

Imagine-then-Plan: Agent Learning from Adaptive Lookahead with World Models

Youwei Liu^{2,1*}, Jian Wang^{1,3†}, Hanlin Wang¹, Beichen Guo¹, Wenjie Li¹

¹ Department of Computing, The Hong Kong Polytechnic University

² Central South University

³ College of Computer Science, Sichuan University

loyiv5477@gmail.com wangjian51@scu.edu.cn

{hanlin-henry.wang, beichen.guo}@connect.polyu.hk

cswjli@comp.polyu.edu.hk

Abstract

Recent advances in world models have shown promise for modeling future dynamics of environmental states, enabling agents to reason and act without accessing real environments. Current methods mainly perform single-step or fixed-horizon rollouts, leaving their potential for complex task planning under-exploited. We propose Imagine-then-Plan (ITP), a unified framework for agent learning via lookahead imagination, where an agent’s policy model interacts with the learned world model, yielding multi-step “imagined” trajectories. Since the imagination horizon may vary by tasks and stages, we introduce a novel adaptive lookahead mechanism by trading off the ultimate goal and task progress. The resulting imagined trajectories provide rich signals about future consequences, such as achieved progress and potential conflicts, which are fused with current observations, formulating a partially *observable* and *imaginable* Markov decision process to guide policy learning. We instantiate ITP with both training-free and reinforcement-trained variants. Extensive experiments across representative agent benchmarks demonstrate that ITP significantly outperforms competitive baselines. Further analyses validate that our adaptive lookahead largely enhances agents’ reasoning capability, providing valuable insights into addressing broader, complex tasks. Our code and data are publicly available at: <https://github.com/loyiv/ITP>.

1 Introduction

The emergence of Large Language Models (LLMs) has sparked a paradigm shift in autonomous agents, enabling them to reason and interact across a wide range of digital and physical environments (Li et al.,

* This work was conducted while Youwei Liu was a remote research assistant at the Hong Kong Polytechnic University.

† Corresponding author. This work was mainly conducted at PolyU, while the author is now at Sichuan University.




Figure 1: Comparison between our ITP framework and conventional agent learning frameworks.

2023; Wang et al., 2024a; Fung et al., 2025). LLM-based agents primarily leverage immediate observations and historical interaction traces to facilitate decision-making (Yao et al., 2023b; Shinn et al., 2023). Despite the impressive performance achieved, most agents remain constrained by *shallow grounding*, a state where they perceive the environment but lack a deep, causal understanding of how their current actions will ultimately reshape the environment. Without the ability to project into the future, agents are prone to catastrophic failures, discovering erroneous actions or state conflicts only after they have been irreversibly executed.

To bridge this grounding gap, world models (LeCun, 2022; Hafner et al., 2023) that focus on future state modeling have emerged as a promising solution, enabling agents to simulate environmental dynamics and “rehearse” actions in a mental sandbox (Wang et al., 2025e). We focus on textual

world models that predict decision-relevant future observations or state consequences, rather than continuous physical dynamics. Contemporary methods typically employ world models for single-step verification or with fixed-horizon rollouts (Hao et al., 2023; Qiao et al., 2024; Zhang et al., 2026). Such rigid strategies are often suboptimal, as they fail to capture long-term dependencies in complex tasks (e.g., household embodied tasks). Furthermore, they are prone to suffer from high computational costs with unnecessary deep rollouts for trivial actions. A truly intelligent agent should be deliberative, allocating deep foresight adaptively to resolve potential state conflicts and account for long-term dependencies in pivotal, high-stakes decisions.

In this paper, we propose **Imagine-then-Plan (ITP)**, a framework that empowers agents to perform task planning with world model-based foresight. The core of ITP is to move beyond passive observation and perform a proactive, deliberative “rehearsal” phase, where decisions are conditioned on both present observations and potential futures. This requires the agent to internally imagine multi-step future trajectories by looking ahead. Unlike previous rigid approaches, ITP introduces an **adaptive lookahead** mechanism that dynamically scales the imagination horizon by trading off the ultimate goal and estimated task progress. More importantly, this shift necessitates a new conceptualization of the agent’s decision-making process. We move beyond the Partially Observable Markov Decision Process (POMDP) (Åström, 1965; Song et al., 2024; Wang et al., 2025b) toward a Partially Observable and Imaginable MDP (POIMDP). As illustrated in Figure 1, the agent’s action policy is optimized over a dual-stream representation: the concrete present (*observable*) and the foresighted future (*imaginable*). These imagined trajectories provide rich signals, such as anticipating goal progress or detecting potential bottlenecks, allowing the agent to close the loop between planning action sequences and estimating their possible consequences. This implicit feedback enables the agent to perform self-correction when necessary, significantly reducing the need for expensive interactions with the real environment.

We instantiate ITP in two variants: a **training-free (ITP_f)** variant that uses reflection as the adaptive lookahead for plug-and-play enhancement of LLM agents, and a **reinforcement-trained (ITP_r)** variant that leverages imagined futures to optimize the agent policy more effectively and more effi-

ciently. Extensive experiments demonstrate that both variants significantly improve task success rates. Further analyses validate the vital role of our adaptive lookahead mechanism.

Our contributions are summarized as follows: 1) We conceptualize the partially observable and imaginable Markov decision process (POIMDP), laying a solid foundation for integrating imagined futures and historical interactions into agent decision-making. 2) We propose Imagine-then-Plan (ITP), a framework that incorporates world model-based imagination with an adaptive lookahead mechanism, which provides deliberative guidance for action policy planning. 3) We demonstrate through training-free and reinforcement-trained variants that ITP significantly improves the success rates of LLM-based agents, providing valuable insights into addressing complex, long-horizon tasks.

2 Preliminaries

Problem Formulation. The reasoning process of LLM agents is often formulated as a Partially Observable Markov Decision Process (POMDP), defined by $(\mathcal{S}, \mathcal{A}, \mathcal{O}, \mathcal{T}, \mathcal{R})$. Here, $\mathcal{S}$ denotes the environment state space, $\mathcal{A}$ the action space, and $\mathcal{O}$ the observation space. $\mathcal{T} : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$ represents the state transition function, and $\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ denotes the reward function that evaluates task performance. At each time step t , the agent receives an observation $o_t \in \mathcal{O}$ and selects an action $a_t \sim \pi_\theta(\cdot|h_t)$, where $h_t = (o_1, a_1, \dots, o_t)$ denotes the interaction history. Executing a_t induces a transition to a new latent state $s_{t+1} \sim \mathcal{T}(s_t, a_t)$, from which the environment emits the next observation o_{t+1} . The interaction terminates when the agent reaches a terminal state or exceeds a predefined maximum number of steps.

LLMs as Textual World Models. A world model is a predictive model of environment dynamics that estimates future states conditioned on actions (LeCun, 2022). In text-based environments, the environment state is typically represented as text and observable to the agent. We treat textual observations as state representations, denoted as s_t at time step t . Under this formulation, LLMs can be interpreted as world models, as they capture transition regularities by predicting the next state given the current interaction context (Zhang et al., 2026; Li et al., 2026). Concretely, a textual world model parameterized by an autoregressive LLM defines the conditional distribution $p_\phi(s_{t+1}|s_t, a_t)$. The




Figure 2: Overview of the proposed Imagine-then-Plan (ITP) framework. It consists of two variants: (a) **ITP_I**, which is training-free and enables LLM agents to learn from the imagination at inference time. (b) **ITP_R**, which leverages an imagined future to optimize the action policy more effectively and more efficiently.

distribution is factorized at the token level, and the model generates the next state sequentially. Such world models serve as powerful proxies for environment dynamics and enable planning in language-based agents.

3 Method

We propose **Imagine-then-Plan (ITP)**, a framework that equips LLM-based agents with adaptive lookahead via learned world models. As illustrated in Figure 2, **ITP** enables agents to condition decisions on both the observable present and imagined future trajectories.

3.1 World Model Training

As shown in Figure 2, we first fine-tune a base LLM on expert demonstrations $\mathcal{D}_{\text{exp}}$ to obtain an initial agent policy π_{θ_0} . This warm-up establishes basic capability to produce executable actions, serving as the foundation for the agent’s exploration. We ask the agent to perform rollouts in the environment, obtaining the rollout trajectories $\mathcal{D}_{\text{roll}}$. As introduced in § 2, we learn a LLM-based world model $\mathcal{M}_\phi$ that approximates the environment dynamics $p_\phi(s'|s, a)$, where s and a denote the current state and action, respectively. s' represents the next-step state. To ensure the world model is grounded and robust to out-of-distribution actions, we train it on a joint dataset $\mathcal{D}_{\text{WM}} = \mathcal{D}_{\text{exp}} \cup \mathcal{D}_{\text{roll}}$.

