

TapeAgents: a Holistic Framework for Agent Development and Optimization

Core contributors¹:

Dzmitry Bahdanau^{†,*} Nicolas Gontier[†] Gabriel Huang[†]
Ehsan Kamalloo[†] Rafael Pardinás[†] Alex Piché[†]
Torsten Scholak[†] Oleh Shliachko^{†,*} Jordan Prince Tremblay^{†,*}

Contributors¹:

Karam Ghanem² Soham Parikh[‡] Mitul Tiwari² Quaizar Vohra[‡]

[†]*ServiceNow Research* [‡]*ServiceNow*

Corresponding authors (*): {dzmitry.bahdanau,oleh.shliachko,jordanprince.t}@servicenow.com

Abstract

We present TapeAgents,³ an agent framework built around a granular, structured log (**tape**) of the agent session that also plays the role of the session’s resumable state. In TapeAgents we leverage tapes to facilitate all stages of the LLM Agent development lifecycle. The agent reasons by processing the tape and the LLM output to produce new thought and action steps and append them to the tape. The environment then reacts to the agent’s actions by likewise appending observation steps to the tape. By virtue of this tape-centred design, TapeAgents can provide AI practitioners with holistic end-to-end support. At the development stage, tapes facilitate session persistence, agent auditing, and step-by-step debugging. Post-deployment, one can reuse tapes for evaluation, fine-tuning, and prompt-tuning; crucially, one can adapt tapes from other agents or use revised historical tapes. In this report, we explain the TapeAgents design in detail. We demonstrate possible applications of TapeAgents with several concrete examples of building monolithic agents and multi-agent teams, of optimizing agent prompts and finetuning the agent’s LLM. We present tooling prototypes and report a case study where we use TapeAgents to finetune a Llama-3.1-8B form-filling assistant to perform as well as GPT-4o while being orders of magnitude cheaper. Lastly, our comparative analysis shows that TapeAgents’s advantages over prior frameworks stem from our novel design of the LLM agent as a resumable, modular state machine with a structured configuration, that generates granular, structured logs and that can transform these logs into training text — a unique combination of features absent in previous work.

Manuscript version: October 16, 2024

1 Introduction

In the coming years, we will likely witness widespread deployments of Large Language Model (LLM) Agents: complex user-facing and background workflows that interleave traditional programming with LLM-based intelligence. This big paradigm shift in software architecture will greatly challenge AI practitioners who put LLM agents to work. The agent developers and applied scientists will have to troubleshoot and improve systems that operate in non-stationary environments and deal with non-deterministic LLM behavior and the LLM’s often fragile instruction following. For the LLM agent adoption to go smoothly and lead to good outcomes, it is crucial that agent developers and applied scientists operate in appropriate frameworks that enable effective tooling. Developers and researchers have recently proposed many agentic frameworks

¹Authors are in alphabetical order.

²Work performed while at *ServiceNow*

³<https://github.com/ServiceNow/TapeAgents>

that support practitioners at different stages of the agent development lifecycle. Several frameworks, like LangChain (Chase, 2022), CrewAI and AutoGen (Wu et al., 2024a), help developers quickly build an agent using low-code paradigms, such as prompt-chaining or multi-agent teams. Others, like LangGraph (Chase, 2023), offer low-level support in achieving resumability, asynchronous execution, concurrency and instrumentation. At the other end of the spectrum are frameworks built by researchers like DSPy (Khattab et al., 2023a), TextGrad (Yuksekonul et al., 2024) and Agents (Zhou et al., 2023a), that usually focus on data-driven optimization of the agent performance with model finetuning and prompt-tuning algorithms, while putting less emphasis on the needs of the agent developers.

In this technical report, we present **TapeAgents** — a new holistic agent framework that supports practitioners at both the agent development and data-driven agent optimization stages. We achieve both objectives by building the framework around a comprehensive, structured, granular, semantic-level log of the agent session that we call a **tape**, a term that also gives the framework its name (see Figure 1 for an illustration). The agents in TapeAgents read the tape to make the LLM prompt and then process the LLM output to append new **steps** to the tape: **thought** steps to express reasoning and **action** steps to request external inputs. The **environment** responds to the action steps at the end of the tape with **observation** steps that it likewise appends to the tape. The **orchestrator** invokes the agent and the environment in an alternate fashion and maintains full control over their interactions. By design, the orchestrator can resume from any intermediate tape, which enables session persistence and step-by-step debugging, both key developer requirements for an agent framework. For data-driven algorithms, tapes record the attribution of each step to the respective part of the agent configuration, which facilitates training, data generation and automatic prompt-tuning. Crucially, for both manual debugging and algorithms, agents can reuse lightly adapted tapes from other agents and revise their own tapes. This allows practitioners to maximally benefit from imperfect historical tapes by earlier versions of the agent, both for evaluating the newer versions and for improving them algorithmically. Last but not least, agents stream their intermediate events to the orchestrator to enable delightful interactive experiences.

We invite the reader to start their TapeAgents journey with the technical presentation of the framework in Section 2. There, we cover the details of agent architecture, agent-environment orchestration, tape content and structure. Section 3 describes three low-code agent-building framework prototypes on top of TapeAgents: one for monolithic agents, another for multi-agent teams and the third one with easy-to-tune function-like prompts. The same section also covers early versions of our Studio toolsuite for development and debugging and our Optimize toolsuite for agent optimization. In Section 4, we present diverse examples of building and optimizing agents using TapeAgents framework and tooling. Section 5 presents a deeper case study of a key practical TapeAgents use case: optimizing the quality of a cost-effective conversational assistant using tapes from an expensive multi-node Teacher agent. After presenting the framework and the examples we offer the reader a detailed comparison of TapeAgents with prior work in Section 6. Lastly, Section 7 discusses possible extensions and applications of TapeAgents.

2 TapeAgents: foundations

Our TapeAgents framework proposes an agent-building paradigm that facilitates all stages of the AI Agent development lifecycle. This section presents the framework in a detailed bottom-up approach. First, we introduce the building blocks: the nodes, the agents, and the environment. Then, we explore how these parts can be composed and orchestrated to build a tape-centered system. In this section, we also describe the tape structure and metadata.

2.1 Nodes and Steps

As outlined in Figure 1, in TapeAgents, one builds the agent from **nodes**: the basic atoms of intelligence. A node describes one LLM call and the classical symbolic processing of the call’s output. The agent will dynamically determine which node to run next based on the tape. Nodes generate new tape entries that we call **steps**: basic atoms of the agent’s memory. Examples of what an agent can do in a step include making a long-term plan, reasoning about how to fulfill the plan or how to use a tool, requesting a tool call. Among these examples, the last one is an **action** step as it requests interaction with or has an impact on the agent’s

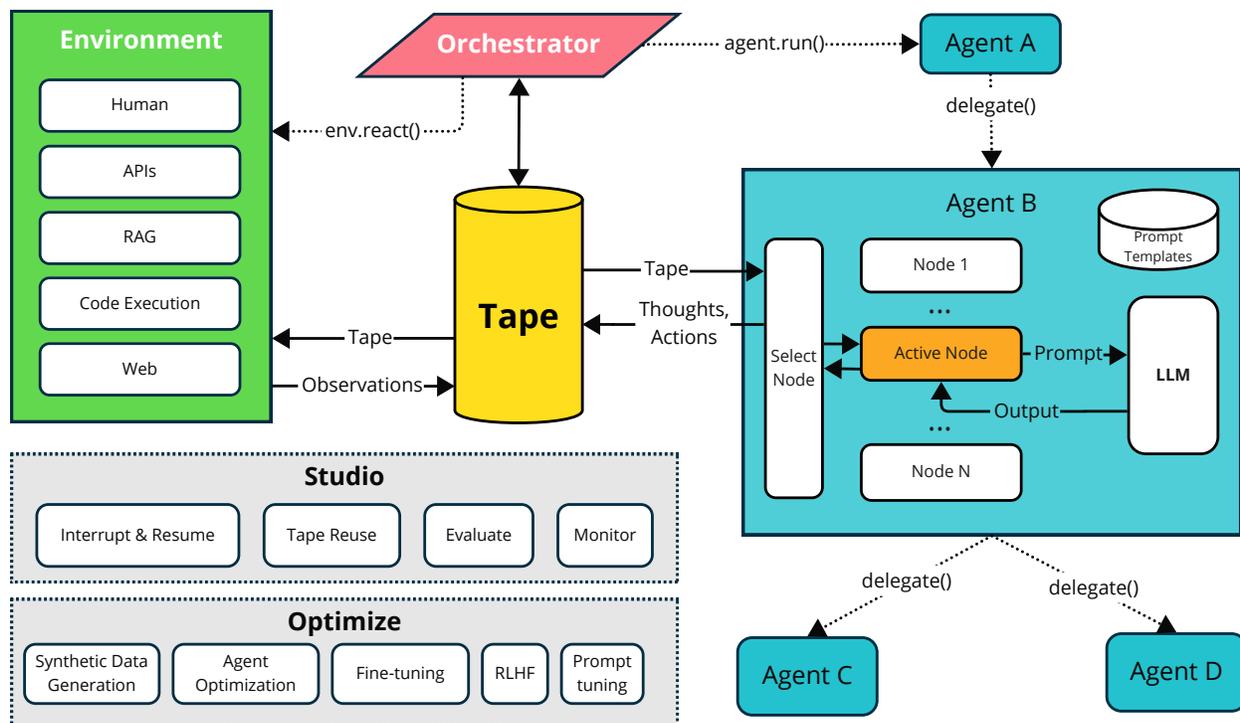


Figure 1: TapeAgents at a glance. The orchestrator alternates between running the agent and the environment interacting with each other via adding steps to the tape: a comprehensive, replayable semantic-level log. Agents are composed of basic reasoning units that we call nodes. The agents are organized in hierarchical teams, with one agent being active at a time. The tape and the agent configurations are highly structured and linked with rich metadata that supports the implementation of broadly usable developer tools (collectively called *Studio*) and optimization methods (collectively called *Optimize*).

environment. The first three examples are **thoughts**: the agent’s inner reasoning steps. The remaining step type in TapeAgents is the **observations** that the agent receives from its environment in response to the agent’s actions. The reader can find an example tape with color-coded actions, thoughts and observations in Figure 3. In TapeAgents we often define a tape type by specifying what specific actions, thoughts, and observations classes it can contain, though all such tapes are currently merely aliases for the one and only `tapeagents.core.Tape` class.

A typical node uses an LLM to generate tape steps. One defines this process with two node methods: `make_prompt` and `generate_steps`. First, the node constructs the LLM prompt through its `make_prompt` method that has the following Python signature:

```
def make_prompt(self, agent, tape) -> Prompt (1)
```

Some nodes perform only the conventional non-neural computation, like taking a branching decision. These nodes can use the default `make_prompt` implementation that produces a null prompt. Note that the node does not call the LLM directly but only makes a prompt. This is a deliberate design decision to keep all node methods pure functions, i.e. deterministic functions with no side effects.

Second, the node generates steps based on the stream of tokens that it receives from the last LLM call. One defines the step-generating behavior of a new node class in its `generate_steps` method:

```
def generate_steps(self, agent, tape, llm_stream) -> Generator[Step | PartialStep] (2)
```

If a node produced a null prompt in its `make_prompt` method above, the `llm_stream` will also be null. All nodes must generate `Step` objects; some can also parse the LLM token stream incrementally to produce partial steps which the agent will pass through to the application without adding them to the tape. Figure 2

shows how the agent runs one node and adds the resulting steps to the tape, along with the relationship between `make_prompt` and `generate_steps`.

By default, the agent calls its nodes sequentially and appends the steps they create to the tape that the agent is asked to continue. A code example of a node implementing these two methods can be seen in Appendix A under `class SearchAgentMainNode`.

2.1.1 Nodes That Can Make Training Data

Some nodes also implement the reverse direction — make the LLM output that would be required to produce the steps at a given index in the tape. The respective node method is

```
def make_llm_output(self, agent, tape, index) -> LLMOutput (3)
```

This method is crucial for making fine-tuning data.

2.2 Agents

Like nodes, a `TapeAgent` agent generates steps and makes a new tape by appending the generated steps to the input tape. Specifically, `agent.run(tape)` runs an iterative reasoning loop that, at every iteration, selects a node, lets it make the prompt and generates the next steps (see Figure 2). By default, the agent will run its nodes sequentially (see next paragraph for more details). The loop continues as long as the nodes only generate thoughts. When a node produces an action, the agent stops and returns a new tape with the generated steps from all iterations appended to it. More precisely, `agent.run(tape)` returns an `AgentStream` object for streaming events like partial tapes and steps, but the final new agent tape is easy to extract from the stream object using `AgentStream.get_final_tape()` method.

An agent may have **subagents** for whom this agent is the **manager**. The subagents can have further subagents, which gives rise to a hierarchical agent organization with a single manager-free root agent on top. Given an input tape, the root agent determines the next active organization member to which delegate the generation of next steps. By default, the root agent makes the delegation decision by looking at the special `Call` and `Respond` thoughts. When an agent A wants the root to delegate to an agent B, A will append `Call(agent_name="B", content=...)` thought to the tape with an optional free-form message in the content field. When B responds by appending `Respond(content=...)`, A becomes active again. Note that both `Call` and `Respond` will affect the delegation logic *at the next agent iteration*. To sum up the delegation description, the root delegates to the agent that was called last and has not responded yet. See Figure 3 for an example of communication between a financial analyst agent and its web search helper. See Appendix A for a listing of the complete code for this example.

