
SYMBOLIC LEARNING ENABLES SELF-EVOLVING AGENTS

Anonymous authors

Paper under double-blind review

ABSTRACT

The AI community has been exploring a pathway to artificial general intelligence (AGI) by developing “language agents”, which are complex large language models (LLMs) workflows involving both prompting techniques and tool usage methods. While language agents have demonstrated impressive capabilities for many real-world tasks, a fundamental limitation of current language agents research is that they are model-centric, or engineering-centric. That’s to say, the progress on prompts, tools, and workflows of language agents requires substantial manual engineering efforts from human experts rather than automatically learning from data. We believe the transition from model-centric, or engineering-centric, to data-centric, i.e., the ability of language agents to autonomously learn and evolve in environments, is the key for them to possibly achieve AGI.

In this work, we introduce *agent symbolic learning*, a systematic framework that enables language agents to optimize themselves on their own in a data-centric way using *symbolic optimizers*. Specifically, we consider agents as symbolic networks where learnable weights are defined by prompts, tools, and the way they are stacked together. Agent symbolic learning is designed to optimize the symbolic network within language agents in a *data-centric* way by mimicking two fundamental algorithms in connectionist learning: back-propagation and gradient descent. Instead of dealing with numeric weights, agent symbolic learning works with text-based weights, loss, and gradients. We conduct proof-of-concept experiments on both standard benchmarks and complex real-world tasks and show substantial improvements over static agent frameworks and simple prompt/tool optimization methods. In addition, agent symbolic learning enables language agents to update themselves after being created and deployed in the wild, resulting in “self-evolving agents”. We will open-source the agent symbolic learning framework to facilitate future research on *data-centric* agent learning.

1 INTRODUCTION

Recent advances in large language models (Radford et al., 2018; 2019; Brown et al., 2020; Ouyang et al., 2022; OpenAI, 2023; Touvron et al., 2023a;b) open the possibility of building language agents that can autonomously solve complex tasks. The common practice for developing AI agents is to decompose complex tasks into LLM workflows where prompts and tools are stacked together (Park et al., 2023; Hong et al., 2023; Zhou et al., 2023b; Chen et al., 2023b; Xie et al., 2023). In a sense, language agents can be viewed as AI systems that connect connectionism AI (i.e., the LLM backbone of agents) and symbolism AI (i.e., the workflow of prompts and tools), which partially explains their effectiveness in real-world problem-solving scenarios.

However, the current state of language agent development is limited by the extensive engineering effort required to build and customize language agent systems for a specific task. Specifically, researchers and developers have to manually decompose complex tasks into subtasks, which we refer to as nodes, that are more tractable for LLMs and then carefully design prompts and tools, including API functions, knowledge bases, memories, etc., for specific nodes. The complexity of this process makes the current landscape of language agent research *model-centric*, or *engineering-centric*. This means it is almost impossible for researchers to manually tune or optimize language agents on datasets on which we can train neural nets in a *data-centric* way. This limits the robustness and

054
055
056
057
058
059
060
061
062
063
064
065
066
067
068
069
070
071
072
073
074
075
076
077
078
079
080
081
082
083
084
085
086
087
088
089
090
091
092
093
094
095
096
097
098
099
100
101
102
103
104
105
106
107

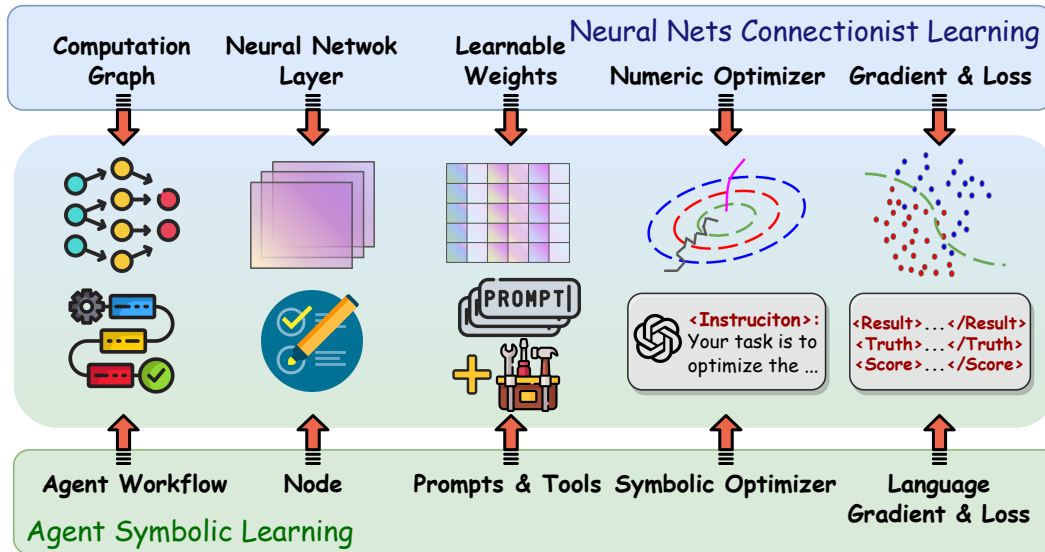


Figure 1: Analogy between agent symbolic learning and neural nets connectionist learning.

versatility of manually coded language agents and requires substantial engineering effort to adapt language agents to new tasks or data distributions. We believe the transition from engineering-centric language agent development to data-centric learning is an important step in language agent research.

To this end, a number of recent efforts have been made on automatic optimization of language agents. For example, Dspy (Khattab et al., 2023) introduces a framework for algorithmically optimizing LLM prompts via bootstrapping or random searching in a combinatorial space of different prompt components and GPTswarm (Zhuge et al., 2024) further proposes to tackle the combinatorial optimization challenge raised in DSPy via an iterative optimization process. Agent-pro (Zhang et al., 2024b) proposes a framework to optimize the components of the prompts corresponding to the agents’ internal policy in competitive environments. AgentOptimizer (Zhang et al., 2024a) proposes a framework to optimize functions with carefully engineered prompts. While effective in some scenarios, these approaches only optimize separate modules in an agent system such as a prompt for a specific node. As a result, these optimization methods are prone to local optimum of isolated prompts, tools, and nodes that lead to compromised performance for the entire agent system. This resembles the early practice in training neural nets (Hinton and Salakhutdinov, 2006) where layers are separately optimized and it now seems trivial that optimizing neural nets as a whole leads to better performance. We believe that this is also the case in agent optimization and joint optimization of all symbolic components within an agent is the key for optimizing agents.

