



HAL
open science

WORLDLLM: IMPROVING LLMS' WORLD MODELING USING CURIOSITY-DRIVEN THEORY-MAKING

Guillaume Levy, Cédric Colas, Pierre-Yves Oudeyer, Thomas Carta, Clément
Romac

► **To cite this version:**

Guillaume Levy, Cédric Colas, Pierre-Yves Oudeyer, Thomas Carta, Clément Romac. WORLDLLM: IMPROVING LLMS' WORLD MODELING USING CURIOSITY-DRIVEN THEORY-MAKING. RLDM 2025 - Multi-disciplinary Conference on Reinforcement Learning and Decision Making, Jun 2025, Dublin, Ireland. ⟨hal-05302378⟩

HAL Id: hal-05302378

<https://inria.hal.science/hal-05302378v1>

Submitted on 8 Oct 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

WORLDLLM: IMPROVING LLMs’ WORLD MODELING USING CURIOSITY-DRIVEN THEORY-MAKING

Guillaume Levy
Inria (Flowers AI & CogSci)
Univ. of Bordeaux, France

Cédric Colas
MIT
Inria (Flowers AI & CogSci)

Pierre-Yves Oudeyer
Inria (Flowers AI & CogSci)
Univ. of Bordeaux, France

Thomas Carta
Inria (Flowers AI & CogSci)
Univ. of Bordeaux, France

Clément Romac
Hugging Face
Inria (Flowers AI & CogSci)
Univ. of Bordeaux, France

ABSTRACT

Large Language Models (LLMs) possess general world knowledge but often struggle to generate precise predictions in structured, domain-specific contexts such as simulations. These limitations arise from their inability to ground their broad, unstructured understanding in specific environments. To address this, we present **WorldLLM**, a framework that enhances LLM-based world modeling by combining Bayesian inference and autonomous active exploration with reinforcement learning. WorldLLM leverages the in-context learning abilities of LLMs to guide an LLM-based world model’s predictions using natural language hypotheses given in its prompt. These hypotheses are iteratively refined through a Bayesian inference framework that leverages a second LLM as the proposal distribution given collected evidence. This evidence is collected using a curiosity-driven reinforcement learning policy that explores the environment to find transitions with a low log-likelihood under our LLM-based predictive model using the current hypotheses. By alternating refining hypotheses and collecting new evidence, our framework autonomously drives continual improvement of the predictions. Our experiments demonstrate the effectiveness of WorldLLM in a textual game environment that requires agents to manipulate and combine objects. The framework not only enhances predictive accuracy, but also generates human-interpretable theories of environment dynamics.

1 INTRODUCTION

Large Language Models (LLMs) possess broad knowledge about the world, but leveraging this knowledge for precise dynamics modeling remains challenging (Vafa et al., 2024). While LLMs can engage in general reasoning, they struggle to make accurate predictions in specific domains with structured observations and dynamics, such as physics simulations or video games. This limitation stems from the gap between their general capabilities and the need for grounded, domain-specific understanding (Carta et al., 2023; Cross et al., 2025; Bender & Koller, 2020).

The challenge of learning world models has been extensively studied in model-based reinforcement learning (RL), as depicted in Moerland et al. (2023). Traditional approaches focus on learning forward dynamics models through direct interactions, typically using neural networks to predict next observations given current states and actions. Recent work has shown promise in using LLMs for world modeling by fine-tuning them on embodied experiences (Xiang et al., 2023). However, such an approach is particularly computationally costly given the large amount of interaction data required as well as the cost of gradient-based finetuning of LLMs. Contrary to this, LLMs are known to possess general knowledge that can be adapted using careful prompting (Zhang et al., 2024).

In a parallel line of work, theory-based RL takes inspiration from humans’ inductive reasoning and intuitive theory-making to suggest focusing on inferring structured and compositional world models (e.g. programs) (Tsvidis et al., 2021). Current approaches either generate programs to capture

environment dynamics directly (Tsividis et al., 2021; Tang et al., 2024a; Li et al., 2024) or use natural language rules as an intermediate representation (Piriyakulkij et al., 2024; Wang et al., 2024). Leveraging the world model’s structure as well as human-provided prior distributions over possible world models, theory-based methods are usually more sample-efficient than gradient-based ones. More importantly, the resulting world model often shows stronger generalization abilities. However, contrary to gradient-based approaches that usually do not require any expert-given knowledge, theory-based RL methods mostly rely on hand-crafted theory spaces. Finally, one major challenge in learning a world model lies in the data collection. Indeed, collecting transitions that cover the space of possible outcomes in the environment is key for learning an accurate world model. Notably, both classic model-based RL approach and theory-based RL methods agree on the importance of active data collection focusing on evidence with low likelihood under the current model (Schmidhuber, 1991a; Pathak et al., 2017; Burda et al., 2019b; Piriyakulkij et al., 2024).

In this paper, we present **WorldLLM**, a framework for autonomous improvement of an LLM’s world modeling abilities. Our approach combines 1) probabilistic theory induction to produce hypotheses that are given in our LLM’s prompt to improve its predictions and 2) curiosity-driven RL to explore the environment and collect transitions poorly predicted with the current hypotheses. Formally, our LLM’s world model is the conditional probability $P(s_{t+1}|s_t, a_t, H)$, where s_t represents a state, a_t an action, and H a set of natural language hypothesized theories. This probability is computed by the LLM by giving it s_t, a_t and H in its prompt and taking the probability of s_{t+1} to follow this prompt. Our key insight is that natural language theories can help ground an LLM’s broad knowledge into precise predictive power by providing domain-specific rules (Zhang et al., 2024). Our approach consists of three interacting components: (1) our LLM that computes $P(s_{t+1}|s_t, a_t, H)$ by conditioning its predictions on both a state-action pair and the current hypotheses, (2) a theory generator that updates natural language hypotheses using Bayesian inference, and (3) a curiosity-driven reinforcement learning agent trained to collect evidence against the current hypotheses. Inspired by how humans, from children to scientists (Piaget, 1954; Gopnik, 1996), actively update their internal world model by performing experiments, our agent’s exploration provides new evidence for hypothesis refinement, creating a virtuous cycle of improvement.

We demonstrate our approach in a video game environment where agents manipulate and combine objects, showing that WorldLLM successfully learns accurate predictive models while generating human-interpretable theories about environment dynamics. This work contributes to a growing body of research on improving LLMs’ world modeling capabilities and grounding their knowledge in specific domains (Carta et al., 2023; Xiang et al., 2023; Patel et al., 2020). By combining ideas from theory-based RL, Bayesian inference, and active exploration, we provide a framework for learning structured, interpretable world models that leverage both the broad knowledge of LLMs and domain-specific experiences without any costly gradient-based learning.

2 IMPROVING LLMs’ WORLD MODEL WITH WORLDLLM

2.1 PROBLEM STATEMENT

We consider an environment framed as a Markov Decision Process $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{T})$ with \mathcal{S} the state space, \mathcal{A} the action space, $\mathcal{R} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \mapsto \mathbb{R}$ the reward function and $\mathcal{T} : \mathcal{S} \times \mathcal{A} \mapsto \mathcal{S}$ the transition function. In particular, as our main objective is to study world modeling abilities, we consider reward-less environments where no particular task must be solved, i.e. the reward function always returns 0. We focus on textual environments where, given a vocabulary \mathcal{V} and a maximum sentence length N , observations $o \in \mathcal{S} \subset \mathcal{V}^N$ and actions $a \in \mathcal{A} \subset \mathcal{V}^N$ are represented in natural language.

In such a setting, we study how to rapidly learn an LLM-based parametric model \mathbb{P}^{LLM} of the transition function \mathcal{T} (also called forward model) given a budget of experiences T allowed to an agent. This parametric model outputs a probability distribution over all possible next states given a state-action pair. The natural approach from the RL literature to improve this forward model uses gradient-based updates (Ha & Schmidhuber, 2018; Hafner et al., 2020; Schrittwieser et al., 2020). However, the scale of modern LLMs makes such an approach particularly costly. LLMs’ natural in-context learning abilities open up an alternative approach: providing natural language hypotheses in the prompt to improve its predictions. This method is less costly, known as more sample efficient (Le Scao & Rush, 2021), and provides human-interpretable updates. Consequently, the parameter we

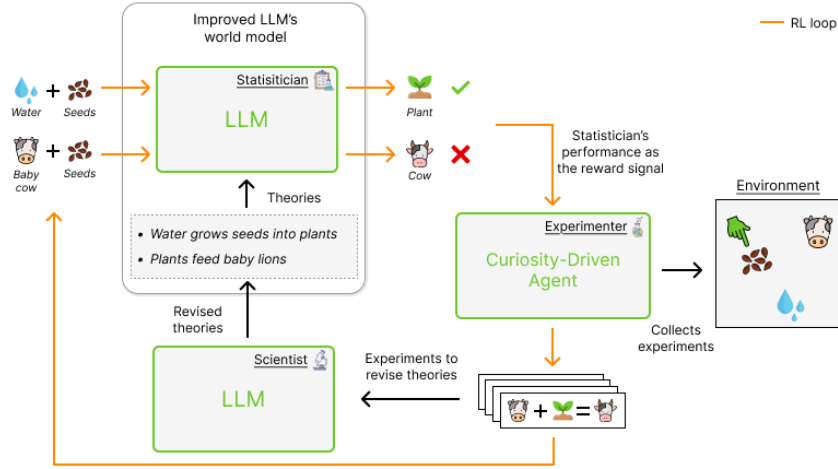


Figure 1: Our framework aims at improving an LLM’s ability to predict the outcome of state-action pairs by augmenting it with natural language hypotheses about the world (*Statistician*). WorldLLM alternates between generating hypotheses (*Scientist*) and collecting evidence to update these hypotheses (*Experimenter*). The former uses Bayesian inference with an LLM as the proposal distribution. The latter uses curiosity-driven RL to train an agent to collect experiences that have a low likelihood under the *Statistician* with the current set of hypotheses.

seek to optimize is a set of hypotheses $\mathcal{H} \subseteq \mathcal{V}^N$ that, when given as input to the forward model (in addition to state-action pairs), maximizes the likelihood of the experienced next states.