2.2.1 Tape Views

In many cases, a middle-level agent A will make its prompts only using a subset of the tape’s steps starting from A’s last `Call` message and excluding the inner steps of the subagents that A called. To keep track of the steps each subagent can see, and to select the active agent and node, most agents compute the **tape view stack** from the tape and delegate to the agent whose view is at the top. For each agent that has not responded yet, the **view** contains the steps that this agent can see and the next node that the agent should

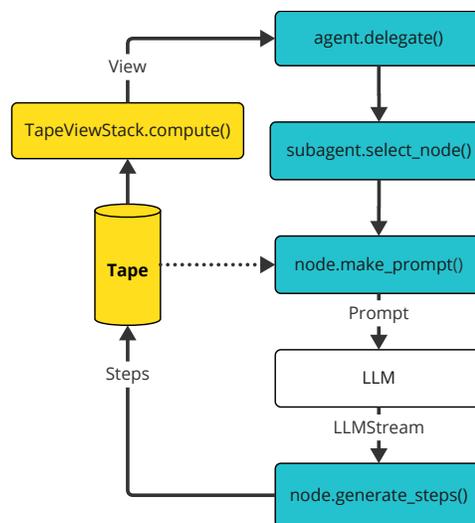


Figure 2: A reasoning loop of an agent in `TapeAgents`. The root agent delegates to a subagent, the subagent selects the node, the node makes the prompt. The subagent calls the LLM with the prompt and lets the node process the resulting stream of tokens (`LLMStream`) that the root agent will then append to the tape.

run when it becomes active (see Figure 3). To track the next node, the `TapeViewStack.compute` function increments the top view’s next node pointer when it sees (from step metadata) that the current node has started running (see more on step metadata in Section 2.5). Exceptionally when `TapeViewStack.compute` encounters a special `SetNextNode(next_node=...)` thought, it changes the next node pointer to the value that the thought carries. A common use-case for `SetNextNode` is to implement looping within an agent.

A reader familiar with how Python interpreter works can find agents similar to Python functions, node similar to lines of Python code, steps similar to Python bytecode instructions, the tape view stack similar to the Python call stack and tape views similar to Python frames.

2.2.2 Optimizable Agents

Agent optimization algorithms tune agent prompts (Khattab et al., 2023a; Pryzant et al., 2023; Zhou et al., 2023b) or alter agent structure (Hu et al., 2024) in order to maximize the agent’s performance. To make such algorithms applicable to as many agents as possible, we standardize the structure of the agent configuration. We achieve this by making `tapeagents.agent.Agent` a Pydantic model⁴ with the following mandatory fields: `.llms` for the LLM configurations, `.templates` for the prompt templates, `.nodes` for the nodes, and `.subagents` for the subagents.

Agents can also make training data for the LLM that they use. An agent’s `agent.make_training_text(tape)` method reconstructs the LLM calls from a given tape, validates the reconstruction by replaying the step generation and returns training text characters. Internally, `agent.make_training_text` uses `node.make_llm_output` method introduced in Section 2.1.1; hence all nodes must implement this method for the agent to be trainable.

2.3 Environment

Just like nodes and agents, the environment in `TapeAgents` makes a new tape by adding steps to an existing tape. The main method of an environment object is:

```
def react(self, tape) -> Tape. (4)
```

The `environment.react` searches for the unfulfilled actions in the tape and adds the corresponding **observation** steps to the tape. Unlike nodes and agents, the environment may be non-deterministic and have side effects. We encourage agent developers to put all the deterministic and pure-function aspects of the system in the agent part, isolating only non-deterministic, computationally heavy or transactional aspects in the environment part.

2.4 Orchestration

To run a `TapeAgent`-based agentic application, one must alternate between running the root agent (which handles the delegation internally) and calling the environment to react to the agent’s actions (see Figure 1). While we provide a default `tapeagents.orchestrator.main_loop` orchestrator for this purpose, we expect many application developers to build their custom orchestrators to closely control the agent-environment communication and ensure safety or enhance iteration logic.

2.4.1 Resumption and Replay

We designed `TapeAgents` with resumption and replay as key priorities. To resume, one can just restart the orchestration from an intermediate tape. For testing purposes, one can run an agent with replayed observations and LLM outputs and verify that this process leads to the same tape or print the diff otherwise. We found the replay tests to be incredibly helpful in our development work. When applicable, one can also replay the tape’s observations (or even some of the agent’s steps) in a new session to evaluate a new agent, though the old observations can be implausible if the new tape deviates too much from the old one.

⁴Pydantic models, from the Python `pydantic` package, provide robust data validation and parsing, ensuring that input data is properly structured and typed. This not only improves reliability by catching errors early but also simplifies code by automatically handling serialization, deserialization, and type conversions.

2.5 Tape Metadata and LLM Call Database

Regardless of the orchestration method, the implementations of `agent.run()` and `environment.react()` ensure that the tape and its steps contain rich metadata, including these fields:

- `tape.metadata.author`: which agent or environment made this tape; either by authoring it, or by adding steps to it, or by making a revision of another tape.
- `tape.metadata.parent_id`: the ID of the parent tape of which the current tape is a continuation (when applicable).
- `step.metadata.agent`: the hierarchical name of the agent that generated the step.
- `step.metadata.node`: the name of the node that generated the step.
- `step.metadata.prompt_id`: the identifier (id) of the prompt that led to the generation of this step, see the explanation below.

When an agent runs a node, the node generates a unique ID for the prompt that it builds at this iteration. The prompt ID thus serves as the unique identifier of a node execution, i.e., of a specific iteration when the node was active. The ID also links the step to the LLM call from the node run so we can trace the origin of each step down to the specific prompt and LLM output. We store the prompt and the output for all LLM calls in an SQLite database. One can view LLM calls as an effective part of the tape in that they are always easily accessible; we don't include them in the tape to keep the latter lightweight.

The metadata is crucial for building the tooling and the algorithms that empower the agent developer. Figure 3 shows a visualization of some metadata fields.

3 TapeAgents: tooling

The TapeAgents foundation that we covered in Section 2 allows the creation of a wide range of reusable agent components, tooling and learning algorithms. What the right building blocks and tooling are often depends on the application area. In our initial release, we provide several prototypes to jump-start future open-source collaborations.

3.1 Low-code Mini-Frameworks

Building agents requires implementing many similar template rendering (`node.make_prompt`) and text parsing (`node.generate_steps`) routines. As a part of TapeAgents, we provide three examples of low-code mini-frameworks for building agents by composing and configuring off-the-shelf components:

1. **MonoAgent** exemplifies the most straightforward way to implement a monolithic agent: make a comprehensive prompt from all the data from the tape and the possible step schemas, then parse the LLM output using the schemas. One creates a `MonoAgent` from `MonoNode` nodes whose prompts are the same except for the final user message instruction. A `MonoAgent` also requires the agent developer to provide Pydantic models for all possible steps that the agent can generate.
2. **TeamAgent** shows how an AutoGen-style agent team can work in one tape. One can create three different kinds of team agents: (a) an initiator that send the first `Call` message, (b) a manager that chooses the next active agent, (c) a worker agent that responds using its system prompt.
3. **LLMFunction** demonstrates how one can build agents using function-style prompt templates, akin to DSPy signatures. These prompt templates are particularly easy to optimize by adding demonstrations.

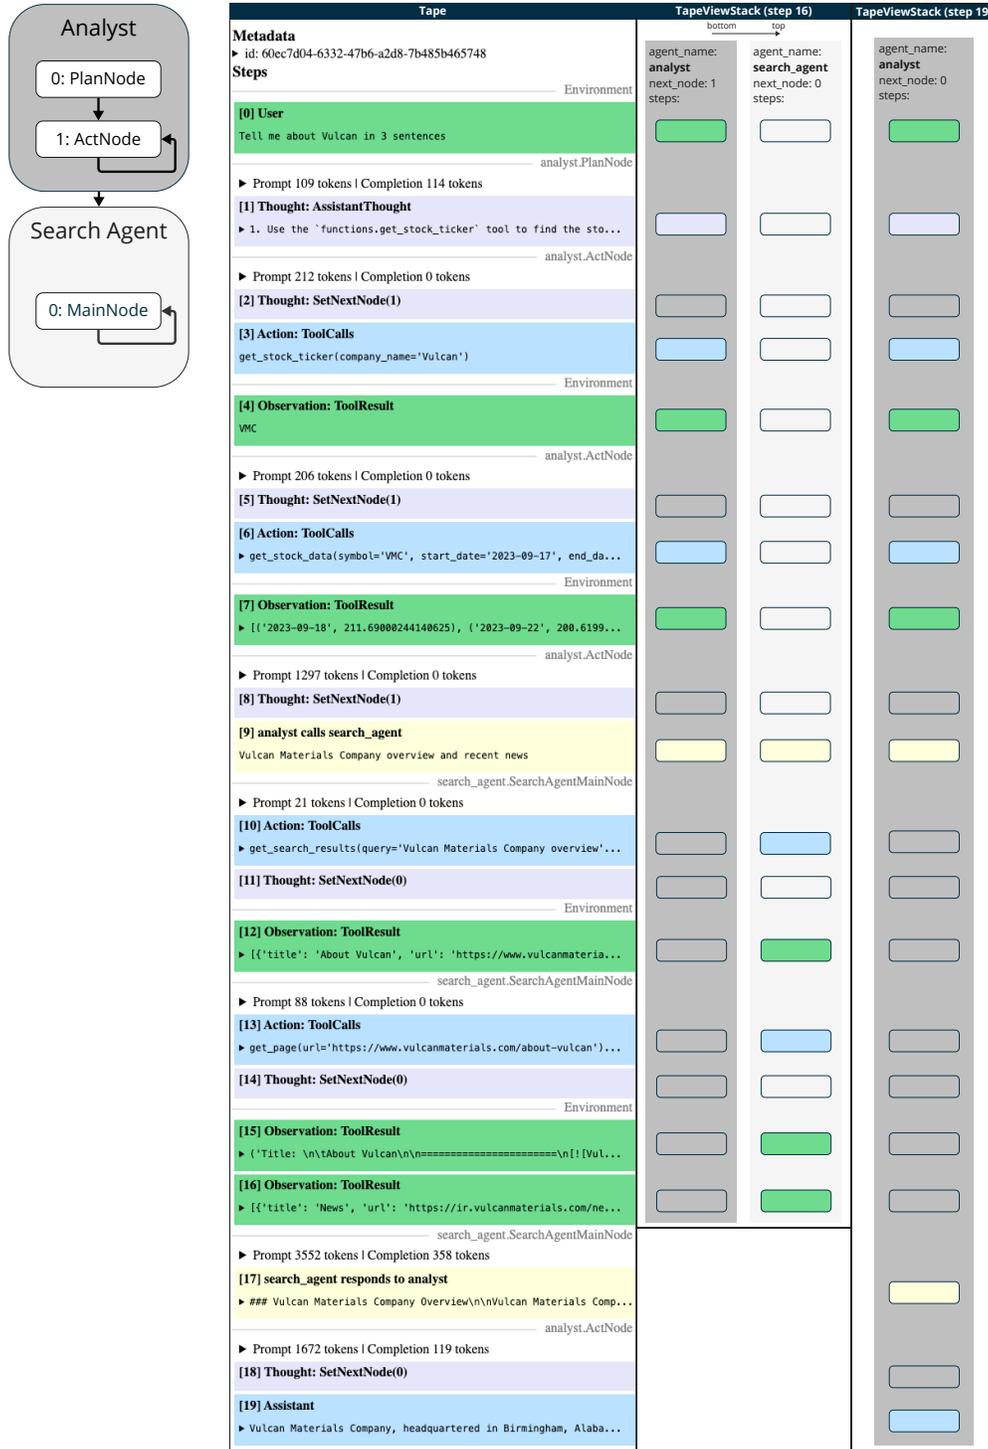


Figure 3: A multi-agent tree structure (left) and a tape resulting from their work (middle) with the TapeViewStack at specific steps (right). At step 16 the stack’s top view is the SearchAgent’s tape view. At step 19, only the Analyst’s view exists. Note how the Analyst’s view does not include the Search Agent’s steps except for its response. Steps are color-coded: yellow for communication thoughts, purple for internal agent thoughts, blue for actions, and green for observations. The step’s author is indicated in grey using the “agent.node” format. Appendix A shows the code implementation that produced this tape.

We include the mini-frameworks mostly for demonstration purposes, as it is hard to offer a high-level programming paradigm without a good knowledge of the intended application domain. The TapeAgents paradigm makes it easy to build such mini-frameworks thanks to the agent’s double compositionality (agents and nodes).

3.2 Tooling

In TapeAgents, the agent configuration and the tape are highly structured and linked with metadata. This allows us to offer developer tooling for a broad range of possible TapeAgents. In the initial release, we include several app prototypes. We offer *TapeAgents Studio* (see Figure 8 in Appendix D), an app to interact with the agent and its tape, *Tape Browser* (Figure 9 in Appendix D), an app to inspect a batch of tapes, and *Tape Diff* (Figure 10 in Appendix D), which compares two batches of tape. Furthermore, for agent optimization, we provide algorithms for auto-prompting, LLM fine-tuning, and a modular Reinforcement Learning orchestrator. The finetuning component uses *Accelerate* (Gugger et al., 2022) and *DeepSpeed* (Rasley et al., 2020) libraries and supports tuning resumption, experiment tracking, reproducibility, LoRA tuning, and distributed training. The above apps and algorithms represent the first steps towards the fully fledged Studio and Optimize modules that we envision in Figure 1.

4 Examples

In an initial set of examples, we demonstrate agents that represent different agent-building paradigms, as well case-studies of using different agent optimization methods.

4.1 Financial Analyst and Their Web Search Helper

To offer an example with the maximal educational value, we have implemented a user-facing financial analyst agent that can delegate searching the web to its subagent. We show the structure of the analyst agent and an example tape in Figure 3. Our introductory hands-on notebook⁵ takes the reader through a journey from TapeAgents basic concepts to building this agent.

For illustrative purposes, we implemented the nodes in this example from scratch, without using mini-frameworks from Section 3.1. We offered the analyst agent an environment with several tools: one to get the company ticker, another to download stock data, as well as several tools to search and browse the web. We inform the analyst and their web search helper of the tools that they can use by including their tools’ schemas in the prompts that the respective agent’s nodes make. The agent operates on a tape type called `DialogTape`, which can only contain two kinds of actions: `ToolCalls` to call one or more tools and `AssistantStep` to respond to the user. The agent uses the same `ToolCalls` step with different content to call different tools. This is the use we intend for `DialogTape`: quick agent prototyping without declaring usecase-specific step schemas, though we believe most TapeAgents users will find it useful to declare their own action and thought types.

