

When Hallucination Costs Millions: Benchmarking AI Agents in High-Stakes Adversarial Financial Markets

Zeshi Dai ^{1,*}, Zimo Peng ^{1,*}, Zerui Cheng ^{2,*}, Ryan Yihe Li ^{1,†}

¹ Surf AI, Cybertino Lab

² Princeton University

Link: [Leaderboard](#) | [Github Repository](#) | [HuggingFace DataSet](#)

We present CAIA, a benchmark exposing a critical blind spot in AI evaluation: the inability of state-of-the-art models to operate in adversarial, high-stakes environments where misinformation is weaponized and errors are irreversible. While existing benchmarks measure task completion in controlled settings, real-world deployment demands resilience against active deception. Using cryptocurrency markets as a natural laboratory, where \$30 billion was lost to exploits in 2024, we evaluate 17 leading models on 178 time-anchored tasks requiring agents to distinguish truth from manipulation, navigate fragmented information landscapes, and make irreversible financial decisions under adversarial pressure.

Our results reveal a fundamental capability gap: without tools, even frontier models achieve only 12-28% accuracy on tasks junior analysts routinely handle. Tool augmentation improves performance but plateaus at 67.4% (GPT-5) versus 80% human baseline, despite unlimited access to professional resources. Most critically, we uncover a systematic tool selection catastrophe: models preferentially choose unreliable web search (55.5% of invocations) over authoritative blockchain data, falling for SEO-optimized misinformation and social media manipulation. This behavior persists even when correct answers are directly accessible through specialized tools, suggesting foundational limitations rather than knowledge gaps.

The implications extend beyond cryptocurrency to any domain where adversaries actively exploit AI weaknesses, e.g. cybersecurity, content moderation, etc. Our finding that Pass@k metrics mask dangerous trial-and-error behavior challenges fundamental assumptions about autonomous deployment. We release CAIA with contamination controls and continuous updates, establishing adversarial robustness as a necessary condition for trustworthy AI autonomy. The benchmark reveals that current models, despite impressive reasoning scores, remain fundamentally unprepared for environments where intelligence must survive active opposition.

* Equal Contributions. † Corresponds to: Ryan Yihe Li, ry@cybertinolab.com.

1. Introduction

The Gap Between Benchmark Performance and Autonomous Agent Deployment. Artificial intelligence benchmarks guide optimization, shape incentives, and define progress in modern AI [13, 26]. Over the past year, foundation models have achieved remarkable milestones: OpenAI models won the International Collegiate Programming Contest [15], and Gemini with DeepThink solved International Mathematical Olympiad problems at gold-medal level [21], surpassing most human experts. These achievements have fueled optimism about deploying autonomous AI agents with minimal human oversight. Yet this optimism rests on a dangerous assumption that high scores translate directly to real-world readiness.

Most benchmarks evaluate models in closed worlds where tools function as expected, information is trustworthy, and other agents cooperate [22, 29, 32, 33]. **They measure competence, not resilience.** Real-world autonomy requires surviving in open systems rife with uncertainty, misinformation, and adversarial incentives. Agents deployed in finance, governance, or infrastructure must distinguish truth from manipulation, avoid catastrophic failure, and act conservatively under uncertainty. Evaluation of autonomous AI agents, where trustworthy deployment is the top priority, should therefore critical capabilities explicitly.

This gap creates a perilous blind spot in measuring AI progress. An agent that excels on challenging reasoning benchmarks may still believe fabricated news, purchase compromised assets, or fall for phishing attacks, because nothing in its evaluation prepared it for deception. As AI agents increasingly interact with untrusted users, real money, and critical infrastructure, this vulnerability represents a safety concern hiding behind impressive scores [4, 16].

We argue for an opinion shift in agent evaluation: **Beyond measuring task completion on curated problems, evaluations should test robust survival in adversarial, high-stakes environments.** Rather than escalating difficulty alone, we should simulate hostile settings where others actively deceive, information is weaponized, and irreversible failures cause substantial loss. We introduce CAIA, the Crypto AI Agent Benchmark, which tests AI agent capabilities under these conditions.