We define the global optimization objective as finding the smallest set of hypothesis $H \subseteq \mathcal{H}$ that maximizes the forward model’s probability of all possible transitions in the environment:

$$\begin{aligned}
 & \max \int_{S \times A \times S} \mathbb{P}^{LLM}(s'|s, a, H) ds' da ds \\
 & \min |H| \\
 & \text{subject to } H \in \mathcal{H}
 \end{aligned} \tag{1}$$

As the full transition space is unknown to our agent, it can only use the collected transitions within the budget of T experiences. Note that it is essential for the agent to efficiently explore the environment within these T experiences, as collecting only similar transitions would lead to hypotheses that predict these transitions well but not the remainder of the transition space. For evaluation, as computing a forward model’s performance on the whole transition space is usually intractable, we define a hidden test set \mathcal{D}_{test} containing representative transitions.

2.2 WORLDLLM

In this section, we present WorldLLM, our methodology for addressing the challenge outlined in the preceding section. Our framework comprises three interconnected components: the Statistician, the Scientist, and the Experimenter. The interactions among these components are illustrated in Figure 1. The Statistician represents our LLM-based forward model (i.e. \mathbb{P}^{LLM}) and is used to evaluate the current hypotheses. The set of hypotheses H the Statistician uses is produced by the Scientist, another LLM. These hypotheses are derived from trajectories collected from interactions with the environment using the Experimenter. WorldLLM alternates from during T iterations between evidence collection from the Experimenter and hypotheses refinement from the Scientist. In the subsequent sections, we elaborate on how these modules contribute to our objective’s optimization.

2.2.1 STATISTICIAN

We first detail the Statistician module whose role is to evaluate a set of hypotheses H_t on collected transitions. We start by renaming, for the sake of clarity, \mathbb{P}^{LLM} by \mathbb{P}^{St} . As described above, the

Statistician is a forward model computing the likelihood $\mathbb{P}^{St}(s'|s, a, H_t)$ of the next state s' given a state-action pair (s, a) and a set of hypotheses H_t . This likelihood is obtained by the LLM by giving it s, a and H_t in the prompt and taking the probability of s' to follow this prompt.

2.2.2 SCIENTIST

To improve the Statistician’s prediction as a forward model, the Scientist’s role is to search through the space of possible sets of hypotheses. To efficiently explore this space, we take inspiration from (Wang et al., 2024; Piriyaakulkij et al., 2024) and use a Bayesian Inference framework. In particular, we rely on another LLM \mathbb{P}^{Sc} as the proposal distribution which generates a candidate set of hypotheses following the distribution $\mathbb{P}^{Sc}(\hat{H}_t^i|\mathcal{D}_t, \hat{H}_t)$ given a data set of collected transitions \mathcal{D}_t and the current best hypotheses candidate \hat{H}_t . We assume a uniform prior (i.e. no particular prior) over the space of hypotheses. For this paper, we chose the Metropolis algorithm, a variant of the Metropolis-Hastings algorithm, which iteratively generates candidate hypotheses for $n_s \text{ steps}$. The full algorithm is displayed in Algorithm 1 and more details are provided in Appendix B.

Algorithm 1 WorldLLM

```

1: Input:  $T$  total number of iterations,  $n_{steps}$  number of Metropolis steps at each iteration,  $\mathbb{P}^{St}$  the
   Statistician’s distribution,  $\mathbb{P}^{Sc}$  the Scientist’s proposal distribution,  $\pi$  the Experimenter
2: Initialize  $H_t \leftarrow null$ 
3: for  $t = 0$  to  $T$  do
4:   Collect trajectories  $\mathcal{D}_t \leftarrow env(\pi)$ 
5:   Evaluate current hypotheses  $current\_score = \sum_{(s,a,s') \sim \mathcal{D}_t} \log \mathbb{P}^{St}(s'|s, a, H_t)$ 
6:   Set  $\hat{H}_t = H_t$ 
7:   for  $i = 0$  to  $n_{steps}$  do ▷ Metropolis step
8:     Generate hypotheses  $\hat{H}_t^i \sim \mathbb{P}^{Sc}(\cdot|\mathcal{D}_t, \hat{H}_t)$ 
9:     Evaluate candidate hypotheses  $score = \sum_{(s,a,s') \sim \mathcal{D}_t} \log \mathbb{P}^{St}(s'|s, a, \hat{H}_t^i)$ 
10:    Generate a random number  $u \sim \mathcal{U}(0, 1)$ 
11:    if  $\log(u) < score - current\_score$  then ▷ Accept or reject
12:      Update hypothesis  $\hat{H}_t = \hat{H}_t^i$ 
13:      Update score  $current\_score = score$ 
14:    end if
15:  end for
16:  Set  $H_t = \hat{H}_t$ 
17:  Update Experimenter  $\pi$ 
18: end for

```

2.2.3 EXPERIMENTER

Given a set of hypotheses produced by the Scientist, the Experimenter must collect transitions that will help assessing and refining the hypotheses (i.e. transitions that have a low likelihood under the Statistician). We explore two approaches for our Experiment: oracle agents, which are hard-coded policies that have access to the environment’s inner dynamics, and curiosity-driven RL agents trained to collect transitions based on an intrinsic signal.

For the oracles, we propose the following policies:

- A policy that collects the same distribution of transitions as the test set’s distribution (O-Ideal)
- An optimal policy that progressively moves from collecting simple transitions to collecting the hardest ones (O-Curriculum)
- An optimal policy that produces the sequence of actions necessary to collect the hardest transitions (O-Hardest)

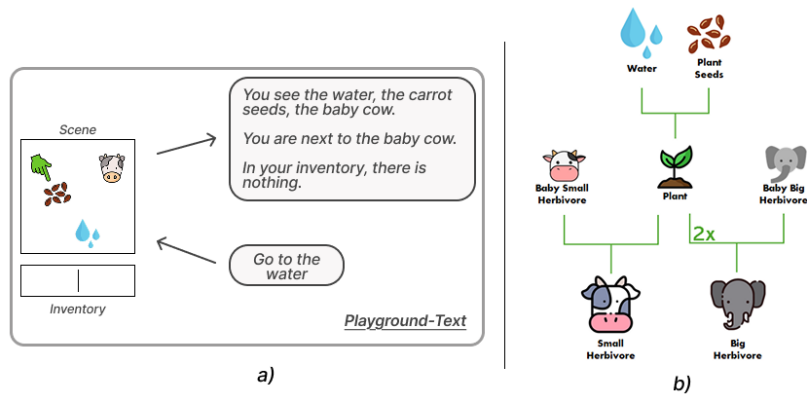


Figure 2: *Our experimental setup.* We use the Playground-Text environment which features objects the agent can interact with along with textual observations and actions (a). We focus on 4 object types from Playground-Text: *Water*, *Plants*, *Small Herbivores*, and *Big Herbivores*. We show in (b) the technology tree for combining objects.

While O-Hardest will collect complex transitions, it may not cover the full transition space. In comparison, O-Ideal matches the test set distribution. However, the iterative process of WorldLLM may require the Experimenter to focus for several iterations on particular subspaces of the transition space for the Scientist to produce accurate hypotheses. While O-Hardest and O-Ideal lack such a developmental trajectory, O-Curriculum focuses on specific transitions for an arbitrary number of collections. We also study the performance of a policy acting randomly (O-Random), which is unlikely to collect complex transitions.

Regarding our RL policies, we propose three agents trained using curiosity-based intrinsic rewards derived from \mathbb{P}^{St} (Oudeyer & Kaplan, 2007). First, we introduce *RL-LogP*, which uses the prediction error or surprisal (Burda et al., 2019a; Pathak et al., 2017) as reward by receiving $r = -\log \mathbb{P}^{St}(s'|s, a, H_t)$ for each collected transition (s, a, s') . However, this method is known to suffer from the "noisy TV problem" (Burda et al., 2019b) where unpredictable transitions give high rewards.

To address this, we inspired from methods relying on Learning Progress (Oudeyer et al., 2007; Lopes et al., 2012; Schmidhuber, 1991a) or information gain over the prediction error (Oudeyer et al., 2007; Lopes et al., 2012; Schmidhuber, 1991a) that reward transitions on which the forward model improves. We thus propose *RL-ALP*, which uses the absolute value of Learning Progress, allowing it to track both progress and forgetting in the forward model (Oudeyer et al., 2007; Baranes & Oudeyer, 2009):

$$r_{alp} = |\log \mathbb{P}^{St}(s'|s, a, H_{t-1}) - \log \mathbb{P}^{St}(s'|s, a, H_t)| \quad (2)$$

Finally, we also propose *RL-ALPEXP*, a variant of RL-ALP which partitions the transition space into subspaces and computes ALP over each subspace. However, the partitions are not automatically inferred (e.g. as in Baranes & Oudeyer (2013)) but hand-designed. It corresponds to a middle ground between the oracles and RL-LogP or RL-ALP as it has access to expert environment knowledge while still relying on RL. Further details are provided in Appendix A.