In this work, we introduce a *agent symbolic learning* framework for training language agents. The agent symbolic learning framework is inspired by the connectionist learning procedure (Hinton, 1990) used for training neural nets. To be specific, we make an analogy between language agents and neural nets: the agent workflow of an agent corresponds to the computational graph of a neural net, a node in the agent workflow corresponds to a layer in the neural net, and the prompts and tools for a node correspond to the weights of a layer. In this way, we are able to implement the main components of connectionist learning, i.e., backward propagation and gradient-based weight update, in the context of agent training using language-based loss, gradients, and weights. We implement the loss function, back-propagation, and weight optimizers in the context of agent training with carefully designed LLM workflows. Specifically, for a training example, our framework first conducts the “forward pass” (agent execution) and stores the input, output, prompts, and tool usage in each node in a “trajectory”. We then use an LLM-based loss function to evaluate the outcome following recent LLM-as-a-judge framework (Zheng et al., 2023), resulting in a text-based loss. Then we back-propagate the text-based loss from the last to the first node along the trajectory, resulting in natural language analysis and reflection for the symbolic components within each node including the prompts and tool descriptions. We refer to these reflections and analyses as “language gradients” since they carry the same role as conventional gradients in the training of neural nets: guide the direction to which optimizers should

108 change the weights so that the overall loss is minimized. Finally, we update all symbolic components
109 in each node, as well as the computational graph consisting of the nodes and their connections,
110 according to the language gradients using LLMs with carefully designed prompts and workflows. Our
111 approach also naturally supports optimizing multi-agent systems by considering nodes as different
112 agents or allowing multiple agents to take actions in one node.

113 The agent symbolic learning framework is an agent learning framework that mimics the standard
114 connectionist learning procedure. In contrast to existing methods that either optimize single prompt
115 or tool in a separate manner, the agent symbolic learning framework jointly optimizes all symbolic
116 components within an agent system, including prompts, tools, and the workflow that stacks them
117 into an agent system. This top-down optimization scheme also enables the agent symbolic learning
118 framework to optimize the agent system “holistically”, avoiding local optimum for each separated
119 component. This makes it possible for language agents targeting complex real-world problems to
120 effectively *learn from data*, opening up the possibility to transform the current state of language agent
121 research from engineering-centric to data-centric.

122 In sum, by learning from LLM-generated critics (language-based loss) and reflections (language-
123 based gradients), the agent symbolic learning framework has the following advantages compared to
124 conventional frameworks for language agents in which the prompts, tools, and workflows are static
125 and require human expert efforts for optimization: first, agent symbolic learning enables the agent
126 system to learn from failure or unstable cases and update the prompts by adding few-shot examples
127 or principles; second, it enables the system to include new nodes (subtasks) and adjust the workflow
128 to improve the overall performance or handle some common failure patterns; third, our approach
129 enables the agent system to update the tool descriptions and implementation or implement new tools
130 for improved performance.

131 Moreover, since the language-based loss function does not require ground-truth when generating the
132 language loss and the optimization framework only requires calling of LLM APIs instead of tons of
133 GPUs, our framework enables language agents to *learn from experience* and *actively* update all their
134 symbolic components after being created and deployed in the wild, enabling “self-evolving agents”¹.
135 We believe this could be very helpful in the pursuit of artificial general intelligence.

136 As a proof-of-concept, we conduct a series of experiments on both standard LLM benchmarks and
137 complex agentic tasks. Our results demonstrate the effectiveness of the proposed agent symbolic
138 learning framework on optimizing and designing prompts and tools, as well as updating the overall
139 agent workflow, by data-centric learning. We will open-source all codes and prompts in the agent
140 symbolic learning framework to facilitate future research on *data-centric* agent learning.

142 2 RELATED WORK

144 2.1 LANGUAGE MODELS, PROMPTS, AND LANGUAGE AGENTS

146 Language model is a family of machine learning model that is trained to evaluate the probability
147 of sequences of words or tokens. Large language models (LLMs) (Radford et al., 2018; 2019;
148 Brown et al., 2020; Ouyang et al., 2022; OpenAI, 2023; Touvron et al., 2023a;b) often refer to
149 language models that adopt the autoregressive probability factorization scheme, parametrized by
150 the Transformer architecture (Vaswani et al., 2017), consists of a large amount of parameters, and
151 trained on large-scale corpus. With scaling of model size, training data, and computation, LLMs have
152 demonstrated remarkable capabilities in generating human-like texts and understanding context.

153 Prompts, on the other hand, is the key for unleashing the capabilities of LLMs. Prompts are critical
154 components in controlling the behavior and output of LLMs and serve as the interface between human
155 and LLMs. The design of prompts significantly impacts the performance of language models and a
156 number of progress have been made on prompt engineering, including in-context learning (Brown
157 et al., 2020), chain-of-thought prompting (Nye et al., 2022; Wei et al., 2022), ReAct (Yao et al., 2022),
158 self-refine (Madaan et al., 2023), self-consistency (Wang et al., 2023), recurrent prompting (Zhou
159 et al., 2023a), etc.

160 ¹Agents can also collect training data in the wild and update the LLM backbone via fine-tuning. In this way,
161 all components in the agent can be updated. We leave this for future work.

162 Language agents further extend the functionality of language models beyond simple prompting by
163 allowing LLMs to use tools (Schick et al., 2023) and integrating LLMs into broader systems capable
164 of executing multi-step tasks (Park et al., 2023; Hong et al., 2023; Zhou et al., 2023b; Chen et al.,
165 2023b; Xie et al., 2023). By stacking prompts and tools into carefully designed workflows, agents are
166 versatile in various applications, from customer service automation to advanced data analysis.

168 2.2 FROM AUTOMATED PROMPT ENGINEERING TO AGENT OPTIMIZATION

169
170 With the increasing popularity of prompt engineering in both academic and industry, a number of
171 recent work investigated methods to automate the prompt engineering process. For example, Pryzant
172 et al. (2020) and Yang et al. (2024) uses carefully designed prompts to unleash LLMs’ ability to do
173 prompt engineering for themselves. On the other hand, Prasad et al. (2023) and Guo et al. (2024)
174 employs different search algorithms such as genetic algorithms for prompt optimization.

175 Since prompts are critical components of agents, the success of automated prompt engineering
176 opens up the possibility of automated agent optimization. Similar to the case in automated prompt
177 engineering, methods for agent optimization can also be categorized into two categories: *prompt-*
178 *based* and *search-based*. For example, Agent-pro (Zhang et al., 2024b) and AgentOptimizer (Zhang
179 et al., 2024a) leverage carefully designed prompts to optimize either the prompts or the tools in a node
180 of the agent workflow. These methods work on isolated components within an agent. Another line of
181 research explored search-based agent optimization algorithms. Sordoni et al. (2023) uses variational
182 inference to optimize stacked LLMs. Dspy (Khattab et al., 2023) uses search algorithms to find the
183 best prompts or nodes in a combinatorial space. GPTSwarm (Zhuge et al., 2024) further improved
184 the search algorithm for the combinatorial optimization problem. These approaches have a few major
185 limitations. First, the search algorithm mainly works when the metric can be defined numerically
186 with equations that can be coded. However, most agentic tasks are real-world complex problems of
187 which the success can not be defined by some equations, such as software development or creative
188 writing. Second, these approaches update each component separately and therefore suffer from the
189 local optimum of each node or component. These approaches also lack the functionality of adding
190 nodes in the workflow or implementing new tools. Our proposed agent symbolic learning framework,
191 on the other hand, is the first agent learning method that optimize the agent system “holistically” and
192 is able to optimize prompts, tools, nodes, as well as the way they are stacked into agents.

193 Furthermore, a number of recent efforts have been done on synthesizing data to fine-tune the LLM
194 backbone of an agent (Chen et al., 2023a; Qiao et al., 2024; Song et al., 2024). This line of research
195 is orthogonal to our work and we believe they can be complementary to each other. ICE (Qian et al.,
196 2024) is also a related work investigating inter-task transfer learning for language agents, which can
197 be complementary with our method for building self-evolving agents.