4.2 Open-domain Question Answering and Web Browsing With Monolithic Agents

To validate TapeAgents quantitatively, we build two agents that target existing benchmarks. The first one is a question-answering (QA) agent that targets the GAIA benchmark (Mialon et al., 2024). The QA agent can search the web, run Python code, read multiple file types. To meet the GAIA evaluation requirements, we prompt the agent to output the precise short answer only. We build the QA agent from `MonoNode` nodes, with two planning nodes and one acting node in which the agent loops (see Figure 4). Table 1 shows that the agent performs quite well for such a simple agent. Our agent beats the more complicated FRIDAY agent (Wu et al., 2024c) which incorporates `Executor` and `Critic` modules and also working memory. For comparison implementing our agent only requires gathering the tools, declaring corresponding action steps (like `ReadDocumentAction` and `UseCalculatorAction`) and declaring usecase-specific thoughts for reasoning (like `ListOfFactsThought` and `NewFactThought`). There is a lot of room for further improvement of the agent,

⁵<https://github.com/ServiceNow/TapeAgents/blob/main/intro.ipynb>

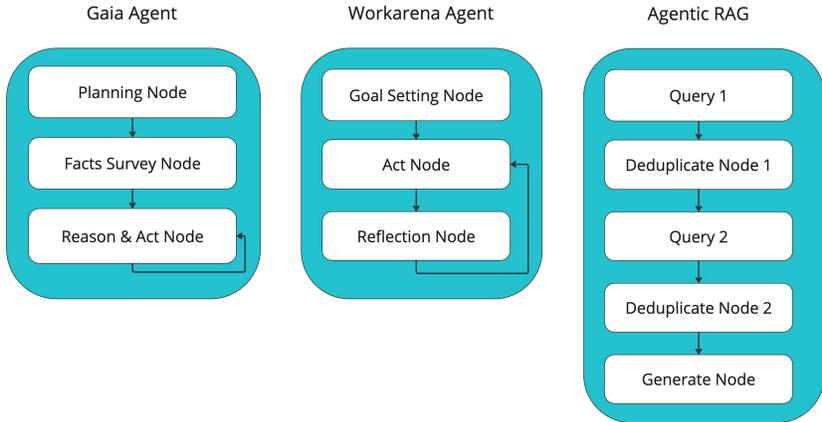


Figure 4: Agent structures for GAIA, WorkArena and Agentic RAG experiments (see sections 4.2 and 4.5 for details).

Table 1: GAIA Agent evaluation results.

Framework & Model	Val accuracy, %	Test accuracy, %
Trase Agent v0.2, Gemini (SOTA)	47.3	39.5
HuggingFace Agents, GPT-4o	44.2	33.3
TapeAgents, GPT-4o	33.9	27.6
FRIDAY, GPT-4-turbo	34.5	24.3
TapeAgents, GPT-4o-mini	25.5	16.6
GPT-4 + manually selected plugins	14.6	14.6

starting from the obviously beneficial majority voting ensembling of runs and up to the multiagents, self-critic techniques, search over the tree of thoughts, etc.

We used a similar approach to build a web-browsing agent that targets the WorkArena benchmark (Drouin et al., 2024). The BrowserGym environment has been used for the evaluation, with a convenient interface to the real browser. We declared action classes for different browser actions like `HoverAction` and `PressAction` and so on. Figure 4 illustrates the exact agent structure, which is following ReAct (Yao et al., 2023)+Planning approach, with goal-setting, reflection and acting nodes. We benchmark our web agent and find that it performs competitively (see Table 2).

4.3 Data Science With a Team of Agents

To demonstrate that TapeAgents natively supports the multi-agent paradigm, we implement a “data science” agent team that consists of the Requestor, Manager, Software Engineer, Code Executor and Asset Reviewer agent. Figure 7 (in Appendix C) shows the team in action as it builds a stock price comparison plot. We drew inspiration from the popular AutoGen framework for the multi-agent communication pattern in this example. Benefits of the TapeAgents implementation of this agent team include that one can easily resume the team from an intermediate tape or use tapes to optimize the entire agent organization algorithmically.

Table 2: Workarena Agent evaluation result on Workarena L1 tasks.

Framework & Model	Accuracy, %
TapeAgents GPT-4o	44.2
Agentlab GPT-4o	42.7
TapeAgents GPT-4o-mini	29.1
Agentlab GPT-4o-mini	23.0

4.4 Finetuning a Cheap Math Agent

We test-drive TapeAgents fine-tuning component with a distillation example. We train a LLAMA3.1-8B-based math agent using tapes by its teacher counterpart with LLAMA3.1-70B under the hood. We equip each agent with a reasoning node and run the environment with a calculator tool. After finetuning on 3,000 samples from 1,000 teacher tapes, the student performance rises significantly from 66.2% to 77.5%, though the teacher’s performance, at 93.1%, remains much higher still.

4.5 Prompt-Tuning for Agentic RAG

In our last example, we show how the tape, the agent configuration and the metadata linking them, can serve as a medium to implement data-driven agent optimization algorithms. In this example, we use `LLMFunction` prompt templates (see examples in Appendix B) that describe the intended behavior of a transformation that the LLM should perform, including the instruction, the input/output format, and optionally, a few demonstrations. We designed `LLMFunction` to make it possible to implement DSPy-like algorithms in TapeAgents. Below we describe how we used TapeAgents components to closely reimplement the DSPy introductory notebook.

We compose a Retrieval-Augmented Generation (RAG) agent. Figure 4 illustrates the structure where the agent performs two rounds of query generation and Wikipedia retrieval and then produces a short factual answer. We build this agent mostly from `LLMFunctionNode` nodes that describe how the input fields in their respective `LLMFunction` templates should be filled with the steps from the tape. The only different kind of a node is a null-prompt node that deduplicates the retrieved paragraphs. We tune the prompts of the resulting 5-node agent by adding demonstrations to the function prompt templates. We obtain demonstrations by running the agent on training examples from HotPotQA and filtering the tapes with the wrong answer or duplicate queries. In this setting, prompt-tuning leads to modest gains but significant gains in retrieval accuracy, raising it from 50% to 56%. The optimized agent is still a TapeAgent that can be resumed from any intermediate tape, unlike a free-form Python program that uses DSPy. Notably, the implementation of the actual demonstration selection algorithm took just 12 lines of code (see Appendix B), highlighting how the TapeAgent structures and metadata facilitate algorithm implementation.

5 Case Study: Building a Cost-Effective Enterprise Form-Filling Assistant

A key use case of TapeAgents is optimizing LLM Agents to offer great quality services at a fraction of the cost. In this section, we present a fleshed-out example of how these goals can be achieved for a conversational assistant that can help fill a request form and submit the request.

5.1 Problem Setting

Employees in large enterprises often fill forms to request resources, assistance or access. A conversational assistant can make the form-filling experience smoother by guiding its user to the right form, by accepting the user’s free-form inputs, and by answering the questions that the user may have in the process. For a great experience, the assistant must also gracefully handle the “unhappy-path” situations, such as when the user’s ask is impossible to fulfill or when the assistant cannot answer the user’s question. In this case study, we show one can use TapeAgents to train a cost-effective assistant that scores high according to a formal metric of user experience that we call the *GREADTH* score. GREADTH stands for Grounded, REsponsive, Accurate, Disciplined, Transparent, and Helpful. We will explain these metrics in Section 5.2.

For simplicity, we consider building a restricted assistant:

- The assistant should answer questions solely based on the form documentation; it does not have to retrieve any additional documents.
- The assistant can only help with one form at a time.

- At the start of the conversation the assistant converses with the user to guide them to the correct form.
- After the form is chosen, the assistant will help the user fill out the form correctly. The agent will not allow the user to switch to a different form after this point.
- During the form-filling process, the assistant maintains the field values that the user has provided so far.
- The interaction ends with either the form submission or the agent exiting the conversation after the user’s confirmation.

The resulting form-filling setup is reminiscent of the Task-Oriented Dialogue setting that has been widely discussed in the literature (Rastogi et al., 2020; Budzianowski et al., 2018). Following this body of work, we will refer to form fields as *slots*.

5.2 Evaluation Criteria: GREADTH Experience

Despite the apparent simplicity of the form-filling setup, it can be non-trivial to develop a form-filler assistant that balances an excellent conversational experience with low hallucination rate and reasonable cost. To balance these desiderata, one must first define them in a measurable way. In our case study, we design our assistant to have maximum **G**READTH: **G**rounded, **R**Esponsive, **A**ccurate, **D**isciplined, **T**ransparent, and **H**elpful. We define these aspects as follows:

- Everything a **G**rounded assistant says must be fully supported by the form documentation, the conversation history and the *grounding statement*. The latter defines the assistant’s identity and purpose and constrains the assistant to form-filling. Small talk is considered ungrounded.
- A **H**elpful assistant must actively take the conversation forward by asking for the user’s intent, requesting the next slot to fill, or asking for confirmation before making the request once all slots have been filled. It should also (a) provide all relevant information regarding a slot when asking for it (default value, allowed values, optionality), (b) answer any user question if the form documentation provides the relevant information, (c) exit the conversation at any time if the user desires so.
- An **A**ccurate assistant must correctly identify the user’s intent and import the relevant form documentation, fill the slots correctly based on user messages, update the slots if the user changes their mind, or skip the slots if relevant.
- A **T**ransparent assistant acknowledges all changes made to the partially filled form. This includes *slot-filling* or *skipping* slots that are optional or have a default value. The summary of slots changes can be concise, yet the user must be able to understand how the slots were affected. While in a mixed modality interaction the user may visually see the form changes, in a purely voice interaction such as talking over the phone the transparent behavior is essential.
- A **D**isciplined assistant must follow its planning thoughts, such as requesting a slot, asking for confirmation, answering a question, rejecting incorrect slot values (as defined by the form documentation), or rejecting an invalid ask.
- We require the assistant to be **R**Esponsive to address a common experience issue with AI assistants: the robotic and opaque behavior when the user goes off the expected conversational path. We wish that AI assistant infers and acknowledges what the user had in mind, while explaining that their request or question is not possible. In particular, we want to cover the following scenarios:
 - if the user tries to fill a slot with an invalid value, the assistant should acknowledge the value and respond that it is invalid;
 - if the user offers information that looks like a value for a nonexistent but plausible slot, the assistant should acknowledge the value and the inferred name of the slot and respond that such slots is not available in this form;

- if at the form-filling stage the user’s ask looks like a request for another form, the agent should acknowledge their ask and say that it can not fulfill it right now (note that in our setup the user must either finish filling the current form or exit);
- if the user asks a plausible question that the form documentation does not answer, the agent should acknowledge the question and say that it can not be answered.

The user may ask other requests that have nothing to do with form-filling (i.e. weather requests). In that case, the agent must politely decline the request and keep moving the conversation towards either submitting or aborting the request. To align our definition of responsiveness with the common sense meaning of this word, we also require a **RE**sponsive assistant to acknowledge all valid slot values and valid questions. Thus we often deem a response that is not **T**ransparent, or not **H**elpful also not **RE**sponsive.

The GREADTH criteria above are binary, a conscious choice that we made to simplify the evaluation and the analysis. We acknowledge that this makes them somewhat crude, as e.g. two assistant answers can be both technically correct but can widely differ in readability and in the choice of the information to present. One can complement these criteria with a preference-based experience evaluation that implicitly covers fluency, verbosity, and other aspects of the assistant’s response.

5.3 Design of a Form Filling TapeAgent

Tape Structure To build an agent that provides a GREADTH experience, we decompose the conversational form-filling task into smaller reasoning steps before each agent message. Having the GREADTH metrics in mind, we define the agent’s thoughts to help it plan its response. The thoughts are used to represent a chain-of-thought of the agent, which includes (a) analyzing the user’s intent (e.g. the form that is requested, the provided slot values, a question being asked); (b) updating the internal state of the conversation (e.g. slot-filling); and (c) planning the next actions (e.g. requesting a specific slot value or requesting a user confirmation). In particular, after each user message or observation, the agent must return:

1. A list of **thoughts**: these include slot-filling related thoughts such as `UpdateFunctionParametersThought`⁶; and message planning thoughts specifying the next slot to request (`RequestFunctionParametersThought`), the need to ask for confirmation before submitting the request (`RequestFunctionCallConfirmationThought`), planning to answer a question (`AnswerFromFunctionSchemaThought`), planning to inform the user that their question cannot be answered (`NoAnswerInFunctionSchemaThought`), refusing unsupported request/behavior/slot values (`RefuseInexistentFunctionThought` / `RefuseToEngageThought` / `RefuseInvalidFunctionParameterValueThought`), etc. The full list of thoughts is described in Appendix E.2.
2. A single **action**: the agent returns a single action, such as searching available forms (`ResolveFunctionAction`), retrieving form documentation (`InspectFunctionAction`), replying to the user (`PromptUserForTextMessageAction`), submitting the request (`CallFunctionAction`), or exiting the conversation (`ExitAction`). Each action results in new observations (available forms, retrieved documentation, or user input), and ends the current agent turn.

Multi-node Teacher Agent We experimented with various combinations of prompting techniques and LLMs to obtain the best performance (see GREADTH metrics in Sections 5.2 and human evaluation procedure in Section 5.4). We found that using Llama-405B with multi-step prompting⁷ (multiple TapeAgent nodes) yielded the most promising results based on human evaluation. For additional details about the

⁶While working on the form-filling case-study we interchangeably used the terms form / request / function and the terms slot / field / parameter. In this report we use the original technical names for the tape’s steps to maximize the coherence between the text and the code.