3 EXPERIMENTS

3.1 EXPERIMENTAL SETUP

Environment To evaluate WorldLLM’s efficiency, we use the Playground-Text environment from (Gaven et al., 2024), a textual environment which produces textual observations and receives textual commands. This environment generates a room that contains multiple objects the agent can interact with by grasping or releasing them. The agent can combine two objects, if an interaction is possible, by releasing one onto the other, as illustrated on Figure 2 b. For instance, releasing *water* on a *carrot seed* transforms them into *grown carrot*. In our experiments, we only consider 4 types of objects:

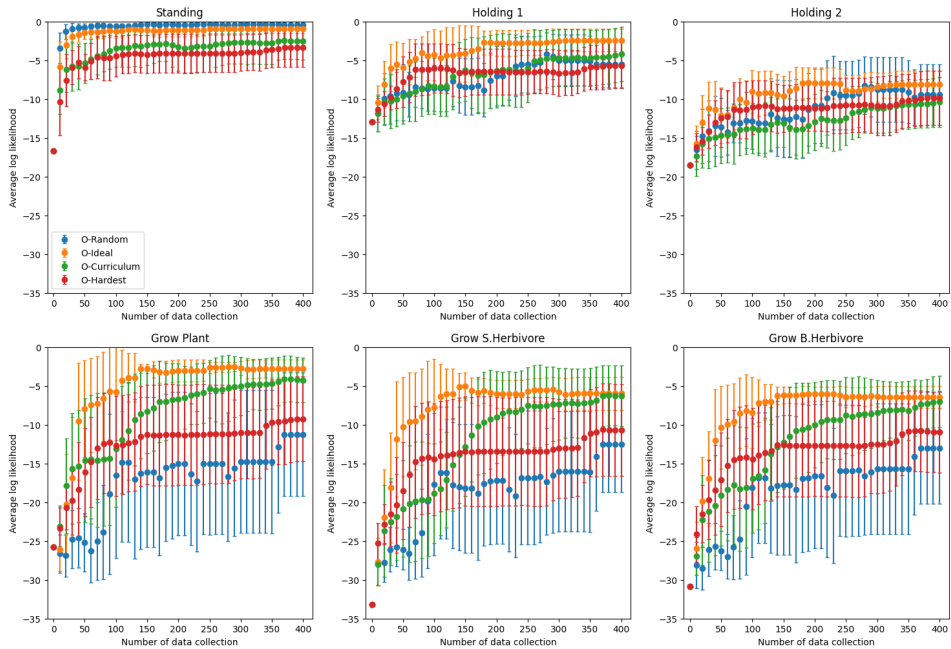


Figure 3: Evolution of log-likelihoods computed by the Statistician (using the Scientist’s last accepted hypotheses) on the test set throughout WorldLLM’s iterations when using different oracles as Experimenters. We average per transition type over 8 random seeds with the standard deviation represented by the error bars. We show at iteration 0 the log-likelihood from the Statistician without any hypotheses given.

Water, *Plants* (seed or grown) that require *water* to grow, *Small Herbivores* (baby or grown) that require a grown *plant* to grow, and *Big Herbivores* (baby or grown) that require 2 grown *plants* to grow. We evaluate WorldLLM’s ability to discover, predict and explain how to grow *Plants*, *Small Herbivores* and *Big Herbivores*. We provide more details on the environment in Appendix A. For analysis, we gather the possible transitions into 6 different types: *Standing* when moving on an object, *Holding 1* when grasping an object, *Holding 2* when grasping two objects at the same time, *Grow Plant* when releasing water on a seed, *Grow S. Herbivore* when releasing a grown plant on a baby small herbivore and *Grow B. Herbivore* when releasing two grown plants on a baby big herbivore. These types are the ones used by RL-ALPEXP.

Optimization We use Phi-3-mini-4k-Instruct¹ for both the Statistician and the Scientist. For our Experimenter, we either use one of our hard-coded oracles or an RL agent. Regarding oracles, O-Ideal collects transitions that match the distribution of the above test set. O-Hardest performs the optimal policy to grow the *Small Herbivore* then the *Big Herbivore* accessible. O-Curriculum plays the optimal policy to grow a *Plant* up to 133 iterations, then plays the optimal policy to grow the accessible *Small Herbivore* up to 266 iterations, and finally grows the *Small Herbivore* then the *Big Herbivore* until the end. See Appendix B.3.1 for more details on the oracles. Concerning the RL agent, its policy is a Multi-Layer Perceptron with 2 hidden layers of 64 units. For simplicity, this policy uses a symbolic representation of Playground-Text’s state as observation. In all experiments below, unless specified otherwise, we use 8 different random seeds. For each seed, we perform 400 iterations of the framework (i.e. T) where 150 transitions are collected per iteration and $n_{steps} = 5$ steps of the Metropolis algorithm are performed by the Scientist at each iteration. When the Experimenter is an RL agent, we collect 3600 transitions per iteration to train it and use PPO to update the policy, while only keeping the last 150 transitions for the Scientist to match the oracles’ hyperparameters. Our test set \mathcal{D}_{test} is composed of 120 *Standing*, 20 *Holding 1*, 7 *Holding 2*, 12 *Grow Plant*, 6 *Grow S. Herbivore* and 3 *Grow B. Herbivore*.

¹<https://huggingface.co/microsoft/Phi-3-mini-4k-instruct>

Methods	Standing	Holding 1	Holding 2	Grow Plant	Grow S.H.	Grow B.H.
O-Random	0.97 ± 0.01	0.45 ± 0.16	0.40 ± 0.15	0.34 ± 0.24	0.44 ± 0.16	0.40 ± 0.17
O-Ideal	0.92 ± 0.03	0.72 ± 0.12	0.50 ± 0.06	0.81 ± 0.06	0.77 ± 0.06	0.74 ± 0.04
O-Curriculum	0.79 ± 0.15	0.48 ± 0.25	0.31 ± 0.15	0.67 ± 0.15	0.63 ± 0.13	0.60 ± 0.14
O-Hardest	0.74 ± 0.13	0.48 ± 0.20	0.39 ± 0.16	0.52 ± 0.21	0.57 ± 0.18	0.56 ± 0.18
RL-ALPEXP	0.95 ± 0.02	0.44 ± 0.16	0.40 ± 0.14	0.42 ± 0.25	0.52 ± 0.19	0.49 ± 0.19
RL-LogP	0.93 ± 0.01	0.52 ± 0.07	0.53 ± 0.07	0.62 ± 0.09	0.61 ± 0.07	0.60 ± 0.07
RL-ALP	0.98 ± 0.01	0.18 ± 0.08	0.17 ± 0.06	-0.07 ± 0.05	0.16 ± 0.04	0.08 ± 0.07

Table 1: This table presents the normalized area under the training curve for each algorithm and transition type. The normalization is performed by dividing the computed area by the area formed by the initial log-likelihoods, i.e. obtained without using hypotheses.

3.2 WORLDLLM WITH AN ORACLE EXPERIMENTER

We initiate the experiments by examining the performance of WorldLLM when using our oracles as Experimenters. In Figure 3, we show the evolution of log-likelihoods outputted by the Statistician using the last accepted set of hypotheses H_t throughout WorldLLMs’ iterations when using different oracles as Experimenters. These log-likelihoods are computed on the test set and averaged over transition types (and seeds).

One can observe that, regardless of the oracle used or transition type, giving the Statistician hypotheses always improves its performance compared to when no hypotheses are given (i.e. the initial point of each curve). When comparing the oracles, one can observe that the data collection strategy of strongly influences WorldLLM’s abilities at producing efficient hypotheses. Indeed, using the same distribution as the one we perform evaluation on, O-Ideal achieves the best overall performance. On the opposite, O-Random obtains the poorest performance on complex transitions (i.e. growing an object) while being the best method on *Standing* transitions. This can easily be explained by the fact that moving actions are predominant in the action space. O-Hardest did not lead to strong performance, highlighting the need of diversity in collected transitions for the Scientist. Finally, O-Curriculum obtains results on par with O-Ideal. Results from Appendix D.2.1 indicate that O-Curriculum particularly helps the Scientist generating improved hypotheses.

Overall, these initial experiments demonstrate that WorldLLM is able to produce hypotheses that improve the predictive capabilities of the Statistician’s LLM. Moreover, the different performance obtained by our oracles highlighted the importance of 1) diversity in the collected transitions and 2) a developmental collection strategy allowing the Scientist to progressively update the hypotheses.

3.3 WORLDLLM WITH A CURIOSITY-DRIVEN RL AGENT AS EXPERIMENTER

We now move to curiosity-driven RL agents as Experimenters. Compared to the hard coded policies, we rely instead on one of our three intrinsic reward signals (RL-ALPEXP, RL-LogP, and RL-ALP), to collect transitions that drive the Scientist towards generating efficient hypotheses.

As in the previous section, we analyze the evolution of the Statistician’s log-likelihoods over the test set. In order to ease the comparison of our three RL-based Experimenters to oracles, we introduce a coverage metric as the area under the evolution of log-likelihoods’ curve of each method. We normalize this area using the area between the Statistician’s log-likelihood without any hypotheses and the best performance achievable (i.e. 0). The complete formula is

$$\text{area} = 1 - \int_0^T \frac{\log \mathbb{P}^{S_t}(\mathcal{D}_{test}|H_t)}{\log \mathbb{P}^{S_t}(\mathcal{D}_{test}|\{\emptyset\})T} dt. \quad (3)$$

The best performing method possible is expected to be very close to 1 while an inept method’s result will be near 0. A method whose result is below 0 means that the hypotheses degrade the initial performance of the Statistician.

From results shown in Table 1, one can observe that RL-ALPEXP is able to leverage its environment knowledge to reach performance significantly better than O-Random but remaining worse O-Ideal.

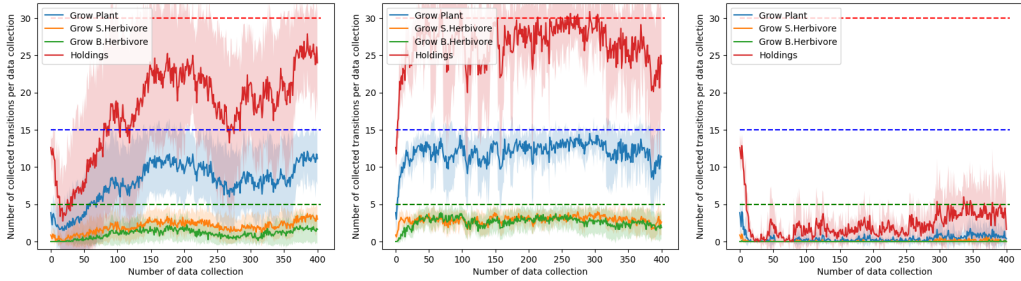


Figure 4: Proportion of collected transitions for RL-ALPEXP (left), RL-LogP (middle) and RL-ALP (right). The colored dashed lines correspond to the amount of transitions collected by O-Hardest.