198 3 AGENT SYMBOLIC LEARNING

199 3.1 PROBLEM FORMULATION

200
201 We first formulate the agent symbolic learning framework by drawing analogies to the components
202 and procedures used in neural network training. We define the key components of the framework and
203 explain the notations used throughout this section.

204
205 The agent symbolic learning framework, as illustrated in Figure 2 and described in Algo 1, is inspired
206 by the connectionist learning procedures used for training neural nets (Hinton, 1990). We first
207 introduce the notations for key concepts by making analogies to that in the connectionist learning
208 framework:

- 209
210
211 • **Agent Workflow \mathcal{A} :** Similar to the computational graph in neural nets that represents the
212 structure of layers and their connections, *agent workflow* represents the sequence of nodes (or
213 steps) through which the agent processes input data. A sequence of nodes $\{\mathcal{N}_1, \mathcal{N}_2, \dots, \mathcal{N}_n\}$
214 that process the input data through various stages. Note that in some agent frameworks,
215 the agent workflow is input-dependent since the nodes are dynamically assigned during
execution, which is similar to the case of dynamic neural nets.

216	Algorithm 1 Agent Symbolic Learning Framework	
217	Require: \mathcal{I}	▷ Input to the agent system
218	Require: \mathcal{A}	▷ Agent workflow with nodes
219	Require: \mathcal{G}	▷ Prompt-based gradient propagation function
220	Require: \mathcal{L}	▷ Prompt-based loss function
221	Ensure: Updated symbolic components in the agent system	
222	1: $\tau \leftarrow \square$	▷ Initialize trajectory
223	2: Forward Pass	
224	3: for each $\mathcal{N} \in \mathcal{A}$ do	
225	4: $\mathcal{I}_n \leftarrow$ Get input for \mathcal{N}	▷ Input to the node
226	5: $\mathcal{O}_n \leftarrow \mathcal{N}(\mathcal{I}_n, \mathcal{P}_n, \mathcal{T}_n)$	▷ Output from the node
227	6: Append $(\mathcal{I}_n, \mathcal{O}_n, \mathcal{P}_n, \mathcal{T}_n)$ to τ	
228	7: end for	
229	8: Loss Computation	
230	9: $\mathcal{L}_{\text{lang}} \leftarrow \mathcal{L}(\tau)$	▷ Compute language loss
231	10: Back-propagation	
232	11: for each $\mathcal{N} \in \text{reverse}(\mathcal{A})$ do	
233	12: $\nabla_{\text{lang}}^n \leftarrow \mathcal{G}(\nabla_{\text{lang}}^{n+1}, \mathcal{I}_n, \mathcal{O}_n, \mathcal{P}_n, \mathcal{T}_n, \mathcal{L}_{\text{lang}})$	▷ $\nabla_{\text{lang}}^{n+1} = \emptyset$ for the last node
234	13: Append ∇_{lang}^n to τ	
235	14: end for	
236	15: Weight Update	
237	16: for each $\mathcal{N} \in \mathcal{A}$ do	
238	17: Update $\mathcal{P}_n, \mathcal{T}_n$ using ∇_{lang}^n	▷ Update prompts and tools
239	18: end for	
240	19: Update \mathcal{A} using $\{\nabla_{\text{lang}}^n\}$	▷ Update the agent workflow
241	20: return $(\mathcal{A}, \mathcal{P}, \mathcal{T})$	▷ Updated agent system

- **Node \mathcal{N} :** An individual step within an agent workflow. The role of a node in an agent is similar to a layer in a neural network. A node \mathcal{N}_n receives **Node Input** \mathcal{I}_n , which are also in natural language form. In general, the input for a node consists of the output of the previous node and (optionally) inputs from the environment (e.g., human input). The node \mathcal{N}_n processes the input \mathcal{I}_n with an LLM using both **prompts** \mathcal{P}_n and **tools** \mathcal{T}_n ². The output \mathcal{O}_n is in natural language and passed to the next node.
- **Trajectory τ :** Similar to the role of computational graph of neural nets, the trajectory stores all information during the forward pass, including the inputs, outputs, prompts, and tools usage for each node, and is responsible for gradient back-propagation.
- **Language Loss $\mathcal{L}_{\text{lang}}$:** Language loss in the agent symbolic learning framework is similar to the loss in neural networks since they both measure the discrepancy between the expected and actual outcomes. The main difference is that the language loss is in textual form and is produced by a natural language loss function implemented by a carefully designed prompt while conventional losses are float numbers computed with loss functions that are numerical equations.
- **Language Gradient ∇_{lang} :** Similar to the role of gradients in connectionist learning, *language gradients* are textual analyses and reflections used for updating each component in the agent with respect to the language loss.

3.2 AGENT SYMBOLIC LEARNING PROCEDURE

After defining the key components, we can summarize the workflow of the agent symbolic learning framework in Algorithm 1. In this section, we describe each step in the agent symbolic learning framework in detail.

Forward Pass The forward pass is almost identical to standard agent execution. The main difference is that we store the input, prompts, tool usage, and the output to the trajectory, which is

² \mathcal{T}_n consists of the input and output for tool usage, and the implementation of the tool itself.

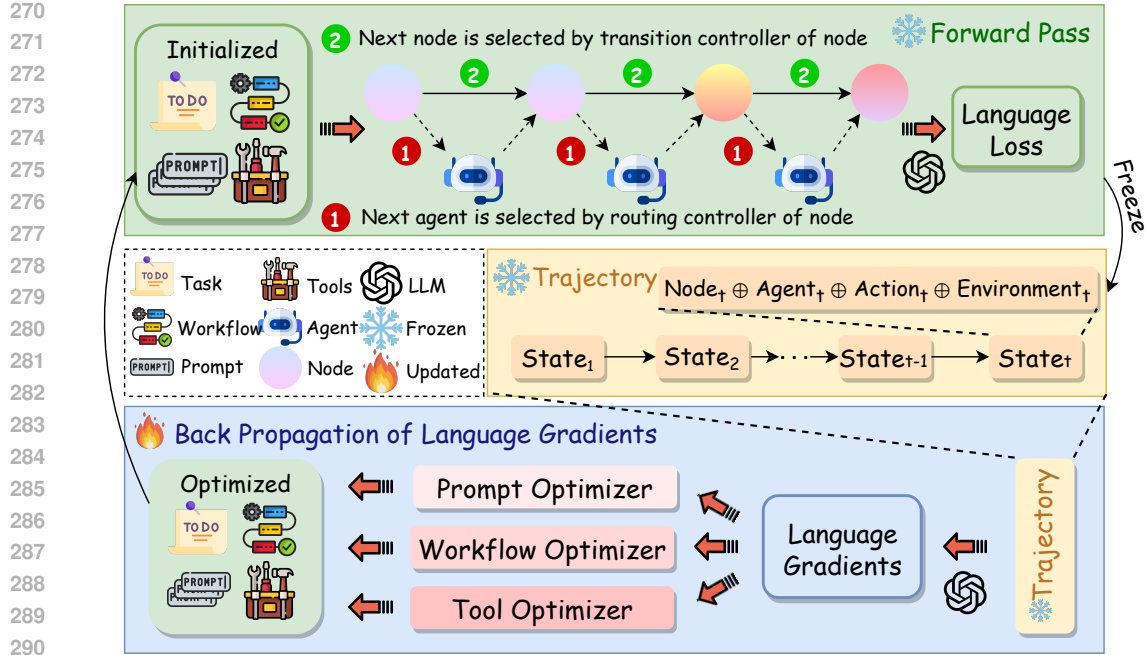


Figure 2: Illustration of the agent symbolic learning framework.

used for language gradient back-propagation. This is similar to deep learning frameworks such as PyTorch (Paszke et al., 2019) and TensorFlow (Abadi et al., 2016) that store the intermediate outputs and activation in the computation graph of the neural network.