⁷Also known as prompt chaining or step-wise prompting. More information can be found in prompt engineering guidelines, e.g. <https://docs.anthropic.com/en/docs/build-with-claude/prompt-engineering/chain-prompts> and <https://platform.openai.com/docs/guides/prompt-engineering/tactic-specify-the-steps-required-to-complete-a-task> (Accessed on Oct 15, 2024)

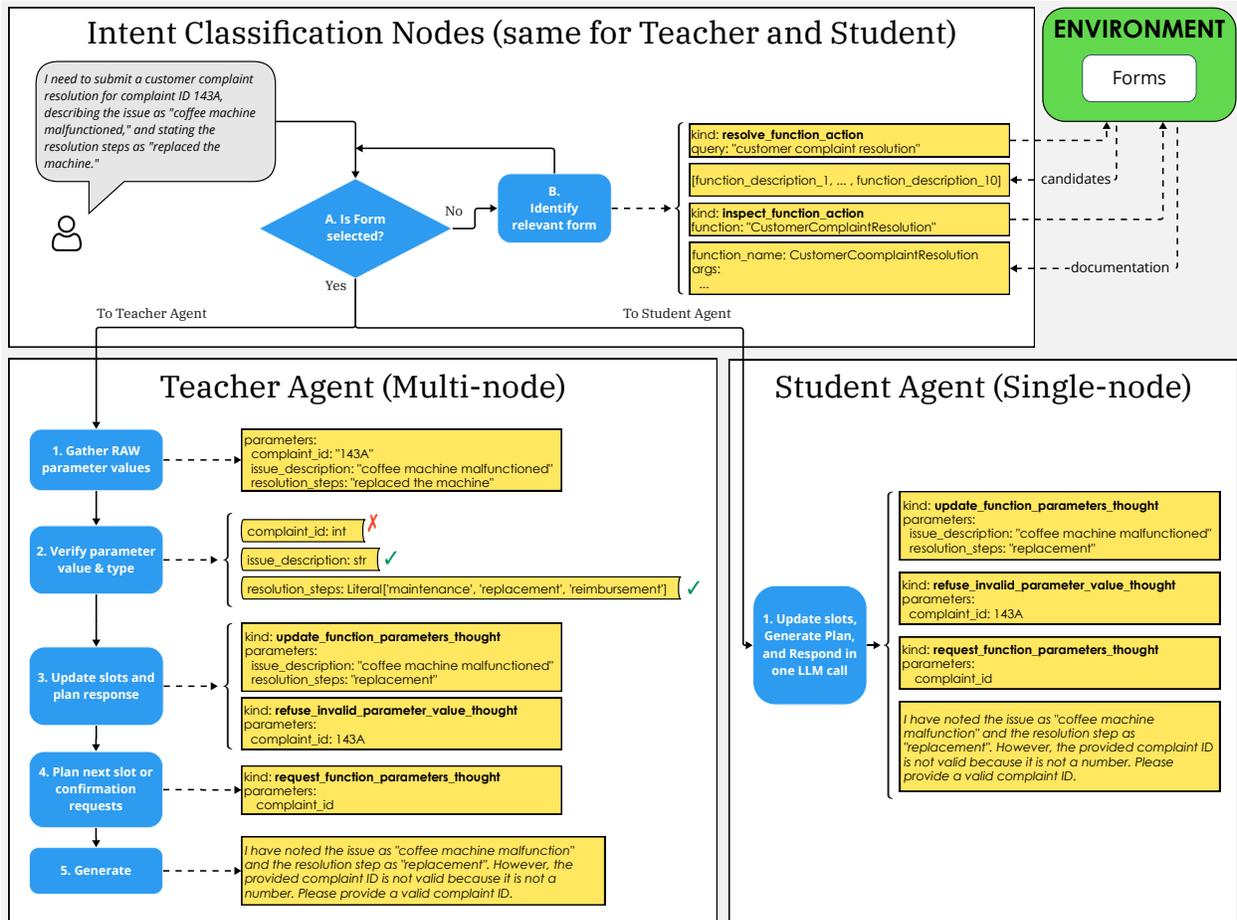


Figure 5: Node structure of the Teacher and the Student agents. The Teacher agent combines intent classification nodes and 5 extra nodes, while the Student agent combines intent classification nodes plus 1 additional node. We represent nodes in blue and produced steps in yellow. Please note that words *intent/form/function* and *slot/field/-parameter* are used interchangeably in this report.

choice of model and agent structure, see Appendix E.4. Our Teacher agent is made of 7 nodes: 2 for intent classification and 5 for slot filling (Figure 5):

- **Is Form Selected?** The initial node checks if we identified the user’s intent. This node does not call any LLM, it simply checks if an `InspectFunctionAction` step is already present in the current Tape. If a form has not yet been selected, the model proceeds to the next step (Step (B) in Figure 5). Otherwise, the model moves directly to the form-filling phase (Step 1 (left) in Figure 5).
- **Intent Discovery.** This node (a) lists the available forms (`ResolveFunctionAction` queries the Environment to return a list of available forms); and (b) finds the relevant one based on the previous user’s message (`InspectFunctionAction` queries the Environment to return the form documentation). If the LLM cannot identify the relevant form, it is prompted to ask again the user for its intent (`RefuseInexistentFunctionThought`; `RequestFunctionThought`; `PromptUserForTextMessageAction`).
- **Gather Raw Parameter Values.** Once the user’s intent is discovered and its form documentation imported, the LLM is prompted to extract all raw slot values present in the user’s message and yield a `GatherValuesThought` step.
- **Verify Parameter Values & Types.** In this node, the LLM is prompted to verify all extracted slot values based on the form documentation imported. The LLM verifies that each value is of the correct type and is valid (in the case of categorical slots) with a `VerifyValuesThought` step.

- **Update Slots and Plan Response.** Once we have identified the correct/incorrect slot values, the LLM is prompted to update the filled slots with a `UpdateFunctionParametersThought` step, refuse invalid slots with `RefuseInvalidFunctionParameterValueThought` steps, and start planning its response to the user with `AnswerFromFunctionSchemaThought` if applicable.
- **Plan Next Slot or Confirmation Requests.** This node prompts the LLM to move the conversation forward by either (i) requesting the next slot (`RequestFunctionParametersThought`), (ii) confirming that the user wants to submit the current request (`RequestFunctionCallConfirmationThought`), or (iii) confirming that the user wants to exit the chat (`RequestExitConfirmationThought`).
- **Generate.** Eventually, the final node of our `Teacher TapeAgent` prompts the LLM to generate the next action. Based on the previous steps/thoughts in the Tape, the agent can either (i) write a message (`PromptUserForTextMessageAction`), (ii) submit the request (`CallFunctionAction`), or (iii) exit the conversation (`ExitAction`).

See Figure 11 of Appendix E.5 for a sample Teacher tape.

Single-node Student Agent We design the Student agent with the goal of minimizing input and output token counts in order to optimize cost. We describe the slot-filling task in one condensed LLM prompt (1 `TapeAgent` node) and use Llama-8B as the LLM to the Student agent to further optimize cost and inference time. While we could get rid of the instructions entirely in the LLM prompt, we still provide enough instructions so that a strong model (e.g. Llama-405B) can still make sense of the task in a zero-shot setting. Our Student agent is made of 3 nodes: 2 for intent classification and 1 for slot filling (Figure 5):

- **Is Form Selected?** Same structure as the Teacher but with fewer and more compact instructions.
- **Intent Discovery.** Same structure as the Teacher but with fewer and more compact instructions.
- **Update Slots, Generate Plan, and Respond.** Using a single short prompt, we ask the agent to generate all its thoughts (e.g. `RefuseInvalidFunctionParameterValueThought`; `UpdateFunctionParametersThought`; `RequestFunctionParametersThought`) and end with an action (e.g. `PromptUserForTextMessageAction`).

We show an example of the Student tape in Figure 12 of Appendix E.5.

5.4 Experiments

We consider the task of training a cost-effective conversational agent to help the user fill and submit a form. We seek the conversational experience to score high in GREADTH metrics (Section 5.2), while requiring less tokens per interaction and less cost per processed token. We attempt to distill a multi-node Teacher agent that uses a large model into a single-node Student agent that uses a small model. For both finetuning and evaluation purposes, we simulate the environment (synthetic forms and user interactions).

Synthetic Companies To simulate an enterprise environment with multiple forms, we prompt a Llama-3-70B-Instruct model to generate the name and descriptions of 6 fictitious companies, which we divide into training domains (`FlyCorp`, `BigBankCorp`, `CoffeeCorp`) and testing domains (`DriveCorp`, `LuxuryCorp`, `ShopCorp`). For each company, we then prompt Llama-3-70B-Instruct to generate 10 plausible request forms based on the company name and description. Each form has a name and a description. Eventually, for each generated form description, we prompt Llama-3-70B-Instruct to generate a `FunctionSchema` for that request form. `FunctionSchemas` are structured data representations including a name, a description, a `JsonSchema` describing slots to fill, and a `JsonSchema` describing the object returned once the form is submitted. Slots to fill have a name, description, type (categorical, date, email, string), optionality, possible values, and default values. The prompts used to generate synthetic companies are described in Appendix E.1. Each of our simulated environments has 10 available forms and identifying which form the user requests is part of the form-filling task we tackle.

Table 3: GREADTH Form Filler experiment results. The Teacher¹ is a multi-node agent with Llama 3.1 405B Instruct FP8 as its LLM. The Student² is a single-node agent with Llama 3.1 8B Instruct as its LLM. We also evaluate the multi-node agent with GPT-4o and with Llama 3.1 8B Instruct as its LLM, as well as the single-node agent with Llama 3.1 405B Instruct for comparison. The metrics are computed over 1524 partial dialogues from the test domains. Read full analysis in Section 5.4.

Agent (LLM+Nodes)	G	Re	A	D	T	H	GREADTH Score (Human Raters)
<i>Reference Comparison (GPT-4o-2024-08-06)</i>							
Multi-node (0-shot)	91.3%	87.1%	91.4%	92.7%	94.3%	87.2%	74.9%
<i>Llama-3.1-405B-Instruct</i>							
Teacher ¹ : Multi-node (0-shot)	89.8%	85.0%	87.9%	91.6%	92.5%	86.5%	75.8%
Single-node (0-shot)	74.2%	72.0%	76.8%	67.3%	78.9%	61.9%	43.2%
<i>Llama-3.1-8B-Instruct</i>							
Multi-node (0-shot)	75.5%	57.7%	72.4%	74.0%	76.3%	60.3%	36.6%
Student ² : Single-node (0-shot)	18.8%	6.2%	10.9%	11.6%	9.4%	12.7%	2.0%
Student ² : Single-node (finetuned)	92.1%	86.4%	90.2%	94.4%	95.1%	87.1%	76.6%

User Agents In the multi-turn dialogue setting, distillation is more complex than running the Teacher on a set of static contexts because of alternating agent and user turns. The user turns need to either be produced by a human or generated. We define 19 different (single-node) User agents to generate the next user message by simulating a variety of user behaviors. Some user behaviors are “easy” such as “*answer the agent’s question*”, while others can be quite adversarial such as “*provide a good value for slot X and a bad value for slot Y*” or “*ask for something unrelated*”. All user agents and their respective behaviors are described in Appendix E.3. Each of these User agents serves as a special instance of the Tape Environment that responds to the assistant after each `PromptUserForTextMessageAction` step.

Synthetic Dialogue Generation We generate a dataset of dialogues by repeatedly and alternatively running the Teacher agent and a User agent to generate the next turn from a set of partial conversations. The minimal partial conversation contains only the “*Hi, how can I help you?*” Assistant message. User agents are sampled randomly after each `PromptUserForTextMessageAction` steps from the Assistant, and yield a user message Observation. The Assistant then continues the conversation until its next message. The conversations end when they reach 18 turns (9 agent messages and 9 user messages), whenever the request is submitted or aborted, or if the agent fails to produce a valid continuation (e.g. trying to fill a nonexistent slot). We show an example of a full conversation between the Teacher agent and the User agent in Figure 11 (Appendix E.5). We structure the dataset as a tree of dialogues, while controlling for the width of the tree (beam search), the diversity of user behaviors, and the diversity of filled forms (requests). We use a beam size of width 500, up to 9 user turns and 9 agent turns. This results in roughly **13k train** and **13k test** agent continuations per synthetic company.

Human Evaluation We perform human evaluation to score each agent turn on the basis of the GREADTH metrics: groundedness, responsiveness, accuracy, discipline, transparency, and helpfulness (Section 5.2). The labeling services were provided by Toloka,⁸ and the labelers were paid above the minimum wage in their respective work locations. Labelers provide the **6 GREADTH binary labels** per agent turn. We also have expert labelers who audit 10% of the labels. We report scores for each metric as well as the **GREADTH score**, which is *equal to 1 only if all 6 metrics are satisfied, and 0 otherwise* (binary AND). If the agent fails to produce a valid continuation – due to unparsable JSON, invalid schema, attempting to fill an invalid slot value, or submitting a form with missing values – then the agent automatically gets 0 for all 6 metrics.

Agent Distillation As described in Section 5.3, we construct the **Teacher agent** by prompting a very large model, Llama-3.1-405B-Instruct-FP8 (runs on 8 H100-80GB GPUs) with sequential and lengthy instructions (**multi-node agent**) in order to obtain the best GREADTH score, with no regards to prompt

⁸<https://toloka.ai/>

Table 4: Agent Cost vs. Performance Tradeoff. On the evaluation set, we report the average number of input and output tokens per agent turn, and the average cost per million agent turns, by multiplying the number of tokens with the retail price of the cheapest provider on Openrouter as of October 3rd 2024 (\$0.055 per million input/output tokens for Llama-3.1-8B-instruct and \$1.79 for Llama-3.1-405B-instruct). The cost per million agent turns is only \$85 for the Student vs. \$28,157 for the Teacher, which represents a **factor of 300** in savings. The cost for the Teacher agent could be mitigated by prefix caching, but most of the multi-node prompts are context-dependent. Read full analysis in Section 5.4.