Crypto: A Natural Laboratory for Adversarial Robustness. Cryptocurrency markets provide a unique environment for evaluating agent robustness under genuinely adversarial conditions. Despite controversy around speculation and fraud, these characteristics create ideal hostile testing conditions for AI agents. Crypto uniquely combines three properties essential for adversarial evaluation:

1. Adversarial Environment with Sophisticated Deception. The cryptocurrency ecosystem operates as a “dark forest” where misinformation is weaponized and adversaries actively hunt victims [25]. Pseudonymous blockchains enable malicious actors to operate without reputation consequences. Potential profits motivate sophisticated attack strategies. Regulatory gaps permit deception tactics illegal in traditional markets. Daily occurrences include honeypot contracts designed to trap victims [17], flash loan exploits manipulating prices within single transactions [6], and coordinated social engineering campaigns [19]. These real adversarial conditions require agents to genuinely distinguish truth from manipulation.

2. High Stakes with Immediate Consequences. Cryptocurrency markets lack traditional financial safeguards. Transactions are irreversible, smart contract executions are final, and no central authority can reverse fraudulent transfers. In 2024 alone, over \$30 billion was lost to exploits and scams [8]. When an AI agent makes a tiny mistake, losses cannot be recovered. This creates genuine high-stakes conditions where errors have immediate, permanent monetary consequences and malicious actors are economically incentivized to exploit weaknesses.

3. Transparent and Verifiable Ground Truth. Despite adversarial chaos, cryptocurrency offers complete transparency and immutability. Every transaction, smart contract interaction, and token transfer is permanently recorded on public blockchains [31]. This enables unique evaluation conditions where: (1) agent decisions can be verified against immutable on-chain records; (2) financial losses trace to specific transactions with cryptographic proof; (3) attack patterns can be analyzed retroactively with perfect information. This transparency in an adversarial environment enables reproducible evaluation with real-world relevance, addressing fundamental limitations of traditional financial benchmarks that must choose between proprietary data or synthetic simulations.

Current AI systems enter this domain fundamentally unprepared. Trained predominantly on centralized, indexed, trustworthy “Web2” data [10, 14], they lack exposure to crypto’s fragmented, rapidly-evolving “Web3” information landscape. Blockchain data spans thousands of nodes without central access points; DeFi protocols update daily without documentation; critical information exists in ephemeral social channels that evade crawlers [31]. Even accessible content is often adversarial, consisting of deliberately misleading information, scams, and market manipulation. This combination makes crypto particularly challenging for AI trained on traditional web data, and hence an ideal testbed for AI agents’ adversarial robustness.

CAIA: Benchmarking Intelligence Under Fire. We present CAIA (Crypto AI Agent Benchmark), the first benchmark explicitly designed to evaluate AI agents in an actively hostile, high-stakes environment. Unlike existing benchmarks measuring task completion in controlled settings, all tasks in CAIA are grounded in crypto, which measure survival and truth-seeking under adversarial pressure.

Our evaluation reveals significant gaps between state-of-the-art large language models and junior human analysts. Models achieve only 12-28% accuracy without tools. Even when equipped with tools providing correct answers, the accuracy at best is 67.4% (GPT-5), while entry-level human analyst baselines reach 80%. Models consistently rely on unreliable web search over domain-specific tools that directly link to the source of truth, suggesting fundamental limitations in tool selection and adversarial reasoning. These patterns reveal that, when users entrust capital to autonomous agents expecting intelligent fund management, agents may effectively be guessing and attempting “trial-and-error”, which is extremely dangerous in high-stakes adversarial scenarios.

Our Contributions. Our work advances agent evaluation through four primary contributions:

Adversarial-First Evaluation: While existing benchmarks assume cooperative environments, CAIA introduces active deception, source validation, and adversarial robustness as core capabilities, reflecting deployment reality where agents face hostile actors, not just noisy data.

Financial Reality Grounding: Using real market tasks where mistakes have monetary consequences creates accountability and real-world transferability absent from synthetic benchmarks.

Temporal Precision Testing: Time-anchored tasks evaluate multi-timescale reasoning and data obsolescence handling required in volatile markets, beyond static benchmark capabilities.

Diagnostic Failure Analysis: Fine-grained diagnostics of evaluation results provide actionable insights about specific failure modes, critical for both model development and deployment decisions.

The implications extend beyond crypto. As AI agents enter other adversarial domains, e.g. cybersecurity, content moderation, medical diagnosis, CAIA’s measured capabilities become universally critical. Crypto represents an extreme adversarial environment characterized by pervasive misinformation, sophisticated scams, and active financial exploits. Success on CAIA therefore provides high

confidence for autonomous deployment in any domain where adversaries actively exploit weaknesses, and establishes a strong foundation for routine deployment in less hostile environments.