Language Loss Computation After the forward pass, we compute the language loss for a training example by feeding the trajectory into an LLM using a carefully designed prompt template $\mathcal{P}_{\text{loss}}$:

$$\mathcal{L}_{\text{lang}} = \text{LLM}(\mathcal{P}_{\text{loss}}(\tau)) \quad (1)$$

The key is the design for the prompt template, which is expected to *holistically* evaluate how the agent performs with respect to the input, environment, and task requirements. To this end, we carefully design a prompt template for language loss computation consisting of the following components: task description, input, trajectory, few-shot demonstrations, principles, and output format control. Among them, task description, input, and trajectory are data-dependent while the few-shot demonstrations, principles, and output format control are fixed for all tasks and training examples. The language loss consists of both natural language comments and a numerical score (also generated via prompting). We can optionally feed the ground-truth label for the input when generating the language loss. We call this scenario *supervised agent learning*. It can also generate language loss without ground-truth by evaluating the output and trajectory according to the task description. In this case, we can say that the agent is doing *unsupervised agent learning*, which enables language agents to self-evolving. We present the detailed implementation of this prompt template in the Appendix.

Back-propagation of Language Gradients In standard connectionist learning, the goal of gradient back-propagation is to calculate the impact of the weights with respect to the overall loss so that the optimizers can update the weights accordingly. Similarly, in our framework, we also design a “back-propagation” algorithm for language gradients. Specifically, we iterate from the last node to the first node and compute the gradient for each node with LLMs using a carefully designed prompt:

$$\nabla_{\text{lang}}^n = \text{LLM}(\mathcal{P}_{\text{gradient}}(\nabla_{\text{lang}}^{n+1}, \mathcal{I}_n, \mathcal{O}_n, \mathcal{P}_n, \mathcal{T}_n, \mathcal{L}_{\text{lang}})) \quad (2)$$

The prompt template $\mathcal{P}_{\text{gradient}}$ is designed to instruct the LLM to generate language gradients that are analyses and reflections for the symbolic components within the node. Inspired by the idea of back-propagation, we give the language gradients of the node executed after the current node, as well as the information on the execution of the current node, which is stored in the trajectory. That’s to

say, when doing analysis and reflection, the LLM not only needs to consider how the prompts and tools suit the subgoal of the current node but also has to consider how they affect the accomplishment of the subgoal of the next node. By chaining from top to bottom, the language gradients for all nodes are relevant and responsible for the overall success of the agent. This method effectively reduces the risk of optimizing toward the local optimum for each isolated prompt and tool, leading to the overall performance of agent systems.

Language Gradient-based Update The final step in the framework is to update the prompts and tools in each node and optimize the overall agent workflow with the help of language gradients. This is accomplished via “symbolic optimizers”. Symbolic optimizers are carefully designed prompt workflows that can optimize the symbolic weights of an agent. We create three types of symbolic optimizers: PromptOptimizer, ToolOptimizer, and WorkflowOptimizer. We present detailed implementation of these prompts in the Appendix.

PromptOptimizer: To facilitate prompt optimization, we split prompts into different components, including task description, few-shot examples, principles, and output format control. We then design separate prompts tailored for the optimization of each prompt component. All prompts share a detailed explanation and demonstration of how the LLM should focus on the language gradients when reasoning about how to edit the original prompt components.

ToolOptimizer: The ToolOptimizer is a workflow of prompts that first instructs the LLM to decide the kind of operation it should use: whether the tools should be improved (by editing the tool description used for function calling), deleted, or new tools need to implement. Then the ToolOptimizer calls different prompts specifically designed for tool editing, deletion, and creation.

WorkflowOptimizer: The goal of the WorkflowOptimizer is to optimize the agent workflow consisting of nodes and their connections. The prompt is designed to first introduce the agent programming language used to define the agent workflow (we use the agent programming language introduced in Zhou et al. (2023b)). Then the prompt describes the definition of a few atomic operations that the LLM can use to update the workflow, including adding, deleting, and moving the nodes. It then instructs the LLM to first analyze how the workflow could be improved and then implement the update using the atomic operations. Detailed descriptions of the agent programming language and the atomic operations used to update the agent workflow are available in the Appendix.

Since all aforementioned optimizers operate in natural language space and some optimization operations need to be done in code space, we use a simple strategy that retries any illegal update up to three times and discards the update if the error persists. We also use a rollback strategy that re-runs the current example after optimization and rolls back to the original agent if the performance evaluated using the language-based loss function drops. Furthermore, we also include a “learning rate” component for each prompts in the optimizers which controls how aggressive the LLM should be when optimizing prompts, tools, and agent workflows.

Batched Training The aforementioned optimization scheme works with one training example at a time, which resembles stochastic gradient descent. Inspired by the fact that mini-batch stochastic gradient descent works better, or more stably, in practice, we also devise a batched training variant for symbolic optimizers. Specifically, we conduct forward pass, loss computation, and back-propagation for each example separately. Then we feed a batch of language gradients for the same node, and prompt the LLM to holistically consider all these language gradients when updating the agent.

Cost and Efficiency Compared to conventional static agent frameworks, agent symbolic learning does not involve additional compute or API costs during inference time. As for training time, for each training example, the agent symbolic learning framework requires roughly 3 to 5 times the API costs (in terms of the number of input and output tokens) compared to that required for inference time.

4 EXPERIMENTS

4.1 SETTINGS

4.1.1 TASKS

We conduct experiments on both standard LLM benchmarks and more complex agentic tasks. We describe the tasks, datasets, and evaluation metrics as follows:

Table 1: Results on standard LLM benchmarks.

Methods	HotPotQA		MATH		HumanEval	
	GPT-3.5	GPT-4	GPT-3.5	GPT-4	GPT-3.5	GPT-4
GPTs	24 / 38.8	33 / 44.3	23.2	53.1	59.2	71.7
Agents	27 / 37.5	39 / 49.8	23.8	56.0	59.5	85.0
Agents w/ AutoPE	29 / 39.8	38 / 50.3	22.5	57.2	63.5	82.3
DSpy	35 / 43.9	40 / 50.5	17.3	48.4	66.7	77.3
Ours	35 / 44.8	41 / 54.0	38.8	60.7	64.5	85.8

Standard Benchmarks We conduct experiments on standard benchmarks for LLMs including HotpotQA (Yang et al., 2018), MATH (Hendrycks et al., 2021), and HumanEval (Chen et al., 2021). HotPotQA is a multi-hop QA task challenging for rich background knowledge. We use the “hard” split in the dataset since we find it to be more challenging for language agents. MATH is a collection of challenging competition mathematics problems. HumanEval is an evaluation set that requires LLMs or agents to synthesize programs from docstrings. As for evaluation metrics, we use F1 and exact match for HotPotQA, accuracy for MATH, and Pass@1 for HumanEval. Tools are disabled in these datasets to ensure the comparison of the results is meaningful with existing literature on these tasks.

