LANGUAGE AND EXPERIENCE: A COMPUTATIONAL MODEL OF SOCIAL LEARNING IN COMPLEX TASKS

Anonymous authorsPaper under double-blind review

ABSTRACT

The ability to combine linguistic guidance from others with direct experience is central to human development, enabling safe and rapid learning in new environments. How do people integrate these two sources of knowledge, and how might AI systems? We present a computational framework that models human social learning as joint probabilistic inference over structured, executable world models given sensorimotor and linguistic data. We make this possible by turning a pretrained language model into a probabilistic model of how humans share advice conditioned on their beliefs, allowing our agents both to generate advice for others and to interpret linguistic input as evidence during Bayesian inference. Using behavioral experiments and simulations across 10 video games, we show how linguistic guidance can shape exploration and accelerate learning by reducing risky interactions and speeding up key discoveries in both humans and models. We further explore how knowledge can accumulate across generations through iterated learning experiments and demonstrate successful knowledge transfer between humans and models — revealing how structured, language-compatible representations might facilitate human-machine collaborative learning.

1 Introduction

Imagine learning to forage mushrooms in autumn woods. Each outing provides direct experience of promising slopes and soil conditions. But an experienced forager's advice can transform your exploration in two crucial ways: "Never touch the red ones with white spots; they're deadly" helps you avoid fatal mistakes, while "Look for chanterelles near oak trees after warm rains" turns random wandering into focused search. This ability to integrate linguistic guidance with direct experience is fundamental to human intelligence, enabling not only safer and more efficient learning, but also the accumulation of knowledge across generations (Tomasello, 2009; Boyd et al., 2011). Yet, we still lack a general computational account of how humans combine these two modes of knowledge acquisition to inform exploration and decision-making in complex tasks.

Current computational models capture only fragments of this dual learning capability. Reinforcement learning (RL) agents, for example, can master complex tasks but require extensive trial-and-error—millions of interaction steps—before achieving proficiency (Sutton and Barto, 2018; Mnih et al., 2015). Theory-based RL mitigates this limitation by combining planning with Bayesian inference over structured world models, achieving human-like sample efficiency, yet remains incapable of social learning (Tsividis et al., 2021; Griffiths et al., 2010; Lake et al., 2017). Attempts to bridge this gap with language-conditioned RL integrate linguistic input into the learning process but rely on massive amounts of paired experience and language, making real-world application impractical (Zhong et al., 2020; Luketina et al., 2019). While large language models (LLMs) excel at processing linguistic guidance flexibly (Brown et al., 2020), they struggle with interactive planning and embodied learning (Valmeekam et al., 2023; Paglieri et al., 2024). Bayesian models of social cognition provide a promising step towards integrating observed behaviors and language to infer goals (Shafto et al., 2014; Jara-Ettinger et al., 2016), yet they typically operate in simple, non-interactive tasks with predefined hypothesis spaces (Ying et al., 2023; Zhi-Xuan et al., 2024).

Our primary aim is to contribute to cognitive science by offering a computational account of how humans integrate language and experience. We propose a Bayesian framework that treats linguistic

¹See more detailed related work in Appendix Section A.

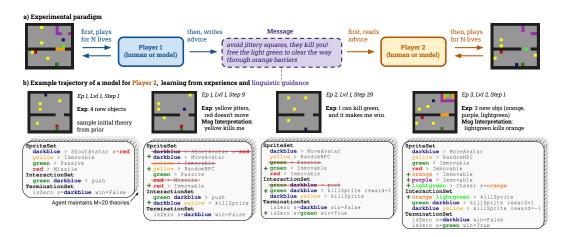


Figure 1: **Overview of the model and experiments** a) *Experimental design*: players (participants or models) are given N=10 lives to learn to play a new video game, either from experience only (Player 1) or from experience and advice written by a previous player (Player 2). b) *Example learning trajectory*: The model maintains beliefs about possible game rules and objectives (programs at bottom), and constantly refines them based on an initial linguistic guidance and new incoming experience.

guidance and direct experience as complementary evidence sources learners can leverage to infer executable, program-like world models. To make this possible, we introduce three key contributions:

- LMs as speaker models: We leverage LMs to approximate the probability that a human with specific world beliefs would produce particular linguistic advice, enabling our model to interpret and generate human-interpretable guidance.
- **Inference from linguistic evidence:** This speaker model is used to evaluate the plausibility of received advice under different candidate world models, allowing linguistic input to shape Bayesian inference alongside experiential learning.
- LM-accelerated inference: We use LMs to transform advice into targeted proposal distributions, guiding Bayesian updates towards the most promising regions of the hypothesis space.