Paper Organization. In the following, Section 2 presents CAIA’s design philosophy and task curation methodology. Section 3 details our experimental framework and quantitative evaluation results across 17 state-of-the-art models. Section 4 analyzes failure modes and derives insights for improving and deploying AI agents. Section 5 discusses future directions and concludes the paper.

2. Benchmark Curation

2.1. Design Principles

CAIA addresses a critical gap in agent evaluation: the absence of benchmarks that capture the worst-case performance under adversarial, high-stakes scenarios, which is exactly the nature of real-world crypto analysis. We identify three core challenges that define this domain:

- irreversible financial consequences where incorrect decisions lead to permanent capital loss (e.g., MEV and execution risks) [11, 24];
- an adversarial information landscape, including coordinated manipulation and pump-and-dump campaigns [5, 28];
- high-density, multi-source data that mixes on-chain traces, social signals, and protocol documentation [22, 33].

Our community-driven curation process, involving over 3,000 contributors including protocol developers, quantitative researchers, and venture capital investors, ensures ecological validity.

To mitigate training-data contamination, a persistent threat to static benchmarks [7, 12], CAIA anchors tasks to recent market events with explicit temporal constraints (block heights, timestamps), following best practices from time-sensitive evaluation [9, 18], and creating an evaluation framework resistant to memorization-based solutions. We will also actively retire out-dated tasks and add new tasks to ensure liveness.

2.2. Quality Assurance

For each task in CAIA, quality is guaranteed through three foundational pillars that mirror expert analytical workflows. This approach moves beyond isolated capability testing to evaluate complete reasoning and acting chains [30], ensuring that successful task completion requires:

Knowledge: Evaluates foundational understanding of crypto-native concepts, from AMM mechanics to governance structures, testing conceptual grasp rather than definitional recall.

Planning: Assesses strategic decomposition of complex questions into executable analytical workflows, requiring agents to specify tool selection and sequencing before execution [27].

Action: Tests real-world execution using production APIs (Etherscan, CoinGecko, DefiLlama) [1–3] dedicated for on-chain detection, evaluating both technical competence and judgment under realistic constraints like rate limiting and data inconsistency.

This progression from understanding through planning to execution reflects established cognitive architectures in complex problem-solving [23] specifically adapted for the crypto domain’s unique requirements. Tasks that satisfy these conditions are naturally suited for testing AI agents’ capabilities, as they precisely mirror the desired approach to tackling complex problems.

2.3. Curation Pipeline

Our dataset originates from more than 10,000 authentic queries we collect from over 3,000 active users spanning different roles, representing the largest and most comprehensive collection of real-world crypto analytical needs to date. Through a rigorous five-stage curation pipeline that operationalizes our design principles and quality metrics (Figure 1), we distill the candidate pool into 178 high-quality CAIA benchmark tasks, as detailed below:




Fig. 1: Data Curation Pipeline.

Stage 1: Automated Filtering: We apply the standard “LLM-as-a-judge” technique, using LLMs to filter out off-topic, ambiguous, non-answerable, or trivial queries while enforcing temporal grounding. After filtration, we ask LLMs to rate each task based on our quality assurance criteria in 2.2, retaining only the top 15%. This reduces the corpus size to approximately 1,000 tasks.

Stage 2: Expert Review: This stage mirrors the traditional paper reviewing process. Our expert team comprises 92 domain specialists, with each assigned 50 tasks to review and grade based on the quality assurance criteria in 2.2. Each task surviving Stage 1 receives at least 4 reviews, and we calculate the final score by averaging all reviews after removing the highest and lowest scores. The top 200 tasks advance to the final pool. After deduplicating similar tasks (i.e., handpicking tasks requiring similar execution logic), we obtain a prototype candidate set of 186 tasks.

Stage 3: Format Standardization: To address inconsistencies arising from different tones and writing styles, we unify the format of each task. This requires explicit anchoring to block numbers or timestamps, enabling objective evaluation and straightforward verification of ground truth answers.

Stage 4: Ground Truth Validation: For each task and its corresponding answer, we verify that a reproducible ground truth toolchain calling scheme exists and associate it with the task. We omit tasks that cannot be reproduced from the benchmark to ensure objectivity. This process provides much more than a single correct answer - It demonstrates the precise methodology that agents should follow to reach the correct solution, ensuring the accuracy, objectivity, and reproducibility of the desired answer. After this validation, we arrive at our final CAIA benchmark of 178 tasks.