Complex Agent Tasks We consider **creative writing** and **software development** as two complex agentic tasks. For the creative writing task, we follow Yao et al. (2023) and give 4 random sentences to the agents and ask them to write a coherent passage with 4 paragraphs that end in the 4 input sentences respectively. Such a task is open-ended and exploratory and challenges creative thinking as well as high-level planning. We use GPT-4 score to evaluate the passages following (Yao et al., 2023). The software development task, on the other hand, requires the agent system to develop an *executable* software given a simple product requirement document (PRD). We evaluate the compared agents according to the *executability* of the generated software, which is quantified by numerical scores ranging from 1 to 4, corresponding to increasing levels of execution capability. Specifically, a score of 1 signifies execution failure, 2 denotes successful code execution, 3 represents conformance to the anticipated workflow, and 4 indicates flawless alignment with expectations.

4.1.2 BASELINES

We compare our proposed method against the following baselines:

- **GPTs**: a simple baseline that uses GPT and a carefully designed prompt following the way OpenAI implements GPTs agents;
- **Agents**: a language agent method implemented using the Agents (Zhou et al., 2023b) framework³ with carefully designed prompts, tools, and workflows;
- **DSpy**: an LLM workflow optimization framework that can search the best combination of prompt components. It is not directly applicable for complex agent tasks where the evaluation metric can not be defined in equation and code;
- **Agents + AutoPE**: a variant where the prompt in each node of the agent workflow is optimized by an LLM following the method described in Yang et al. (2024). Compared

³We have tested with other agent frameworks such as OpenAgents and AgentVerse and got similar results.

with our approach, this baseline does not involve language gradient back-propagation and language gradient-based optimization.

We conduct the experiments with both GPT-3.5 and GPT-4. We use the `gpt-3.5-turbo-0125` endpoint for GPT-3.5 and the `gpt-4-turbo-0409` endpoint for GPT-4. As for our approach, we start with the **Agents** baseline and then conduct agent symbolic learning on top of it. All agent systems included in the experiments are implemented and optimized with the best efforts from the same group of engineers with good proficiency on agent development.

4.2 RESULTS

Table 2: Results on software development.

Task	GPTs	Agents	Ours
Flappy bird	2	2	3
Tank battle game	1	2	4
2048 game	1	2	4
Snake game	2	3	4
Brick breaker game	2	3	4
Average score	1.6	2.4	3.8

Table 3: Results on creative writing.

Methods	GPT-3.5	GPT-4
GPTs	4.0	6.0
Agents	4.2	6.0
Agents w/ AutoPE	4.4	6.5
ToT	3.8	6.8
Ours	6.9	7.4

Results on LLM Benchmarks The results on standard LLM benchmarks are shown in Table 1. We can see that the proposed agent symbolic learning framework consistently improves over all compared methods. The performance improvement on MATH, a competition-level benchmark, is especially large. In contrast, the conventional LLM-based prompt optimization method (Agents w/ AutoPE) and the search-based prompt optimization approach (DSPy) are not as stable: they results in good performance improvements in some cases but lead to significant performance degradation in some other cases. This suggests that the agent symbolic learning framework is more robust and can optimize the overall performance of language agents more effectively.

Results on Complex Tasks We present the results on software development and creative writing in Table 2 & 3, respectively. We can see that our approach significantly outperforms all compared baselines on both tasks with an even larger performance gap compared to that on conventional LLM benchmarks. Interestingly, our approach even outperforms tree-of-thought, a carefully designed prompt engineering and inference algorithm, on the creative writing task. We find that our approach successfully finds a “plan, write, and revision” workflow for professional creative writing, and the prompts are very well optimized in each step. We also find that the agent symbolic learning framework recovers a similar standard operation procedure developed in MetaGPT (Hong et al., 2023), an agent framework specifically designed for software development. This confirms the effectiveness of the proposed agent symbolic learning framework on real-world tasks where there is no ground truth and the overall performance cannot be calculated by equations or codes, as contrary to search-based algorithms such as DSPy.

4.3 CASE STUDY & ANALYSIS

We then show a case study for the optimization dynamics of the agent symbolic learning framework in Figure 3. We can see that our approach can effectively do prompt engineering and designing of the agent workflow in the way a human expert develops language agents. Specifically, agent symbolic learning successfully adds an “edit” or “revision” node in the workflow of a creative writing agent and substantially improves the design of the prompts.

Moreover, we find that the initialization of the agent system has non-negligible impacts on the final performance, just as the initialization of neural nets is important for training. In general, we find that it is generally helpful to initialize the agent in the simplest way and let the symbolic optimizers to do the optimization. In contrast, the performance tends to become unstable if the initial agent system is over-engineered. A natural extension of this observation is that maybe we can do some kind of pre-training on large-scale and diverse tasks as a versatile initialization for general-purpose agents and then adapt it to specialized tasks with agent symbolic learning. We also find that the success of

486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539

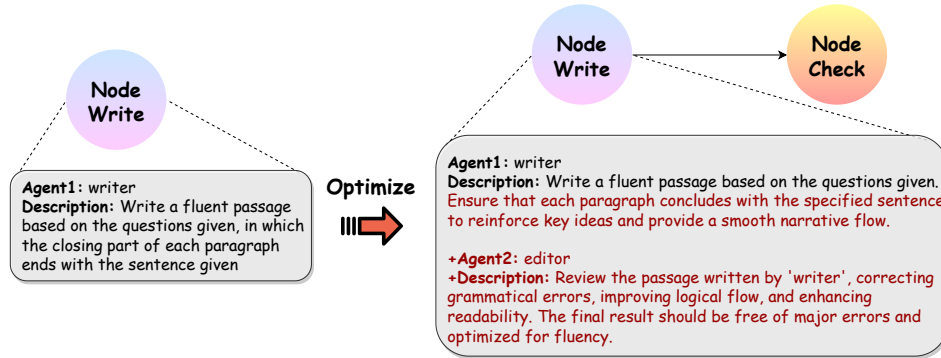


Figure 3: A case study conducted on creative writing task.

our approach is more significant and stable on complex real-world tasks compared to that on standard benchmarks where the performance is evaluated by traditional metrics such as accuracy or F1. This suggests that future research on agent learning should focus more on real-world tasks, and the agent research community should work on building a benchmark focusing on agent learning evaluation that consists of diverse complex agentic tasks and investigating robust approaches to measure progress.

5 CONCLUSION

This paper introduces agent symbolic learning, a framework for agent learning that jointly optimizes all symbolic components within an agent system. The agent symbolic learning framework draws inspiration from standard connectionist learning procedure to do symbolic learning. It uses language-based loss, gradients, and optimizers to optimize prompts, tools, and the agent workflow with respect to the overall performance of the agent system. The proposed framework is among the first attempts to optimize agents that can solve complex real-world tasks using sophisticated workflows. Our frameworks enables language agents to “learn from data” and perform “self-evolve” after being created and deployed in the wild. We conduct several proof-of-concept experiments and show that the agent symbolic learning framework can effectively optimize agents across different task complexity. We believe this transition from model-centric to data-centric agent research is a meaningful step towards approaching artificial general intelligence and open-source the codes and prompts for the agent symbolic learning framework to accelerate this transition.

