COMPUTER SCIENCE / PREPRINT / 25 NOVEMBER 2026

# A Framework for Evaluating Autonomous AI Agents

## A reproducible synthetic pilot for task-execution policies

**Vladimir Chernega (PhD)**
Contact: vchernega@bestagent-ai.com

## Abstract

Agent evaluations frequently conflate an answer with task completion. This preprint defines a measurement protocol that separates verified completion, unsafe actions, intervention, latency, and full execution cost. A synthetic, paired pilot compares three execution policies on 12 task templates and 2,400 shared scenarios. Validation with at most one retry improved verified success by 10.0 percentage points over retrying only on a tool error (95% paired bootstrap interval 8.7 to 11.2). Cost per verified completion also increased in the simulator. The results describe the configured simulator and do not estimate the performance of deployed language models.

**Keywords.** AI agents, task evaluation, reliability, tool use, safety, reproducibility.

## 1. The unit of evaluation

The unit is one attempted business task with a specified initial state and an independently checkable final state. For a catalogue update, for example, a correct response must correspond to the expected change in the store. A plausible description of the change does not receive completion credit. Each task fixture records allowable tools, permissions, expected side effects, time and cost limits, and a validator.

AgentBench compares agents across interactive environments [1]. GAIA tests tasks that require reasoning and the use of tools [2]. SWE-bench checks whether software patches resolve real repository issues [3]. These benchmarks motivate verifiable outcomes. The protocol here focuses on the additional operational measurements needed to choose an execution policy for a specific workflow.

## 2. Evaluation protocol

Freeze the task fixture, reference data, permission boundary, model and prompt versions, tool versions, retry budget, and environment snapshot before evaluating a policy. Reset the environment for each run. Give competing policies the same scenario and initial state. Keep development tasks separate from the final task set. Log version identifiers and timestamps without recording sensitive payloads.

### Measures to report

**Verified completion.** Share of all attempted tasks whose observable end state satisfies every required postcondition. Repeated attempts on one task count as one denominator unit.

**Reliability.** Report the completion rate by task family and shift, the distribution across repeated runs, and the fraction of tasks that exhaust their retry budget.

**Safety.** Count committed actions that violate an explicit permission, destination, data, or spending constraint. Record blocked attempts separately. Require review for any irreversible external action.

**Autonomy.** Report completion without human intervention separately from assisted completion. An escalation required by policy is an observable outcome and should remain in the denominator.

**Latency.** Measure from receipt to terminal outcome, including tool time, validation and retries. Report the median and 95th percentile, with a separate tally for timeouts.

**Cost.** Sum model, tool, validation, retry and human-review costs over all attempts. Divide this total by the count of verified completions. Specify currencies, prices and the observation window.

### Example pre-registration

Primary outcome: paired difference in verified completion on a locked task set. Report safety and cost alongside the primary outcome and specify operating thresholds before looking at results. If a critical violation is observed, analyze the trace and stop that arm. The NIST Generative AI Profile provides a complementary risk management reference [4].

## 3. Synthetic pilot

Four templates were specified in each of three families: extraction, tool use and policy checks. For each template, the program generated 200 pairs of potential execution attempts. Each attempt sampled transient failure, invalid output and an authorization violation from independently specified Bernoulli probabilities. The same draws were reused for all policies. These are constructed scenarios; the template names describe their intended role in the simulator.

The single-attempt policy accepts any tool result. The retry policy repeats once only after an observable tool failure. The validation policy blocks all violations represented by the simulator’s authorization rule, detects 75% of simulated invalid outputs and flags 3% of valid outputs for retry. These are input assumptions. Every execution costs one nominal unit, and validation of a tool result adds 0.12 units. No price, latency, or human review was measured.

## 4. Results

| Policy | Verified completion | Unsafe commits | Cost per attempt | Cost per success |
|---|---|---|---|---|
| single | 1,653/2,400 (68.9%) | 190 (7.9%) | 1.000 | 1.452 |
| retry on tool error | 1,832/2,400 (76.3%) | 209 (8.7%) | 1.108 | 1.451 |
| validate and retry | 2,072/2,400 (86.3%) | 0 (0.0%) | 1.432 | 1.658 |

Table 1. Trial counts and nominal cost units. Each row represents the same 2,400 scenarios.

| Task family | Single | Retry on tool error | Validate and retry |
|---|---|---|---|
| extraction | 76.5% | 80.9% | 90.0% |
| tool use | 68.9% | 80.8% | 87.1% |
| policy checks | 61.2% | 67.4% | 81.9% |

Table 2. Verified completion by family. Each family contains 800 scenarios per policy.

The interval resamples paired scenario outcomes within each task 2,000 times, keeping task weights fixed. It estimates variability conditional on the 12 configured templates. An organization can use this same comparison to decide whether a validation stage justifies its cost, provided that the task set and outcomes represent the intended deployment.

## 5. Limits and deployment implications

The absence of unsafe commits under validation is a consequence of the simulated rule. It does not estimate real security. Real failures can be correlated across attempts, safety checks can miss harmful actions, and tools can return ambiguous signals. The pilot has no human baseline, no measured latency, no deployment traffic, and no proprietary user records. Domain-specific validation and a held-out set of actual tasks are needed to estimate operating performance.

For a live evaluation, use de-identified or authorized fixtures, retain a versioned trace of tool calls, and document retries, cascades and human approvals. Start with read-only scenarios or mocks for irreversible side effects. Assess the inter-rater agreement of human adjudicators where automated checks are unavailable. Re-run the locked evaluation after changing a model, policy, or tool contract.

### Reproduction, disclosure and license

The PDF embeds **agent_evaluation.py** and **results.json**. Extract both attachments in a PDF reader, then run **python agent_evaluation.py**. Python 3.11+ and the standard library are sufficient. The files contain the task probabilities, both seeds, per-family summaries and the paired interval.

The work is associated with BestAgent AI, which has a commercial interest in AI agent evaluation. The pilot uses synthetic inputs and evaluates no BestAgent AI product or third-party model. Manuscript licensed under CC BY 4.0.

## References

[1] Xiao Liu et al. AgentBench: Evaluating LLMs as Agents. ICLR 2024. arXiv:2308.03688.

[2] Grégoire Mialon et al. GAIA: a benchmark for General AI Assistants. ICLR 2024. arXiv:2311.12983.

[3] Carlos E. Jimenez et al. SWE-bench: Can Language Models Resolve Real-World GitHub Issues? ICLR 2024. arXiv:2310.06770.

[4] NIST. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. NIST AI 600-1, 2024. doi:10.6028/NIST.AI.600-1.