Agent Harness ???????????? 12 ???

> [illegible] Akshay Pachaar[illegible] The Anatomy of an Agent Harness[illegible] 2026-04-06[illegible]
>
> [illegible] https://x.com/akshay_pachaar/status/2041146899319971922 [illegible] `Agent/design/harness.pdf`

????? Agent Harness

[illegible] chatbot[illegible] ReAct [illegible] ——demo [illegible]

> [illegible]

LangChain [illegible] TerminalBench 2.0 [illegible] "[illegible]"[illegible] 30 [illegible] 5[illegible] LLM [illegible] pass rate [illegible] 76.4% [illegible]

[illegible] **Agent Harness**[illegible]

1.1 ??

> **Harness** [illegible] **LLM** [illegible]

- Anthropic [illegible] Claude Code [illegible] SDK [illegible] "the agent harness that powers Claude Code"
- OpenAI [illegible] Codex [illegible] "agent" [illegible] "harness" [illegible]
- LangChain [illegible] Vivek Trivedy [illegible]

> **"If you're not the model, you're the harness."** □□□□□□□□□□□□□□
> harness□□

1.2 ?????"Agent" vs "Harness"

- **Agent**□□□□□□□□ "□□□□ "——□□□□□□□□□□□□□□□□□□□□
- **Harness**□□□□□□□□ □□

> ❝ □□□□ "□□□□ agent"□□□□□□ □□□ **harness**□□□□□□□□□ □

1.3 ????·??????Beren Millidge, 2023?

□ LLM = □□□□ RAM□□□□□□□□ I/O □ CPU□

□□□	**LLM Agent**
RAM□□□□□	Context window
□□□□□□□	□□□□□
□□□□	Tool □□
□□□□	**Harness**

??????

□□□□□□□□□□□□

1. **Prompt Engineering**□□□□□□□□□□
2. **Context Engineering**□□□□□□□□□□□□
3. **Harness Engineering**□□□□□□□□□□□□ □□□□□□□□□ ——
   □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□

> ❝ Harness □□ prompt □□□□□□□□□ agent □□□□□□□ □□□□ □

????? Harness ? 12 ???

> □□ Anthropic□ OpenAI□ LangChain □□□□□□□□□□□□□□□

---

1. Orchestration Loop??????

□□□□□□□ **TAO □□□** **Thought-Action-Observation**□ □□□ ReAct loop□

```
□□ prompt → □ LLM → □□□□ → □□ tool calls → □□□□ → □□□□□
```

□□□□□ □□□□ **while □□** □□□□□□□□□□□□□□□□□□□□□□□□

> Anthropic □□□ runtime □□ **"dumb loop"**□□□□□□□□
> □□□□□□□□□□ □ harness □□□□□

---

2. Tools????

□□□ agent □ "□ "□□ Schema □□□□□ name□ description□ parameter types□□□□□ LLM □□□□□□□□□□□□□□□ Tool □□□□

- □□□ registration□
- Schema □□□ validation□
- □□□□□ argument extraction□
- □□□□□ sandboxed execution□
- □□□□□ result capture□
- □□□□□□□□□□□ observation

□□□□ □

- **Claude Code**□ 6 □□□□ ——□□□□□□□□□□ Web □□□□□□□□ subagent □□
- **OpenAI Agents SDK**□ function tools□ `@function_tool`□ + hosted tools□ WebSearch / CodeInterpreter / FileSearch□ + MCP server tools

---

3. Memory????

□□□□□□□□□□

- **Short-term memory**□□□□□□□□□

- **Long-term memory**□□□□□□□
  - Anthropic□ `CLAUDE.md` □□□□ + □□□□□ `MEMORY.md`
  - LangGraph□ namespace □□□ JSON Stores
  - OpenAI□ SQLite/Redis □□□ Sessions