Agent (LLM+Nodes)	Input Tokens /turn	Output Tokens /turn	Cost /1M turns	GREADTH score
<i>Reference Comparison (GPT-4o-2024-08-06)</i>				
Multi-node (0-shot)	15,431	526	\$43,837	74.9%
<i>Llama-3.1-405B-Instruct</i>				
Teacher: Multi-node (0-shot)	15,189	541	\$28,157	75.8%
Single-node (0-shot)	1,435	97	\$2,742	43.2%
<i>Llama-3.1-8B-Instruct</i>				
Multi-node (0-shot)	15,189	541	\$865	36.6%
Student: Single-node (0-shot)	1,437	208	\$90	2.0%
Student: Single-node (finetuned)	1,441	110	\$85	76.6%

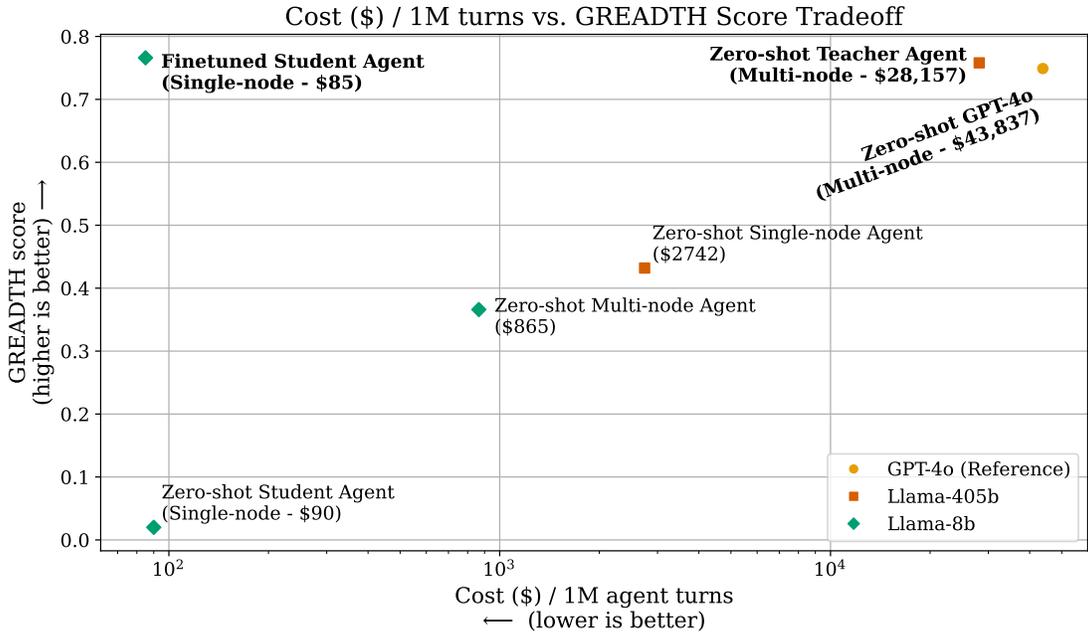


Figure 6: Cost per 1M Agent Turns vs. GREADTH Score Tradeoff. The finetuned Student agent (top-left) performs on par with the Teacher and reference agents (top-right) for a fraction of the cost. We also provide ablations for various combinations of LLMs and node flows (single -vs- multi), which shows that finetuning is instrumental in getting the desired performance when using a short and simple LLM prompt (single-node). See Section 5.4 for the full discussion.

length optimization or latency. From a qualitative perspective, our interaction with the Teacher agent reveals a generally satisfactory user experience and validates its choice as a Teacher agent. We construct the **Student agent** by combining a much smaller model, Llama-3.1-8B-Instruct model (runs on a single A100-80GB GPU), with a single-node (**single-node agent**) optimized for length. We distill the Teacher agent into the Student agent by finetuning it over **13k teacher agent turns** (continuations) on the training domains. We perform a single epoch of LoRA (Hu et al., 2021) optimization using AdamW (Loshchilov and Hutter, 2017) with learning rate 1e-5 and batch size 32. More epochs did not seem to help.

Results We evaluate the Teacher and Student agents by generating a single agent turn (continuation) over a set of **1524 partial dialogues** subsampled from the testing domains generated previously. We control for the diversity of user behaviors (19 mostly uniform) and for the diversity of requested forms (30 forms split over 3 domains + dialogues where no form has been selected yet). Here is what we observed:

1. The Teacher agent achieves a GREADTH Score of **75.8%** (Table 3) at the expensive cost of **28,157\$/1M** agent turns (Table 4), due to using a large model and multiple nodes with lengthy instructions. In comparison, the Student agent achieves an honorable **76.6%** GREADTH score comparable to the Teacher agent but costing only **85\$/1M** agent turns, a **factor 300x cheaper**, thanks to shorter instructions and smaller model.
2. We also evaluate single-node agents (same as the Student) but with different LMs in a 0-shot setting, with no fine-tuning. We achieve GREADTH scores of only 43.2% with Llama-3.1-405B and 2.0% with Llama-3.1-8B. This is not particularly surprising given that the single-node Student agent prompt is designed for cost optimization over task success, though it is interesting to observe that the big gap between the two LLMs (43.2%-2.0%=41.2% points) is entirely closed by the finetuning process (the finetuned Student is 0.8% points above the Teacher).
3. Similarly, we evaluate a multi-node agent (same as the Teacher) but with a small LM (Llama-3.1-8B-Instruct) in a 0-shot setting, with no fine-tuning. We observe a large difference of 75.8%-36.6%=39.2% points, which mostly confirms that while using the largest models is most crucial in the zero-shot regime, finetuning smaller models can close the gap and generalize across domains.

6 Related Work

Among the developer-oriented frameworks most similar to TapeAgents are LangGraph (Chase, 2023) and AutoGen (Wu et al., 2024a). In LangGraph, one builds the agent at a very low-level as a concurrent graph-based state machine. AutoGen offers a high-level paradigm to build multi-agent teams. TapeAgents combines the best of these two worlds, as it allows both the low-level control and the implementation of higher-level low-code paradigms like AutoGen. Neither LangGraph nor AutoGen are designed with agent optimization in mind.

On the other side of the spectrum are AI frameworks that have recently demonstrated techniques to automatically optimize prompts and tweak other aspects of agent configuration, such as the flow or the assignment of tools to the agents. Solutions using DSPy (Khatab et al., 2023b) and TextGrad (Yuksekonul et al., 2024) attain higher performance compared to human prompt engineering. AgentOptimizer (Zhang et al., 2024) and agent symbolic learning (Zhou et al., 2024) enable improving agent tools and pipelines. Meta-Agent Search (Hu et al., 2024) aims to create the best multi-agent architecture. The rich metadata that links the tape with the agent configuration in TapeAgents provides a perfect medium for implementing agent optimization algorithms like the ones from the above works. Notably, DSPy and TextGrad implement control flow in pure Python, which created challenges for resuming the agent from a persistent session state. In TapeAgents, developers can freely use Python within one node, but between-node control flow is handled by adding steps to the tape, which makes the tape perfect for session persistence.

6.1 Detailed comparison with LangGraph, AutoGen, DSPy

To show that TapeAgents is uniquely holistic in targeting all stages of LLM agent lifecycle, we present a more detailed side-by-side comparison of TapeAgents with three particularly popular frameworks: LangGraph, AutoGen⁹ and DSPy. We focus on the following seven axis that are particularly helpful to differentiate these frameworks:

- 1. Building from Components while Allowing Finegrained Flow Control** All the frameworks we compare here allow building agentic systems from components, which are called agents in TapeAgents and AutoGen, subgraphs in LangGraph, predictors in DSPy. TapeAgents and LangGraph additionally offer fine-grained control flow inside each module via nodes and transition between them. To have the same level of control in AutoGen, one must write long non-resumable blocks of Python code.
- 2. Native Streaming Support** LangGraph and TapeAgents natively support response streaming by propagating intermediate event loops through the orchestration loops. Streaming support for AutoGen agents requires substantial change to the agent implementation.
- 3. Concurrent LLM Calls** LangGraph state-machine natively supports concurrent node execution, which allows one to make concurrent LLM Calls. At the moment the agent in TapeAgents can only run one node of one agent at a time. We have a plan on how to address this limitation, see Section 7 for details. We believe AutoGen also has no intra-agent concurrency at the framework level. DSPy users orchestrate DSPy modules with Python code, hence they can use Python multithreading to run multiple modules concurrently.
- 4. Resumable State Machine Agents** Developer-friendly frameworks like LangGraph and TapeAgents share the following design pattern: (a) there is a notion of the agent state (tape) that is serializable and that fully determines the agent’s behavior (b) one implements an agent by defining how the agent makes the new state from the old state in response to external inputs, such as LLM outputs or API responses. This *state machine agent* pattern should be contrasted with implementing the agent in pure Python code as advocated by DSPy. In the latter case the agent state is entangled in the non-serializable state of the Python interpreter. The advantage of having a well-defined state machine for the agent is that it can resume from frequent state checkpoints, such as intermediate tapes in case of TapeAgents, and it can be stopped at any time. Our understanding is that AutoGen provides less control over resumption and agent execution, e.g. one can not just rerun the next speaker selection or restart the agent from a pre-recorded speaker selection.
- 5. Log Reuse Across Agents** A practitioner that uses TapeAgents can reuse tapes from one agent to evaluate or improve another agent. We show an example of this pattern in our form-filling case-study in Section 5. There, the tapes could be reused as is, though we believe often minor modifications can suffice for adapting the tape from e.g. a monolithic agent to an agent team. Our understanding is that from the frameworks in comparison only LangGraph can be modified to support this pattern.
- 6. Structured Logs and Agent Configurations for Data-Driven Agent Optimization** In TapeAgents the agent configuration and the tape are highly structured and linked with metadata, helping the implementation of agent optimization algorithms. LangGraph does not currently position its event sequences as a medium to express algorithms. We likewise found the message history and the agent configuration in AutoGen were not designed with agent optimization in mind. On the contrary, the hierarchical structure of DSPy predictors and the traces that DSPy programs enabled implementation of numerous effective algorithms.
- 7. Making Training Text From Semantic-Level Logs** Tape offers a semantic-level representation of the agent session that the agent can convert into low-level training text for LLM finetuning. This is a unique trait of our framework which we believe LangGraph and AutoGen do not share. We believe with some effort a similar pattern could be implemented in DSPy.

⁹A new AutoGen 0.4 version with major changes came out when we were writing this report. The comparison below is based on a prior version.

Method	Development				Optimization			
	Building from Components while Allowing Finegrained Flow Control	Native Streaming Support	Concurrent LLM Calls	Resumable State Machine Agents	Log Reuse Across Agents	Structured and Agent Configurations for Data-Driven Agent Optimization	Logs for Agent	Making Training Text From Semantic-Level Logs
DSPy	✓	✗	✓	✗	✗	✓		▲
LangGraph	✓	✓	✓	✓	▲	▲		✗
AutoGen	▲	▲	✗	▲	✗	▲		✗
TapeAgents (Ours)	✓	✓	✗	✓	✓	✓		✓

Table 5: TapeAgents vs Other Frameworks. TapeAgents stands out in features it offers to the practitioner to the support them throughout the LLM Agent development cycle. In this figure, we use the cross sign (✗) to indicate that major core changes would be required for the framework support the feature. Triangle sign (▲) indicates partial support of a feature, meaning that practitioner would have to do extra effort or accept associated limitations to achieve the respective functionality. Check sign (✓) indicates that the framework natively supports a feature. TapeAgents’s only weakness in this table is the lack of Concurrent LLM Calls, see Section 7 for a discuss of how we intend to tackle it.

We refer the reader to Table 5 for a tabular summary of the above analysis. One can see that TapeAgents uniquely helps practitioners to both develop the agent and optimize it in a data-driven way.

6.2 Observability Platforms vs TapeAgents

Agent observability software such as LangSmith¹⁰ and Langfuse¹¹ adds visibility to agent execution. They allow one to incrementally instrument agent code to track specific components. In TapeAgents, the tape offers complete observability by design, but beyond that it can also be used for point-in-time resumption and agent optimization.

7 Discussion and Future Work

We have presented TapeAgents, a holistic framework that targets all stages of the LLM Agent lifecycle. We believe the tape-centered approach of our framework can facilitate responsible deployment and continual improvement of LLM agents. Initial tapes will help debugging and testing at the development stages, historical tapes from production agent sessions will serve as a machine-readable source of evaluation and training data. Red-teaming algorithms can use historical tapes as seed data for testing the agent on potential attacks or business-critical dangerous situations. The practitioner can also use historical tapes to seed the simulation that they use for testing the agent. These are but a few benefits that TapeAgents can bring to practitioners.

7.1 Immediate Next Steps

TapeAgents is still in early development stages. A key next step for TapeAgents is adding coroutine implementations for the agent loop and for agent-environment orchestration. This will enable both running many agent-environment loops in parallel on their respective tapes and running members of the same agent team on their shared tape. The latter will require changes in the tape view computation to ignore steps of the agents running in parallel, but we believe that the main framework concepts that we introduced in this paper will stay the same.

On the optimization front, we will soon release an online Reinforcement Learning (RL) trainer for TapeAgents, which will improve the assistant agent using the rewards that the annotator agent computes. Another welcome optimizer addition would be implementing in TapeAgents a text-based feedback-propagation algorithm like the one in TextGrad. Tapes steps are a perfect medium to attach feedback to.

¹⁰<https://www.langchain.com/langsmith>

¹¹<https://github.com/langfuse/langfuse>

7.2 Agent as an Optimizable Workflow

Stepping back from the immediate future plans, we believe it is worth reflecting on what should be called an LLM Agent and what should be called “just” a program, a workflow or software. In the TapeAgents context, this philosophical question is what the developer asks themselves when they build an agent—whether they should implement as agent nodes, and what should go in the application that uses the agent. Our current recommendation is that one should treat and implement the parts of the system that they intend to optimize with data-driven algorithms as LLM Agents. Frameworks provide the structure that the algorithms require to identify the issue, propose a change, and test the change’s outcome. In TapeAgents, this process is particularly clear: the algorithm will identify an issue in the tape, attribute it to a root cause step, propose a change to the agent configuration and test this change by resuming the agent from intermediate points in the tape. Thus the gains from algorithmic improvement will compensate for the overhead of respecting the TapeAgents engineering constraints. To sum up, our recommendation is to implement *optimizable workflows* as LLM Agents and use other appropriate tools for the software that will not be subject to data-driven improvement.