When given no expert knowledge, RL-ALP achieves poor performance while RL-LogP obtains results close or even better than oracle baselines.

To better understand the effect of the RL-based Experimenters on the Statistician’s predictions, we studied the evolving distribution of collected transitions. We grouped *Holding 1* and *Holding 2* for simplicity and ignored *Standing* transitions which are predominant by nature. We show these results in Figure 4.

They show that RL-LogP manages to explore the environment and cover the whole transition space. The resulting collected distribution is near O-Hardest’s one (which is showed in dashed lines). This performance is notably explained by the syntax used by Playground-Text for observations and actions that makes complex transitions naturally harder to predict for an LLM (e.g. moving to a *Big Herbivore* involves only this entity whereas growing a *Big Herbivore* implies an action where two *Plants* are released on a baby *Big Herbivore*). This natural ordering appears in the rewards obtained by the RL policy and favors complex transitions. RL-ALPEXP, which does not leverage such a reward ordering, progressively converges to collecting more complex transitions. When looking at RL-ALP, most of the collected transitions are *Standing* ones. As we have seen in the previous section, the diversity of the transitions given to the Scientist is key in WorldLLM. Consequently, the hypotheses produced with RL-ALP focus solely on explaining *Standing* transitions. A more detailed analysis of why RL-ALP failed is available in Appendix D.4.1.

This sections’ results show that WorldLLM successfully leads to improvement of the LLM’s forward modeling abilities when using curiosity-driven RL-based Experimenters. Moreover, we show that rewarding Experimenters for collecting transitions that help refine the hypotheses not only leads to efficient hypotheses but also leads to policies which solve our environment’s full technology tree.

3.4 EVALUATING THE PRODUCED WORLD MODEL

We now focus on evaluating the usefulness of the produced world model (i.e. the Statistician equipped with the last accepted hypotheses). While previous sections indicate that WorldLLM succeeds in increasing the likelihood of the observed next states, these results do not guarantee that the observed next states receive the highest likelihood compared to other possible next states. We thus examine our LLM-based world model’s capacity to correctly predict (i.e. generate) the next state on our test. In particular, for each state-action pair in \mathcal{D}_{test} , we used our Statistician’s LLM and constrained decoding to only generate the valid next states, given the current state, in Playground-Text. We then report in Table 2 the percentage of next states that belonged to the top-3 generated sequences.

The results show that no method performs best on all transition types. However, similarly to what was observed in previous sections, *O-Ideal* , *O-Curriculum* , *RL-ALPEXP* and *RL-LogP* appear to be the best performing Experimenters. More importantly, these results show a significant improvement over the performance without any hypotheses.

3.5 GENERATING HYPOTHESES WITH WORLDLLM VS FINE-TUNING

WorldLLM improves an LLM’s forward modeling abilities by providing it natural language hypotheses about the world. In this section, we study how our approach compares to a more classic

Methods	Standing	Holding 1	Holding 2	Grow Plant	Grow S.H.	Grow B.H.
No Hypo.	0.82 ± 0.00	1.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
O-Random	1.00 ± 0.00	0.48 ± 0.26	0.07 ± 0.19	0.18 ± 0.31	0.00 ± 0.00	0.04 ± 0.11
O-Ideal	1.00 ± 0.00	0.94 ± 0.15	0.00 ± 0.00	0.94 ± 0.17	0.25 ± 0.26	0.46 ± 0.37
O-Curriculum	1.00 ± 0.00	0.88 ± 0.31	0.00 ± 0.00	0.85 ± 0.28	0.60 ± 0.41	0.54 ± 0.44
O-Hardest	1.00 ± 0.00	0.86 ± 0.21	0.11 ± 0.28	0.29 ± 0.34	0.27 ± 0.37	0.21 ± 0.37
RL-ALPEXP	1.00 ± 0.00	0.71 ± 0.29	0.34 ± 0.36	0.40 ± 0.43	0.02 ± 0.06	0.08 ± 0.14
RL-LogP	1.00 ± 0.00	0.67 ± 0.30	0.34 ± 0.25	0.74 ± 0.32	0.02 ± 0.06	0.29 ± 0.42
RL-ALP	1.00 ± 0.00	0.20 ± 0.26	0.00 ± 0.00	0.03 ± 0.08	0.00 ± 0.00	0.00 ± 0.00

Table 2: Top 3 performance of Constrained Decoding for each algorithm by transition type. We also report the performance without any hypotheses given (No Hypo.).

Methods	Standing	Holding 1	Holding 2	Grow Plant	Grow S.H.	Grow B.H.
No Hypo.	0.28 ± 0.00	0.85 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
O-Random	0.69 ± 0.16	0.45 ± 0.26	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
O-Ideal	0.73 ± 0.23	0.46 ± 0.19	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
O-Hardest	0.33 ± 0.33	0.71 ± 0.33	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
O-Curriculum	0.46 ± 0.34	0.72 ± 0.29	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
RL-ALPEXP	0.65 ± 0.16	0.42 ± 0.30	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
RL-LogP	0.43 ± 0.16	0.30 ± 0.23	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
RL-ALP	0.76 ± 0.18	0.61 ± 0.31	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
F-LogP	0.88 ± 0.17	0.24 ± 0.25	0.00 ± 0.00	0.07 ± 0.13	0.21 ± 0.32	0.67 ± 0.44

Table 3: Top 3 performance on Constrained Decoding for each algorithm by transition type on the test set using the generalization environment. We also report the performance without any hypotheses given (No Hypo.).

gradient-based approach where the LLM is directly fine-tuned on collected transitions. In addition to avoiding gradient computation through the LLM and being human interpretable, the theories produced by WorldLLM are expected to better generalize to changes in the environment.

We use the transitions collected by RL-LogP and fine-tune our Statistician’s LLM on them using a Causal Language Modeling objective (i.e. we maximize the likelihood of the next state’s tokens given the state-action pair). We name this new approach *F-LogP*. When evaluated on its top-3 constrained decoding accuracy over the test set as in Section 3.4, F-LogP obtains perfect results (i.e. accuracy of 1) over all transition types. In the following sections, we study the generalization abilities of both methods and provide a deeper analysis of their impact on log-likelihoods.

3.5.1 GENERALIZATION TO CHANGES IN THE ENVIRONMENT

We now assess whether using natural language hypotheses in the prompt produces better generalization abilities for the LLM-based world model than fine-tuning it on collected transitions. For this, we create a generalization environment in which the syntax of observations is changed (see Appendix 10). We use this new environment to report the top-3 constrained decoding accuracy on our test set as in previous sections.

Results from Table 3 show an important deterioration of the performance for all methods. F-LogP manages to maintain a non-zero accuracy on complex transition types while WorldLLM-based world models do not perform better than the Statistician without any hypotheses. These results hints that WorldLLM may struggle in finding abstract hypotheses that generalize. The next section provides additional insights on these results.

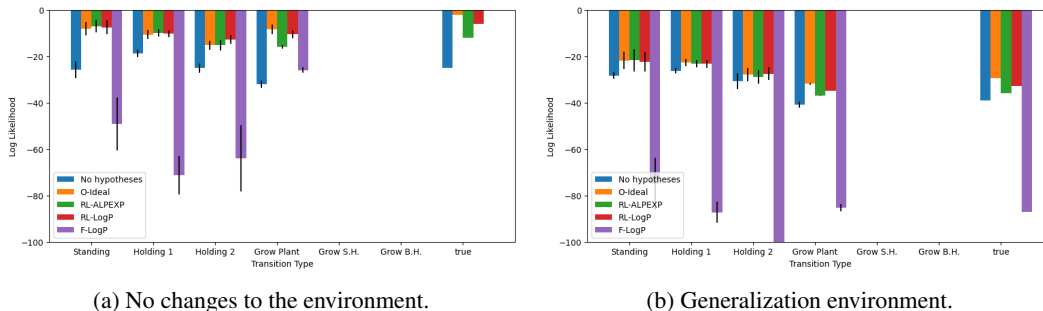


Figure 5: Log-likelihoods (averaged by transition type) produced by the Statistician when facing an instance of a *Grow Plant* transition. We show how using the generalization environment from Section 3.5.1 affects the log-likelihoods (b).

3.5.2 EFFECT ON LOG-LIKELIHOODS

We now propose to study how both methods affect the Statistician’s log-likelihoods. We show in Figure 5 the log-likelihoods (averaged by transition type) produced by the Statistician when facing an instance of a *Grow Plant* transition.

Looking at results when using the classic environment (a), WorldLLM increases the log-likelihood of all the possible transitions (no matter the Experimenter) while maintaining the correct transition more likely. In comparison, the Causal Language Modeling objective used for F-LogP increases the log-likelihood of the correct transition while significantly decreasing the log-likelihood of all the other transitions. These results illustrate the difference in the two optimization processes: while fine-tuning aims to increase the distance in log-likelihoods between the correct transition and the other ones, WorldLLM focuses on increasing the log-likelihood of the correct transition, regardless of the others.

When focusing on performance with the generalization environment, WorldLLM leads to very few changes compared to when no hypotheses are given. On its side, F-LogP produces very low log-likelihoods for all transitions and keeps the correct one slightly above the others. Overall these results show that both approaches generalize poorly.

One explanation for WorldLLM’s poor generalization performance lies in the hypotheses produced. We provided samples in Appendix D.3 which show that our Scientist notably used examples of transitions in the hypotheses (even though some hypotheses did mention abstract categories or even mentioned objects not in our environment). Such examples no longer align with the generalization environment’s syntax. Multiple explanations can be given for this. First, our current experiments rely on a Scientist’s LLM with limited capabilities. Second, while our objective in 1 minimizes the number of hypotheses, the current implementation of WorldLLM only focuses on maximizing the Statistician’s log-likelihood. We argue that favoring short set of hypotheses could help abstract theories emerge. We leave these investigations for future work.