Claude Code □□□□□ □

1. □□□□□□□ ~150 □□□ □□□□ □
2. □□□□□□□□□□□□
3. □□ transcript□ □□□□□□ □

> □□□□ □ **agent** □□□□□□□□□ "□□ "□□□□□□□□□□□□□□□□

---

4. Context Management???????

□□ agent □□□ □□□□ □□□□□□ **Context Rot**□

- □□□□□□□□□□□□□□□□□□ **30%+**□ Chroma □□□□ Stanford "Lost in the Middle" □□□
- □□ **token** □□ □□□□□□□□□□□□□□□□□

□□□□ □

□□	□□
Compaction	□□□□□□□□□ Claude Code □□□□□□□□ bug □□□□□□□□
Observation Masking	□□□□□□□□□□□□□□□□□ JetBrains Junie□
Just-in-time Retrieval	□□□□□□□□□□□□□ Claude Code □ grep / glob / head / tail□□□□□□□□□
Sub-agent Delegation	□ agent □□□□□□□□ 1k–2k token □□□□

> Anthropic context engineering □□□□□□ □□□□□□□□□□□□□□ "□□□□□ **token** □□ "□

---

5. Prompt Construction?Prompt ???

□□□□□□□□□□□□ □□□□ □

```
System Prompt → Tool Definitions → Memory Files → Conversation History → Current User Message
```

OpenAI Codex □□ □□□□□□ □□□□□□□□□□ □

- Server-controlled system message□□□□
- Tool definitions
- Developer instructions
- User instructions□□□□ AGENTS.md□ **32 KiB** □□ □
- Conversation history

---

6. Output Parsing??????

□□ harness □□□□ □□ **tool calling**□□□□□□□□□□□□ `tool_calls` □□□ □□□□□□□□□□□ □

```
□ tool calls□ → □□ + □□
□ tool calls□ → □□□□□□□□
```

□□□□□ □ OpenAI □ LangChain □□□ Pydantic □□□ schema □□□

□□□ `RetryWithErrorOutputParser`□□□ prompt + □□□□ + □□□□□□□□□□□□□□□□□□□□□□

---

7. State Management??????

□□□□□□□□

□□	□□
LangGraph	□□□□□□□□□□□□□□ reducer □□□□□ super-step □□□ checkpoint□□□□□□□□□□□□□
OpenAI	□□ □□ □□□□□□□□ / SDK Sessions / □□□ Conversations API / □□□ `previous_response_id` □□□□
Claude Code	□ **git commit** □ **checkpoint**□ **progress** □□□□□□□□□

---

8. Error Handling??????

□□□□□□□□□□ □

> “ 10 □□□□□□□□□□ 99% → □□□□□ **~90.4%**□ □□□□□□□□ □

LangGraph □□□□□ □

□□	□□
Transient□□□□	retry + backoff
LLM-recoverable□□□□□□□	□□□ `ToolMessage` □□□□□□□□□□□□
User-fixable□□□□□□□	□□□ interrupt□□□□□□□
Unexpected□□□□	□□□□□□□□□

- Anthropic□□ tool handler □□□□□□□□□ error result □□□ □□□□□□□□
- Stripe □□ harness□ □□□□ **2** □

---

9. Guardrails and Safety???????

OpenAI SDK □□□ □

- **Input guardrails**□□□□□ agent □□□
- **Output guardrails**□□□□□□□□□
- **Tool guardrails**□□□□□□□□□□
- **Tripwire** □□□□□□□□□□

Anthropic □□□□□ □ □□□□□□□□□□□□□□□□ □

- □□□□ □□□□
- □□□□□□ □□□□□

Claude Code □□ **40** □□□□□□□ □□□□□□□□□□

1. □□□□□□□□□□ trust establishment□
2. □□□□□□□□□□ permission check□
3. □□□□□□□□□□□□□