The world model $\mathcal{M}_\phi$ is optimized by minimizing the negative log-likelihood as follows:

$$\mathcal{L}_{\text{WM}}(\phi) = -\mathbb{E}_{(s', s, a) \sim \mathcal{D}_{\text{WM}}} [\log p_\phi(s'|s, a)]. \quad (1)$$

We provide detailed training setup, dataset scale, and computational costs in Appendix C.1.

3.2 Lookahead Imagination and POIMDP

To integrate world-model foresight into decision making, we extend the standard POMDP (as introduced in §2) to a **Partially Observable and Imaginable Markov Decision Process (POIMDP)**. Under this formulation, the agent is endowed with a lookahead imagination operator induced by the learned world model. POIMDP does not replace the underlying POMDP or modify the environment dynamics. Instead, it augments the agent-side information state with an internally imagined trajectory before action execution. When the lookahead horizon is selected as zero, the imagined trajectory is empty, and the decision reduces to the standard POMDP policy.

Formally, given the current observed state s_t and a lookahead horizon $K_t \in \{0, 1, \dots, K_{\text{max}}\}$ at each time step t , the agent policy π_θ and the world model $\mathcal{M}_\phi$ interact for K_t steps within a “mental sandbox”. This imagination process is given by:

$$\hat{a}_t \xrightarrow{\mathcal{M}_\phi} \hat{s}_{t+1} \xrightarrow{\pi_\theta} \dots \xrightarrow{\pi_\theta} \hat{a}_{t+K_t-1} \xrightarrow{\mathcal{M}_\phi} \hat{s}_{t+K_t}. \quad (2)$$

This yields an imagined future trajectory $\hat{\tau}_t^{(K_t)} = \{(\hat{a}_{t+i}, \hat{s}_{t+i+1})\}_{i=0}^{K_t-1}$, where $\hat{a}_{t+i} \sim \pi_\theta(\cdot | s_t, \hat{\tau}_t^{(i)})$ and $\hat{s}_{t+i+1} \sim \mathcal{M}_\phi(\cdot | \hat{s}_{t+i}, \hat{a}_{t+i})$. The objective of the agent policy is to yield an appropriate action a_t conditioned on both the observable state s_t and the imagined future $\hat{\tau}_t^{(K_t)}$. As such, our POIMDP formulates the policy decision as follows:

$$a_t \sim \pi_\theta(\cdot | s_t, \hat{\tau}_t^{(K_t)}). \quad (3)$$

This allows the agent to anticipate goal progress or detect potential bottlenecks before generating

the next action, enabling the agent to perform self-correction when necessary.

3.3 Planning with Adaptive Lookahead

For effective task planning, a key challenge is determining how far the agent should imagine. Short-horizon imagination may miss long-term dependencies, while excessive rollouts can amplify model errors and incur unnecessary computation. To resolve this, we aim to adaptively select the imagination horizon K_t based on the estimated task progress against the ultimate goal. We instantiate **ITP** via two distinct variants to provide both flexibility and optimization: (i) **ITP_I**, which is an inference-time method that learns from the imagination, and (ii) **ITP_R**, which is a reinforcement-trained method that jointly optimizes the lookahead horizon selection and action policy.

3.3.1 In-Imagination Learning (ITP_I)

ITP_I is a training-free variant that improves LLM agents at inference time. Both the policy and world model remain frozen, and the agent performs deliberative reasoning based on the current state s_t . At each step t , it follows a three-stage “Imagine-then-Plan” procedure: 1) *Adaptive horizon selection*: The agent analyzes the instruction and state s_t to choose an imagination horizon $K_t \in \{0, 1, \dots, K_{\max}\}$, assigning deeper foresight to critical decisions while avoiding unnecessary computation on trivial ones. 2) *World-model imagination*: The agent interacts with $\mathcal{M}_\phi$ for K_t steps to obtain a future trajectory $\hat{\tau}_t^{(K_t)}$. 3) *Reflective policy generation*: Instead of directly executing the first imagined action, the agent treats $\hat{\tau}_t^{(K_t)}$ as implicit feedback for self-refinement.

Specifically, the agent reflects on the imagined trajectory to assess progress toward the goal and identify potential conflicts, bottlenecks, or catastrophic failures before acting in the real environment. It then refines its reflection and selects the next optimal action a_t as follows:

$$a_t \sim \pi_\theta(\cdot \mid s_t, \text{Reflect}(\hat{\tau}_t^{(K_t)})). \quad (4)$$

By grounding decisions in the imagined future, ITP_I turns passive observation into proactive, deliberative learning, improving task success without additional training.

3.3.2 Reinforced Training (ITP_R)

ITP_R aims to explicitly learn when and how long to imagine. We augment the agent with a lightweight

K -head predictor $P_\theta(K_t \mid s_t)$, which is a linear layer built on top of a backbone LLM that predicts distributions over imagination horizons. The action policy and predictor are optimized jointly with the following three stages.

Stage 1: Pseudo-Labeling Lookahead Horizon.

A key obstacle for learning adaptive lookahead is that expert trajectories provide only (s_t, a_t^*) pairs but do not specify the “right” lookahead horizon. We therefore construct pseudo labels using the world model $\mathcal{M}_\phi$ and the initial agent policy π_{θ_0} . Specifically, we use teacher-forced expert actions to rollout on the frozen world model $\mathcal{M}_\phi$ by looking ahead one step and obtain future states, from which we derive lookahead-conditioned trajectories $\{\hat{\tau}_t^{(k)}\}_{k=0}^{K_{\max}}$. We then score each candidate step by the log-likelihood of the expert action under π_{θ_0} , and select the optimal lookahead step $\tilde{K}_t$ according to:

$$\arg \max_{0 \leq k \leq K_{\max}} \left[\log p_{\theta_0}(a_t^* \mid s_t, \hat{\tau}_t^{(k)}) - \lambda_K k \right], \quad (5)$$

where λ_K is a hyperparameter controlling the lookahead penalty. Based on the selection criteria in Eq. (5), we obtain a dataset $\mathcal{D}_K$ containing the pseudo-labels of the optimal lookahead step for each action in the expert trajectory.

Stage 2: Warm-Up Training. Starting from the initial agent policy π_{θ_0} , we further fine-tune the agent to jointly (i) act under lookahead-conditioned pseudo-trajectories and (ii) predict the required lookahead step. Specifically, given labeled tuples $\{(s_t, a_t^*, \tilde{K}_t)\}$, we condition the agent policy on $\hat{\tau}_t^{(\tilde{K}_t)}$ and train $\pi_\theta(a_t \mid s_t, \hat{\tau}_t^{(\tilde{K}_t)})$ to imitate the expert action a_t^* , with a standard negative log-likelihood loss $\mathcal{L}_\pi(\theta)$. Meanwhile, we train the K -head predictor to estimate the pseudo label $\tilde{K}_t$, with a similar negative log-likelihood loss $\mathcal{L}_K(\theta)$. Our warm-up training is given by:

$$\mathcal{L}_{\text{WT}}(\theta) = \mathcal{L}_\pi(\theta) + \eta \mathcal{L}_K(\theta). \quad (6)$$

where η is a weighted coefficient. This yields (i) a competent agent policy that can generate reliable actions and (ii) an adaptive lookahead horizon predictor that approximates lookahead steps, providing a stable initialization for the subsequent online reinforcement optimization.

Stage 3: Online Optimization. To balance task performance and imagination cost, we further refine the agent policy through online reinforcement

learning. At each step t , the agent samples a lookahead step K_t from the K -head predictor, invokes the frozen world model $\mathcal{M}_\phi$ to generate a K_t -step imagined trajectory $\hat{\tau}_t^{(K_t)}$, and subsequently samples an action $a_t \sim \pi_\theta(\cdot | s_t, \hat{\tau}_t^{(K_t)})$. As illustrated in Figure 2, we employ a reward function that augments the environment reward r_{env} with explicit penalties for computational and interaction overhead, given by:

$$r_{t+1} = r_{\text{env}} - \lambda_K K_t - \lambda_{\text{step}}, \quad (7)$$

where λ_K is the lookahead penalty coefficient, and λ_{step} is a factor discouraging reasoning redundancy.