4 RELATED WORK

Building forward models Building forward models is a long-established research direction (Ha & Schmidhuber, 2018; Moerland et al., 2023). In addition to producing more interpretable agents, forward models were particularly used in RL to improve sample efficiency allowing the agent to train on imagined transitions or perform multi-step planning (Hafner et al., 2020; 2021; 2025; Schrittwieser et al., 2020). While these works learned the forward model at the same time as the RL policy, a line of work also studied how learning the forward model could be used as an intrinsic reward to foster the policy’s exploration (Schmidhuber, 1991b; Pathak et al., 2017; Burda et al., 2019a; Oudeyer et al., 2007). WorldLLM employs a similar curiosity-driven scheme where our RL-based Experimenters use reward signals from prior work to promote exploration and collect transitions helping to improve the forward model. However, our method differs in the way the model is improved: instead of using a gradient-based approach to fine-tune the LLM-based forward model, as done in (Xiang et al., 2023),

WorldLLM searches for natural language theories that improve the forward model when given in the prompt. While recent work studied how a forward model could leverage external textual information such as a manual explaining the environment (Zhang et al., 2024; Dainese et al., 2023), our approach iteratively generates and refines its own explanations.

Inductive reasoning for world modeling The importance of theories in shaping one’s representations about the world was argued as central, in both children development and adults (Piaget, 1954; Gopnik, 1996). Building on this, the literature of theory-based RL has been advocating for augmenting artificial agents with inductive reasoning (Tsividis et al., 2021). In particular, (Tsividis et al., 2021; Tang et al., 2024a; Li et al., 2024) proposed approaches based on LLMs which generate programs representing the theories about an environment’s dynamics. The use of programs allowed the authors to execute and compare how the proposed theories match empirical evidence. Parallel to this, (Piriyakulkij et al., 2024; Wang et al., 2024) proposed to generate theories using natural language. While the former constrained the generated theories to programs which can be easily evaluated in their environment, the latter used the generated theories to condition a program generator to simulate the environment’s dynamics. WorldLLM also builds on the theory-based RL literature and closely relates to (Piriyakulkij et al., 2024). Notably, WorldLLM also leverages Bayesian Inference with an LLM as the proposal distribution to generate the theories. However, the major difference lies in how the produced theories are used: while (Piriyakulkij et al., 2024) constrain the format of theories such that they can be directly tested in the environment, our work focuses on giving these theories in another LLM’s prompt to improve its forward modeling capabilities. *In other words, our work strongly differs in its perspective of producing theories that improve our LLM’s world modeling abilities instead of producing correct theories.* Nonetheless, the Bayesian Inference framework we used in our Scientist is largely inspired by (Piriyakulkij et al., 2024). Both our work and theirs also consider an active learning paradigm to choose which experiments to perform for theories’ refinement. While their approach computes exact Information Gain of all possible experiments, ours relies on curiosity-driven RL agents and Information Gain approximates as rewards to overcome the intractability of exact Information Gain in most environments.

5 DISCUSSION

We propose WorldLLM, a method to improve LLMs’ world modeling abilities by generating a set of natural language hypotheses that are given to an LLM when predicting the outcome of a state-action pair. Our method improves the LLM using human-interpretable theories with natural language while being less computationally costly than directly fine-tuning large-scale LLMs over collected data. WorldLLM is also based on an architecture that does not require any hand-crafted theory space. Our results showed that the adversarial framework between the Scientist and simple curiosity-driven RL Experimenters fuels an exploratory dynamic that covers the environment and allows the discovery of the underlying rules.

Nonetheless, our results show several limitations. First, the inner dynamics of Playground-Text are arguably simple. Moreover, the natural ordering created by the hardest transitions being longer and harder to predict is not expected in more complex environments. This could significantly affect RL-LogP’s results. Studying WorldLLM’s scaling to more complex environments (e.g. Mars (Tang et al., 2024b)) is therefore an important step. Secondly, our results were performed using a single LLM. This model showed its limitations on Playground-Text, especially in its capacity to generate high-level and abstract hypotheses when used in the Scientist. We believe that using a more efficient LLM in the Scientist in addition to taking into account the hypotheses length (e.g. adding a non-uniform prior in our proposal) should improve WorldLLM’s efficiency. Then, our paper studied three simple curiosity-driven RL Experimenters. Studying more advanced methods such as autotelic approaches (Colas et al., 2022) could improve the quality of the collected transitions, which is key in our framework. Finally, our online Metropolis approach for the Scientist suffers from several limitations. Notably, it only considers the transitions collected at time t , creating potential forgetting when the Experimenter is non-stationary. This choice is motivated by the fact that 1) our LLM used as the proposal has a fixed and limited context window and thus cannot consider all transitions collected up to time t , and 2) computing the likelihood on a large set of transitions is costly. This could be improved by either keeping track of past data (e.g. using a buffer containing transitions from different timesteps) or considering more advanced Bayesian inference methods more suited for

online data such as particle filters. The latter would also overcome another current limitation of our approach: a single particle is used and may lead our Scientist to get stuck in local optima (as hinted by Appendix D).

In conclusion, WorldLLM represents a significant advancement, opening up promising new avenues for developing more grounded LLMs that can efficiently predict and explain their operational environments. While the current iteration has certain limitations, these also present opportunities for substantial enhancements and future innovations in the methodology.

ACKNOWLEDGMENTS

This work was granted access to the HPC resources of IDRIS under the allocation A0171011996 made by GENCI.

REFERENCES

- Adrien Baranes and Pierre-Yves Oudeyer. R-IAC: Robust Intrinsically Motivated Exploration and Active Learning. *IEEE Transactions on Autonomous Mental Development*, 1(3):155–169, October 2009. ISSN 1943-0612. doi: 10.1109/TAMD.2009.2037513. URL <https://ieeexplore.ieee.org/document/5342516/?arnumber=5342516&tag=1>. Conference Name: IEEE Transactions on Autonomous Mental Development.
- Adrien Baranes and Pierre-Yves Oudeyer. Active learning of inverse models with intrinsically motivated goal exploration in robots. *Robotics and Autonomous Systems*, 61(1):49–73, January 2013. ISSN 0921-8890. doi: 10.1016/j.robot.2012.05.008. URL <https://www.sciencedirect.com/science/article/pii/S0921889012000644>.
- Emily M. Bender and Alexander Koller. Climbing towards NLU: On Meaning, Form, and Understanding in the Age of Data. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 5185–5198, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.463. URL <https://aclanthology.org/2020.acl-main.463>.
- Yuri Burda, Harri Edwards, Deepak Pathak, Amos Storkey, Trevor Darrell, and Alexei A. Efros. Large-scale study of curiosity-driven learning. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *7th International Conference on Learning Representations (ICLR 2019)*. Curran Associates, Inc., 2019a.
- Yuri Burda, Harrison Edwards, Amos Storkey, and Oleg Klimov. Exploration by random network distillation, 2019b.
- Thomas Carta, Clément Romac, Thomas Wolf, Sylvain Lamprier, Olivier Sigaud, and Pierre-Yves Oudeyer. Grounding large language models in interactive environments with online reinforcement learning. In *Proceedings of the 40th International Conference on Machine Learning*, pp. 3676–3713. PMLR, 2023. URL <https://proceedings.mlr.press/v202/carta23a.html>. ISSN: 2640-3498.
- Cédric Colas, Tristan Karch, Nicolas Lair, Jean-Michel Dussoux, Clément Moulin-Frier, Peter Dominey, and Pierre-Yves Oudeyer. Language as a Cognitive Tool to Imagine Goals in Curiosity Driven Exploration. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 3761–3774. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/274e6fcf4a583de4a81c6376f17673e7-Paper.pdf.
- Cédric Colas, Tristan Karch, Olivier Sigaud, and Pierre-Yves Oudeyer. Autotelic agents with intrinsically motivated goal-conditioned reinforcement learning: a short survey. *Journal of Artificial Intelligence Research*, 74:1159–1199, 2022. URL <https://www.jair.org/index.php/jair/article/view/13554>.
- Logan Cross, Violet Xiang, Agam Bhatia, Daniel LK Yamins, and Nick Haber. Hypothetical Minds: Scaffolding Theory of Mind for Multi-Agent Tasks with Large Language Models. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *13th International Conference on Learning Representations (ICLR 2025)*. Curran Associates, Inc., 2025. URL <https://openreview.net/forum?id=otW0TJOUYF>.
- Nicola Dainese, Pekka Marttinen, and Alexander Ilin. Reader: Model-based language-instructed reinforcement learning. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 16583–16599, Singapore, December 2023. Association for Computational Linguistics. URL <https://aclanthology.org/2023.emnlp-main.1032>.
- Loris Gaven, Clement Romac, Thomas Carta, Sylvain Lamprier, Olivier Sigaud, and Pierre-Yves Oudeyer. Sac-glam: Improving online rl for llm agents with soft actor-critic and hindsight relabeling, 2024. URL <https://arxiv.org/abs/2410.12481>.