These mechanisms allow our computational model to learn efficiently from naturalistic linguistic input, update beliefs in real-time, and share discoveries with others — forming the basis for more structured and communicative learning systems.

We validate our framework through human experiments and computational simulations, showing that linguistic guidance accelerates learning, shapes exploration strategies, and supports knowledge transfer across generations as well as between humans and models. These results reveal how language drives cultural transmission and illustrate how structured, language-compatible representations can facilitate human–AI collaborative learning.

2 VIDEO GAMES AS LEARNING ENVIRONMENTS

Video games provide an ideal experimental paradigm for studying social and individual learning (Allen et al., 2024): they offer rich causal environments for systematic exploration, create natural pressure to learn efficiently through costs like lost lives, and enable the study of how complex mechanics can be communicated through language. Drawing inspiration from earlier work modeling human causal learning (Tsividis et al., 2021; Tomov et al., 2023), and cultural knowledge transmission (Tessler et al., 2021), we reuse a suite of 10 games defined in the Video Game Description Language (VGDL) (Schaul, 2013).

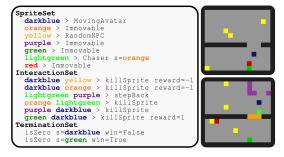


Figure 2: **Game example (beesAndBirds).** Players must discover rules and objectives on their own, sometimes helped with advice from others, to solve 4 game levels (right: levels 1 and 2/4).

VGDL lets us define games by specifying object types, collision effects, rewards, and win/loss conditions (e.g. yellow objects are deadly, victory requires eliminating green objects; see Figure 2 and DSL description in Appendix Section C). This formal representation enables precise experimental control and provides a structured language for our model to represent and update hypotheses about game dynamics. Players, however, encounter the games as grids of colored squares, having to discover game dynamics and objectives through exploration. Our games were designed to test diverse learning capabilities: some require spatial reasoning to navigate hazards (*pushBoulders*) or use teleportation mechanics (*portals*); others demand quick tactical decisions like shooting threats (*aliens*) or defending resources (*plaqueAttack*). Most challenging are games requiring systematic experimentation to discover novel objects through object combinations (*relational*). Play them here.

3 Computational model of integrated learning

We frame game learning as a problem of sequential decision-making under uncertainty (Kaelbling et al., 1998). At each time step t, agents select actions a_t to maximize expected cumulative rewards:

$$a_t = \arg\max_{a} \mathbb{E}_{s_{i+1} \sim P(s_{i+1} \mid s_i, a)} \left[\sum_{i=t}^{\infty} \gamma^{i-t} r_i \right],$$

where s_i and r_i are the agent's state and reward at time i, and γ is a discount factor. The main challenge lies in the agent's uncertainty about the transition function $P(s_{i+1} | s_i, a)$, which is governed by unknown game dynamics.

Our approach extends the *theory-based RL* framework by proposing to infer causal world models *jointly from experience and linguistic guidance* to support goal-directed planning and strategic exploration (Tsividis et al., 2021). The model alternates between three phases: 1) it constructs and infers a posterior over probabilistic world models given both experience and linguistic guidance, 2) it plans action sequences by identifying and pursuing high-value interactions that balance exploration and exploitation, and (3) it executes these actions in the game. The following sections detail each component of this learning loop.

3.1 Inference of structured, causal world models

Our agent models its environment with a distribution over structured world models, each represented as a probabilistic program specifying game rules and objectives. These beliefs are continually updated as the agent gathers new evidence from gameplay experience E and linguistic guidance L. We formalize these beliefs update as a Bayesian inference over possible world theories T:

$$P(T \mid E, L) \propto P(E \mid T) \times P(L \mid T) \times P(T),$$

where P(T) encodes prior beliefs over plausible world models, while $P(E \,|\, T)$ and $P(L \,|\, T)$ measure the consistency of theory T with experiential and linguistic evidence, respectively. As illustrated in Figure 1b, the agent continuously refines its beliefs, incrementally integrating new data from experience and reinterpreting language to narrow down its hypothesis space. In the following subsections, we detail: 1) the search space of possible world models, 2) the likelihood functions that quantify the fit of experience and linguistic guidance, and 3) the inference algorithm that efficiently approximates the posterior distribution over possible worlds.