Specifically, we utilize the Advantage Actor-Critic (A2C) algorithm (Mnih et al., 2016) to jointly optimize the action policy and the K -head parameters. The objective is decomposed into three components: (i) an actor term $\mathcal{L}_{\text{act}}(\theta) \triangleq -\mathbb{E}\left[A_t \left(\log p_\theta(K_t | s_t) + \log \pi_\theta(a_t | s_t, \hat{\tau}_t^{(K_t)})\right)\right]$ that jointly updates the lookahead predictor and the agent policy; (ii) a Critic regression term $\mathcal{L}_{\text{value}}(\theta) \triangleq \mathbb{E}\left[(V_\theta(s_t) - \hat{V}_t)^2\right]$ that trains a Value-head to match the TD learning target; and (iii) an entropy regularizer $\mathcal{L}_{\text{ent}}(\theta) \triangleq -\mathbb{E}[\mathcal{H}(p_\theta(K_t | s_t))]$ that encourages sufficient exploration over the lookahead steps to prevent premature convergence to sub-optimal horizons. The final training objective is:

$$\mathcal{L}_{\text{A2C}}(\theta) = \mathcal{L}_{\text{act}}(\theta) + \alpha \mathcal{L}_{\text{value}}(\theta) + \beta \mathcal{L}_{\text{ent}}(\theta), \quad (8)$$

where α and β are hyperparameters balancing value estimation and exploration. Algorithm 1 shows the pseudocode of ITPR’s training process.

At inference time, ITP_R utilizes the learned K -head to perform adaptive lookahead, following a similar deliberative procedure as ITP_I. We provide a stage-wise analysis on the computational cost of ITP_R in Appendix C.4.

4 Experiments

4.1 Experimental Settings

Benchmarks. We evaluate ITP on four representative agent benchmarks: **ALFWorld** (Shridhar et al., 2021) for embodied household tasks, **ScienceWorld** (Wang et al., 2022) for simulated science experiments, **WebShop** (Yao et al., 2022) for long-horizon web navigation, and **StableToolBench** (Guo et al., 2024) for multi-turn tool use. Appendix A provides details of these benchmarks.

Backbone Models. For a fair comparison, all methods use the same backbone suite: Qwen2.5-7B (Yang et al., 2024), Qwen3-8B (Yang et al., 2025), and Llama-3.1-8B-Instruct (Grattafiori et al., 2024). Unless otherwise stated, we use Qwen3-8B to train the world model, and evaluate other world-model backbones in § 4.4. All models are prompted using their official chat templates.

Baseline Methods. We compare ITP against two categories of baselines. (1) *Prompting-based* methods: **CoT** (Wei et al., 2022) elicits reasoning capabilities by prompting the agent with step-by-step rationales. **ReAct** (Yao et al., 2023b) interleaves reasoning and action to solve interactive tasks. **RAP** (Hao et al., 2023) leverages the LLM as both a world model and a policy model, employing Monte Carlo Tree Search to perform planning. (2) *Training-based* methods: **SFT** (Chen et al., 2023) conducts behavioral cloning on expert trajectories. **WKM** (Qiao et al., 2024) trains a parametric world knowledge model that provides global task priors and local dynamic state knowledge to guide planning. **IWM** (Zhang et al., 2026) augments imitation learning with an implicit world-modeling objective to encourage the policy to internalize environment dynamics. All baselines are evaluated using the same benchmark splits, executable action space, state-action serialization, backbone family, and evaluation metric as ITP. We additionally consider stronger search-based, world-model-based, and trajectory-optimization methods. Their complete results and reproduction details are provided in Appendix C.6.

Evaluation Metrics. We adopt success rate (**SR**) as the evaluation metric, defined as the percentage of episodes that successfully achieve the task goal (Feng et al., 2025; Wang et al., 2025a). For ALFWorld, we report SR on each task, as well as the overall SR. For ScienceWorld, we report SR on both the *seen* and *unseen* test splits. Following Guo et al. (2024), we report solvable pass rate (**SoPR**) and solvable win rate (**SoWR**) on the solvable subset for StableToolBench, with the detailed definitions provided in Appendix A.

Implementation Details. To ensure a rigorous and fair comparison, we maintain a consistent training protocol across ITP and all baseline methods. The maximum lookahead horizon is configured based on task completion steps per benchmark. For more details, please refer to Appendix C.

Type	Method	ALFWorld							ScienceWorld		WebShop
		PICK	CLEAN	HEAT	COOL	LOOK	PICK2	Overall	Seen	Unseen	Total
Backbone: Qwen2.5-7B 🦙											
Prompting	CoT	17.14	18.52	18.75	16.00	15.38	0.00	14.29	3.09	4.63	5.10
	ReAct	20.00	22.22	18.75	20.00	23.08	0.00	17.14	8.24	9.93	15.28
	RAP	40.00	33.33	6.25	32.00	15.38	20.83	27.86	10.30	16.55	11.28
	ITP _I (Ours)	65.71	25.93	25.00	24.00	30.77	25.00	35.71	16.49	17.88	20.10
Training	SFT	85.71	66.67	56.25	68.00	38.46	66.67	67.86	55.67	49.00	51.60
	WKM	77.14	77.78	75.00	76.00	76.92	75.00	76.43	54.12	56.29	58.80
	IWM	90.60	85.20	88.20	84.20	42.90	76.90	82.80	60.82	57.61	56.20
	ITP _R (Ours)	94.29	88.89	87.50	53.84	76.00	91.67	85.07	62.58	58.94	60.20
Backbone: Qwen3-8B 🦙											
Prompting	CoT	14.29	14.81	12.50	12.00	15.38	12.50	13.57	2.44	1.99	6.32
	ReAct	25.71	22.22	12.50	12.00	7.69	25.00	19.29	9.79	8.61	18.62
	RAP	42.86	37.04	37.50	16.00	15.38	4.17	28.57	15.46	27.14	12.40
	ITP _I (Ours)	82.86	25.93	12.50	16.00	23.08	54.17	41.43	20.61	19.86	25.25
Training	SFT	71.43	70.37	68.75	72.00	69.23	70.83	70.71	56.70	49.67	52.30
	WKM	80.00	77.78	81.25	80.00	76.92	79.17	79.29	60.31	47.68	61.15
	IWM	85.71	85.19	87.50	84.00	46.15	87.50	82.14	59.27	54.30	57.30
	ITP _R (Ours)	97.14	88.88	93.75	88.00	76.92	79.17	88.57	61.85	56.95	68.10
Backbone: Llama-3.1-8B-Instruct 🦙											
Prompting	CoT	17.14	14.81	18.75	16.00	15.38	8.33	15.00	3.09	3.31	8.20
	ReAct	22.86	22.22	6.25	24.00	23.08	25.00	21.43	9.27	13.24	19.32
	RAP	25.71	25.93	6.25	32.00	7.69	25.00	22.86	11.34	17.21	16.05
	ITP _I (Ours)	57.14	37.04	31.25	28.00	23.08	33.33	37.86	19.58	19.20	24.30
Training	SFT	85.71	85.20	82.40	89.50	85.70	53.80	79.28	57.21	50.33	47.30
	WKM	85.71	85.19	75.00	80.00	38.46	79.17	77.86	61.34	54.96	65.58
	IWM	87.50	88.90	82.40	94.70	85.90	84.60	85.90	57.56	56.29	58.60
	ITP _R (Ours)	88.57	92.59	93.75	92.00	46.15	91.67	87.14	63.91	57.61	67.50

Table 1: Evaluation of task success rates (%) across ALFWorld, ScienceWorld, and WebShop benchmarks. **Bold** and underlined values denote the best and second-best performance within each backbone model group, respectively.

Type	Method	SoPR (%)	SoWR (%)
Prompting	ReAct	26	22
	RAP	<u>28</u>	<u>26</u>
	ITP_I (Ours)	35 (+7)	28 (+2)
Training	SFT	42	36
	IWM	44	36
	ITP_R (Ours)	68 (+24)	54 (+18)

Table 2: Evaluation of solvable pass rate (SoPR) and solvable win rate (SoWR) on StableToolBench using Qwen3-8B 🦙 as the backbone model.

4.2 Main Results

How does ITP perform across different benchmarks and backbone models? As shown in Table 1, without any additional training, **ITP_I** substantially enhances zero-shot performance compared to strong prompting baselines, such as ReAct and RAP. The trained variant **ITP_R** consistently surpasses all training-based baselines, achieving the highest success rate in every backbone group (e.g.,





Figure 3: Ablation results of ITP_R on ALFWorld and ScienceWorld benchmarks.

88.57% with Qwen3-8B on ALFWorld). These results verify that **ITP** achieves consistent gains across tested benchmarks and backbone models.