Stage 5: Categorization: For diagnostics of model capabilities, we identify 6 fine-grained categories encompassing all tasks and carefully categorize each task accordingly, as shown in Table 1. This step enables detailed assessment beyond aggregate metrics, supporting our analysis in Section 4.

By design, CAIA addresses weaknesses noted in prior evaluations: contamination in static datasets [12], lack of ecological validity in synthetic tasks [22, 33], and single-metric reporting that masks capability gaps [20]. Grounding in real-world high-stakes needs, operating under an adversarial environment with false information, verifiable with objective and immutable answers, CAIA provides a durable foundation for measuring autonomous agentic intelligence in adversarial financial markets.

Category	N	%	Focus	Validation Method
On-Chain Analysis	77	43.3	Transaction patterns, MEV, fund flows	Transaction hash verification
Project Discovery	49	27.5	Protocol evaluation, security analysis	Documentation cross-reference
Tokenomics	23	12.9	Incentive design, value accrual	Mathematical proof
Overlap	14	7.9	Multi-domain synthesis	Composite verification
Trend Analysis	8	4.5	Temporal patterns, adoption metrics	Statistical validation
General Knowledge	7	3.9	Foundational concepts	Canonical reference

Table 1: Distribution of 178 benchmark tasks across 6 analytical categories.

3. Evaluation Results

3.1. Experimental Setup

We conduct a comprehensive evaluation of 17 state-of-the-art large language models on the CAIA benchmark, encompassing leading proprietary models (GPT-4.1, GPT-5, Claude, Gemini, Grok, Kimi) and prominent open-source flagships (Llama, Qwen, DeepSeek, GPT-OSS).

Tool Augmentation. We evaluate each model under two distinct conditions that mirror complementary aspects of real-world deployment. The **without-tools** condition functions as a closed-book examination, testing models’ internalized knowledge and reasoning capabilities when forced to rely solely on parametric memory. This reveals their fundamental understanding of concepts, market dynamics, and analytical reasoning without external assistance. Conversely, the **with-tools** condition resembles an open-book examination and tests agentic abilities where models gain access to 23 specialized tools spanning web search APIs, blockchain analytics platforms, market data feeds, and computational interpreters. Crucially, our data curation process in 2 ensures that correct answers are always accessible through appropriate tool use, and thus the challenge lies not in information availability but in tool selection and synthesis. This design choice deliberately isolates the agent’s tool orchestration capabilities from knowledge limitations, providing a pure test of whether agents can identify and invoke the right resources when given unlimited access to professional instruments.

Agentic Framework. When equipped with tools, each model operates within a standardized ReAct-style [30] agentic framework that handles tool dispatch, result parsing, and iterative reasoning. This ensures that our evaluation result is not affected by implementation variations.

Human Baseline. To establish human performance benchmarks, we recruited 16 participants from university blockchain clubs and early-stage blockchain companies, representing entry-level analyst expertise. These participants completed a stratified 10% sample of our benchmark, carefully balanced across all six analytical domains. Their averaged performance of 80% accuracy provides a critical baseline—notably, these junior analysts achieved this score in the open-book equivalent condition with full tool access, establishing the minimum bar for professional competence in the domain.

3.2. Quantitative Performance Analysis

To ensure robust evaluation, we employ multiple complementary metrics. Our primary measure is average accuracy via majority voting across five independent runs, which mitigates the substantial variance inherent in single-run evaluations of stochastic language models. We additionally report standard Pass@1 and Pass@5 metrics to capture both first-attempt performance and eventual success rates through exploration. However, as we will discuss in Section 4, the traditional Pass@k is misleading in

high-stakes adversarial contexts where trial-and-error carries unacceptable risks.

Beyond performance metrics, we track computational costs by logging token consumption for each query and computing the associated monetary expense. This enables us to derive cost efficiency (cost-per-accuracy-point), revealing critical trade-offs between model capability and economic viability. This analysis proves particularly illuminating when comparing proprietary APIs against open-source alternatives, where we observe up to 100-fold differences in cost for comparable performance.

Our results reveal a stark performance landscape that challenges fundamental assumptions about tool-augmented language models. As illustrated in Figure 2 and detailed in Tables 2 and 3, model performance exhibits a bimodal distribution heavily dependent on tool availability.




Fig. 2: Average accuracy across five evaluation runs using majority voting. The dashed line indicates 80% human baseline performance. Without tools, all models perform near random (12–28%); with tools, performance improves but plateaus below human capability.