-
- Alison Gopnik. The scientist as child. 63(4):485–514, 1996. ISSN 0031-8248. URL <https://www.jstor.org/stable/188064>. Publisher: [Cambridge University Press, The University of Chicago Press, Philosophy of Science Association].
- David Ha and Jürgen Schmidhuber. World models, 2018. URL <http://arxiv.org/abs/1803.10122>.
- Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to Control: Learning Behaviors by Latent Imagination. In *8th International Conference on Learning Representations (ICLR 2020)*. Curran Associates, Inc., 2020.
- Danijar Hafner, Timothy Lillicrap, Mohammad Norouzi, and Jimmy Ba. Mastering Atari with Discrete World Models, 2021. URL <http://arxiv.org/abs/2010.02193>.
- Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. Mastering diverse control tasks through world models. *Nature*, 640(8059):647–653, April 2025. ISSN 1476-4687. doi: 10.1038/s41586-025-08744-2. URL <https://www.nature.com/articles/s41586-025-08744-2>. Publisher: Nature Publishing Group.
- Teven Le Scao and Alexander Rush. How many data points is a prompt worth? In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 2627–2636, Online, June 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.naacl-main.208. URL <https://aclanthology.org/2021.naacl-main.208>.
- Michael Y. Li, Emily B. Fox, and Noah D. Goodman. Automated statistical model discovery with language models. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *ICML'24*, pp. 27791–27807, Vienna, Austria, July 2024. JMLR.org.
- Manuel Lopes, Tobias Lang, Marc Toussaint, and Pierre-yves Oudeyer. Exploration in model-based reinforcement learning by empirically estimating learning progress. In *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc., 2012. URL <https://proceedings.neurips.cc/paper/2012/hash/a0a080f42e6f13b3a2df133f073095dd-Abstract.html>.
- Laurens van der Maaten and Geoffrey Hinton. Visualizing Data using t-SNE. *Journal of Machine Learning Research*, 9(86):2579–2605, 2008. ISSN 1533-7928. URL <http://jmlr.org/papers/v9/vandermaaten08a.html>.
- Thomas M. Moerland, Joost Broekens, Aske Plaat, and Catholijn M. Jonker. Model-based Reinforcement Learning: A Survey. *Foundations and Trends® in Machine Learning*, 16(1):1–118, January 2023. ISSN 1935-8237, 1935-8245. doi: 10.1561/22000000086. URL <https://www.nowpublishers.com/article/Details/MAL-086>. Publisher: Now Publishers, Inc.
- Pierre-Yves Oudeyer and Frederic Kaplan. What is intrinsic motivation? a typology of computational approaches. 1, 2007. ISSN 1662-5218. doi: 10.3389/neuro.12.006.2007. URL <https://www.frontiersin.org/journals/neurorobotics/articles/10.3389/neuro.12.006.2007/full>. Publisher: Frontiers.
- Pierre-Yves Oudeyer, Frdric Kaplan, and Verena V. Hafner. Intrinsic motivation systems for autonomous mental development. 11(2):265–286, 2007. ISSN 1941-0026. doi: 10.1109/TEVC.2006.890271. URL <https://ieeexplore.ieee.org/abstract/document/4141061>. Conference Name: IEEE Transactions on Evolutionary Computation.
- Roma Patel, Ellie Pavlick, and Stefanie Tellex. Grounding Language to Non-Markovian Tasks with No Supervision of Task Specifications. *Robotics: Science and Systems XVI*, July 2020. doi: 10.15607/RSS.2020.XVI.016. URL <http://www.roboticsproceedings.org/rss16/p016.pdf>. Conference Name: Robotics: Science and Systems 2020 ISBN: 9780992374761 Publisher: Robotics: Science and Systems Foundation.

-
- Deepak Pathak, Pulkit Agrawal, Alexei A. Efros, and Trevor Darrell. Curiosity-driven exploration by self-supervised prediction. In *Proceedings of the 34th International Conference on Machine Learning*, pp. 2778–2787. PMLR, 2017. URL <https://proceedings.mlr.press/v70/pathak17a.html>. ISSN: 2640-3498.
- Jean Piaget. *The construction of reality in the child*. The construction of reality in the child. Basic Books/Hachette Book Group, New York, NY, US, 1954. doi: 10.1037/11168-000. Pages: xiii, 386.
- Wasu Top Piriyaikulij, Cassidy Langenfeld, Tuan Anh Le, and Kevin Ellis. Doing Experiments and Revising Rules with Natural Language and Probabilistic Reasoning. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 53102–53137. Curran Associates, Inc., 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/5f1b350fc0c2affd56f465faa36be343-Paper-Conference.pdf.
- J. Schmidhuber. Curious model-building control systems. In *[Proceedings] 1991 IEEE International Joint Conference on Neural Networks*, pp. 1458–1463 vol.2. IEEE, 1991a. ISBN 978-0-7803-0227-3. doi: 10.1109/IJCNN.1991.170605. URL <https://ieeexplore.ieee.org/document/170605/>.
- Juergen Schmidhuber. A possibility for implementing curiosity and boredom in model-building neural controllers. 1991b. doi: 10.7551/mitpress/3115.003.0030. URL <https://direct.mit.edu/books/edited-volume/3865/chapter/162771/A-Possibility-for-Implementing-Curiosity-and>.
- Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Simon Schmitt, Arthur Guez, Edward Lockhart, Demis Hassabis, Thore Graepel, Timothy Lillicrap, and David Silver. Mastering Atari, Go, Chess and Shogi by Planning with a Learned Model. *Nature*, 588(7839):604–609, December 2020. ISSN 0028-0836, 1476-4687. doi: 10.1038/s41586-020-03051-4. URL <http://arxiv.org/abs/1911.08265>. arXiv:1911.08265 [cs].
- Hao Tang, Darren Key, and Kevin Ellis. WorldCoder, a Model-Based LLM Agent: Building World Models by Writing Code and Interacting with the Environment. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 70148–70212. Curran Associates, Inc., 2024a. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/820c61a0cd419163ccbd2c33b268816e-Paper-Conference.pdf.
- Xiaojuan Tang, Jiaqi Li, Yitao Liang, Song-chun Zhu, Muhan Zhang, and Zilong Zheng. Mars: Situated Inductive Reasoning in an Open-World Environment. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 17830–17869. Curran Associates, Inc., 2024b. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/1fb6d0b52f5e41b11392841a66dbfe7d-Paper-Datasets_and_Benchmarks_Track.pdf.
- Pedro A. Tsividis, Joao Loula, Jake Burga, Nathan Foss, Andres Campero, Thomas Pouncy, Samuel J. Gershman, and Joshua B. Tenenbaum. Human-level reinforcement learning through theory-based modeling, exploration, and planning, 2021. URL <http://arxiv.org/abs/2107.12544>.
- Keyon Vafa, Justin Y. Chen, Ashesh Rambachan, Jon Kleinberg, and Sendhil Mullainathan. Evaluating the World Model Implicit in a Generative Model. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 26941–26975. Curran Associates, Inc., 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/2f6a6317bada76b26a4f61bb70a7db59-Paper-Conference.pdf.
- Ruocheng Wang, Eric Zelikman, Gabriel Poesia, Yewen Pu, Nick Haber, and Noah D. Goodman. Hypothesis search: Inductive reasoning with language models. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *12th International Conference on Learning Representations (ICLR 2024)*. Curran Associates, Inc., 2024.

Jiannan Xiang, Tianhua Tao, Yi Gu, Tianmin Shu, Zirui Wang, Zichao Yang, and Zhiting Hu. Language models meet world models: embodied experiences enhance language models. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, pp. 75392–75412. Curran Associates Inc., December 2023.

Alex Zhang, Khanh Nguyen, Jens Tuyls, Albert Lin, and Karthik Narasimhan. Language-guided World Models: A Model-based Approach to AI Control. In Parisa Kordjamshidi, Xin Eric Wang, Yue Zhang, Ziqiao Ma, and Mert Inan (eds.), *Proceedings of the 4th Workshop on Spatial Language Understanding and Grounded Communication for Robotics (SpLU-RoboNLP 2024)*, pp. 1–16, Bangkok, Thailand, August 2024. Association for Computational Linguistics. URL <https://aclanthology.org/2024.splurobonlp-1.1/>.

A THE PLAYGROUND-TEXT ENVIRONMENT

We use the Playground-Text environment from (Gaven et al., 2024), itself being a textual adaptation of the Playground environment introduced by (Colas et al., 2020). This textual environment returns, at each time step, a description of its state and possible actions. In our experiments, we used a modified version of Playground-Text where the current state still describes the whole environment’s state but the next state only describes how the performed action changed the environment (see Figure 6).

Playground-Text features 4 different classes of objects with different varieties: *Water*, *Plants* (carrot, potato, beet, berry and pea), *Small Herbivores* (pig, cow and sheep) and *Big Herbivores* (elephant, giraffe, rhinoceros). As depicted in Figure 2, all objects except *Water* possess two states: young (seed for plants and baby for animals) and grown. These objects transition from their young state to their grown one by consuming other objects.

Three different types of actions are possible: *Go to <object>*, to move to the selected object, *Grasp*, to put the object the agent is standing on in its inventory (which has two slots) and *Release <object>/all* to release one or all objects from the inventory on the object the agent is currently standing on. The resulting transitions can be categorized into 6 different types: *Standing*, *Holding 1*, *Holding 2*, *Grow Plant*, *Grow Small Herbivore* or *Grow Big Herbivore*. Growing transitions are triggered when the appropriate objects are released on a young one (see the technology tree in Figure 2). All actions involving non-existent objects in the current scene, as well as release actions that do not result in a transformation (e.g. releasing water onto water) are considered illegal and masked.

State: You see the baby cow, the water, the potato seed, the baby rhinoceros, the water, the pea seed, the water, the potato seed. You are next to the potato seed. In your inventory, there is nothing.

Action: You go to the water.

Change: You are standing on the water.

State: You see the baby cow, the water, the potato seed, the baby rhinoceros, the water, the pea seed, the water, the potato seed. You are next to the water. In your inventory, there is nothing.

Action: You pick up the object.

Change: In your inventory, there is the water.

State: You see the baby cow, the potato seed, the baby rhinoceros, the water, the pea seed, the water, the potato seed. In your inventory, there is the water.

Action: You go to the potato seed.

Change: You are standing on the potato seed.

State: You see the baby cow, the potato seed, the baby rhinoceros, the water, the pea seed, the water, the potato seed. You are next to the potato seed. In your inventory, there is the water.

Action: You give the water.