Do ITP's performance gains stem from lookahead with world models? The effectiveness of our approach is twofold. First, **ITP_I** uses the same backbone LLM as the prompting baselines, so its gains come solely from conditioning actions on world-model rollouts at inference time, isolating



(a) Adaptive lookahead vs. fixed- k lookahead. Left: success rate. Right: normalized budget. (b) Our ITP_I vs. ReAct + Random Lookahead (c) Our ITP_R vs. SFT + Random Lookahead

Figure 4: Performance-Cost comparison across lookahead strategies. (a): Our **adaptive** lookahead outperforms **fixed** lookahead with higher success rates and lower computational cost. (b) and (c): Both ITP_I and ITP_R with the adaptive lookahead surpass baselines (ReAct and SFT) with a **random** lookahead strategy.

the value of explicit lookahead. Second, ITP_R significantly outperforms training-based alternatives such as IWM and WKM. Its strongest results (e.g., 63.91% on ScienceWorld test-seen with Llama3.1-8B) indicate that combining policy learning with selective adaptive lookahead yields benefits beyond static or implicit methods.

Ablation Study. To assess the contribution of ITP_R , we remove its online reinforced training (**w/o RT**) for an ablation study. In this setting, the model relies solely on initial training and does not learn when to invoke the world model. We use Qwen-3-8B for both the agent policy and world model, and evaluate on ALFWorld and ScienceWorld. As shown in Figure 3, removing RT substantially degrades performance: ALFWorld drops from 88.57% to 71.42%, and ScienceWorld from 59.70% to 46.00%. This confirms that the online reinforced training is a core component rather than a minor training detail. The gap further suggests that supervised learning provides only a basic capability, while reinforcement optimization is crucial for learning when to “imagine” and use the world model efficiently on complex tasks.

4.3 Benefits of Adaptive Lookahead

Superior performance-efficiency trade-off over Fixed Lookahead. We first compare ITP with fixed- k lookahead, where the agent always imagines a constant horizon k . We measure computational cost by the total episode tokens $T = \sum_t (T^{(t)}(\pi_\theta) + T^{(t)}(\mathcal{M}_\phi))$, and define a **Normalized Budget (NB)** metric as follows:

$$\text{NB}(k) = \frac{\bar{T}(k) - \bar{T}(0)}{\bar{T}(K_{\max}) - \bar{T}(0)},$$

which rescales the average token cost $\bar{T}(k)$ to $[0, 1]$. Appendix C reports the setting of $K_{\max}$.

As shown in Figure 4a, fixed- k lookahead is brittle: success rate peaks at a moderate k and then declines, while cost rises sharply with k . In contrast, ITP ’s adaptive lookahead achieves higher success rates with a substantially lower budget, avoiding both the need for global horizon tuning and the high cost of large lookahead steps.

Higher success rates with lower cost than Random Lookahead. Beyond fixed-horizon lookahead, we compare ITP with *random lookahead*, which samples K_t independently at each step. This isolates the effect of adaptive horizon selection from simply using a varying horizon. We evaluate the performance-cost trade-off on ALFWorld, using Qwen3-8B as both the policy and world model. We run 140 tasks grouped into 14 folds and report fold-averaged SR and NB, where each point in Figure 4b and Figure 4c corresponds to one fold.

Our adaptive lookahead consistently outperforms the random strategy, achieving higher SR with lower and more stable budgets. This shows that the gains come from state-conditioned allocation of lookahead, rather than horizon variability.

4.4 Impact of World Models

Does the choice of world-model backbones affect ultimate performance? We evaluate the sensitivity of ITP to the underlying world models by varying the world-model backbones among Qwen3-8B, Llama-3.1-8B, and DeepSeek-V3.2 (DeepSeek-AI et al., 2025). We fix the agent policy to Qwen3-8B and evaluate on the ALFWorld benchmark. As illustrated in Figure 5, the choice of backbone is most consequential in the training-free (ITP_I) setting. Specifically, despite its superior scale, the zero-shot DeepSeek-V3.2 exhibits a noticeable performance deficit, likely due to a lack of domain-specific alignment with ALFWorld’s state-transition dynamics. In contrast, Qwen and Llama




Figure 5: Impact of different world-model backbones. We report success rates of (a) ITP₁ and (b) ITP_R across six task types on ALFWorld.




Figure 6: World model’s reliability and lookahead horizon statistics on ALFWorld with Qwen3-8B. (a) Longer lookahead horizon reduces *Fact-F1* of imagined states, while success rate peaks at $k=3$. (b) Longer-horizon tasks induce larger lookahead steps.

backbones maintain robust success rates, suggesting better inherent compatibility. Crucially, ITP_R substantially narrows this gap. After optimizing the use of the DeepSeek-V3.2 world model, ITP’s performance becomes highly competitive. These results demonstrate that ITP is model-agnostic and can effectively distill world-modeling capabilities from various architectures. Furthermore, targeted world-modeling training is essential in specific interactive environments.

How reliable is the world model across lookahead horizons? We evaluate adaptive lookahead from two perspectives: world-model reliability and the selection of the lookahead horizon (k). We use *Fact-F1* as the reliability metric, computed by comparing canonicalized atomic facts extracted from predicted versus ground-truth states. As shown in Figure 6a, *Fact-F1* degrades as the horizon increases due to compounding errors. Interestingly, the agent’s success rate follows a bell curve, peaking at $k=3$, which suggests a trade-off between foresight and prediction accuracy. Figure 6b further demonstrates that ITP’s learned k distributions

adapt to task complexity: longer-horizon environments (e.g., ScienceWorld) favor larger lookahead values. These results validate ITP’s core design: effective planning requires dynamically anchoring the lookahead horizon within the world model’s reliable predictive range.

5 Related Work

LLM-based Agents. LLM-based agents use language models as policies that map instructions and partial observations to actions. A significant line of work formulates agent learning as trajectory- or step-level optimization. For example, ETO (Song et al., 2024) frames learning as exploration-based trajectory optimization, while IPR (Xiong et al., 2024) and E²CL (Wang et al., 2024b) refine agent behavior via iterative revision and correction signals. More recent post-training methods further improve robustness and generalization through reflective updates, such as Agent-R (Yuan et al., 2025), STeCa (Wang et al., 2025b), and AgentRefine (Fu et al., 2025). However, these methods primarily focus on learning from historical traces, often leaving the agent in a state of “shallow grounding” where it lacks active foresight to anticipate future environmental shifts before execution.

World Models for Planning. World models provide a “mental sandbox” for agents, enabling model-based decision-making by predicting environment dynamics (LeCun, 2022; Xiang et al., 2023; Wang et al., 2025e; Liu et al., 2026). Recent research has increasingly positioned LLMs as either implicit or explicit world models for task planning. For instance, RAP (Hao et al., 2023) treats LLM reasoning as a planning process over an implicit state space, while other works explicitly construct world models to simulate future states and verify plan feasibility (Guan et al., 2023). To improve simulation fidelity, methods like WKM (Qiao et al., 2024), DMWM (Wang et al., 2025d), and IWM (Zhang et al., 2026) incorporate world knowledge or historical experience to guide rollouts, while D2PO (Wang et al., 2025e) and internalizing strategies (Chen et al., 2025) align planning behavior with environmental feedback. Further specialized work, such as WebEvolver (Fang et al., 2025), explores world model-based imagination in web agents. Unlike existing methods constrained by rigid, fixed-horizon planning, our work introduces an adaptive lookahead mechanism that dynamically adjusts the imagination horizon, enabling the agent

to optimally navigate the trade-off between foresight reliability and computational overhead.

6 Conclusion

In this paper, we propose Imagine-then-Plan (ITP), an agent learning framework that equips LLM-based agents with adaptive lookahead with world models. By extending the conventional Partially Observable MDP to a Partially Observable and Imaginable MDP, ITP enables agents to explicitly reason over both the present and foresighted trajectories, addressing the shallow grounding limitation of reactive decision-making. We instantiate ITP into both an inference-time variant and a reinforcement-trained variant. Extensive experiments demonstrate that our approach significantly improves task success and robustness across domains. Further analyses validate the critical role of the adaptive control of the imagination horizon. We believe that ITP provides a principled advancement toward more deliberative utilization of world models for autonomous agent learning.

Limitations

While our approach demonstrates superior performance compared to baseline methods, it is important to acknowledge the limitations as follows: (1) Current evaluation primarily focuses on interactive text-based benchmarks. While these environments provide a rigorous test of long-horizon reasoning, they do not fully capture the challenges of multi-modal or real-world physical environments. The transition from textual state descriptions to visual or sensorimotor observations may introduce additional noise that could affect the stability of the adaptive lookahead mechanism. (2) Although our adaptive lookahead mechanism is designed to optimize efficiency by scaling the imagination horizon, the use of world models inherently introduces higher inference-time overhead compared to purely reactive agents. While higher success rates in high-stakes tasks often justify this trade-off, further optimization, e.g., via speculative decoding (Xia et al., 2023) or balancing exploration and exploitation (Wei and Yang, 2013), is needed for real-time applications. We will leave these directions as our future work.