In the absence of external tools, every evaluated model, including frontier systems, demonstrates catastrophic failure, achieving merely 12–28% accuracy. This represents performance scarcely above random guessing for many tasks. Even GPT-5, the strongest model, has only 27.5% accuracy, indicating how poorly current parametric knowledge transfers to specialized adversarial domains.

On the other hand, tool augmentation yields substantial improvements, yet even our best-performing model GPT-5 achieves only 67.4% accuracy, falling significantly short of the 80% human baseline established by junior analysts. This performance ceiling persists despite unlimited access to professional-grade tools and comprehensive documentation, suggesting fundamental architectural limitations of the state-of-the-art LLMs today, rather than simple knowledge gaps.

The cost-efficiency analysis reveals a striking economic disparity across model families. While proprietary systems like Claude Opus 4 incur costs exceeding \$1 per problem, open-source alternatives such as GPT-OSS 120B achieve competitive accuracy at under \$0.01 per query, which is a remarkable 100-fold improvement in cost efficiency. Even more compelling, GPT-OSS 120B actually *outperforms* several proprietary models while maintaining this dramatic cost advantage. This economic reality has profound implications for deployment at scale: organizations processing thousands of queries daily could achieve near-frontier performance at a fraction of the cost, fundamentally challenging the assumed superiority of commercial APIs in specialized domains, as illustrated in 3.

Crypto AI Agent Benchmark




Fig. 3: Cost-accuracy tradeoff reveals GPT-OSS 120B and Grok 4 Fast as Pareto-optimal choices, achieving near-frontier performance at minimal cost.

Model	Majority Vote	Pass@1 (%)	Pass@5 (%)	Avg Cost (\$)	Cost/Score
Claude Sonnet 4	0.118	12.9	18.0	0.0070	0.0593
Claude Opus 4	0.135	13.5	16.9	0.0334	0.2481
Claude Opus 4.1	0.135	15.2	17.4	0.0356	0.2642
DeepSeek R1	0.208	21.9	35.4	0.0038	0.0184
DeepSeek V3.1	0.157	15.7	29.2	0.0005	0.0030
Gemini 2.5 Flash	0.185	20.2	21.9	0.0012	0.0062
Gemini 2.5 Pro	0.225	20.2	29.8	0.0051	0.0226
GPT-4.1	0.197	20.8	24.2	0.0025	0.0126
GPT-4o	0.169	19.1	20.8	0.0016	0.0098
GPT-5	0.275	28.1	42.7	0.0207	0.0753
GPT-o3	0.208	22.5	29.2	0.0085	0.0407
GPT-OSS 120B	0.146	18.5	21.3	0.0003	0.0022
Grok 4	0.202	20.2	24.2	0.0345	0.1705
Grok 4 Fast Reasoning	0.191	21.3	23.6	0.0006	0.0029
Kimi K2 Instruct	0.185	17.4	25.3	0.0006	0.0033
Llama 4 Maverick	0.174	16.9	24.7	0.0003	0.0015
Qwen 3 235B	0.163	14.6	18.0	0.0010	0.0061

Table 2: Performance *without tools*: accuracy, Pass@k, cost, and cost efficiency across all models.

Model	Majority Vote	Pass@1 (%)	Pass@5 (%)	Avg Cost (\$)	Cost/Score
Claude Sonnet 4	0.567	57.9	66.9	0.2291	0.4037
Claude Opus 4	0.573	59.6	71.9	1.1139	1.9439
Claude Opus 4.1	0.563	56.3	69.0	0.9357	1.6614
DeepSeek R1	0.174	26.4	54.5	0.0121	0.0695
DeepSeek V3.1	0.492	55.9	71.2	0.0216	0.0438
Gemini 2.5 Flash	0.371	39.3	62.4	0.0070	0.0190
Gemini 2.5 Pro	0.449	49.4	61.2	0.0407	0.0906
GPT-4.1	0.466	51.7	60.7	0.0913	0.1958
GPT-4o	0.303	50.0	55.6	0.0909	0.2997
GPT-5	0.674	70.2	77.0	0.1154	0.1712
GPT-o3	0.534	59.6	73.6	0.1047	0.1962
GPT-OSS 120B	0.629	56.2	72.5	0.0066	0.0104
Grok 4	0.449	52.2	66.9	0.1980	0.4405
Grok 4 Fast Reasoning	0.612	57.9	71.9	0.0098	0.0160
Kimi K2 Instruct	0.472	46.6	64.6	0.0273	0.0579
Llama 4 Maverick	0.242	30.3	64.6	0.0031	0.0129
Qwen 3 235B	0.367	38.4	61.0	0.0062	0.0170

Table 3: Performance *with tools*: accuracy, Pass@k, cost, and cost efficiency across all models.