6 LIMITATIONS & BOARDER IMPACT

The scope of experiments in this paper is not super large enough to cover most agentic tasks in the real world. They are rather proof-of-concept experiments showcasing the effectiveness of the proposed method. We believe the community on agent research should work on a standard evaluation procedure to facilitate future research. Another limitation is that the experiments are done with text-only models and tasks, while experiments with multi-modal agents and tasks would be very interesting.

As for the boarder impact, we would like to point out that enabling language agents to self-evolve in the wild poses certain safety risks. We believe it is important to reveal these potential risks to the agent research & development community and we need to discuss methods for effective regulation.

REFERENCES

- Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. {TensorFlow}: a system for {Large-Scale} machine learning. In *12th USENIX symposium on operating systems design and implementation (OSDI 16)*, pages 265–283, 2016.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel

540 Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler,
541 Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Ben-
542 jamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and
543 Dario Amodei. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell,
544 M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33,
545 pages 1877–1901. Curran Associates, Inc., 2020. URL [https://proceedings.neurips.
546 cc/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf](https://proceedings.neurips.cc/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf).

547 Baian Chen, Chang Shu, Ehsan Shareghi, Nigel Collier, Karthik Narasimhan, and Shunyu Yao.
548 Fireact: Toward language agent fine-tuning, 2023a.

549

550 Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Pondé de Oliveira Pinto, Jared
551 Kaplan, Harrison Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri,
552 Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan,
553 Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian,
554 Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios
555 Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino,
556 Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders,
557 Christopher Hesse, Andrew N. Carr, Jan Leike, Joshua Achiam, Vedant Misra, Evan Morikawa,
558 Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob
559 McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. Evaluating
560 large language models trained on code. *CoRR*, abs/2107.03374, 2021. URL [https://arxiv.
561 org/abs/2107.03374](https://arxiv.org/abs/2107.03374).

562 Weize Chen, Yusheng Su, Jingwei Zuo, Cheng Yang, Chenfei Yuan, Chen Qian, Chi-Min Chan, Yujia
563 Qin, Yaxi Lu, Ruobing Xie, et al. Agentverse: Facilitating multi-agent collaboration and exploring
564 emergent behaviors in agents. *arXiv preprint arXiv:2308.10848*, 2023b.

565 Qingyan Guo, Rui Wang, Junliang Guo, Bei Li, Kaitao Song, Xu Tan, Guoqing Liu, Jiang Bian,
566 and Yujiu Yang. Connecting large language models with evolutionary algorithms yields powerful
567 prompt optimizers. In *The Twelfth International Conference on Learning Representations*, 2024.
568 URL <https://openreview.net/forum?id=ZG3RaNIso8>.

569 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,
570 and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *NeurIPS*,
571 2021.

572 Geoffrey E Hinton. Connectionist learning procedures. In *Machine learning*, pages 555–610. Elsevier,
573 1990.

574

575 Geoffrey E Hinton and Ruslan R Salakhutdinov. Reducing the dimensionality of data with neural
576 networks. *science*, 313(5786):504–507, 2006.

577

578 Sirui Hong, Mingchen Zhuge, Jonathan Chen, Xiawu Zheng, Yuheng Cheng, Ceyao Zhang, Jin-
579 lin Wang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, Liyang Zhou, Chenyu Ran, Lingfeng
580 Xiao, Chenglin Wu, and Jürgen Schmidhuber. Metagpt: Meta programming for a multi-agent
581 collaborative framework, 2023.

582 Omar Khattab, Arnav Singhvi, Paridhi Maheshwari, Zhiyuan Zhang, Keshav Santhanam, Sri Vard-
583 hamanan, Saiful Haq, Ashutosh Sharma, Thomas T. Joshi, Hanna Moazam, Heather Miller,
584 Matei Zaharia, and Christopher Potts. Dspy: Compiling declarative language model calls into
585 self-improving pipelines. *arXiv preprint arXiv:2310.03714*, 2023.

586

587 Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon,
588 Nouha Dziri, Shrimai Prabhunoye, Yiming Yang, Sean Welleck, Bodhisattwa Prasad Majumder,
589 Shashank Gupta, Amir Yazdanbakhsh, and Peter Clark. Self-refine: Iterative refinement with
590 self-feedback, 2023.

591 Maxwell Nye, Anders Johan Andreassen, Guy Gur-Ari, Henryk Michalewski, Jacob Austin, David
592 Bieber, David Dohan, Aitor Lewkowycz, Maarten Bosma, David Luan, Charles Sutton, and
593 Augustus Odena. Show your work: Scratchpads for intermediate computation with language
models, 2022. URL <https://openreview.net/forum?id=iedYJm92o0a>.

594 OpenAI. GPT-4 technical report, 2023.
595

596 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong
597 Zhang, Sandhini Agarwal, Katarina Slama, Alex Gray, John Schulman, Jacob Hilton, Fraser Kelton,
598 Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and
599 Ryan Lowe. Training language models to follow instructions with human feedback. In Alice H. Oh,
600 Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information
601 Processing Systems*, 2022. URL <https://openreview.net/forum?id=TG8KACxEON>.

602 Joon Sung Park, Joseph C. O’Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and
603 Michael S. Bernstein. Generative agents: Interactive simulacra of human behavior, 2023.
604

605 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor
606 Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style,
607 high-performance deep learning library. *Advances in neural information processing systems*, 32,
608 2019.

609 Archiki Prasad, Peter Hase, Xiang Zhou, and Mohit Bansal. GrIPS: Gradient-free, edit-based
610 instruction search for prompting large language models. In Andreas Vlachos and Isabelle
611 Augenstein, editors, *Proceedings of the 17th Conference of the European Chapter of the As-
612 sociation for Computational Linguistics*, pages 3845–3864, Dubrovnik, Croatia, May 2023.
613 Association for Computational Linguistics. doi: 10.18653/v1/2023.eacl-main.277. URL
614 <https://aclanthology.org/2023.eacl-main.277>.

615
616 Reid Pryzant, Richard Diehl Martinez, Nathan Dass, Sadao Kurohashi, Dan Jurafsky, and Diyi
617 Yang. Automatically neutralizing subjective bias in text. In *The Thirty-Fourth AAAI Conference
618 on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial
619 Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in
620 Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 480–489.
621 AAAI Press, 2020. doi: 10.1609/AAAI.V34I01.5385. URL [https://doi.org/10.1609/
622 aai.v34i01.5385](https://doi.org/10.1609/aaai.v34i01.5385).

623 Cheng Qian, Shihao Liang, Yujia Qin, Yining Ye, Xin Cong, Yankai Lin, Yesai Wu, Zhiyuan
624 Liu, and Maosong Sun. Investigate-consolidate-exploit: A general strategy for inter-task agent
625 self-evolution, 2024.
626

627 Shuofei Qiao, Ningyu Zhang, Runnan Fang, Yujie Luo, Wangchunshu Zhou, Yuchen Eleanor Jiang,
628 Chengfei Lv, and Huajun Chen. AUTOACT: automatic agent learning from scratch via self-
629 planning. *CoRR*, abs/2401.05268, 2024. doi: 10.48550/ARXIV.2401.05268. URL [https:
630 //doi.org/10.48550/arXiv.2401.05268](https://doi.org/10.48550/arXiv.2401.05268).