Ethics Statement

We strictly follow the protocols governing the academic use of all LLMs. Our study is conducted

in simulated, text-based benchmarks and involves no human subjects or personally identifying information. We cite and comply with the licenses of all models, datasets, and software used. We acknowledge that world-model-based lookahead may introduce potential risks if transferred beyond our simulated evaluation settings: prediction errors and hallucinated rollouts could lead to unsafe or unintended actions in open-world tool-use systems or robots, and the generation of imagined trajectories increases inference-time computation and energy cost. Additionally, while AI assistants (e.g., Cursor and ChatGPT) were partially utilized for coding and linguistic refinement, we affirm that all core content and findings in this paper are the original work of the authors.

Acknowledgements

This work was supported by the Fundamental Research Funds for the Central Universities (YJ202620 and 2026SCUQJTX022B), the General Research Fund (GRF) of the Research Grants Council of Hong Kong (PolyU 15205325), and was supported in part by the PolyU Postdoc Matching Fund (4-W40Z). The authors would like to thank the anonymous reviewers for their valuable feedback and constructive suggestions.

References

- Karl Johan Åström. 1965. [Optimal control of markov processes with incomplete state information I](#). *Journal of Mathematical Analysis and Applications*, 10:174–205.
- Baian Chen, Chang Shu, Ehsan Shareghi, Nigel Collier, Karthik Narasimhan, and Shunyu Yao. 2023. FireAct: Toward language agent fine-tuning. *arXiv preprint arXiv:2310.05915*.
- Shiqi Chen, Tongyao Zhu, Zian Wang, Jinghan Zhang, Kangrui Wang, Siyang Gao, Teng Xiao, Yee Whye Teh, Junxian He, and Manling Li. 2025. Internalizing world models via self-play finetuning for agentic RL. *arXiv preprint arXiv:2510.15047*.
- DeepSeek-AI, Aixin Liu, Aoxue Mei, Bangcai Lin, Bing Xue, Bingxuan Wang, Bingzheng Xu, Bochao Wu, Bowei Zhang, Chaofan Lin, Chen Dong, and 1 others. 2025. [DeepSeek-V3.2: Pushing the frontier of open large language models](#). *arXiv preprint arXiv:2512.02556*.
- Tianqing Fang, Hongming Zhang, Zhisong Zhang, Kaixin Ma, Wenhao Yu, Haitao Mi, and Dong Yu.

2025. [WebEvolver: Enhancing web agent self-improvement with co-evolving world model](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 8959–8975, Suzhou, China. Association for Computational Linguistics.
- Lang Feng, Zhenghai Xue, Tingcong Liu, and Bo An. 2025. [Group-in-group policy optimization for LLM agent training](#). In *Advances in Neural Information Processing Systems*, volume 38, pages 46375–46408. Curran Associates, Inc.
- Dayuan Fu, Keqing He, Yejie Wang, Wentao Hong, Zhuoma GongQue, Weihao Zeng, Wei Wang, Jingang Wang, Xunliang Cai, and Weiran Xu. 2025. AgentRefine: Enhancing agent generalization through refinement tuning. In *The Thirteenth International Conference on Learning Representations (ICLR)*.
- Pascale Fung, Yoram Bachrach, Asli Celikyilmaz, Kamalika Chaudhuri, Delong Chen, Willy Chung, Emmanuel Dupoux, Hongyu Gong, Hervé Jégou, Alessandro Lazaric, and 1 others. 2025. Embodied ai agents: Modeling the world. *arXiv preprint arXiv:2506.22355*.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. [The Llama 3 herd of models](#). *arXiv preprint arXiv:2407.21783*.
- Lin Guan, Karthik Valmeekam, Sarath Sreedharan, and Subbarao Kambhampati. 2023. Leveraging pre-trained large language models to construct and utilize world models for model-based task planning. In *Advances in Neural Information Processing Systems*, volume 36, pages 79081–79094.
- Zhicheng Guo, Sijie Cheng, Hao Wang, Shihao Liang, Yujia Qin, Peng Li, Zhiyuan Liu, Maosong Sun, and Yang Liu. 2024. [StableToolBench: Towards stable large-scale benchmarking on tool learning of large language models](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 11143–11156, Bangkok, Thailand. Association for Computational Linguistics.
- Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. 2023. Mastering diverse domains through world models. *arXiv preprint arXiv:2301.04104*.
- Shibo Hao, Yi Gu, Haodi Ma, Joshua Hong, Zhen Wang, Daisy Wang, and Zhiting Hu. 2023. [Reasoning with language model is planning with world model](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 8154–8173, Singapore. Association for Computational Linguistics.
- Levente Kocsis and Csaba Szepesvári. 2006. [Bandit based monte-carlo planning](#). In *Machine Learning: ECML 2006*, volume 4212 of *Lecture Notes in Computer Science*, pages 282–293. Springer.
- Yann LeCun. 2022. [A path towards autonomous machine intelligence](#). OpenReview. Version 0.9.2, 2022-06-27.
- Guohao Li, Hasan Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. 2023. [CAMEL: Communicative agents for “mind” exploration of large language model society](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 51991–52008. Curran Associates, Inc.
- Yixia Li, Hongru Wang, Jiahao Qiu, Zhenfei Yin, Dongdong Zhang, Cheng Qian, Zeping Li, Xiaoteng Ma, Guanhua Chen, and Heng Ji. 2026. [From word to world: Can large language models be implicit text-based world models?](#) In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8084–8111, San Diego, California, United States. Association for Computational Linguistics.
- Youwei Liu, Jian Wang, Hanlin Wang, and Wenjie Li. 2026. [COMAP: Co-evolving world models and agent policies for LLM agents](#). *arXiv preprint arXiv:2606.02372*.
- Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. 2016. [Asynchronous methods for deep reinforcement learning](#). In *International conference on machine learning*, volume 48, pages 1928–1937. PMLR.
- Shuofei Qiao, Runnan Fang, Ningyu Zhang, Yuqi Zhu, Xiang Chen, Shumin Deng, Yong Jiang, Pengjun Xie, Fei Huang, and Huajun Chen. 2024. Agent planning with world knowledge model. In *Advances in Neural Information Processing Systems*, volume 37, pages 114843–114871.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. Reflexion: Language agents with verbal reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 36, pages 8634–8652.
- Mohit Shridhar, Xingdi Yuan, Marc-Alexandre Côté, Yonatan Bisk, Adam Trischler, and Matthew Hausknecht. 2021. [ALFWorld: Aligning text and embodied environments for interactive learning](#). In *International Conference on Learning Representations*.
- Yifan Song, Da Yin, Xiang Yue, Jie Huang, Sujian Li, and Bill Yuchen Lin. 2024. [Trial and error: Exploration-based trajectory optimization of LLM agents](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7584–7600, Bangkok, Thailand. Association for Computational Linguistics.
- Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. 2024a. [Voyager: An open-ended embodied agent with large language models](#). *Transactions on Machine Learning Research*.