A space of possible worlds. Each candidate theory T is represented as a VGDL program that specifies object types for each object color (e.g. missile, shooting avatar), interaction effects between objects (e.g. collision with yellow kills avatar), reward functions (e.g. +1 when avatar kills green), and win/loss conditions (e.g. kill all green), see example in Figure 2 and complete list of VGDL primitives in Appendix Section C. We define a simplicity-biased prior P(T) over the search space, favoring theories with fewer rules: 1) object types are uniformly distributed, 2) any object pair has a p=0.25 chance of interacting, with interaction types sampled uniformly, and 3) each object's death has a p=0.1 chance of contributing to win or loss conditions. Sampling a theory from this prior involves generating object types, object pair interactions, and win/loss conditions accordingly. These theories are executable: they can be compiled into playable games that lets the agent simulate trajectories internally. Play them yourself on our demo website.

Likelihood from experience. To estimate the likelihood $P(E \mid T)$, we first decompose the agent's experience E—a sequence of symbolic state transitions—into a sequence of discrete events e_i (object movements, appearance or disappearance, rewards, and win/loss events). We assume that these local events are conditionally independent given the theory T, which lets us factorize the likelihood as $P(E \mid T) = \prod_{i=1}^n P(e_i \mid T)$. Because candidate world models are executable, we can estimate $P(e_i \mid T)$ through simulations. Specifically, we replay the agent's actions under T by initializing the game engine to the agent's previous state and executing its chosen action. We then track the occurrence of each observed event e_i across 20 independent simulations, using their frequencies as empirical estimates of $P(e_i \mid T)$.

Likelihood from language. Linguistic advice received from other agents serves as evidence for evaluating candidate theories T, modeled through $Bayesian\ Theory-of-Mind\ (Baker\ et\ al.,\ 2011).$ We formalize $P(L\,|\,T)$ as the probability that a speaker, believing T to be true, would produce the observed message L. To approximate this, we use a language model (LLaMA-3.1-70B) as a probabilistic speaker model. Given a description of T, the LM is prompted to generate advice for a future player, see prompt in Appendix J. The likelihood is then estimated as the LM's probability of producing the exact message L: $P(L\,|\,T)\approx P_{LM}(L\,|\,\mathrm{prompt}(T))$. This approximation measures how well T explains the speaker's linguistic behavior: e.g. if L contains the advice "avoid yellow at all cost!", then $P(L\,|\,T)$ will be higher if T contains the rule "yellow kills avatar," than if it does not.

Inference algorithm. The space of possible theories is vast — exceeding 10^{20} configurations for games with just five objects — making exact Bayesian inference intractable. To approximate the posterior distribution P(T | E, L), we use a particle filter with Metropolis rejuvenations (Metropolis et al., 1953; Chopin, 2002). We maintain a population of M=20 candidate theories and iteratively refine them by: (1) resampling theories proportional to their posterior probability, and (2) proposing local modifications guided by observed events and linguistic guidance (see details in Appendix Section E). These modify exactly one rule at a time: an object's type, the interaction between a pair of objects, a win condition, a loss condition, or a reward function. They are accepted in proportion to the ratio of posterior probabilities between modified and original theories: p_{accept} min(1, P(T' | E, L)/P(T | E, L)) (Metropolis et al., 1953). This process efficiently approximates the posterior distribution P(T | E, L) over candidate theories, with each particle T_i assigned a weight w_i proportional to its posterior probability given experience and linguistic guidance. The resulting distribution represents the agent's belief over possible game dynamics and objectives. The agent executes one inference step (1 resampling step + 5 rejuvenation steps) every 20 environment steps, and 20 inference steps every time a new kind of object appears in the scene, or when the agent dies; see full pseudo-code of the inference algorithm in Appendix Section D.