---

10. Verification Loops??????

□□□□□ **demo** □□□ **agent** □□□□□□□

Anthropic □□□□□

- **Rules-based feedback**□□□□ Linter□□□□□
- **Visual feedback**□ UI □□□ Playwright □□□□
- **LLM-as-judge**□□□□ subagent □□□□□

> Claude Code □□ **Boris Cherny**□ □□□□□□□□□□□□□□□□□□□□ **2** □ **3** □□

---

11. Subagent Orchestration?? agent ???

Claude Code □□□□□□□ □

□□	□□
Fork	□□□□□□□□□□□□□□
Teammate	□□□□□□ + □□□□□□
Worktree	□□□ git worktree□□□ agent □□□□

- **OpenAI SDK**□ agents-as-tools□□□□□□□□□□□ + handoffs□□□□□□□□
- **LangGraph**□ subagent □□□□□ state graph

---

12. Long-Horizon Execution????? / Ralph Loop ?

> □□□□□ 12 □□□□□□□□□□□ 11□□□□ "□□□□□□□□□□□ "□□□ 12 □ ——□□ PDF □□□□□□ **"Ralph Loop"** □□□□□

□□ □□□□□□□□ 1–2 □□□□□□ □□□□□□□□□□□□□□□□ —— □□□□□□□□□□□□□

Anthropic □□□□ **Ralph Loop**□

1. **Initializer Agent**□□□□□
   - init □□
   - progress □□
   - feature □□
   - □□ git commit

2. **Coding Agent**
   - git log + progress
   - feature
   - → commit →

> “

/ tool call / / token / guardrail tripwire / / /

??????? 7 ???

12

1. Prompt Assembly	system prompt + tool schemas + memory files + conversation history + **prompt** Lost in the Middle
2. LLM Inference	API text / tool call requests /
3. Output Classification	tool call → tool call → handoff → agent
4. Tool Execution	→ → →
5. Result Packaging	LLM **error result**
6. Context Update	compaction
7. Loop	1

> “ 1–2 tool call

?????????

Anthropic Claude Agent SDK	□□ `query()` □□□□□□□□□□□□	"dumb loop"□ Gather-Act-Verify □□□ gather context → take action → verify results → repeat□
OpenAI Agents SDK	`Runner` □□ async/sync/streamed□	**Code-first**□□□□ Python □□□□□□□ graph DSL
OpenAI Codex Harness	□□□□	Codex Core□ agent + runtime□ + App Server□□□ JSON-RPC□ + □□□□ CLI / VS Code / Web□□ **"Codex □□□ Codex** □□□□□□ "□□□□□□□ **harness**
LangGraph	□□ state graph□ `llm_call` □ `tool_node` + □□□	□□□□□ `AgentExecutor` □□□□□□□□□□ agent
LangChain Deep Agents	□□□□ "agent harness" □□	□□□□ + □□□ `write_todos`□ + □□□□ + subagent + □□□□
CrewAI	Agent / Task / Crew + Flows	□□□□ role/goal/backstory/tools□□ Flows □□ "intelligence where it matters"□□□□□□
AutoGen → Microsoft Agent Framework	Core / AgentChat / Extensions □□	□□□□□□□ sequential / concurrent (fan-out/fan-in) / group chat / handoff / **magentic**□ manager agent □□□□□□□□□□□□

?????? Harness ? 7 ?????

#	□□	□□ **trade-off**
1	□ **agent vs** □ **agent**	Anthropic & OpenAI □□□ □□□ **agent** □□□□ □□ agent □□□□□□□ LLM □□□ handoff □□□□□□□□ □□□□□□□□□ **10** □□□□□□□□□□□□□□□□□ □□
2	**ReAct vs Plan-and-Execute**	ReAct □□□□□□□□□□□□□□□□□□ □□□ Plan-and-Execute □□□□□□□□ **LLMCompiler** □□□□□ **ReAct** □ **3.6×**
3	□□□□□□□□□	□□□ time-based clearing / summarization / observation masking / structured note-taking / sub-agent delegation□ **ACON** □□□ **26–54% token** □□ **+ 95%+** □□□□□ □□□□ reasoning trace □□□ raw tool output□