7.3 Synthetic Data Generation with Worlds of TapeAgents

In addition to helping practitioners with their solution-specific challenges, we envision synthetic data generation as another application area where TapeAgents can make an impact. A key trend in the data-making trade is building modular pipelines with many agent-like modules, such as prompt-generating workflows (Mitra et al., 2024), judges (Bai et al., 2022), meta-judges (Wu et al., 2024b), process supervisors (Lightman et al., 2023; Uesato et al., 2022), annotator augmented with tools (Wei et al., 2024) among other examples. We believe TapeAgents is a great foundation for the continual improvement of such multi-agent pipelines with human feedback, as implementing all pipeline modules as TapeAgents immediately makes them optimizable.

Acknowledgements

We would like to thank Nicolas Chapados, Chris Manning, Chris Pal, Sebastien Paquet, Siva Reddy, Arkil Patel and David Vazquez for their feedback and ideas.

References

- Bai, Y., Kadavath, S., Kundu, S., Askill, A., Kernion, J., Jones, A., Chen, A., Goldie, A., Mirhoseini, A., McKinnon, C., Chen, C., Olsson, C., Olah, C., Hernandez, D., Drain, D., Ganguli, D., Li, D., Tran-Johnson, E., Perez, E., Kerr, J., Mueller, J., Ladish, J., Landau, J., Ndousse, K., Lukosuite, K., Lovitt, L., Sellitto, M., Elhage, N., Schiefer, N., Mercado, N., DasSarma, N., Lasenby, R., Larson, R., Ringer, S., Johnston, S., Kravec, S., Showk, S. E., Fort, S., Lanham, T., Telleen-Lawton, T., Conerly, T., Henighan, T., Hume, T., Bowman, S. R., Hatfield-Dodds, Z., Mann, B., Amodei, D., Joseph, N., McCandlish, S., Brown, T., and Kaplan, J. (2022). Constitutional ai: Harmlessness from ai feedback.
- Budzianowski, P., Wen, T.-H., Tseng, B.-H., Casanueva, I., Stefan, U., Osman, R., and Gašić, M. (2018). Multiwoz - a large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Chase, H. (2022). LangChain.
- Chase, H. (2023). LangGraph.
- Drouin, A., Gasse, M., Caccia, M., Laradji, I. H., Verme, M. D., Marty, T., Vazquez, D., Chapados, N., and Lacoste, A. (2024). Workarena: How capable are web agents at solving common knowledge work tasks? In *Forty-first International Conference on Machine Learning*.
- Gugger, S., Debut, L., Wolf, T., Schmid, P., Mueller, Z., Mangrulkar, S., Sun, M., and Bossan, B. (2022). Accelerate: Training and inference at scale made simple, efficient and adaptable. <https://github.com/huggingface/accelerate>.

- Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., and Chen, W. (2021). Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv: 2106.09685*.
- Hu, S., Lu, C., and Clune, J. (2024). Automated design of agentic systems.
- Khattab, O., Singhvi, A., Maheshwari, P., Zhang, Z., Santhanam, K., Vardhamanan, S., Haq, S., Sharma, A., Joshi, T. T., Moazam, H., Miller, H., Zaharia, M., and Potts, C. (2023a). DSPy: Compiling declarative language model calls into self-improving pipelines. *arXiv preprint arXiv:2310.03714*.
- Khattab, O., Singhvi, A., Maheshwari, P., Zhang, Z., Santhanam, K., Vardhamanan, S., Haq, S., Sharma, A., Joshi, T. T., Moazam, H., Miller, H., Zaharia, M., and Potts, C. (2023b). DSPy: Compiling declarative language model calls into self-improving pipelines.
- Lightman, H., Kosaraju, V., Burda, Y., Edwards, H., Baker, B., Lee, T., Leike, J., Schulman, J., Sutskever, I., and Cobbe, K. (2023). Let’s verify step by step. *arXiv preprint arXiv:2305.20050*.
- Loshchilov, I. and Hutter, F. (2017). Decoupled weight decay regularization. *International Conference on Learning Representations*.
- Mialon, G., Fourrier, C., Wolf, T., LeCun, Y., and Scialom, T. (2024). GAIA: a benchmark for general AI assistants. In *The Twelfth International Conference on Learning Representations*.
- Mitra, A., Corro, L. D., Zheng, G., Mahajan, S., Rouhana, D., Codas, A., Lu, Y., ge Chen, W., Vrousos, O., Rosset, C., Silva, F., Khanpour, H., Lara, Y., and Awadallah, A. (2024). Agentinstruct: Toward generative teaching with agentic flows.
- Pryzant, R., Iyer, D., Li, J., Lee, Y., Zhu, C., and Zeng, M. (2023). Automatic prompt optimization with “gradient descent” and beam search. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7957–7968, Singapore. Association for Computational Linguistics.
- Rasley, J., Rajbhandari, S., Ruwase, O., and He, Y. (2020). Deepspeed: System optimizations enable training deep learning models with over 100 billion parameters. *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*.
- Rastogi, A., Zang, X., Sunkara, S., Gupta, R., and Khaitan, P. (2020). Towards scalable multi-domain conversational agents: The schema-guided dialogue dataset. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8689–8696.
- Uesato, J., Kushman, N., Kumar, R., Song, F., Siegel, N., Wang, L., Creswell, A., Irving, G., and Higgins, I. (2022). Solving math word problems with process- and outcome-based feedback.
- Wei, J., Yang, C., Song, X., Lu, Y., Hu, N., Huang, J., Tran, D., Peng, D., Liu, R., Huang, D., Du, C., and Le, Q. V. (2024). Long-form factuality in large language models.
- Wu, Q., Bansal, G., Zhang, J., Wu, Y., Li, B., Zhu, E., Jiang, L., Zhang, X., Zhang, S., Liu, J., Awadallah, A. H., White, R. W., Burger, D., and Wang, C. (2024a). Autogen: Enabling next-gen llm applications via multi-agent conversation framework. In *COLM*.
- Wu, T., Yuan, W., Golovneva, O., Xu, J., Tian, Y., Jiao, J., Weston, J., and Sukhbaatar, S. (2024b). Meta-rewarding language models: Self-improving alignment with LLM-as-a-meta-judge. *arXiv preprint arXiv:2407.19594*.
- Wu, Z., Han, C., Ding, Z., Weng, Z., Liu, Z., Yao, S., Yu, T., and Kong, L. (2024c). Os-copilot: Towards generalist computer agents with self-improvement.
- Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., and Cao, Y. (2023). ReAct: Synergizing reasoning and acting in language models.
- Yuksekgonul, M., Bianchi, F., Boen, J., Liu, S., Huang, Z., Guestrin, C., and Zou, J. (2024). Textgrad: Automatic “differentiation” via text.

- Zhang, S., Zhang, J., Liu, J., Song, L., Wang, C., Krishna, R., and Wu, Q. (2024). Training language model agents without modifying language models. *ICML'24*.
- Zhou, W., Jiang, Y. E., Li, L., Wu, J., Wang, T., Qiu, S., Zhang, J., Chen, J., Wu, R., Wang, S., Zhu, S., Chen, J., Zhang, W., Tang, X., Zhang, N., Chen, H., Cui, P., and Sachan, M. (2023a). Agents: An open-source framework for autonomous language agents.
- Zhou, W., Ou, Y., Ding, S., Li, L., Wu, J., Wang, T., Chen, J., Wang, S., Xu, X., Zhang, N., Chen, H., and Jiang, Y. E. (2024). Symbolic learning enables self-evolving agents.
- Zhou, Y., Muresanu, A. I., Han, Z., Paster, K., Pitis, S., Chan, H., and Ba, J. (2023b). Large language models are human-level prompt engineers. In *The Eleventh International Conference on Learning Representations*.

A Multi-Agent Code Example

```

import json
from tapeagents.agent import Agent, Node
from tapeagents.core import Prompt, SetNextNode, Call, Respond
from tapeagents.dialog_tape import DialogTape, AssistantThought, ToolCalls, UserStep,
    AssistantStep
from tapeagents.llms import LLMStream, LiteLLM
from tapeagents.prompting import view_to_messages
from tapeagents.orchestrator import main_loop
from tapeagents.tools.simple_browser import SimpleTextBrowser
from tapeagents.tools.stock import get_stock_data, get_stock_ticker
from tapeagents.environment import ToolEnvironment
# to visualize Tapes in Python notebooks
from tapeagents.rendering import render_tape_with_prompts
from tapeagents.renderers.camera_ready_renderer import CameraReadyRenderer
from IPython.display import HTML, clear_output

##### Search Agent #####

# define environment of search agent
browser = SimpleTextBrowser()
search_agent_env = ToolEnvironment([
    browser.get_search_results, browser.get_page, browser.get_next_page
])

# define the main node of the search agent
class SearchAgentMainNode(Node):
    def make_prompt(self, agent, tape: DialogTape) -> Prompt:
        view = agent.compute_view(tape)
        search_system_message = {
            "role": "system",
            "content": ""Use at most 5 tool calls to search the request info on on the web.""
        }
        return Prompt(
            messages=[search_system_message] + view_to_messages(view.top, agent),
            tools=search_agent_env.get_tool_schema_dicts()
        )

    def generate_steps(self, agent, tape, llm_stream: LLMStream):
        m = llm_stream.get_message()
        if m.content:
            # if the LLM responds, yield Respond(..) as your last step. This special step tells the
            # orchestrator that this agent is done. The next active agent will be the one that `Call`
            # ed it (the Analyst Agent in this case).
            yield Respond(content=m.content)
        elif m.tool_calls:
            # while the LLM suggests tool calls, yield them as action steps
            yield ToolCalls.from_llm_output(m)
            yield SetNextNode(next_node=0) # set yourself as next node for when you get called again
        else:
            raise ValueError()

# create the search agent

```

```

search_agent = Agent.create(name="search_agent", llms=LiteLLM(model_name="gpt-4o"), nodes=[
    SearchAgentMainNode()])

##### Analyst Agent #####

# define environment of analyst agent

# We will use the tool choice mechanism to let the main agent call its search specialist agent.
# To this end, we create a mock tool that represents calling the search agent.
def call_search_agent(query: str):
    """Use this tool to ask a fellow AI agent to search for information on the web."""
    pass

main_agent_env = ToolEnvironment([
    get_stock_ticker, get_stock_data, call_search_agent
])

today = "2024-09-17" # fixed date for reproducible tests
system_message = {"role": "system", "content": f"""
You will help the user to learn about financials of companies. For general user queries, include
some info about stock price changes during the last year, as well as some general
information on the company. Today is {today}.
"""}

# define the two nodes of the main agent
class PlanNode(Node):
    def make_prompt(self, agent, tape) -> Prompt:
        view = agent.compute_view(tape)
        guidance_message = {
            "role": "user",
            "content": """Write a natural language plan on how to use tools help the user. Output a
list of numbered items, like 1., 2., 3., etc."""
        }
        return Prompt(
            messages=[system_message] + view_to_messages(view.top, agent) + [guidance_message],
            tools=main_agent_env.get_tool_schema_dicts(),
        )

    def generate_steps(self, agent, dialog, llm_stream: LLMStream):
        # the PlanNode should only yield thoughts based on the llm output
        if content := llm_stream.get_message().content:
            yield AssistantThought(content=content)
        else:
            raise ValueError()

class ActNode(Node):
    def make_prompt(self, agent, tape: DialogTape) -> Prompt:
        view = agent.compute_view(tape)
        guidance_message = {
            "role": "user",
            "content": """Follow the plan you created to earlier. When you are done, respond to the
user."""
        }
        return Prompt(

```

```

    messages=[system_message] + view_to_messages(view.top, agent) + [guidance_message],
    tools=main_agent_env.get_tool_schema_dicts(),
)

def generate_steps(self, agent, dialog, llm_stream: LLMStream):
    m = llm_stream.get_message()
    if m.content:
        # show the llm output as your response and give back control to plan node
        yield SetNextNode(next_node=0) # set next node to Plan node
        yield AssistantStep(content=m.content) # show the llm output
    elif m.tool_calls:
        yield SetNextNode(next_node=1) # set next node to Act node (self) until we get tool_calls
        # only keep the tool calls before the call to another agent
        agent_call = None
        for i, tc in enumerate(m.tool_calls):
            if tc.function.name == "call_search_agent":
                agent_call = tc
                m.tool_calls = m.tool_calls[:i]
                break
        # either produce the ToolCalls action OR call another agent
        if m.tool_calls:
            yield ToolCalls.from_llm_output(m)
        else:
            assert agent_call and agent_call.function.name == "call_search_agent"
            yield Call(
                agent_name="search_agent",
                content=json.loads(agent_call.function.arguments)["query"]
            )
        else:
            raise ValueError()

# define the main (root) agent
multi_agent_analyst = Agent.create(name="analyst", subagents=[search_agent.clone()], llms=LiteLLM
    (model_name="gpt-4o"), nodes=[PlanNode(), ActNode()])

# define the starting point: a DialogTape with only 1 UserStep
start_tape = DialogTape(steps=[UserStep(content="Tell_me_about_Vulcan_in_3_sentences")])

# define the whole environment as the combination of all agents & subagents environments
whole_env = ToolEnvironment([
    get_stock_ticker, get_stock_data, browser.get_search_results, browser.get_page, browser.
    get_next_page
])

# Main loop executing the `multi_agent_analyst` on the `start_tape` in the `whole_env`
environment.
for event in main_loop(multi_agent_analyst, start_tape, whole_env):
    # `event`s are all types of steps yielded by the `multi_agent_analyst` when running on the `
    start_tape` in the `whole_env` environment.
    if new_tape := event.agent_tape or event.env_tape:
        # show a fresh render every time when the environment finishes reacting with new actions
        clear_output()
        display(HTML(render_tape_with_prompts(new_tape, CameraReadyRenderer()))))