- Hanlin Wang, Chak Tou Leong, Jian Wang, and Wenjie Li. 2024b. [E²CL: Exploration-based error correction learning for embodied agents](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 7626–7639, Miami, Florida, USA. Association for Computational Linguistics.
- Hanlin Wang, Chak Tou Leong, Jiashuo Wang, Jian Wang, and Wenjie Li. 2025a. SPA-RL: Reinforcing LLM agents via stepwise progress attribution. *arXiv preprint arXiv:2505.20732*.
- Hanlin Wang, Jian Wang, Chak Tou Leong, and Wenjie Li. 2025b. [STeCa: Step-level trajectory calibration for LLM agent learning](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 11597–11614, Vienna, Austria. Association for Computational Linguistics.
- Kangrui Wang, Pingyue Zhang, Zihan Wang, Yaning Gao, Linjie Li, Qineng Wang, Hanyang Chen, Yiping Lu, Zhengyuan Yang, Lijuan Wang, Ranjay Krishna, Jiajun Wu, Fei-Fei Li, Yejin Choi, and Manling Li. 2025c. [VAGEN: Reinforcing world model reasoning for multi-turn VLM agents](#). In *Advances in Neural Information Processing Systems*, volume 38, pages 172871–172933. Curran Associates, Inc.
- Lingyi Wang, Rashed Shelim, Walid Saad, and Naren Ramakrishnan. 2025d. [DMWM: Dual-mind world model with long-term imagination](#). In *Advances in Neural Information Processing Systems*, volume 38, pages 4911–4945. Curran Associates, Inc.
- Ruoyao Wang, Peter Jansen, Marc-Alexandre Côté, and Prithviraj Ammanabrolu. 2022. [ScienceWorld: Is your agent smarter than a 5th grader?](#) In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11279–11298, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Siyin Wang, Zhaoye Fei, Qinyuan Cheng, Shiduo Zhang, Panpan Cai, Jinlan Fu, and Xipeng Qiu. 2025e. [World modeling makes a better planner: Dual preference optimization for embodied task planning](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 21518–21537, Vienna, Austria. Association for Computational Linguistics.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 24824–24837. Curran Associates, Inc.
- Xiao-Yong Wei and Zhen-Qun Yang. 2013. [Coaching the exploration and exploitation in active learning for interactive video retrieval](#). *IEEE Transactions on Image Processing*, 22(3):955–968.
- Heming Xia, Tao Ge, Peiyi Wang, Si-Qing Chen, Furu Wei, and Zhifang Sui. 2023. [Speculative decoding: Exploiting speculative execution for accelerating seq2seq generation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 3909–3925, Singapore. Association for Computational Linguistics.
- Jiannan Xiang, Tianhua Tao, Yi Gu, Tianmin Shu, Zirui Wang, Zichao Yang, and Zhiting Hu. 2023. [Language models meet world models: Embodied experiences enhance language models](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 75392–75412. Curran Associates, Inc.
- Weimin Xiong, Yifan Song, Xiutian Zhao, Wenhao Wu, Xun Wang, Ke Wang, Cheng Li, Wei Peng, and Sujian Li. 2024. [Watch every step! LLM agent learning via iterative step-level process refinement](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 1556–1572, Miami, Florida, USA. Association for Computational Linguistics.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, and 22 others. 2024. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*.
- Shunyu Yao, Howard Chen, John Yang, and Karthik Narasimhan. 2022. WebShop: Towards scalable real-world web interaction with grounded language agents. In *Advances in Neural Information Processing Systems*, volume 35, pages 20744–20757.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik Narasimhan. 2023a. Tree of thoughts: Deliberate problem solving with large language models. In *Advances in Neural Information Processing Systems*, volume 36, pages 11809–11822.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2023b. [ReAct: Synergizing reasoning and acting in language models](#). In *The Eleventh International Conference on Learning Representations*.
- Siyu Yuan, Zehui Chen, Zhiheng Xi, Junjie Ye, Zhengyin Du, and Jiecao Chen. 2025. Agent-R: Training language model agents to reflect via iterative self-training. *arXiv preprint arXiv:2501.11425*.
- Kai Zhang, Xiangchao Chen, Bo Liu, Tianci Xue, Zeyi Liao, Zhihan Liu, Xiyao Wang, Yuting Ning,

Zhaorun Chen, Xiaohan Fu, Jian Xie, Yuxuan Sun, Boyu Gou, Qi Qi, Zihang Meng, Jianwei Yang, Ning Zhang, Xian Li, Ashish Shah, and 11 others. 2026. [Agent learning via early experience](#). In *Proceedings of the 43rd International Conference on Machine Learning*, volume 306 of *Proceedings of Machine Learning Research*, Seoul, South Korea. PMLR.

A Datasets and Preprocessing

A.1 Datasets

We evaluate our method on four representative agent benchmarks: ALFWorld¹ (Shridhar et al., 2021), ScienceWorld² (Wang et al., 2022), WebShop³ (Yao et al., 2022), and StableToolBench⁴ (Guo et al., 2024). Table 3 presents task descriptions and data statistics.

ALFWorld: ALFWorld is a text-based household embodied benchmark where an agent interacts with a simulated environment via natural-language observations and admissible text actions. Each episode specifies a goal instruction that can be instantiated from compositional templates, and success requires planning over multi-step action sequences under partial observability.

- **PICK:** to find an object of the desired type, pick it up, navigate to the correct location/receptacle, and place it there.
- **CLEAN:** to find the target object, pick it up, go to a sink/basin, wash it by turning on the faucet, then navigate to the target receptacle and place it.
- **HEAT:** to find the target object, pick it up, go to a microwave, heat it by turning on the microwave, then place it at the specified location.
- **COOL:** to find the target object, pick it up, go to a fridge, cool it by placing it inside the fridge, then return and place it at the specified location.
- **LOOK:** to find the target object, locate a light source, turn it on, and examine the object with the light while holding it.
- **PICK2:** to find the first target object and place it at the destination, then find a second object of the same type, return to the destination, and place it together with the first one.

ScienceWorld: ScienceWorld is a text-based interactive science environment that evaluates an

agent’s ability to solve procedural and reasoning-intensive tasks grounded in everyday scientific phenomena. Compared to household tasks, ScienceWorld typically involves longer horizons and requires the agent to combine information gathering, tool use, and multi-step experimentation.

WebShop: WebShop is a text-based web shopping benchmark in which an agent interacts with simulated e-commerce webpages to search for products that satisfy a user instruction. Compared with embodied environments, WebShop requires the agent to reason over longer and noisier textual observations such as search results, product titles, attributes, and descriptions, and to perform multi-step browsing, comparison, and selection before making a final purchase decision.

StableToolBench: StableToolBench is a multi-turn tool-use benchmark that evaluates an agent’s ability to solve tasks by invoking executable tools under realistic interaction constraints. In contrast to text navigation environments, StableToolBench emphasizes the correctness and robustness of tool calling, intermediate execution feedback handling, and solvability-aware evaluation. Although episodes are typically shorter, each decision is more sensitive because an incorrect tool invocation may immediately derail downstream progress.

We employ Qwen3-8B as the backbone model for different methods. Following existing studies, we adopt solvable pass rate (**SoPR**) and solvable win rate (**SoWR**) as the evaluation metrics. SoPR measures the average task success on the “solvable” subset (i.e., queries that are answerable given the available tools/KB, where each instance is judged as solved/unsolved/uncertain and mapped to 1/0/0.5). SoWR is a pairwise metric on the same solvable subset, reporting how often our method wins against a fixed baseline. In particular, SoWR is computed on the solvable subset as the average head-to-head win rate against the baseline set (ReAct, RAP, SFT, and IWM), using the official evaluation prompts and tool API.

A.2 Data Preprocessing

We use the provided expert trajectories as supervision to warm-start the agent policy. For the world model, we repurpose the same expert rollouts into transition-level supervision for a text world model by emitting one record per environment step, containing: (i) a compact state string (goal + current

¹<https://github.com/alfworld/alfworld>

²<https://github.com/allenai/ScienceWorld>

³<https://github.com/princeton-nlp/WebShop>

⁴<https://github.com/THUDM/StableToolBench>

Domain	Task Description	Dataset Statistics
ALFWorld		
Household, text-based embodied tasks	Six compositional task families: PICK, CLEAN, HEAT, COOL, LOOK, PICK2.	Training: 3,119 Test: 140
ScienceWorld		
Elementary science curriculum in an interactive text environment	30 subtasks with many variations (entities, initial conditions, distractors, room layouts), partitioned following the benchmark protocol.	Training: 1,483 Test-Seen: 194 Test-Unseen: 151
WebShop		
Goal-driven web shopping in a text-based browsing environment	Product search, browsing, attribute comparison, and final item selection based on user instructions over noisy webpage observations.	Training: 1,824 Test: 210
StableToolBench		
Multi-turn tool-use tasks with executable tool interactions	Tool invocation, intermediate execution feedback handling, and solvability-aware evaluation under realistic function-calling constraints.	Training: 1,972 Test: 169

Table 3: Dataset statistics. We report dataset splits following the standard benchmark protocol.

observation/page state/tool context + optional inventory or execution feedback), (ii) the executed expert action, and (iii) the next observation, optionally augmented with scalar signals such as reward and done. All records are stored in JSONL format and serialized as dialogue-style causal-LM input-output pairs, following standard SFT practice for text-based world models. The state serialization is benchmark-specific but follows a unified transition-level format. For ALFWorld and ScienceWorld, the state mainly consists of the task instruction, current observation, and optional inventory. For WebShop, the state additionally includes webpage content, product attributes, and browsing context. For StableToolBench, the state further incorporates tool descriptions, intermediate tool calls, and execution feedback. This unified serialization allows the

same world-model training and adaptive-lookahead pipeline to be applied across embodied, web, and tool-use settings.

To annotate the number of lookahead steps (i.e., K) in ITP_R , we precompute imagined rollouts using the trained world model for a small discrete set of horizons ($k_candidates$) at each expert step. Each imagined rollout is summarized into a short lookahead text snippet (`lookahead_summary`). We then compute a per-horizon score that balances improved expert-action likelihood against deeper rollout cost, and store the resulting pseudo label (k_label) along with all scores. This adaptive-horizon annotation procedure is applied consistently across all four benchmarks under their respective state/action serialization formats.