#	□□	□□ **trade-off**
4	□□□□□□	Computational□□□□ linter□ = □□□ ground truth□ Inferential□ LLM-as-judge□ = □□□□□□□□□□ Martin Fowler / Thoughtworks □□□□□ **guides**□ feedforward □□□□□□□ vs **sensors**□ feedback □□□□□□□
5	□□□□□□□	Permissive □□□□□□□□□□□□□□ vs Restrictive □□□□□□□□□□□□□□□□ □□□□□□□
6	□□□□□□	□□□□□□□□□□□□□□□ □ Vercel □ v0 □□ 80% □□□□□□□□□□ Claude Code □□□□□□ 95% □□□□□□ □□□□□□□□□□□□□□□□□ □
7	**Harness □□**	□□□□□□ harness □□□□□□□□□ **Anthropic □□□ harness** □□□□□ □□□□□□□□□□□ **Anthropic □□□□□□□□□□□ Claude Code □ harness** □□□□□□□ □□□□□□□□□□□

??????????

7.1 ???????????????

- □□□□□ = □□□□□□ □□□□□□□□□□□□□□□□
- □ □□□□□□ □□ □□□□□□□□□□□□
- □□□□□□□□□□□□

> “□□□□□□□□□ → **harness □□□□□□□□**

Manus 6 □□□□□ 5 □□ □□□□□□□□□□ □

- □□□□□□□ → □□ shell □□
- "Management agents" → □□□□□□ handoff

7.2 ??????

[illegible] [illegible] harness [illegible] post-train [illegible]

- Claude Code [illegible] harness
- [illegible]

7.3 "?????"??

> **[illegible] scale up [illegible] harness [illegible]**

??Harness ???

[illegible] harness [illegible] [illegible] [illegible] TerminalBench [illegible] [illegible] **harness[illegible] agent [illegible] 20+ [illegible]**

Harness [illegible] [illegible]

- [illegible]
- [illegible]
- [illegible]
- [illegible] "[illegible]" [illegible] "[illegible]" [illegible]

> **[illegible] harness [illegible] ——[illegible] [illegible]**
> [illegible]
> [illegible]

> **[illegible] agent [illegible] harness[illegible]**

??????AllData????

□□□□□□□□ Agent □□□□□□ `□□□□□□□□□□□□□□.md` □ `□□□□□□□□□□□□□□□□□□□□.md` □□□ WeaveAgent □□□□□□□□□□

Harness □□	□□□□□□□□
Memory □□□□	□ MindStore □□□□□□□□ pyramid + MarkdownLogic □□□□□□□□□□ "□□□□ /□□ /□□ "□□□□□□□ MindStore □
Context Management — Compaction & Tool Output Offloading	□□□□ `project_chronicle_grain.md` □ WeaveAgent Chronicle □□ "□□ + □□ "□□□□□□□□ ——□□□ PDF □□□ **tool output offloading**
Verification Loops□ rules / visual / LLM-as-judge□	□□ Agent □□□□□□□□□□□□ Qwen3-Embedding □□□□□□□□□□□□ LLM-as-judge
Long-Horizon Execution / Ralph Loop	Auto/ □□□□ PDF→MD→Logic→Embedding→MindStore □□□□□□□□□□□□□ progress □□ + git checkpoint □ Ralph Loop □□
Harness □□□□	□□ Qwen3-Embedding checkpoint □□□□□□□□□ Agent □ harness □ □□□□□□ ——□□□□□□□□□
□□□□ — "□□□□ "	Vercel □ 80% □□□□□□□□□□□ Auto/ □□ `gen_*.py` □□□□□□□□□□□□□□□□□

Revision #1
Created 2026-05-18 08:20:55 UTC by Colin
Updated 2026-05-18 08:20:55 UTC by Colin