Language-guided proposals. We use the same LM to bias the proposal of game rules — e.g. after receiving the message "yellow kills you," the LM may propose rules capturing this lethal interaction:

$$P(r_{\text{new}} | r_{\text{old}}, E, L, T) \propto ((P_0(r_{\text{new}} | E, T) + P_{\text{LM}}(r_{\text{new}} | \text{prompt}(L, T)))/2$$

where $r_{\rm new}$ is a candidate rule (e.g. "yellow kills avatar"), $r_{\rm old}$ is the current rule, P_0 is the base proposal distribution, and $P_{\rm LM}$ is the probability the language model assigns to that rule given the message. This is implemented by prompting the LM with the received message L and instructing it to answer multiple-choice questions about specific VGDL rules: e.g. does the yellow object: 1) kills the avatar, 2) steps back against the avatar, etc.; see detailed prompt in Appendix Section J. This process biases inference towards theories containing rules most compatible with received advice, resulting in faster convergence.

3.2 GOAL-DIRECTED PLANNING AND STRATEGIC EXPLORATION

To maximize its expected long-term utility, the agent must balance exploration — gathering information to refine its world model — and exploitation — leveraging its current understanding to achieve game objectives. Planning is guided by the maximum a posteriori (MAP) theory $T_{\rm MAP}$, inferred during Bayesian updates: the agent's best estimate of game rules and objectives. Based on $T_{\rm MAP}$ simulations, the agent selects high-level goals, and plans action sequences to achieve them.

Goal sampling. Based on $T_{\rm MAP}$, the agent defines a space of high-level goals as object-object interactions that it can cause in the environment: collisions between the agent, something it can push or shoot, and any other object. Each goal is assigned two values: 1) an *exploitation value* reflecting its contribution to game objectives (e.g. higher if it is thought to trigger a reward or a win),

and 2) an *exploration value* representing its potential to reduce model uncertainty, measured as the disagreement about what would happen across the M=20 candidate world models. Subgoals are sampled in proportion to their combined value, balancing both learning and game progress.

Action planning. To achieve these goals, the agent optimizes 10-step action sequences using $T_{\rm MAP}$ for simulation. Initial action plans are refined through a simple genetic algorithm to maximize both game rewards and progress toward goals. To prevent catastrophic errors, the agent performs ten 3-step lookaheads to detect possible deaths or major deviations from the expected reward, triggering replanning when necessary. More details about planning can be found in Appendix Section F.

3.3 GENERATING LINGUISTIC GUIDANCE FOR OTHERS

The agent generates linguistic advice for future players by sampling from the same speaker model used to evaluate language likelihood during inference — effectively translating its MAP theory $T_{\rm MAP}$ into natural language $L_{\rm generated} \sim P_{\rm LM}(L,|,{\rm prompt}(T_{\rm MAP}))$. This provides optimal speaker modeling when message emitters are computational agents and approximate modeling when they are human. By using the LM both to interpret linguistic guidance and to generate it, the agent captures key aspects of Bayesian Theory of Mind — modeling how humans communicate their beliefs and how they interpret the beliefs of others through language.

3.4 BASELINE MODELS

We compare our approach to three baselines: 1) *Oracle*: a model that plans to solve the game using ground-truth game rules, 2) *Deep RL*: a Double Deep Q-Network agent implementing pure trial-and-error learning without structured representations (Mnih et al., 2015; Van Hasselt et al., 2016), and 3) *pure LM*: an LM agent (*LLaMA-3.1-70B*) prompted to reflect, plan and save notes in a scratchpad before each action (Yao et al., 2022; Nye et al., 2021); see detailed prompt in Appendix Section J. These baselines test the importance of structured representations and goal-directed planning. Note that we do not compare to language-conditioned RL baselines (Luketina et al., 2019) as these methods require training across thousands of episodes to learn language-policy mappings, while our setting involves single-shot linguistic guidance per game.

4 EXPERIMENTAL PARADIGM

To investigate how humans and computational models integrate experiential and linguistic evidence during learning, we conducted a series of IRB-approved experiments comparing three learning conditions (see Figure 1a):

- 1. **Experience**: Players learn solely through direct interaction with the game.
- 2. Experience + human message: Players receive additional advice from previous human players.
- 3. Experience + model message: Players receive additional advice from previous model players.

This design allows us to examine both the effectiveness of linguistic guidance and potential asymmetries in human-model knowledge transfer. In each condition, players had 15 lives to solve four levels of each game, advancing only after completing the current level.

