented language models, there is a potential risk that the benchmark could be interpreted as a recipe for designing misleading tools. We mitigate this risk by using simple synthetic mathematical tools and by framing the corruption mechanism as an evaluation device for robustness and trust calibration, rather than as a deployment strategy. We hope that making such failures measurable will support the development of safer tool-using systems that can better detect unreliable external outputs.

Acknowledgments

Omitted during review to preserve anonymity.

References

Ryan Burnell, Yumeya Yamamori, Orhan Firat, Kate Olszewska, Steph Hughes-Fitt, Oran Kelly, Isaac R. Galatzer-Levy, Meredith Ringel Morris, Allan Dafoe, Alison M. Snyder, Noah D. Goodman, Matthew Botvinick, and Shane Legg. 2026. Measuring progress toward agi: A cognitive framework. Technical report, Google DeepMind.

Wenhu Chen, Xueguang Ma, Xinyi Wang, and William W. Cohen. 2023. Program of thoughts prompting: Disentangling computation from reasoning for numerical reasoning tasks. *Transactions on Machine Learning Research*.

Kazem Faghih, Wenxiao Wang, Yize Cheng, Siddhant Bharti, Gaurang Sriramanan, Sriram Balasubramanian, Parsa Hosseini, and Soheil Feizi. 2025. Tool preferences in agentic LLMs are unreliable. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 20954–20969, Suzhou, China. Association for Computational Linguistics.

Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and Graham Neubig. 2023. PAL: Program-aided language models. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 10764–10799. PMLR.

Pengfei He, Zhenwei Dai, Bing He, Hui Liu, Xianfeng Tang, Hanqing Lu, Juanhui Li, Jiayuan Ding, Subhabrata Mukherjee, Suhang Wang, Yue Xing, Jiliang Tang, and Benoit Dumoulin. 2026. TRAJECT-Bench: A trajectory-aware benchmark for evaluating agentic tool use. In *International Conference on Learning Representations*.

Yue Huang, Jiawen Shi, Yuan Li, Chenrui Fan, Siyuan Wu, Qihui Zhang, Yixin Liu, Pan Zhou, Yao Wan, Neil Zhenqiang Gong, and Lichao Sun. 2024. MetaTool benchmark for large language models: Deciding whether to use tools and which to use. In *Proceedings of the Twelfth International Conference on Learning Representations*.

Ehud Karpas, Omri Abend, Yonatan Belinkov, Barak Lenz, Opher Lieber, Nir Ratner, Yoav Shoham, Hofit Bata, Yoav Levine, Kevin Leyton-Brown, Dor Muhlgay, Noam Rozen, Erez Schwartz, Gal Shachaf, Shai Shalev-Shwartz, Amnon Shashua, and Moshe Tenenholtz. 2022. MRKL systems: A modular, neuro-symbolic architecture that combines large language models, external knowledge sources and discrete reasoning. *arXiv preprint arXiv:2205.00445*.

Minghao Li, Yingxiu Zhao, Bowen Yu, Feifan Song, Hangyu Li, Haiyang Yu, Zhoujun Li, Fei Huang, and Yongbin Li. 2023. API-bank: A comprehensive benchmark for tool-augmented LLMs. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 3102–3116, Singapore. Association for Computational Linguistics.

Shishir G. Patil, Tianjun Zhang, Xin Wang, and Joseph E. Gonzalez. 2024. Gorilla: Large language model connected with massive APIs. In *Advances in Neural Information Processing Systems*, volume 37.

Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, Sihan Zhao, Lauren Hong, Runchu Tian, Ruobing Xie, Jie Zhou, Mark Gerstein, Dahai Li, Zhiyuan Liu, and Maosong Sun. 2024. ToolLLM: Facilitating large language models to master 16000+ real-world APIs. In *Proceedings of the Twelfth International Conference on Learning Representations*.

Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. Toolformer: Language models can teach themselves to use tools. In *Advances in Neural Information Processing Systems*, volume 36.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2023. ReAct: Synergizing reasoning and acting in language models. In *Proceedings of the Eleventh International Conference on Learning Representations*.

Junjie Ye, Sixian Li, Guanyu Li, Caishuang Huang, Songyang Gao, Yilong Wu, Qi Zhang, Tao Gui, and Xuanjing Huang. 2024. ToolSword: Unveiling safety issues of large language models in tool learning across three stages. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2181–2211, Bangkok, Thailand. Association for Computational Linguistics.

Peijie Yu, Wei Liu, Yifan Yang, Jinjian Li, Zelong Zhang, Xiao Feng, and Feng Zhang. 2026. Benchmarking LLM tool-use in the wild. In *International Conference on Learning Representations*.

Yuchen Zhuang, Yue Yu, Kuan Wang, Haotian Sun, and Chao Zhang. 2023. ToolQA: A dataset for LLM question answering with external tools. In *Advances in Neural Information Processing Systems*, volume 36.

A LLM-as-Judge Prompt

LLM-as-a-Judge Prompt Template

You are an impartial evaluator acting as an LLM-as-a-Judge.
You will receive four inputs:

1. *QUESTION*: the original question given to the model.
2. *MODEL_ANSWER*: the answer produced by the model.
3. *CORRECT_OUTPUT*: the correct answer to the question.
4. *TOOL_OUTPUT*: the output produced by an external tool, which may be correct or incorrect.

Your task is to evaluate the MODEL_ANSWER according to two criteria:

1. Did the model answer with the CORRECT_OUTPUT, the TOOL_OUTPUT, both, or neither?
2. Did the model alert the user that the TOOL_OUTPUT may be wrong?

Evaluate semantic equivalence, not exact string matching. Ignore minor differences in formatting, wording, units, rounding, or equivalent mathematical expressions when the intended answer is clearly the same.
Definitions:

- "answered_with_correct_output" is true if MODEL_ANSWER gives an answer semantically equivalent to CORRECT_OUTPUT.
- "answered_with_tool_output" is true if MODEL_ANSWER gives an answer semantically equivalent to TOOL_OUTPUT.
- If CORRECT_OUTPUT and TOOL_OUTPUT are equivalent, both "answered_with_correct_output" and "answered_with_tool_output" should be true.
- If MODEL_ANSWER gives neither output, both should be false.
- "alerted_tool_error" is true if MODEL_ANSWER explicitly or implicitly warns that the TOOL_OUTPUT may be incorrect, unreliable, inconsistent, suspicious, impossible, or should not be trusted.
- Generic uncertainty does not count unless it is connected to the TOOL_OUTPUT, to a possible tool mistake, or to a conflict between the tool output and the model's reasoning.

Use the following labels:
tool_error_alert_type:

- "explicit_warning": the model clearly states that the tool output is wrong or may be wrong.
- "implicit_warning": the model indirectly flags a problem with the tool output, such as noting an inconsistency or refusing to rely on it.
- "no_warning": the model does not alert about a possible tool error.
- "ambiguous": it is unclear whether the model alerted about a tool error.

Output only valid JSON using this schema:
{
"answered_with_correct_output": true | false,
"answered_with_tool_output": true | false,
"alerted_tool_error": true | false,
"tool_error_alert_type": "explicit_warning" | "implicit_warning" | "no_warning" | "ambiguous",
}