Change: The objects transform into the potato.

Figure 6: An example of a trajectory in Playground-Text.

B WORLDLLM DETAILS

In this section, we provide details about our WorldLLM implementation. All our LLMs were loaded and prompted using the *transformers* library from Hugging Face. We also provide our code for reproduction on the following repository: <https://github.com/flowersteam/WorldLLM>.

B.1 STATISTICIAN

The LLM used for the Statistician is *Phi-3-mini-4k-instruct* quantized to 4 bits. The prompt used is shown in Figure 7.

System Prompt:
You like doing a lot of puzzles. Please answer with a brief answer and be as precise as you can.

User Prompt:
You are in an environment that contains multiple objects. It can contain water, plant seeds(carrot, porato, beet, berry and pea seeds), small herbivores(pig, cow and ship) and large herbivores(elephant, giraffe, rhinoceros). You can move objects, like water, plants or herbivores and release them on another object to make them interact and transform into a new object. You know that:
<hypothesis>
Your objective is to predict the next change in the environment given the state of the environment and the action taken.
The last state was: <state>
The last action was: <action>
The change is: **Assistant Prompt:**
<change>

Figure 7: Prompt used for the Statistician. The <hypothesis> corresponds to the hypotheses to test. The <state>, <action> and <change> represent the current state, the action taken, and the resulting change, respectively, for the given transition.

B.2 SCIENTIST

The LLM used for the Scientist is the same as the one used for the Statistician. We show its prompt in Figure 8.

B.3 EXPERIMENTER

For the Experimenter, we provide further details for both our oracles and RL-based policies.

B.3.1 ORACLES

We show in Figure 9 the distribution of transitions collected by each of our four oracles.

B.3.2 RL AGENTS

For our RL agents we used Masked PPO from [stable-baselines3 contributions](https://stable-baselines3.readthedocs.io/en/master/modules/ppo_mask.html)². To accelerate training, we used 15 parallel environments and collected 8 trajectories per environment. We show the hyperparameters used in Table 4.

Concerning the reward signals, we provide in Algorithm 2 the details on how RL-ALPEXP compute its ALP reward. RL-ALPEXP introduces an additional hyperparameter α which regulates the reward’s smoothing of the reward. For each transition type, α corresponds to the weight of the previous rewards

²https://sb3-contrib.readthedocs.io/en/master/modules/ppo_mask.html

System Prompt:
 You like doing a lot of puzzles. Please answer with a brief answer and be as precise as you can.

User Prompt:
 You are in an environment that contains multiple objects. It can contain water, plant seeds(carrot, porato, beet, berry and pea seeds), small herbivores(pig, cow and ship) and large herbivores(elephant, giraffe, rhinoceros). You can move objects, like water, plants or herbivores and release them on another object to make them interact and transform into a new object. Your previous experiences were:
 <trajectories>
 Can you find a set of easily understandable and concise hypotheses to predict how the environment will change based on these trajectories? You can take inspiration from the previous rules:
 <previous hypothesis>
 You also know that the previous set of rules failed the most on those trajectories:
 <worst trajectories>
 Answer with just the hypothesis.

Figure 8: Prompt used for the Scientist. <trajectories> correspond to the trajectories collected by the Experimenter, <previous hypothesis> corresponds to the set of hypotheses that have been accepted at the last iteration and <worst trajectories> to the trajectories where the previous hypotheses obtained the worst log-likelihood.

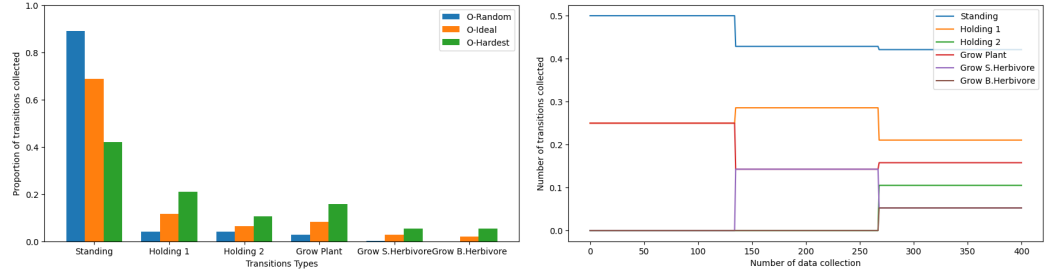


Figure 9: Proportion of collected transition types for each oracle. We show on the left the distribution of our stationary oracles and on the right the evolving distribution of O-Curriculum.

obtained for this transition type. Its role is highlighted in our study from Appendix D.4.1 but α can be interpreted as an adjustment factor determining the extent to which rewards from previous steps are retained. A value of 0 for α implies that only the reward from the most recent step is considered, which corresponds to RL-ALP, while a value of 1 indicates that the rewards are accumulated from the beginning. After a manual hyperparameter search (see Appendix D.4.2), we chose $\alpha = 0.9$ across all transition types.

hyperparameter	value
gamma	0.99
lr	$5e - 4$
vf coef	0.5
gae λ	0.95
entropy	0.01
epochs	20
architecture	64×64

Table 4: PPO hyperparameters used for our RL agents

Algorithm 2 RL-ALPEXP

1: **Input:** α moving average coefficient, \mathcal{I} set of transition types, T total number of iterations, n_{steps} number of Metropolis steps at each iteration, f function returning the type of a transition given a transition, π the Experimenter
2: **Initialize** $(m_i)_{i \in \mathcal{I}} \leftarrow 0$ ▷ Initialize moving average for each type
3: **for** $i = 0$ to n **do**
4: Collect trajectories $\mathcal{D}_t \leftarrow \text{env}(\pi)$
5: Perform n_{steps} Metropolis steps
6: **for** $\tau \in \mathcal{D}_t$ **do**
7: Compute reward $(r_{alp})_\tau$ ▷ Compute ALP reward using equation 2
8: $(r_{alpexp})_\tau = \alpha * m_{f(\tau)} + (1 - \alpha) * (r_{alp})_\tau$ ▷ Compute ALPEXP reward
9: **end for**
10: **for** $i \in \mathcal{I}$ **do**
11: $\tilde{m}_\beta = \frac{\sum_{\tau \in \mathcal{D}_t} \mathbf{1}_{f(\tau)=\beta} (r_{alp})_\tau}{\sum_{\tau \in \mathcal{D}_t} \mathbf{1}_{f(\tau)=\beta}}$ ▷ Compute mean over transitions
12: $m_\beta = \alpha * m_\beta + (1 - \alpha) * \tilde{m}_\beta$ ▷ Update moving average
13: **end for**
14: Update Experimenter π
15: **end for**

C GENERALIZATION TESTS

We show in Figure 10 how the syntax of observations was changed in our generalization tests. In particular, we modified the words associated to the transition (e.g. standing, transform) and inverted the sentence. These changes were also applied to the prompts used by the Scientist and the Statistician.

Standing: You are standing on the $\langle object \rangle$. \rightarrow The $\langle object \rangle$ is beneath you.
Holding 1: In your inventory, there is the $\langle object \rangle$. \rightarrow The $\langle object \rangle$ is in your grasp.
Holding 2: In your inventory, there are the $\langle object \rangle$ and the $\langle object \rangle$. \rightarrow The $\langle object \rangle$ and the $\langle object \rangle$ are in your grasp.
Grow Plant: The objects transform into the $\langle object \rangle$. \rightarrow The $\langle object \rangle$ results from combining the objects.
Grow Small Herbivore: The objects transform into the $\langle object \rangle$. \rightarrow The $\langle object \rangle$ results from combining the objects.
Grow Big Herbivore: The objects transform into the $\langle object \rangle$. \rightarrow The $\langle object \rangle$ results from combining the objects.

Figure 10: Changes in the observations for the generalization environment.

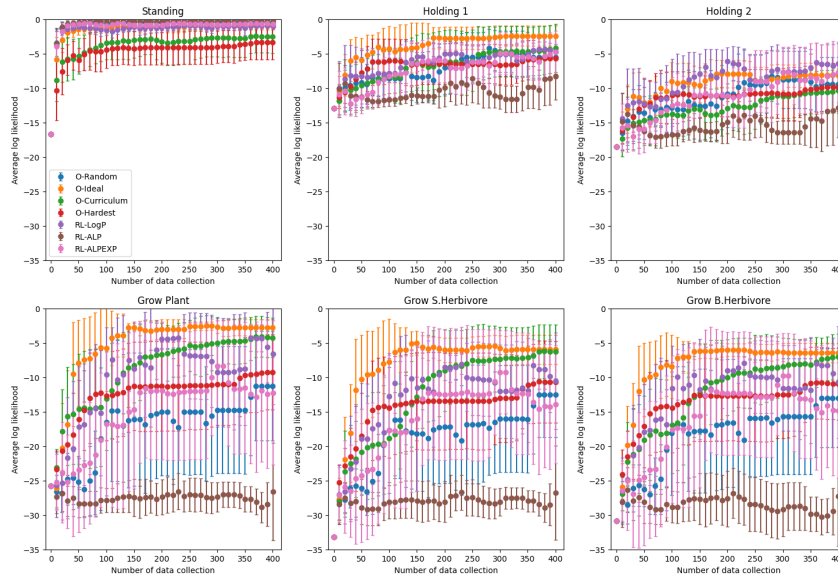


Figure 11: Evolution of log-likelihoods on the test set for all methods.

D ADDITIONAL RESULTS

In this section, we first provide the complete evolution of log-likelihoods on our test set from Section 3.3. Then, we provide additional studies and experiments to improve one’s comprehension of WorldLLM’s behavior. These investigations notably explore the reasons behind the failure of RL-ALP and its differences with RL-ALPEXP.

D.1 EVOLUTION OF LIKELIHOODS

We report in Figure 11 the evolution of log-likelihoods on our test set for all the evaluated methods.