Participants. We recruited 122 participants through Prolific to play 5 randomly-assigned games. To ensure task engagement while maintaining a representative sample, we excluded participants who failed to complete at least one level in ≥ 3 games (final N=120). Participants were randomly assigned to one of the 3 conditions (N=40 each). In social conditions, each participant received advice from a randomly-selected previous player (either human or model) who had completed the game in the experience-only condition (1-to-1 mapping).

Procedure. All participants first completed a brief tutorial game to familiarize themselves with the interface and basic game mechanics. In social conditions, participants read advice from previous players before starting each new game, and during gameplay. After completing each game (either by winning or depleting lives), participants wrote advice "to help future players who have not yet played the game." This prompt encouraged participants to distill their learned knowledge into linguistic guidance, see full instructions in Appendix Section G.

Analysis approach. To evaluate learning efficiency, we tracked both the number of lives required to complete each level and the total proportion of levels completed. We use a normalized area-under-curve (nAUC) metric ($\in [0,1]$) to integrate both proficiency and learning efficiency. To analyze message effectiveness, we manually coded advice content along four dimensions: the fraction of useful information about game dynamics, about loss conditions, about win conditions, and the presence of incorrect information. This coding scheme lets us examine how specific types of linguistic guidance shape exploration and learning outcomes.

Computational simulations. We conducted 20 simulation runs per condition, matching the human sample size to enable direct comparison. The model received the same information as human participants: in social conditions, it processed the same messages (human- or model-generated) that humans received, while in the experience-only condition, it learned purely through interaction. We also ran iterated learning experiments where a sequence of 10 agents, each given two lives, played the game and passed a message to the next agent. This design tests whether partial knowledge can accumulate incrementally across generations, mirroring human cultural learning (Tessler et al., 2021).

5 RESULTS

We first examine how humans and models learn novel games from pure experience, before analyzing how linguistic guidance shapes this process. We then look at human–model transfer, before exploring how our model can accumulate partial knowledge across generations. Our results reveal both striking similarities and systematic differences in learning strategies across humans and models.

5.1 LEARNING FROM EXPERIENCE

How efficiently can humans and models learn novel game dynamics through pure experience? Both demonstrated remarkable sample efficiency, with median participants solving 9 of 10 games and our model solving all 10 games within a 10-life budget (Figure 3). However, systematic differences emerged in games requiring specific cognitive capabilities (see per-game plots in Appendix Figure 7). In *relational*, which demands systematic exploration of object combinations, humans showed a bimodal pattern: 25% achieved model-like efficiency by systematically testing interactions, while 40% failed to solve even two out of

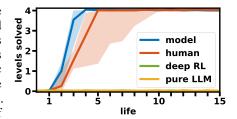


Figure 3: **Learning from experience.** Median across games (N=10, IQ range).

four levels. This split suggests that while humans can perform systematic experimentation, not everyone defaults to it—unlike our model which explicitly reasons about information gain. Conversely, in *avoidGeorge*, which requires rapid planning to protect allies, models consistently outperformed humans (median levels: 4 vs 0), likely due to their capacity for accurate short-term planning.

The importance of structured reasoning becomes clear when comparing against baselines. Pure deep RL (double DQN) failed to solve any level within 10 lives, while pure LM agents never solved more than one level per game, often solving none (7/10 games). The stark difference between these baselines and both human and model performance underscores the value of structured theories of game dynamics in supporting efficient exploration and decision-making.

5.2 LEARNING FROM EXPERIENCE AND HUMAN LANGUAGE

Having established baseline learning capabilities, we next examine how linguistic guidance from previous human players shapes exploration and learning outcomes. Our results demonstrate substantial benefits from social learning while revealing key patterns in effective knowledge transmission.

Benefits of linguistic guidance. How does linguistic guidance shape learning? Both humans and our model showed significantly faster learning when provided with human-written advice, reducing median attempts needed by 1.75 for humans (4 \rightarrow 2.25) and 1.25 for models (2.5 \rightarrow 1.25) (see Figure 4). To quantify these benefits in terms of learning speed, we computed a normalized area under the learning curve (nAUC). A fixed-effect model controlling for game difficulty revealed significant improvements from linguistic guidance for both humans (Δ (nAUC) = 0.12, p = 2.2 × 10⁻⁴) and