631 Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language
632 understanding by generative pre-training. 2018.
633

634 Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language
635 models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
636

637 Timo Schick, Jane Dwivedi-Yu, Roberto Dessi, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke
638 Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language models can teach
639 themselves to use tools. In *Thirty-seventh Conference on Neural Information Processing Systems*,
640 2023. URL <https://openreview.net/forum?id=Yacmpz84TH>.

641 Yifan Song, Da Yin, Xiang Yue, Jie Huang, Sujian Li, and Bill Yuchen Lin. Trial and error:
642 Exploration-based trajectory optimization for LLM agents. *CoRR*, abs/2403.02502, 2024. doi: 10.
643 48550/ARXIV.2403.02502. URL <https://doi.org/10.48550/arXiv.2403.02502>.

644
645 Alessandro Sordani, Xingdi Yuan, Marc-Alexandre Côté, Matheus Pereira, Adam Trischler, Ziang
646 Xiao, Arian Hosseini, Friederike Niedtner, and Nicolas Le Roux. Joint prompt optimization of
647 stacked LLMs using variational inference. In *Thirty-seventh Conference on Neural Information
Processing Systems*, 2023. URL <https://openreview.net/forum?id=iImnbUVhok>.

-
- 648 Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée
649 Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand
650 Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language
651 models, 2023a.
- 652
- 653 Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay
654 Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cris-
655 tian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu,
656 Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn,
657 Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel
658 Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee,
659 Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra,
660 Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi,
661 Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh
662 Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen
663 Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic,
664 Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models,
665 2023b.
- 666 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz
667 Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing
668 systems*, 30, 2017.
- 669 Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H. Chi, Sharan Narang, Aakanksha
670 Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language
671 models. In *The Eleventh International Conference on Learning Representations*, 2023. URL
672 <https://openreview.net/forum?id=1PL1NIMMrw>.
- 673
- 674 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny
675 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in
676 neural information processing systems*, 35:24824–24837, 2022.
- 677 Tianbao Xie, Fan Zhou, Zhoujun Cheng, Peng Shi, Luoxuan Weng, Yitao Liu, Toh Jing Hua, Junning
678 Zhao, Qian Liu, Che Liu, Leo Z. Liu, Yiheng Xu, Hongjin Su, Dongchan Shin, Caiming Xiong,
679 and Tao Yu. Openagents: An open platform for language agents in the wild, 2023.
- 680
- 681 Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc V Le, Denny Zhou, and Xinyun
682 Chen. Large language models as optimizers. In *The Twelfth International Conference on Learning
683 Representations*, 2024. URL <https://openreview.net/forum?id=Bb4VGOWELI>.
- 684
- 685 Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and
686 Christopher D. Manning. HotpotQA: A dataset for diverse, explainable multi-hop question answer-
687 ing. In Ellen Riloff, David Chiang, Julia Hockenmaier, and Jun’ichi Tsujii, editors, *Proceedings of
688 the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380,
689 Brussels, Belgium, October–November 2018. Association for Computational Linguistics. doi:
690 10.18653/v1/D18-1259. URL <https://aclanthology.org/D18-1259>.
- 691 Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao.
692 React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629*,
693 2022.
- 694
- 695 Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik
696 Narasimhan. Tree of thoughts: Deliberate problem solving with large language models, 2023.
- 697
- 698 Shaokun Zhang, Jieyu Zhang, Jiale Liu, Linxin Song, Chi Wang, Ranjay Krishna, and Qingyun Wu.
699 Offline training of language model agents with functions as learnable weights, 2024a.
- 700
- 701 Wenqi Zhang, Ke Tang, Hai Wu, Mengna Wang, Yongliang Shen, Guiyang Hou, Zeqi Tan, Peng Li,
Yueting Zhuang, and Weiming Lu. Agent-pro: Learning to evolve via policy-level reflection and
optimization, 2024b.

702 Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang,
703 Zi Lin, Zhuohan Li, Dacheng Li, Eric. P Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica.
704 Judging llm-as-a-judge with mt-bench and chatbot arena, 2023.
705

706 Wangchunshu Zhou, Yuchen Eleanor Jiang, Peng Cui, Tiannan Wang, Zhenxin Xiao, Yifan Hou,
707 Ryan Cotterell, and Mrinmaya Sachan. Recurrentgpt: Interactive generation of (arbitrarily) long
708 text, 2023a.

709 Wangchunshu Zhou, Yuchen Eleanor Jiang, Long Li, Jialong Wu, Tiannan Wang, Shi Qiu, Jintian
710 Zhang, Jing Chen, Ruipu Wu, Shuai Wang, Shiding Zhu, Jiyu Chen, Wentao Zhang, Ningyu
711 Zhang, Huajun Chen, Peng Cui, and Mrinmaya Sachan. Agents: An open-source framework for
712 autonomous language agents, 2023b.

713 Mingchen Zhuge, Wenyi Wang, Louis Kirsch, Francesco Faccio, Dmitrii Khizbullin, and Jurgen
714 Schmidhuber. Language agents as optimizable graphs. *arXiv preprint arXiv:2402.16823*, 2024.
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755

756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809

A IMPLEMENTATION DETAILS

We adopt the agent programming language and framework introduced in Agents (Zhou et al., 2023b), a language agent framework that enables developers to build language agents that stacks prompts and tools together into complex pipelines. The main advantage of the Agents framework is that it enables developers to use a config file to define the agent system, which makes it easier for the symbolic optimizers in the agent symbolic learning framework to perform update operations on the agent system.

B PROMPT TEMPLATES

Prompt Template for Language Loss Function

Loss with ground truth:

You are a fine-tuner of a large model. I will provide you with some output results from the model and the expected correct results. You need to evaluate these data and provide a score out of 10, please wrap the score using `<score></score>`. Additionally, please provide some suggestions for modifying the model’s output, using `<suggestion></suggestion>` to wrap your suggestions.

Here is the model’s output:

`<result>result</result>`;

The expected result is:

`<ground_truth>ground_truth</ground_truth>`

Please note:

1. Ensure that the output is wrapped with `<score></score>` and `<suggestion></suggestion>` respectively.
2. The output should be as consistent as possible with the expected result while being correct. For example, if the expected result is “BUST”, and the model’s output is “The women’s lifestyle magazine is ‘BUST’ magazine.”, even though the answer is correct, you should advise the model to be more concise.
3. The standard for a score of 10 is that the model’s output is exactly the same as the expected result in a case-insensitive manner, and without any unnecessary content. Even if the model’s output is semantically correct, if it includes superfluous content, points should be deducted.

Loss with ground truth and score:

You are a large language model fine-tuner. I will provide you with a model’s output and the expected correct result. You need to evaluate it and suggest modifications to the model’s output. Please use ‘`<suggestion></suggestion>`’ to enclose your feedback.