```

B Agentic RAG Code Examples

Listing 1: *add_demos* function

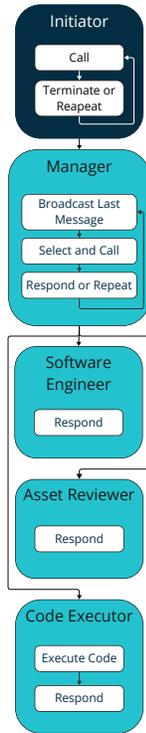
```
def add_demos(agent: Agent, tapes: list[Tape], max_n_demos: int, seed: int = 1):
    """Extract demos for function templates from the given tapes.
    When there are too many demos, select random ones.
    """
    demos = {template_name: [] for template_name in agent.templates}
    for tape in tapes:
        for node, index in agent.get_node_runs(tape):
            if isinstance(node, LLMFunctionNode):
                demos[node.template_name].append(node.extract_demo(agent, tape, index))
    rng = random.Random(seed)
    agent_copy = agent.model_copy(deep=True)
    for template_name, template in agent_copy.templates.items():
        k = min(max_n_demos, len(demos[template_name]))
        template.demos = rng.sample(demos[template_name], k)
    return agent_copy
```

Listing 2: *LLMFunction* templates

```
# Create a template for answering questions
def make_answer_template() -> LLMFunctionTemplate:
    return LLMFunctionTemplate(
        desc="Answer_questions_with_short_factoid_answers.",
        inputs=[
            ContextInput(name="context", desc="may_contain_relevant_facts", separator="\n"),
            # The question to be answered
            Input(name="question"),
        ],
        outputs=[
            # Rationale for the answer
            RationaleOutput.for_output("answer"),
            # The actual answer
            AssistantOutput(name="answer", desc="often_between_1_and_5_words")
        ]
    )

# Create a template for generating search queries with a retrieval tool
def make_query_template() -> LLMFunctionTemplate:
    return LLMFunctionTemplate(
        desc="Write_a_simple_search_query_that_will_help_answer_a_complex_question.",
        inputs=[
            ContextInput(name="context", desc="may_contain_relevant_facts", separator="\n"),
            # The question for which a search query is needed
            Input(name="question"),
        ],
        outputs=[
            # Rationale for the generated query
            RationaleOutput.for_output("query"),
            # The generated search query, to be used with a retrieval tool
            ToolCallOutput(name="query", tool_name="retrieve", arg_name="query")
        ]
    )
```

C Agent Tree and Tape



Metadata
 ▶ id: 5b2e59c8-7ca9-4bc8-b597-c73ee81c2e53

Steps

[0] Initiator calls Manager
 ▶ Make a plot comparing the stocks of ServiceNow and Salesforce...

[1] Manager broadcasts
 to: SoftwareEngineer, CodeExecutor, AssetReviewer

[2] Manager calls SoftwareEngineer
 ▶ Prompt 113 tokens | Completion 2 tokens

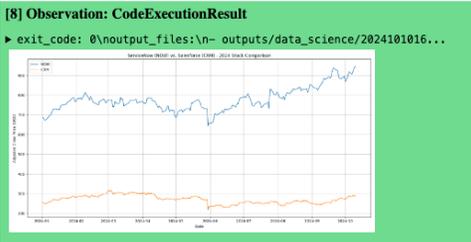
[3] SoftwareEngineer responds to Manager
 ▶ To create a plot comparing the stocks of ServiceNow and Sale...

[4] Thought: SetNextNode(0)

[5] Manager broadcasts
 to: CodeExecutor, AssetReviewer

[6] Manager calls CodeExecutor
 ▶ Prompt 522 tokens | Completion 2 tokens

[7] Action: ExecuteCode
 ▶ ```sh\n# filename: install.sh\n#!/bin/sh\n\npip install matplotlib...

[8] Observation: CodeExecutionResult
 ▶ exit_code: 0\noutput_files:\n- outputs/data_science/2024101016...


[9] CodeExecutor responds to Manager
 ▶ I ran the code `outputs/data_science/202410165546/install...`

[10] Thought: SetNextNode(0)

[11] Manager broadcasts
 to: SoftwareEngineer, AssetReviewer

[12] Manager calls AssetReviewer
 ▶ Prompt 612 tokens | Completion 2 tokens

[13] AssetReviewer responds to Manager
 ▶ Prompt 619 tokens | Completion 86 tokens
 ▶ - Currently, the plot directly charts the absolute adjusted ...

[14] Thought: SetNextNode(0)

[15] Manager broadcasts
 to: SoftwareEngineer, CodeExecutor

[16] Manager calls SoftwareEngineer
 ▶ Prompt 709 tokens | Completion 2 tokens

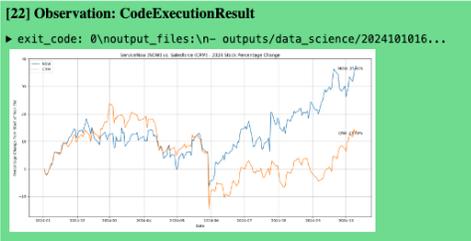
[17] SoftwareEngineer responds to Manager
 ▶ Let's enhance the plot to include a relative performance bas...

[18] Thought: SetNextNode(0)

[19] Manager broadcasts
 to: CodeExecutor, AssetReviewer

[20] Manager calls CodeExecutor
 ▶ Prompt 1167 tokens | Completion 2 tokens

[21] Action: ExecuteCode
 ▶ ```python\n# filename: compare_stocks_percentage.py\nimport yf...

[22] Observation: CodeExecutionResult
 ▶ exit_code: 0\noutput_files:\n- outputs/data_science/2024101016...


[23] CodeExecutor responds to Manager
 ▶ I ran the code `outputs/data_science/202410165546/compare_...`

[24] Thought: SetNextNode(0)

[25] Manager broadcasts
 to: SoftwareEngineer, AssetReviewer

[26] Manager calls AssetReviewer
 ▶ Prompt 1248 tokens | Completion 2 tokens

[27] AssetReviewer responds to Manager
 TERMINATE

[28] Manager responds to Initiator

[29] Action: FinalStep
 reason: Termination message received

Figure 7: A multi-agent tree configuration showing nodes (left) and a tape resulting from their collaboration (right) with color-coded steps: yellow for external agent thoughts (enabling collaboration), purple for internal agent thoughts, blue for actions, and green for observations. The step’s author is indicated in grey using the “Agent.node” format.

D Tape Tools

Raw Tape content
Press Enter to rerender the tape

```

metadata:
id: 67103d7e-a354-4c3b-9c72-080019bf07fb
parent_id: c0e7a438-1a40-47f6-b7cb-1dbe00fd96e7
author: ""
author_tape_id: null
n_added_steps: 7
error: null
result: null
context: null
steps:
- prompt_id: 8dca5402-bee0-4f1d-8b30-8e35016107f1
  task: call
  by: UserProxy
  role: call
  content: Make a plot comparing the stocks of ServiceNow and Salesforce since beginning of 2024. Save it to a PNG file.
  agent_name: GroupChatManager
- prompt_id: 5bb44b5d-3bee-4a18-86dd-73c2dcd17779
  task: broadcast_last_message
  by: UserProxy/GroupChatManager
  content: Make a plot comparing the stocks of ServiceNow and Salesforce since beginning of 2024. Save it to a PNG file.
            
```

Pop N last steps
Press Enter to proceed

Keep N first steps
Press Enter to proceed

Load tape by id
Press Enter to load tape

Choose the tape renderer

full

Metadata
► Show / Hide

Steps

[1] **UserProxy calls GroupChatManager**
Make a plot comparing the stocks of ServiceNow and Salesforce since beginning of 2024. Save it to a PNG file.

[2] **Thought: Broadcast**
by: UserProxy/GroupChatManager
content: Make a plot comparing the stocks of ServiceNow and Salesforce since beginning of 2024. Save it to a PNG file.
from_: UserProxy
to:
- CodeExecutor
- SoftwareEngineer

► Prompt 494 characters
► Completion

[3] **GroupChatManager calls SoftwareEngineer**

► Prompt 2403 characters
► Completion

[4] **SoftwareEngineer responds to GroupChatManager**
To help you plot and compare the stock prices of ServiceNow (ticker: NOW) and Salesforce (ticker: CRM) since the beginning of 2024, we need to:

1. Gather the stock price data for both companies starting from January 1, 2024.
2. Use Python (with pandas, yfinance, and matplotlib libraries) to process the data and create the plot.
3. Save the plot to a PNG file.

We will start by writing a Python script to accomplish these steps.

```

...python
# filename: stock_comparison.py

import yfinance as yf
import pandas as pd
import matplotlib.pyplot as plt

# Define the stock tickers and date range
tickers = ['NOW', 'CRM']
start_date = '2024-01-01'
end_date = pd.Timestamp.today().strftime("%Y-%m-%d")
            
```

Agent Org Chart

```

UserProxy
- GroupChatManager
- SoftwareEngineer
- CodeExecutor
            
```

Agent configuration
Press Enter to update the agent

```

name: UserProxy
llms: {}
subagents:
- name: GroupChatManager
  llms:
  default:
  model_name: gpt-4o
  parameters:
  timeout: 15.0
  context_size: 32000
  tokenizer_name: ""
  use_cache: false
subagents: []
- name: SoftwareEngineer
  llms:
  default:
  model_name: gpt-4o
  parameters:
  timeout: 15.0
  context_size: 32000
  tokenizer_name: ""
  use_cache: false
subagents: []
templates:
system: "You are a helpful AI assistant.\nSolve tasks using your coding and\n\ language skills.\n\nIn the following cases, suggest python code (in a\n python\n\ coding block) or shell script (in a sh coding block) for the user to\n execute.\n\n\ 1. When you need to collect info, use the code to output the info\n you\n\ need, for example, browse or search the web, download/read a file,\n print()
            
```

Figure 8: *TapeAgents Studio: Application to help AI Admin to edit Tape, resume and debug Agentic Systems*

Context

Steps

Question:
If Eliud Kipchoge could maintain his record-making marathon pace indefinitely, how many thousand hours would it take him to run the distance between the Earth and the Moon its closest approach? Please use the minimum perigee value on the Wikipedia page for the Moon when carrying out your calculation. Round your result to the nearest 1000 hours and do not use any comma separators if necessary.

▶ **Prompt 2064 tokens**

Agent Thought: Plan
plan:
- Look up Eliud Kipchoge's record marathon pace in minutes per mile or kilometers per hour.
- Look up the minimum perigee value (closest approach) of the Moon from its Wikipedia page.
- Convert Eliud Kipchoge's marathon pace to a consistent unit (e.g., kilometers per hour).
- Calculate the time it would take to run the distance between the Earth and the Moon at the closest approach using the converted pace.
- Convert the time from hours to thousand hours.
- Round the result to the nearest 1000 hours.
- Provide the final answer.

▶ **Prompt 2167 tokens**

Agent Thought: List Of Facts

▶ **Prompt 2310 tokens**

Agent Action: Search
source: web_search
query: Eliud Kipchoge record marathon pace

▶ **Observation: Search Results**

▶ **Prompt 2901 tokens**

Agent Thought: New Fact
fact_name: eliud_kipchoge_marathon_pace
unit: min/km
value: 2.5

▶ **Prompt 2951 tokens**

Agent Action: Search
source: wikipedia
query: minimum perigee value of the Moon

▶ **Observation: Search Results**

▶ **Prompt 3466 tokens**

Agent Thought: New Fact
fact_name: minimum_perigee_moon
unit: km
value: 362600

▶ **Prompt 3510 tokens**

Agent Action: Convert Fact

Observation: Calculation Result
name: eliud_kipchoge_speed
result: 24.0

▶ **Prompt 3689 tokens**

Agent Action: Use Calculator

Observation: Calculation Result
name: time_to_run_to_moon
result: 15108.333333333334

▶ **Prompt 3841 tokens**

Agent Action: Use Calculator

Observation: Calculation Result
name: time_to_run_to_moon_thousand_hours
result: 15000.0

▶ **Prompt 4018 tokens**

Agent Action: Plan Finished
success: True
overview: Calculated the time it would take Eliud Kipchoge to run the distance between the Earth and the Moon at its closest approach.
answer_unit: thousand hours
answer: 15000

Answer:
answer: 15000
success: True
overview: Calculated the time it would take Eliud Kipchoge to run the distance between the Earth and the Moon at its closest approach.

Reload Tapes

File
ll_gpt-4o-2024-05-13_2024-07-02_shortpwn

Accuracy 50.94%, 27 out of 53

Tape
LL1: If Eliud Kipchoge could maintain

Tape Result
Golden Answer: 17
Agent Answer: 15000
Steps: 17
Failures: 0
Finished successfully: True
LLM Calls: 10
Overview:
Calculated the time it would take Eliud Kipchoge to run the distance between the Earth and the Moon at its closest approach.

Figure 9: *Tape Browser: Application to inspect a batch of tapes result. This tape is a GAIA task where the Agent did not provide the right answer. No step failed during the session.*

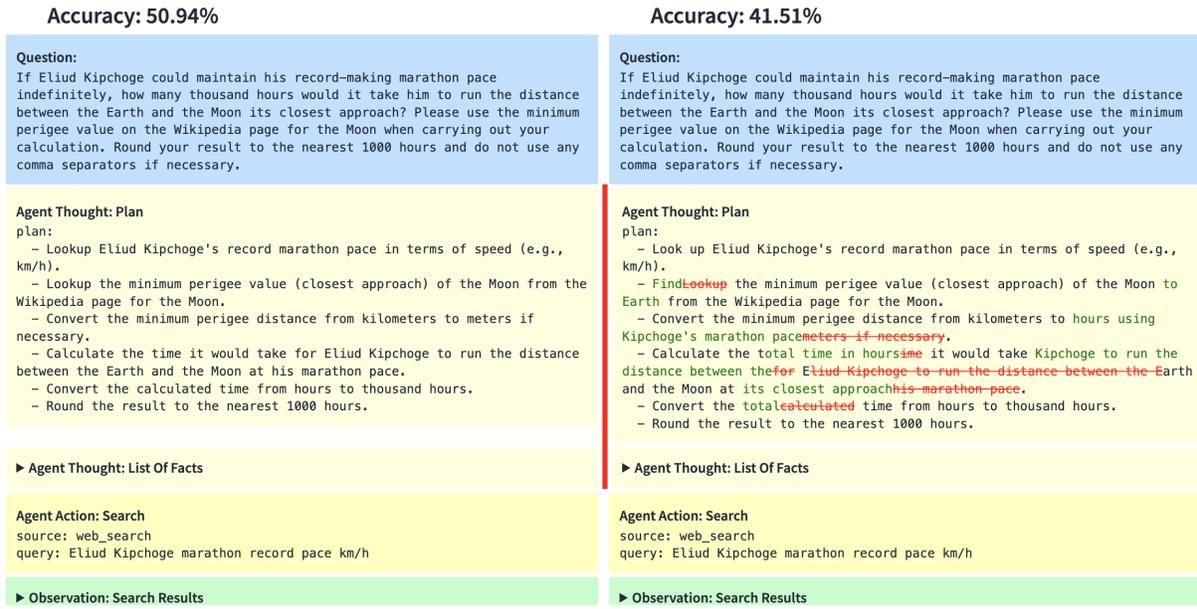


Figure 10: Tape Diff: Application to compare two batches of tape.