B Additional Baseline Comparisons

B.1 Tree Search and RL-from-World-Model Methods

We additionally compare ITP against two baseline families that represent natural alternatives to our formulation. For inference-time planning, we include **ToT** (Yao et al., 2023a) and **MCTS** (Kocsis and Szepesvári, 2006), implemented under the same backbone/world-model interface and evaluated on ALFWorld with **Success Rate (SR)** and **Normalized Budget (NB)**. For training-based comparison, we include **D2PO** (Wang et al., 2025e) and **VaGen** (Wang et al., 2025c). Since D2PO and VaGen were originally developed for VLM settings, we adapt their training objectives to our text-only environment using the same state/action serialization, benchmark interface, and evaluation protocol as our method. All methods use the same executable action space and are compared under matched backbone and benchmark settings.

Table 4 reports the results. We obtain the following findings: (1) **Tree search alone is insufficient**. Both ToT and MCTS underperform ITP_I and are less budget-efficient, while ITP_I achieves the best trade-off (41.43 SR / 0.25 NB versus RAP at 28.57 / 0.42). (2) **RL-from-world-model training is helpful but not sufficient**. D2PO and VaGen consistently improve over Base SFT (70.71 $\rightarrow$ 76.28/78.89), but ITP_R further advances the frontier to 88.57 SR with only a modest NB increase (0.21). Taken together, these results suggest that the key benefit of our framework lies not in search or world-model usage alone, but in integrating world-model-based learning with *adaptive lookahead* as

Tree Search-based			World-model-based Training		
Method	SR (%) $\uparrow$	NB $\downarrow$	Method	SR (%) $\uparrow$	NB $\downarrow$
RAP	28.57	0.42	SFT	70.71	—
ToT	16.29	0.58	D2PO	76.28	0.29
MCTS	25.00	0.39	VaGen	78.89	0.18
ITP_I (Ours)	41.43	0.25	ITP_R (Ours)	88.57	0.21

Table 4: Comparison with additional baselines on ALFWorld. Left: tree-search-based methods under the same backbone/world-model interface. Right: RL-from-world-model baselines adapted to our text-only setting using the same state/action serialization and evaluation protocol. ITP consistently achieves the strongest overall performance, and for prompting-based methods also the best SR–budget trade-off.

Method	World Model	Prompting-based SR	Training-based SR
Self-Refine	✗	22.59	68.89
Multi-Attempt	✗	20.48	59.68
K-candidate Selection	✗	36.42	78.14
ITP_I (Ours)	✓	41.43 _{+5.01}	—
ITP_R (Ours)	✓	—	88.57 _{+10.43}

Table 5: Compute-matched comparison with no-world-model alternatives. Subscripts denote the gain over the strongest no-world-model baseline in each setting.

an explicit budget/error control mechanism.

B.2 Compute-Matched Comparisons with No-World-Model Alternatives

To test whether **ITP** gains are merely due to extra compute, we compare against *compute-matched* no-world-model baselines: **Self-Refine**, **Multi-Attempt**, and **K-candidate Selection**. These methods spend a similar budget on repeated direct action generation or refinement, but do not perform future-state imagination.

Table 5 shows that no-world-model alternatives still lag behind **ITP** under matched compute. **ITP_I** improves over the strongest prompting-based baseline by 5.01 points, and **ITP_R** improves over the strongest training-based baseline by 10.43 points. This suggests that the gains come not from extra compute alone, but from allocating it to **state-conditioned adaptive lookahead over imagined future states**.

C Implementation Details

For fine-tuning, we employed several open-source models, including Qwen3-8B (Yang et al., 2025)⁵, Qwen2.5-7B (Yang et al., 2024)⁶, and Llama3.1-

8B (Grattafiori et al., 2024)⁷. All experiments were conducted on a computational cluster equipped with $2 \times$ NVIDIA A100 80GB GPUs. We report additional implementation details for reproducibility, including world-model training details, deployment profiling, stage-wise training cost of **ITP_R**, and the prompting templates.

C.1 World Model Training Details

Our text world model is trained to predict the next textual observation conditioned on the current textual state and executed action, i.e., to model $p_\phi(s_{t+1} \mid s_t, a_t)$. Following the transition serialization described in Appendix A, each training instance is formatted as a dialogue-style causal-LM example, where the input contains the task instruction, current observation, optional inventory or tool feedback, and action, and the target is the next observation.

We further construct adaptive-horizon supervision by scoring imagined rollouts over a small discrete set of horizon candidates at each expert step, producing one K -label record per step. Figure 7 visualizes the resulting data scale across benchmarks, while Table 6 reports the optimization setup and convergence diagnostics for representative world-model training runs.

⁵<https://huggingface.co/Qwen/Qwen3-8B/blob/main/LICENSE>

⁶<https://huggingface.co/Qwen/Qwen2.5-7B-Instruct/blob/main/LICENSE>

⁷<https://huggingface.co/meta-llama/Llama-3.1-8B-Instruct/blob/main/LICENSE>




Figure 7: Visualization of world-model training data scale across benchmarks. For each benchmark, we show the number of expert episodes (D_{exp}) and rollout episodes (D_{roll}) used for training data construction.

Item	Value
Objective	NLL / CE on $p_\phi(s_{t+1} s_t, a_t)$
Training epochs / steps	3 epochs / 3,450 steps
Wall-clock time	8.69 h
Throughput	0.11 steps/s ; 1.764 samples/s
Held-out validation split	5% of D_{WM} , fixed seed
Best / final validation loss	0.052 / 0.055
Effective batch size (incl. accum.)	16
Max sequence length / truncation	2048
Optimizer + LR schedule	AdamW + cosine decay with warmup
Precision + hardware	bf16 on 2× NVIDIA A100 GPUs
Rollout collection cost	9.5 GPU-hours on 1× A100 GPU

Table 6: World-model training diagnostics for Qwen3-8B as the backbone. We report optimization settings, convergence indicators, and rollout collection cost.

C.2 Deployment Profiling

Normalized Budget (NB) captures algorithmic token efficiency, but does not directly reflect deployment-facing runtime characteristics. We therefore additionally profile interactive inference and report wall-clock latency, throughput, and GPU memory usage under representative deployment settings.

C.3 ITP_R Training Algorithm

Algorithm 1 presents the full training procedure of ITP_R, including pseudo-label generation for adaptive horizon supervision, warm-up training with lookahead labels, and online actor-critic optimization.

C.4 Training Cost of ITP_R

To improve training-cost transparency, we report the stage-wise runtime of ITP_R. The training pipeline consists of three stages: (i) pseudo-label generation for adaptive horizon supervision, (ii) supervised warm-up of the policy and K -head, and (iii) online actor-critic optimization.

Metric	p50	p95	p99
Decision latency per step (s)	14.2	31.8	44.7
Policy time per step (s)	5.1	11.6	17.4
Imagination time per step (s)	7.8	18.9	28.5
Tokens / second	43.5	31.2	24.0

Episodes / hour: 11.8

Peak allocated / reserved VRAM (policy GPU): 31.0 / 38.4 GB

Peak allocated / reserved VRAM (WM GPU): 28.7 / 35.9 GB

Table 7: Deployment-oriented profiling of ITP_R. We report interactive latency, throughput, and peak GPU memory usage.

Algorithm 1 ITP_R: Reinforced Training with Adaptive Lookahead

Input: Dataset $\mathcal{D} = \{(s_t, a_t^*)\}$; Agent policy π_{θ_0} ; World model $\mathcal{M}_\phi$; Parameters $K_{\text{max}}, \lambda_K, \eta, \alpha, \beta$.

Output: Agent policy π_θ ; Lookahead predictor P_θ .

```

// Stage 1: Pseudo-Labeling Lookahead Horizon
1: for each episode in  $\mathcal{D}$  do
2:   Cache  $\{\hat{\tau}_t^{(k)}\}_{k=0}^{K_{\text{max}}}$  via  $\mathcal{M}_\phi$  & teacher-forced  $a_t^*$ .
3:   for each step  $t$  do
4:      $S_t(k) \triangleq \log p_{\theta_0}(a_t^* | s_t, \hat{\tau}_t^{(k)})$ .
5:      $\tilde{K}_t \leftarrow \arg \max [S_t(k) - \lambda_K k]$ .
6:     Store  $(s_t, a_t^*, \tilde{K}_t)$  in  $\mathcal{D}_K$ .
7:   end for
8: end for

// Stage 2: Warm-Up Training with Lookahead
9: for each episode in  $\mathcal{D}_K$  do
10:  Sample  $(s_t, a_t^*, \tilde{K}_t) \sim \mathcal{D}_K$ .
11:  Update  $\theta$  via Eq. (6).
12: end for

// Stage 3: Online Actor-Critic Optimization
13: for each episode in  $\mathcal{D}$  do
14:  Sample  $K_t \sim P_\theta(K | s_t)$ , query  $\mathcal{M}_\phi$  for  $\hat{\tau}_t^{(K_t)}$ .
15:  Update  $\theta$  via Eq. (8).
16: end for
17: return  $\pi_\theta$  and  $P_\theta$ .