Below is the model’s output:

`<result>result</result>`

The expected result is:

`<ground_truth>ground_truth</ground_truth>`

Here is the evaluation score for the model. Your goal is to optimize this score:

`<score>score</score>`

The relevant information about this score is as follows:

`<evaluation_info>score_info</evaluation_info>`

Note:

1. Ensure that ‘`<suggestion></suggestion>`’ exists and appears once.
 2. If the model’s output is satisfactory, you can output `<suggestion>The output is satisfactory, no additional requirements</suggestion>`.
 3. The output should be as close to the expected result as possible while ensuring correctness. For example, if the expected result is "BUST" and the model’s output is "The women’s lifestyle magazine is 'BUST' magazine.", even though this answer is correct, you should remind the model to be concise.
-

Table 4: Prompt Template for Language Loss Function

810	Prompt Template for Gradient Back-propagation
811	
812	Prompt-Level
813	You are now a prompt fine-tuner for a large language model. You are tasked with providing suggestions for optimizing the prompt template.
814	Please enclose your suggestions using <code><suggestion></suggestion></code> , for example, <code><suggestion>it could be made shorter</suggestion></code> .
815	The task is divided into multiple steps; I will provide you with the output from the previous step, the requirement proposed by the next step for the current output, the current output itself, and the prompt template. You need to suggest improvements for the current step's prompt template.
816	
817	
818	
819	- The prompt template that needs optimization is: <code><prompt_template>prompt_template</prompt_template></code>
820	- The output from the previous step is: <code><previous_output>previous_output</previous_output></code>
821	- The current output is: <code><output>response</output></code>
822	- The requirement proposed by the next step for the current output is: <code><requirement>suggestion</requirement></code>
823	
824	In addition to suggesting modifications for the current prompt template, you also need to propose requirements for the output of the previous step. Please wrap these using <code><suggestion></suggestion></code> , for example: <code><suggestion>the analysis should include a comparison of original data</suggestion></code> .
825	
826	
827	Note:
828	1. Ensure that the results are wrapped with <code><suggestion></suggestion></code> and <code><requirement></requirement></code> , and each tag appears only once.
829	2. If you are the first node, you can state within <code><suggestion></suggestion></code> "This is the first node."
830	3. Please note that during your analysis, remember that this prompt template will be applied to multiple different datasets, so your suggestions should be general and not solely focused on the examples provided here.
831	4. Please analyze step by step.
832	
833	
834	Node-Level
835	You are a large model fine-tuner. Now you need to try to optimize the information of a node. For a complex task, it has been divided into multiple nodes, each of which contains multiple roles that work together to complete the task of this node. Each role is backed by an LLM Agent, and you need to optimize the configuration information of one of the nodes.
836	
837	
838	
839	Here are the relevant explanations for the Node configuration:
840	- The fields in the "controller" indicate the scheduling method of the model. If there is only one role, this item does not need to be optimized:
841	- "route_type" indicates the scheduling method, which has three values: "random" means random scheduling, "order" means sequential scheduling, and "llm" means scheduling determined by the LLM model.
842	- "route_system_prompt" and "route_last_prompt" are used when "route_type" is "llm" and are respectively the system prompt and last prompt given to the LLM model responsible for scheduling.
843	- "begin_role" is a string indicating the name of the starting role of this node.
844	- "roles" is a dictionary where the key is the role name, and the value is the prompt used by this role.
845	
846	
847	You need to decide how to optimize the configuration of this node. Specifically, you need to try to provide suggestions in the following aspects:
848	
849	1. Update the node description field. This field describes the function of the node and is also an important indicator to measure the performance of a node.
850	2. Update the scheduling method of the role. Note that if there is only one role, no optimization is needed.
851	3. Add a new role, and you need to clearly describe the function of this role.
852	4. Delete a role, and you need to clearly describe the reason for deleting this role.
853	5. Update a role, and you need to indicate how to update the description of this role.
854	
855	Next, I will give you a Node configuration, and you need to provide optimization suggestions based on the current Node configuration. Please use <code><suggestion>[put your suggestion here]</suggestion></code> to enclose your suggestions.
856	
857	
858	<code>## Current Node Config</code>
859	<code>{node_config}</code>
860	
861	You need to first provide your analysis process, then give your optimized result. Please use <code><analyse></analyse></code> to enclose the analysis process. Please use <code><suggestion></suggestion></code> to enclose the optimization suggestions for the current node. Please use <code><requirement></requirement></code> to enclose the requirements for the previous node.
862	
863	Note: The suggestions provided need to be in one or more of the five aspects mentioned above.

Table 5: Prompt Template for Gradient Back-propagation

864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917

Prompt Template for Optimizers

Prompt Optimizer:

You are now a prompt fine-tuner for a large language model. I will provide you with a prompt template along with its corresponding input and output information.

Please modify the prompt based on the provided data:

- The current prompt template is: `prompt_template`.

Here is some information about the model when using this template:

Example index

- Output result: `<output>response</output>`

- Suggestion: `<suggestion>suggestion</suggestion>`

You need to analyze the content above and input the optimized prompt result. Please wrap your analysis in `<analyse></analyse>` and the new prompt in `<new_prompt></new_prompt>`.

Please note:

1. When actually using the prompt template, the Python `format()` method is employed to fill variables into the prompt. Therefore, please ensure that the content enclosed in `in` in both the new and old prompts remains the same, with no variables added or removed.
2. Ensure that your new prompt template can be directly converted to a dictionary using the `json.loads()` method. Therefore, you need to be careful to use double quotes and escape characters properly.
3. Ensure that `<analyse></analyse>` and `<new_prompt></new_prompt>` each appear only once.
4. If you believe that the current prompt template performs sufficiently well, leave `<new_prompt></new_prompt>` empty.

Node Optimizer:

You are a large model fine-tuner. Now you need to try to optimize the information of a node. For a complex task, it has been divided into multiple nodes, each containing multiple roles that work together to complete the task of this node. Each role is backed by an LLM Agent, and you need to optimize the configuration information of one of the nodes.

Here are the relevant explanations for the Node configuration:

- The fields in the "controller" indicate the scheduling method of the model. If there is only one role, this item does not need to be optimized:

- "route_type" indicates the scheduling method, which has three values: "random" means random scheduling, "order" means sequential scheduling, and "llm" means scheduling determined by the LLM model.

- "route_system_prompt" and "route_last_prompt" are used when "route_type" is "llm" and are respectively the system prompt and last prompt given to the LLM model responsible for scheduling.

- "begin_role" is a string indicating the name of the starting role of this node.

- "roles" is a dictionary where the key is the role name, and the value is the prompt used by this role.

Next, I will give you a Node configuration and several modification suggestions. You need to modify the Node configuration based on the suggestions:

```
## Current Node Config
```

```
{node_config}
```

```
## Suggestions
```

```
{suggestions}
```

When providing the modification plan, you need to give the optimized result in the following format. It is a list, each element is a dict, and the dict contains an action field indicating the operation on the Node.

Your optimized result should be enclosed in `<result></result>`, that is, the content inside `<result></result>` should be a JSON-formatted list, which should be able to be directly loaded by `json.loads()`.

Note:

1. If you think the current configuration is already excellent and does not need modification, you can directly output an empty list.
2. The format of `<result>[optimization method]</result>` needs to strictly follow the given format, otherwise, it will be judged as incorrect.

Table 6: Prompt Template for Optimizers