4. Analysis and Discussion

4.1. The Illusion of Competence: Why Pass@k Metrics Mislead

A critical finding emerges from the stark divergence between Pass@1 and Pass@5 metrics shown in Table 3. While Pass@1 accuracy hovers around 50–60% for top models except for 70.2% of GPT-5, Pass@5 consistently exceeds 60%, with GPT-5 reaching 77%. This improvement might initially suggest robust problem-solving capabilities. However, this interpretation fundamentally misunderstands the nature of real-world deployment, particularly in high-stakes financial contexts.

Consider the implications: Gemini 2.5 Flash’s jump from 39.3% (Pass@1) to 62.4% (Pass@5) indicates that the model essentially *guesses* correctly through repeated trials rather than reasoning strategically on first attempt. In cryptocurrency markets, where a single incorrect transaction can result in irreversible financial loss, this trial-and-error approach represents an unacceptable risk profile. **Real-world financial decisions do not offer multiple attempts.** When users entrust capital to autonomous agents, expecting them to manage funds “cleverly” through clear reasoning, it is unacceptable and extremely dangerous if the agent is effectively guessing its next action.

The minimal improvement (and sometimes decrease) from Pass@1 to majority voting further underscores this concern. If models were exhibiting genuine understanding with occasional errors, we would expect majority voting to substantially improve accuracy. Instead, the modest gains suggest that errors stem from fundamental reasoning failures rather than stochastic variations. This pattern is particularly alarming given that automated agents are increasingly deployed in financial contexts where users may place unwarranted trust in their recommendations without human oversight.

4.2. Domain-Specific Tool Usage and Performance Patterns




Fig. 4: Tool usage frequency vs. accuracy.




Fig. 5: Performance on different categories.

Figure 4 reveals a positive correlation between tool usage and accuracy with a Pearson coefficient of 0.82, showing that statistically more tool calls can help iterate and refine the response. On the other hand, the improvement by more tool calls is not apparent, which suggests that effective tool use depends on strategic selection rather than quantity. Models making numerous unfocused tool calls may perform worse than those making fewer, well-targeted queries to appropriate tools.

Figure 5 illustrates how this tool effectiveness varies dramatically across CAIA’s six analytical categories. In **general knowledge** tasks, tools provide minimal benefit since this stable, universal knowledge is already well-represented in pre-training corpora, and models perform consistently well with or without tool access. Conversely, **on-chain analysis** and **trend analysis** show the largest

performance gaps between tool-assisted and non-tool scenarios, yet remain the lowest-performing categories overall. These domains demand not just tool access but sophisticated reasoning about which tools to deploy and how to interpret their outputs within dynamic environments.

The intermediate categories, **tokenomics**, **project discovery**, and **overlap**, demonstrate the most successful tool integration, with substantial performance improvements when external resources are available. These domains benefit from tools because they are underrepresented in pre-training data, allowing targeted information retrieval to compensate for knowledge gaps.

However, as we reveal in the following subsection, these improvements are largely driven by generic web search rather than specialized tools. LLMs essentially get lucky on these topics because malicious actors have less economic incentive to manipulate information in these relatively obscure domains.

4.3. The Tool Selection Catastrophe

From our evaluation, we identify the systematic failure of models to select appropriate tools, even when optimal choices are unambiguous. Table 4 reveals that models default to generic web search for 55.5% of all tool invocations (combining Google and Twitter searches), despite having access to specialized blockchain analytics tools that provide authoritative data and direct answer.

Tool Category	Invocations	Percentage
Google search	11,626	49.6%
Specialized blockchain tools	8,351	35.6%
URL fetching	1,743	7.4%
Twitter search	1,388	5.9%
Code execution	355	1.5%
Total	23,463	100.0%

Table 4: Tool usage distribution reveals heavy reliance on generic search over domain-specific tools.

This behavior pattern is not merely suboptimal: it is dangerous. In an adversarial and manipulated environment, web search returns manipulated social media posts, coordinated shilling campaigns, and deliberately false information. Meanwhile, blockchain data provides immutable, verifiable ground truth. Yet models consistently choose the unreliable source over the authoritative one.