```

C.5 Parameter Settings

Table 9 summarizes the hyperparameters used for training and inference. Unless otherwise specified, we apply the same configuration across different backbone models. During exploration, the agent samples actions with temperature 0.7. Due to the difference in task trajectory lengths and planning horizons across benchmarks (e.g., ALFWorld has an average of 8 steps, ScienceWorld 15, WebShop 11, and StableToolBench 4), during ITP_R training, we set K_{max} to 5 for ALFWorld, 8 for ScienceWorld, 7 for WebShop, and 3 for StableToolBench.

C.6 Comparison with Strong Planning and Search Baselines

We further compare ITP with recent search-based, world-model-based, and trajectory-optimization

Stage	Wall-clock	GPU-hours
I: Lookahead Pseudo-Labeling	6.63 h	13.27
II: Warm-Up Training	6.38 h	12.77
III: Online Optimization	4.78 h	9.57
Total	17.80 h	35.60

Table 8: Stage-wise training cost of ITP_R . Stage I corresponds to pseudo-label generation, Stage II to warm-up supervised training, and Stage III to online actor-critic optimization.

Name	Value
<i>Warm-Up Training</i>	
cutoff_len	2048
epochs	3
per_device_train_batch_size	1
gradient_accumulation_steps	16
learning_rate	2×10^{-5}
warmup_ratio	0.03
lr_scheduler	cosine
fp16 / bf16	False / True
gradient_checkpointing	False
lora_r	8
lora_alpha	16
lora_dropout	0.05
merge_lora	True
<i>Online A2C Optimization</i>	
γ (discount)	0.99
rl_learning_rate	5×10^{-6}
λ_K (lookahead penalty)	0.2
λ_{step} (step cost)	0.01
success_bonus	0.01
invalid_action_penalty	-0.1
η	0.5
α	1.0
β	0.01
max_grad_norm	1.0
<i>Inference Stage</i>	
do_sample (exploration)	True
temperature (exploration)	0.7
top_p (exploration)	0.9
action_max_new_tokens	16
imagine_action_max_new_tokens	12
wm_max_new_tokens	192
K_{max} (ALFWorld)	5
K_{max} (ScienceWorld)	8
K_{max} (WebShop)	7
K_{max} (StableToolBench)	3

Table 9: Hyperparameter setup.

methods on ALFWorld. All methods use the same benchmark split, executable action space, backbone

Method	ALFWorld SR (%) $\uparrow$
CoT	13.57
ReAct	19.29
ToT	16.29
MCTS	25.00
RAP	28.57
ITP_I	41.43

Table 10: Comparison with training-free planning and search methods on ALFWorld.

Method	ALFWorld SR (%) $\uparrow$
SFT	70.71
ETO	70.50
D2PO	76.28
VAGEN	78.89
WKM	79.29
QLASS	80.35
IWM	82.14
ITP_R	88.57

Table 11: Comparison with training-based planning methods on ALFWorld.

family, and success-rate metric as ITP .

ITP_I outperforms RAP, the strongest training-free baseline, by 12.86 percentage points. ITP_R improves over IWM, the strongest training-based baseline, by 6.43 points. These results show that the advantage of ITP cannot be explained solely by search, additional rollout computation, or world-model-based policy training. Instead, ITP benefits from adaptively determining when imagination should be invoked and how deeply the model should look ahead.

C.7 Prompting Templates

We provide the prompt templates for the agent policy and for the adaptive K -step lookahead inference procedure of ITP_I . Figure 8 shows the prompting template for the ReAct baseline. Figure 9 illustrates the adaptive horizon selector, where the policy model maps the task instruction and the *current* textual state s_t to a single discrete lookahead depth $K \in [0, K_{\text{max}}]$. Figure 10 presents the world-model foresight generator, which conditions only on s_t and produces a concise K -step imagined trajectory enclosed by $\langle \text{Foresight} \rangle \dots \langle / \text{Foresight} \rangle$. Figure 11 depicts the foresight-conditioned agent prompt, where the policy model consumes the task, the current state s_t , and the generated foresight to produce a ReAct-style response, and outputs an admissible environment action via exact copying.

Prompt Template for Base Agent Policy

Interact with a household to solve a task. You are an intelligent agent in a household environment. Your goal is to perform valid actions to complete the task goal.

At each step t , you will be given the **task goal**, the **current state** (observation and optional inventory), and the **previous step context** (the previous action and its resulting environment feedback/observation). You must follow the **ReAct** paradigm: **Reason** $\rightarrow$ **Thought** $\rightarrow$ **Action**.

You need to process the information in a specific order:

1. **Reason:** Briefly interpret the current state and the latest environment feedback. Identify what has been achieved and what remains.
2. **Thought:** Plan the next few steps to make progress toward the task goal based on the reasoning.
3. **Action:** Output **exactly one** next action that is valid in environment.

After each turn, the environment will provide immediate feedback (a new observation, and optionally inventory updates). If the environment outputs “Nothing happened”, the previous action is invalid; revise your plan and try a different valid action.

Your response must use the following format:

```
Reason: <brief interpretation of state/feedback>
Thought: <short plan for next steps>
Action: <EXACTLY ONE valid action line>
```

Inputs at step t :

```
Task goal:
{task_goal}

Current state (observation + optional inventory):
{state}

Previous action (optional):
{prev_action}

Latest environment feedback / observation:
{feedback}
```

Figure 8: Prompt template used for base agent policy on ALFWorld and ScienceWorld benchmarks.

Adaptive Selection of K (PolicyModel.decide_k)

System prompt: You are a planning assistant. Your job is to decide how many steps of look-ahead are needed right now. Given a task instruction and the dialogue/action history, output a single integer K in the range $[0, K_{\max}]$

Task instruction: task

History trajectory (thoughts, actions, observations so far): history

Question: Output a single integer K in $[0, K_{\max}]$.

Your response should be:

```
K
```

Figure 9: Prompt used to adaptively select the lookahead horizon K from the task instruction and trajectory history. The output must be a single integer within the range $[0, K_{\max}]$.

K -Step Foresight Generation (WorldModel.imagine)

System prompt: You are a world model for the ALFWorld environment. Given an action/observation history, imagine the next few steps, describing likely observations and key objects.

User prompt: Predict the next k step(s). Return a concise plan inside `<Foresight>...</Foresight>` with numbered steps.

User prompt: History so far: `<history>`

Your response should be:

```
Foresight state: <Foresight>...</Foresight>
```

Figure 10: Prompt used by the world model to generate a K -step foresight trajectory.

Foresight-Conditioned Action (PolicyModel.reflect_and_act)

You are an agent that **first imagines and then acts**. At each step, you will be given the task instruction, the **current state**, and a K -step foresight trajectory imagined by a world model. Use the foresight to reflect on progress and bottlenecks, then decide the next admissible action.

You will be given:

1. **Task instruction** ({task})
2. **Current state** ({state})
3. K -step foresight trajectory from the world model ({foresight})
4. **Admissible actions (instance-level)** ({admissible_actions_joined_by_newlines})

After each turn, the environment will give you immediate feedback based on which you plan your next few steps. If the environment output “Nothing happened”, that means the previous action is invalid and you should try more options.

You must follow this order:

1. **Reflection:** Briefly assess whether the foresight indicates progress, contradictions, or missing subgoals.
2. **Thought:** Provide a short plan for the next few steps based on the current state and foresight.
3. **Action:** Output **exactly one** admissible action by **exact copying**.

Hard constraint on the action: You MUST choose the Action by copying **EXACTLY one line** from the provided **admissible actions** list. Do NOT paraphrase. Do NOT add extra tokens.

User prompt:

```
Task instruction:
{task}

Current state:
{state}

 $K$ -step foresight trajectory from the world model:
{foresight}

Admissible actions (copy exactly one line for Action):
{admissible_actions_joined_by_newlines}
```

Your response should use the following format:

```
Reflection: <Reflection> ... </Reflection>
Thought: <Thought> ... </Thought>
Action: <Action> ... </Action>
```

Figure 11: Prompt used to generate a **ITP**-style response action conditioned on world-model foresight on ALFWorld and ScienceWorld benchmarks.