E GREADTH Form Filler

E.1 Virtual Companies Prompts

To generate synthetic environments/companies with 10 request forms available in each we use a three-step prompting method with Llama-3-70B.

1. The first step is to generate a description of a fake company. For this, we use the following prompt:

```
messages = [  
  {  
    "role": "system",  
    "content": "You are a helpful assistant."  
  },  
  {  
    "role": "user",  
    "content": f"Give me a description of {real_name} but replace all occurrences of  
      `{real_name}` by `{fake_name}`."  
  }  
]
```

with `real_name` and `fake_name` being set to different company names (e.g. “Starbucks” and “CoffeeCorp” respectively). This step produces a `DESCRIPTION` variable that will be used in the next step.

2. The second step is to generate a list of 10 request forms for each fake company. For now, we only ask the model to generate a name and a description for each form with the following system and user prompts:

```
SYSTEM_MESSAGE = """  
You are a helpful enterprise assistant who is very well integrated into the  
internal system of [ENTERPRISE_NAME]. You have access to the database of  
[ENTERPRISE_NAME], which contains REQUEST_FORMS. REQUEST_FORMS can be used  
by employees and clients to submit requests and trigger automations.  
REQUEST_FORMS have the following data format:  
```python  
class FunctionSchema(BaseModel):
 name: FunctionName
 description: str = Field(description="The description of the function.")
 parameters: JsonSchema = Field(default=None, description="The JSON schema
 of the function's parameters.")
 return_value: JsonSchema = Field(default=None, description="The JSON
 schema of the function's return value.")
...
"""
messages = [
 {"role": "system", "content": SYSTEM_MESSAGE},
 {"role": "user", "content": f"""
 [ENTERPRISE_NAME] is '{fake_name}'.
 {DESCRIPTION}
 Give me the name and description of 10 REQUEST_FORMS that are often used
 at {fake_name}. Make sure the REQUEST_FORMS are specific to {fake_name}
 use cases. Format the output as a JSON list of JSON dictionaries with
 'name' and 'description' keys. Make sure the output is JSON parsable.
 """}],
]
```

with `fake_name` and `DESCRIPTION` defined in the previous step. The output is a list of 10 request form names and descriptions. We programmatically verify that the output is JSON parsable.

3. Eventually, for each request form dictionary (`FORM`) generated previously, we prompt llama3-70B to generate the `FunctionSchema` like this:

```
messages = [
 {"role": "system", "content": SYSTEM_MESSAGE},
 {"role": "user", "content": f"""
 [ENTERPRISE_NAME] is '{fake_name}'.
 {DESCRIPTION}
 Show me the FunctionSchema of the following REQUEST_FORM:
    ```json
    {FORM}
    ```
 """
}
```

First, rephrase the description to be more specific. The description must include rules, facts, and policies at `{fake_name}` about `{FORM['name']}`, such as when to fill this request, who will process the request, etc... The description must also include information about the parameters to help the user fill the form. For multiple choice parameters (enum), the description must explain the difference between each possible value.

Once the description is written, remove all colon ( : ) characters from it.

Second, write the `yaml` format of this `FunctionSchema` by replacing the description with your detailed version. Parameters cannot be nested objects, but return values can. The `FunctionSchema` must contain both required and optional parameters.

Write the output in `yaml` format. The output must be parsable as a `FunctionSchema`.

```
"""},
]
```

with `SYSTEM_MESSAGE`, `fake_name`, `DESCRIPTION`, and `FORM` defined in the previous steps. We verify programmatically that the output is parsable as a `FunctionSchema`. An example of generated `FunctionSchema` can be seen in Step [5] of the Tapes in Figures 11 & 12 (Appendix E.5).

## E.2 List of Form-Filler Agent Thoughts and Actions

**Form-filler thoughts.** Below, we list the available thoughts for our Form-Filler TapeAgent and explain when they should be used:

- `AnswerFromFunctionSchemaThought`: This thought premeditates answering the user's question that the agent knows how to answer based on the form description. Attributes:
  - `function`: The name of the function to answer a question about.
- `NoAnswerInFunctionSchemaThought`: This thought premeditates informing the user that the agent does not know the answer to their question. Attributes:
  - `function`: The name of the function.
- `RefuseInexistentFunctionThought`: A thought that indicates that the query could not be resolved to a function.
- `RefuseInvalidFunctionParameterValueThought`: This thought premeditates informing the user that a function parameter value is invalid. Attributes:
  - `function`: The name of the function.

- parameter: The name of the parameter the user tried to set a value for.
- parameter\_value: The value the user tried to set for the parameter.
- **RefuseSkippingParameterThought:** This thought premeditates informing the user that a function parameter cannot be skipped because it is required. Attributes:
  - function: The name of the function.
  - parameter: The name of the required parameter that the user tried to skip.
- **GatherValuesThought:** This thought records extracted parameters from the user’s message and is mainly used in the multi-node Teacher agent. Attributes:
  - function: The name of the function.
  - parameters: Dictionary mapping parameter names to their extracted values.
- **VerifyValuesThought:** This thought records the validity status of parameter values extracted from the user’s message and is mainly used in the multi-node Teacher agent. Attributes:
  - function: The name of the function.
  - parameters: Dictionary mapping parameter names to their value, their validity, and the explanation when invalid.
- **UpdateFunctionParametersThought:** This thought updates the values of the parameters of a function based only on the user’s last message. It is used when the user provides NEW information about the parameters of a function or when the user wants to skip an optional parameter. Attributes:
  - function: The name of the function.
  - assign: The dictionary assignment of parameter names to their NEW values (optional, default: {}). If no NEW values are provided, set ‘assign’ to {}.
  - skip: The list of NEW optional parameter names to skip (optional, default: []). If no NEW parameters are skipped, set ‘skip’ to [].

The ‘assign’ attribute sets or updates parameter values based on the last user message. It is only applied to parameters that are newly filled or have updated values. Parameters that remain unchanged will not have the ‘assign’ attribute. The ‘skip’ attribute is set for optional parameters when the user wants to ignore or skip an optional parameter. It is only applied to new optional parameters that the user wants to skip and will not be used for parameters that were already skipped.

- **RequestFunctionThought:** This thought premeditates requesting the user to select a function. Attributes:
  - functions: The list of available functions the user can select.
- **RequestFunctionParametersThought:** This thought premeditates requesting the user to provide a value for one parameter of a given function. Attributes:
  - function: The name of the function to request parameters for.
  - parameters: A list containing only one value which is the parameter to request.
- **RequestFunctionCallConfirmationThought:** This thought is a preliminary step before calling a function. This step is ONLY used once ALL parameters are filled or skipped. Attributes:
  - function: The name of the function.
- **RequestExitConfirmationThought:** This thought is a preliminary step before exiting the dialogue. The message MUST be short and concise. The message should contain clear and explicit confirmation of the acceptance of the last value provided by the user.

**Form-filler actions.** Here are the available actions our Form-Filler TapeAgent can take.

- **ResolveFunctionAction:** An action that resolves which function candidates are compatible with a given user query. Attributes:
  - query: The query to resolve the function with, by default the user’s last message.
  - result: A list of candidate functions that are compatible with the query.
- **InspectFunctionAction** An action that inspects the schema of a function given its name. This action inspects and returns the schema of the function, which contains its description and the names and types of the function’s parameters. Attributes:
  - function: The name of the function.
  - result: The schema of the function.
- **PromptUserForTextMessageAction:** An action that asks the user to enter a message in a text input field. This action is used to capture free-form text input from the user through a standard text input field. This action must be used in every single agent’s response. Attributes:
  - prompt: The message from the agent to the user.
  - result: User message
- **CallFunctionAction:** This action calls a function with all available parameter values. Attributes:
  - function: The name of the function to call.
  - result: Result of the function call
- **ExitAction:** With this action the agent indicates that the dialogue has ended. No further steps will be executed after this step. Attributes:
  - text: The message from the agent to inform the user that the dialogue has ended.

### E.3 User Agents

**Table 6:** The list of user agents implemented for all domains. The User LLM is a prompted Llama-3.1-405B-Instruct model.

Agent Name	Behavior
UserInitMessageAmazing	At the beginning of the conversation, the user should request a specific intent and provide values for some of the parameters in that request.
UserInitMessageShort	At the beginning of the conversation, the user should request a specific intent in a short message.
UserInitMessageAsk	At the beginning of the conversation, the user should ask the agent what it can do.
UserBadInitMessage	At the beginning of the conversation, the user should ask to do something impossible.
UserHappyPath	The user correctly answers the agent question.
UserMultislotInstruct	The user replies to the agent question and provides values for additional parameters in the request.
UserMultislotFuture	The user provides values for additional parameters in the request.
UserChangesSlot	The user changes the value to a previously set request parameter.
UserAdjectiveOrPosition	For categorical parameters, the user answers the agent but instead of giving the choice name, uses an adjective or a positional number that defines its choice.
UserSkipsOptional	For optional parameters, the user asks to skip it.
UserAsksAboutDocs	The user does not answer the agent, and instead asks a question that can be answered based on the request documentation.
UserAsksAboutFilledParameters	The user does not answer the agent, and instead asks a question about the parameters filled, or still missing.
UserAsksForHelp	For categorical parameters, the user ask the agent for help in choosing a value based on its problem.
UserCorrectValueInvalidValue	The user answers the agent correctly, but also gives an invalid value for a categorical parameter not yet filled.
UserInvalidValue	For categorical parameters, the user gives an impossible choice.
UserChangesTopic	The user does not answer the agent’s question and change the topic of conversation instead.
UserSkipsUnskipable	For required parameters, the user asks to skip it.
UserIncorrectDate	For date/time parameters, the user gives an invalid date.
UserIncorrectWeb	For technical parameters such as url/email, the user gives a invalid value.

#### E.4 Choosing the teacher model

The Teacher agent uses a Llama-405B-Instruct model<sup>12</sup> that is carefully prompted to perform the form-filling task according to all our metrics (Section 5.2). We experimented with various language models such as GPT4o, and Mistral-Large, but according to human evaluation, Llama 405B was on par or better than other models. We explored two different prompting techniques. The first (referred to as ‘long-node’ in Table 7) consists of crafting one long and detailed prompt explaining how the agent should make a thought plan and answer the user (single node TapeAgent). The other prompting technique explored (referred to as ‘multi-node’ in Table 7) consists in decomposing the task into multiple smaller subtasks, each handled by a different node. The multi-node agent is first tasked to identify potential parameter values provided by the last user message, then it is prompted to verify whether the provided values are correct or not, eventually, the agent is prompted to note valid information, refuse invalid information, and ask the user for the next parameter to fill. We generated agent messages as a continuation to 1016 unfinished conversations where the last user message behaved according to specific behaviors and sent these messages to human labelers. Results show that the multi-node approach yields better performance overall, so we decided to keep this prompting strategy as our final Teacher agent in Section 5.3.

**Table 7:** Evaluation of the Teacher agent with two prompting methods (long-node & multi-node) across all user behaviors on 1016 test domain conversations.

Author of last user message	Count	GREADTH metric	
		long-node	multi-node
UserInitMessageAmazing	46	60.87%	86.96%
UserInitMessageShort	48	87.50%	89.58%
UserInitMessageAsk	46	91.30%	84.78%
UserBadInitMessage	46	97.83%	100.00%
UserHappyPath	186	72.58%	89.25%
UserMultislotInstruct	48	68.75%	83.33%
UserMultislotFuture	46	58.70%	76.09%
UserChangesSlot	46	71.74%	82.61%
UserAdjectiveOrPosition	46	58.70%	76.09%
UserSkipsOptional	46	71.74%	71.74%
UserAsksAboutDocs	46	84.78%	93.48%
UserAsksAboutFilledParameters	48	83.33%	81.25%
UserAsksForHelp	45	40.00%	42.22%
UserCorrectValueInvalidValue	45	62.22%	77.78%
UserInvalidValue	46	63.04%	91.30%
UserChangesTopic	46	76.09%	86.96%
UserSkipsUnskipable	48	87.50%	95.83%
UserIncorrectDate	46	47.83%	19.57%
UserIncorrectWeb	42	16.67%	23.81%
<b>ALL_BEHAVIORS</b>	<b>1016</b>	<b>69.39%</b>	<b>78.54%</b>

<sup>12</sup><https://huggingface.co/meta-llama/Llama-3.1-405B-Instruct-FP8>

## E.5 Teacher and Student Tapes



Figure 11: Sample Tape between the Teacher agent and User agents. The user is requesting the creation of a digital coupon for their online store. Steps are color-coded: purple for internal agent thoughts, blue for actions, and green for observations. The step's author is indicated in grey using the "Agent.node" format. User agent names are described in Appendix E.3.



Figure 12: Sample tape between the Student agent and User agents. The user is requesting the creation of a digital coupon for their online store. Steps are color-coded: purple for internal agent thoughts, blue for actions, and green for observations. The step's author is indicated in grey using the "Agent.node" format. User agent names are described in Appendix E.3.