Our analysis reveals that certain tools require orchestration to be effective. Twitter search accuracy plummets from 40.7% when used in combination to 6.6% when used alone, indicating that social sentiment tools need market context to provide value. Conversely, direct blockchain queries (e.g., ERC-20 token info) maintain high accuracy in isolation. Models fail to recognize these compositional requirements, treating all tools as functionally equivalent, and problematically have a preference towards generic search tools, which may deliver second-hand manipulated information, over domain-specific tools that directly provide the source of truth for each query.

A Case Study: When Simple Tasks Become Impossible. Task 49 in CAIA epitomizes the depth of model failure in tool selection. The task requires retrieving monthly token launch counts from Pump.fun. The data is readily available through a single blockchain analytics API call, and the ground truth solution is trivial: `DEFILLAMA_PUMP_STATS(MONTH="2025-01", METRIC="LAUNCHES")`

Yet across all 17 evaluated models, **not a single one succeeded**. Instead, we observed a consistent pattern of cascading failure where models fall for misinformation:

1. Initial web searches return SEO-optimized but outdated blog posts
2. Refined searches for specific months yield social media speculation rather than data
3. Desperation leads to Twitter searches, surfacing coordinated misinformation
4. Models synthesize incorrect answers from these unreliable sources

The models never attempt to use DeFiLlama, Dune Analytics, or any blockchain-specific tool, despite these being explicitly documented and available. This represents not just a failed execution but a fundamental inability to recognize when specialized tools are necessary and identify source of truth.

5. Conclusion

We introduce CAIA, the first benchmark evaluating AI agents in high-stakes, adversarial environments. Our evaluation of 17 state-of-the-art models reveals critical gaps: leading models achieve only 67.4% accuracy with tools versus 80% human baseline, consistently preferring unreliable web search over specialized blockchain tools. The key obstacle is not tool access but fundamental lack of skeptical reasoning. Agents are easily misled by manipulation and confidently hallucinate critical data.

These vulnerabilities extend beyond crypto to any adversarial domain where misinformation is weaponized. Current models remain dangerously unreliable when stakes are high and adversaries are present. For trustworthy autonomy, future work should prioritize adversarial robustness over task completion metrics, and CAIA provides a vital testbed for building truly reliable autonomous agents.

References

- [1] Coingecko api: Common errors & rate limit. <https://docs.coingecko.com/docs/common-errors-rate-limit>. Accessed 2025-09-24.
- [2] Defillama api docs (pro api and limits). <https://api-docs.defillama.com/>. Accessed 2025-09-24.
- [3] Etherscan api: Rate limits and pro tiers. <https://docs.etherscan.io/etherscan-v2/rate-limits>. Accessed 2025-09-24.
- [4] Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. Concrete problems in ai safety. In *arXiv preprint arXiv:1606.06565*, 2016.
- [5] David Ardia and Keven Bluteau. The role of twitter in cryptocurrency pump-and-dumps. *arXiv preprint arXiv:2306.02148*, 2023.
- [6] Yixin Cao, Chuanwei Zou, and Xianfeng Cheng. Flashot: A snapshot of flash loan attacks on the DeFi ecosystem. *arXiv preprint arXiv:2102.00626*, 2021.
- [7] Nicholas Carlini et al. Extracting training data from large language models. *USENIX Security Symposium*, 2021.
- [8] Chainalysis Team. Crypto crime report 2025. Chainalysis Report, Jan. 2025. <https://go.chainalysis.com/2025-crypto-crime-report>, 2025.
- [9] Wenhui Chen et al. A dataset for answering time-sensitive questions. In *NeurIPS Datasets and Benchmarks*, 2021.
- [10] Common Crawl Foundation. Common crawl. <https://commoncrawl.org/>, 2024.
- [11] Philip Daian, Steven Goldfeder, Tyler Kell, Yuan Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. Flash boys 2.0: Frontrunning, transaction reordering, and consensus instability in decentralized exchanges. *arXiv preprint arXiv:1904.05234*, 2019.
- [12] Chunyuan Deng, Yilun Zhao, Xiangru Tang, Mark Gerstein, and Arman Cohan. Investigating data contamination in modern benchmarks for large language models. *NAACL*, pages 8706–8719, 2024. URL <https://aclanthology.org/2024.naacl-long.482/>.
- [13] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255, 2009.
- [14] Jesse Dodge, Maarten Sap, Ana Marasović, William Agnew, Gabriel Ilharco, Dirk Groeneveld, Margaret Mitchell, and Matt Gardner. Documenting large webtext corpora: A case study on the colossal clean crawled corpus. *arXiv preprint arXiv:2104.08758*, 2021.
- [15] ICPC Foundation. Openai took a historic step by joining the inaugural experiment to incorporate ai development tools into the 49th annual icpc world finals. <https://worldfinals.icpc.global/2025/openai.html>, 2025. Accessed: 2025-09-24.