D.2 ANALYSIS OF THE METROPOLIS ALGORITHM IN THE SCIENTIST

D.2.1 HYPOTHESES RETENTION RATE

First, we analyze how the Metropolis algorithm performs when using an LLM as both the proposal and target distribution. In particular, we study the evolution of accepted hypotheses throughout WorldLLM’s iterations when using different Experimenters (Figure 12).

Our results show a high percentage of accepted hypotheses (around 20%) at the beginning of the experiment. However, as iterations progress, it becomes increasingly challenging for the Scientist to generate more effective hypotheses, leading to a decline in the retention rate (it reaches 2%). As highlighted in our manuscript, a developmental strategy is key in WorldLLM. This is here illustrated by O-Curriculum’s ability to maintain a higher retention rate than all other oracles. Additionally, preliminary experiments performed with Phi4 suggested that employing a larger LLM for the Scientist improves hypotheses generation and leads to higher retention rates for the same Statistician model.

D.2.2 EMBEDDINGS OF THE HYPOTHESES

In order to compare the hypotheses generated by our Scientist, we used a sentence transformer, **all-mpnet-base-v2**³, to project hypotheses in a latent space. We then used T-SNE(Maaten & Hinton, 2008) to obtain a 2D representation of the embeddings. As the number of generated hypotheses is high, we plotted only 1/10 of them for each Experimenter and seed. We show these embeddings in Figure 13.

³<https://huggingface.co/sentence-transformers/all-mpnet-base-v2>

Percentage of hypotheses kept as a function of the Metropolis algorithm inference.

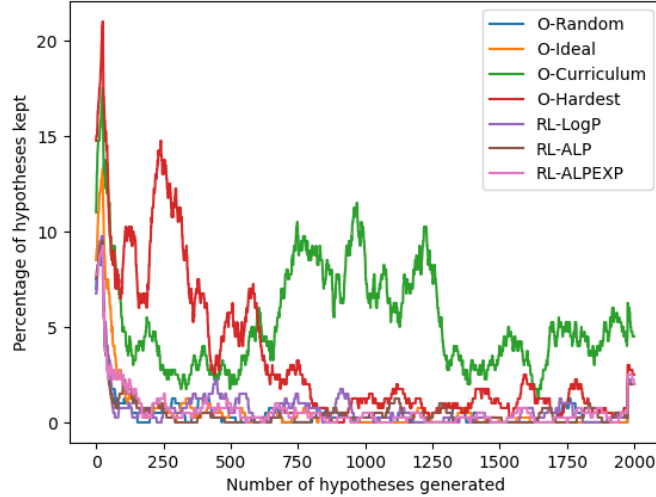



Figure 12: Evolution of the hypotheses retention rate for the different Experimenters.

The results show that O-Curriculum and O-Hardest mostly produce similar hypotheses. This indicates a Scientist’s tendency to refine previous hypotheses rather than generating brand new ones. This corroborates our results on retention rate from previous section. We also compute an inertia score across all seeds for each Experimenter:

$$\forall \beta \in \mathcal{M}, \quad I = \sum_{i=1}^n (x_i - C_\beta)^2 \quad \text{with} \quad C_\beta = \sum_{i=1}^n \frac{x_i}{n} \quad (4)$$

with \mathcal{M} the set of Experimenters, x_i the 2D projection from T-SNE for each set of hypotheses and n the total number of hypotheses.

Results indicate that the clusters produced by O-Curriculum and O-Hardest are each produced by a different seed. This highlights the current limitation of our Bayesian Inference approach where a single particle is used and never reset. As a result, our Scientist may easily get stuck in local optima.



Methods	Inertia Score
O-Random	4.12e6
O-Ideal	4.38e6
O-Curriculum	2.77e6
O-Hardest	2.88e6
RL-LogP	4.52e6
RL-ALP	3.62e6
RL-ALPEXP	4.92e6

Table 5: Mean of the inertia score on the T-SNE embeddings across the different seeds for each Experimenter.

Figure 13: T-SNE embeddings obtained using the *all-mpnet-base-v2* sentence transformer on the hypotheses generated by different Experimenters.

D.3 ANALYSIS OF GENERATED HYPOTHESES

In this section, we propose a qualitative study of hypotheses produced by WorldLLM when using O-Curriculum, O-Ideal and RL-LogP as Experimenters (see Figure 14).

One can observe that each set of hypotheses covers multiple transition types. This can be explicit, as seen in the hypotheses from O-Ideal, or more abstract, as the ones from O-Curriculum, which replace specific animals and objects with abstract representations. A closer examination of O-Curriculum’s hypotheses shows that new entities that do not exist in Playground-Text such as "bear", "mouse" or "wolf" are introduced. This aligns with our argument that natural language hypotheses could help the world model generalize. However, our poor generalization results can be explained by the fact that most hypotheses contain examples of transitions from the environment. Such examples tremendously help our Statistician predict outcomes but also significantly impair its generalization capabilities.

D.4 A DEEPER STUDY RL-ALP

In this section, we provide a deeper analysis of RL-ALP’s failure. While expected to be more robust than RL-LogP (e.g. to the "noisy TV" problem), our results showed RL-ALP was unable to collect other transitions than *Standing*, leading to very poor hypotheses. We begin this section by studying how directly using the per-transition ALP as reward for an RL policy is challenging. Then, show how RL-ALPEXP differs and helps stabilizing training.

D.4.1 FAILURE OF RL-ALP AS THE REWARD

To get a better grasp of the RL policy’s behavior when using RL-ALP, we plot the evolution of rewards obtained by the policy in Figure 15. We observe that most rewards equal 0. Higher rewards are only obtained in the first data collections when *Standing* transitions are not well predicted by the Statistician. After gathering such transitions, the policy eventually converges to a random one. This absence of rewards can be explained by our reward definition from Eq. 2.

A positive reward is obtained only if a set of hypotheses are accepted between the previous and the current iteration. However, as discussed in Appendix D.2.1, the number of accepted hypotheses tends to quickly drop and eventually reach a very low percentage. As a result, very few improvements (or degradation) are made to the Statistician’s log-likelihoods, leading to zero rewards.

D.4.2 ABLATION STUDY ON RL-ALPEXP

In Appendix B.3.2, we detailed RL-ALPEXP, which includes a new hyperparameter α . This hyperparameter governs the weighting of smoothing applied to the reward computed on various transition types over time. In this section, we explore the impact of this hyperparameter.

To assess how α influences RL-ALPEXP, we conducted qualitative experiments with α values set to $\{0.6, 0.75, 0.9, 0.95\}$. To gain insights into the underlying behavior, we analyze the transitions collected by the policy. The results can be found in Figure 16.

From these results, we observe two behaviors depending on the values of α . On one hand, when α is smaller than 0.9, the RL policy demonstrates a periodic behavior. This phenomenon becomes more pronounced as α decreases, leading to instability that hinders WorldLLM’s ability to improve consistently.

It seems that the reward diminishes to zero before the Scientist and Statistician can refine the existing hypothesis. This decline ultimately leads to a collapse in the RL agent’s policy. As a result, the distribution of the collected data begins to resemble that of a random agent, as in the first iterations of WorldLLM. In response, the Scientist generates hypotheses that cause the framework to revert to an earlier stage.

On the other hand, as α gets closer to 1, the reward takes increasingly longer to fade. This can be problematic as the reward function used is an estimation of the learning progress. For example, the high rewards collected at the beginning of training are no longer relevant after multiple iterations.

This hyperparameter search explicitly shows that RL-ALPEXP has not fully resolved the sparse reward problem inherent to RL-ALP. If α is too small, the agent’s policy collapses to random. If it is too close to 1, the old rewards obtained require too much time to fade. One potential solution could be to add a small fraction of RL-LogP’s rewards. This addition would prevent the agent’s policy from collapsing but would require a new hyperparameter to balance the reward from RL-ALP and RL-LogP.

O-Curriculum:

1. If the action is 'You go to' and the object is a seed or gentle creature (cow, elephant, sheep, giraffe, bear, mouse, wolf), predict "You are standing on the [object]."
2. If you are standing on an animal's seed and the action is 'You give the water', predict "The objects transform into the [animal]."
3. If the object is a seed or gentle creature (cow, elephant, sheep, giraffe, bear, mouse, or wolf) and the action is 'You give the water', predict "The objects transform into the [animal]."
4. If you are standing on the water and the action is 'You pick up', predict "In your inventory, there is the water."
5. If you are standing on the water and the action is 'You pick up the object', predict "In your inventory"

O-Ideal:

1. The action was: You go to the water, and the change was: You are standing on the water.
2. The action was: You pick up the object, and the change was: In your inventory, there is the water.
3. The action was: You go to the beet seed, and the change was: You are standing on the beet seed.
4. The action was: You give the water, and the change was: The objects transform into the beet.
5. The action was: You pick up the object, and the change was: In your inventory, there is the beet.
6. The action was: You go to the baby cow, and the change was: You are standing on the baby cow.
7. The action was: You give the beet, and the change was: The objects transform into the cow.
8. The action was: You go to the pe'

RL-LogP:

1. Go to the water.
Change: You are standing on the water.
2. Pick up the object.
Change: In your inventory, there is the water.
3. Go to the water.
Change: You are standing on the water.
4. Pick up the object.
Change: In your inventory, there are the water and the water.
5. Go to the berry seed.
Change: You are standing on the berry seed.
6. Give the water.
Change: The objects transform into the berry.
7. Go to the water.
Change: You are standing on the water.
8. Pick up the object.
Change: In your inventory, there are the water and the water.
9. Go to the pea seed.
Change: You are standing on the pea seed.
10. Give the water.
Change: The objects transform into

Figure 14: Examples of accepted hypotheses.




Figure 15: The evolution of rewards obtained by the RL agent when using RL-ALP for a single seed. We plot the sum of rewards obtained for each PPO update.




Figure 16: Transitions collected by the Experimenter when using RL-ALPEXP for $\alpha = \{0.6, 0.75, 0.9, 0.95\}$ respectively.