- [16] Dan Hendrycks, Nicholas Carlini, John Schulman, and Jacob Steinhardt. Unsolved problems in ml safety. *arXiv preprint arXiv:2109.13916*, 2021.
- [17] Ting Hu, Shangwang Li, Yin Wu, Ming Fan, Wuxia Jin, and Ting Liu. Machine-learning approach using solidity bytecode for smart-contract honeypot detection in ethereum. In *Proc. IEEE Intl. Conference on Software Quality, Reliability and Security Companion (QRS-C)*, pages 652–659, 2021.
- [18] Jungo Kasai et al. Realtime qa: What’s the answer right now? In *NeurIPS Datasets and Benchmarks*, 2023.
- [19] Xigao Li, Anurag Yepuri, and Nick Nikiforakis. Double and nothing: Understanding and detecting cryptocurrency giveaway scams. In *Proc. Network and Distributed System Security Symposium (NDSS)*, 2023.
- [20] Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, et al. Holistic evaluation of language models. *arXiv preprint arXiv:2211.09110*, 2022.
- [21] Thang Luong and Edward Lockhart. Advanced version of gemini with deep think officially achieves gold-medal standard at the international mathematical olympiad. [https://deepmind.google/discover/blog/advanced-version-of-gemini-with-deep-think-officially-achieves-gold-medal-standard-at-the-](https://deepmind.google/discover/blog/advanced-version-of-gemini-with-deep-think-officially-achieves-gold-medal-standard-at-the-2025) 2025. Accessed: 2025-09-24.
- [22] Grégoire Mialon, Clémentine Fourrier, Craig Swift, Thomas Wolf, Yann LeCun, and Thomas Scialom. GAIA: a benchmark for general AI assistants. *arXiv preprint arXiv:2311.12983*, 2023.
- [23] Allen Newell. *Unified Theories of Cognition*. Harvard University Press, 1990.
- [24] Kaihua Qin, Liyi Zhou, and Arthur Gervais. Quantifying blockchain extractable value: How dark is the forest? *arXiv preprint arXiv:2106.07337*, 2021.
- [25] Dan Robinson and Georgios Konstantopoulos. Ethereum is a dark forest. <https://www.paradigm.xyz/2020/08/ethereum-is-a-dark-forest>, 2020.
- [26] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. Glue: A multi-task benchmark and analysis platform for natural language understanding. In *Proc. International Conference on Learning Representations (ICLR)*, 2019.
- [27] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [28] Jiahua Xu, Benjamin Livshits, and Aviv Zohar. The anatomy of a cryptocurrency pump-and-dump scheme. In *USENIX Security Symposium*, 2019.
- [29] Jianzhu Yao, Kevin Wang, Ryan Hsieh, Haisu Zhou, Tianqing Zou, Zerui Cheng, Zhangyang Wang, and Pramod Viswanath. Spin-bench: How well do llms plan strategically and reason socially? *arXiv preprint arXiv:2503.12349*, 2025.
- [30] Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629*, 2022.

- [31] Rui Zhang, Rui Xue, and Ling Liu. Challenges and opportunities in blockchain data management. *IEEE Transactions on Knowledge and Data Engineering*, 33(4):1451–1468, 2021.
- [32] Zihan Zheng, Zerui Cheng, Zeyu Shen, Shang Zhou, Kaiyuan Liu, Hansen He, Dongruixuan Li, Stanley Wei, Hangyi Hao, Jianzhu Yao, Peiyao Sheng, Zixuan Wang, Wenhao Chai, Aleksandra Korolova, Peter Henderson, Sanjeev Arora, Pramod Viswanath, Jingbo Shang, and Saining Xie. LiveCodeBench Pro: How do olympiad medalists judge LLMs in competitive programming? *arXiv preprint arXiv:2506.11928*, 2025.
- [33] Shuyan Zhou, Frank F. Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Tianyue Ou, Yonatan Bisk, Daniel Fried, Uri Alon, and Graham Neubig. Webarena: A realistic web environment for building autonomous agents. *arXiv preprint arXiv:2307.13854*, 2023.