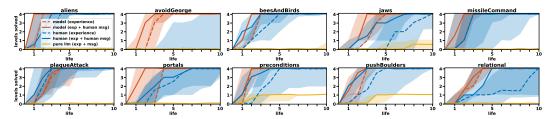


Figure 4: **Learning from experience and human advice.** In most games, both humans (blue) and our model (orange) learn significantly faster with human advice (plain lines) than without (dashed lines), see median in Figure 5a. The *pure LLM* baseline does not learn significantly better with advice, solving the first level in only 3/10 games (N=20, median \pm interquartile range).

models ($\Delta(\text{nAUC}) = 0.04$, $p = 3.3 \times 10^{-2}$). These benefits were most pronounced in games with opaque mechanics like *relational*, or multiple hazards like *portals* and *jaws*, where guidance could directly communicate critical interactions. Benefits were smallest in games requiring primarily motor skills like *aliens*, *missileCommand* or *plaqueAttack*.

What makes effective guidance? Analysis of message content revealed systematic patterns in communication and message effectiveness. Most messages (88%) contained information about game dynamics ("[avoidGeorge] light blue can transform green into them"), with many also including information about win (64%, e.g. "[jaws] The objective is to stay alive") and loss conditions (74%, e.g. "[avoidGeorge] you lose if all squares get turned purple"). A small but notable fraction (11%) contained errors (e.g. "[relational] Push all the blue into the orange," when blue should contact yellow). Longer messages and those containing detailed win conditions proved particularly beneficial to both models and humans($p = 6.3 \times 10^{-3}$ and p = 0.031 respectively). Importantly, messages that helped humans also helped models, shown by significant correlations in performance gains across nAUC (r = 0.17), lives to first level (r = 0.24), and lives to second level (r = 0.16, all p < 0.02)—suggesting shared mechanisms for integrating linguistic and experiential knowledge.

How language shapes exploration. Linguistic guidance systematically altered how players explored game mechanics. Messages warning about dangers significantly reduced costly mistakes: in average, humans and models receiving such warnings experienced between 37% and 67% fewer deaths in avoidGeorge, relational, jaws, and aliens. Messages about key mechanics accelerated their discovery: in relational, players informed about tool creation discovered essential combinations 43% to 83% faster, while in plaqueAttack they learned to revive allies 43% to 62% faster when told about this possibility. However, incorrect advice could also mislead: in avoidGeorge, the human player and the model who were wrongly warned about the danger of green blocks (in fact harmless) completely avoided them for the first three episodes, while others players interacted with them in average 5 times in the same period — demonstrating how unhelpful linguistic advice can also shape direct experience.

5.3 From cognitive models to learning partners?

Our model demonstrated the ability to efficiently learn from human-generated advice. But can it also help humans in return? Can it generate guidance that is useful to other models and, more importantly, to humans? Leveraging its LM-based advice generation capability (Section 3.3), our model produced detailed, pedagogical advice that rivaled human teaching. For example, in *preconditions*:

"Control the darkblue square with arrow keys. Your goal is to kill all gold objects by touching them, earning points along the way. Watch out for green objects—touching them will kill you unless you have white resources to protect yourself. Collect white resources to safeguard against green, and use them to kill green objects if needed, but be aware that each kill will cost a resource."

Model-generated guidance significantly improved learning for both humans (Δ (nAUC)= 0.15, $p < 10^{-5}$) and models (Δ (nAUC)= 0.088, $p < 10^{-8}$) compared to experience alone. Interestingly, models learned better from model- than human-generated advice (Δ (nAUC)= 0.052, $p < 10^{-10}$), while these gave humans only a modest advantage (Δ (nAUC)= 0.035, p = 0.26, n.s.).

These asymmetries reveal important differences in human and model communication. When we had our model generate advice based on human players' trajectories, models still learned better from this model-generated guidance than from human-written messages (Δ (nAUC)= 0.21, $p < 10^{-8}$). This suggests the asymmetry stems not from differences in knowledge, but from communication style. Human messages often included metacognitive strategies ("take time to look for safe patterns"), analogies ("orange is like terminator"), or emotional content ("[I] was left very confused") — aspects that human learners readily use but our model finds harder to interpret. These differences highlight the potential and challenges for language-mediated human-machine collaborative learning. We show more examples of human- and model-generated messages in Appendix Section I and on our website.

Ablating language-guided proposals in our model leads to a significant performance drop compared to the full version ($\Delta(\text{nAUC}) = -0.058$, $p = 7.2 \times 10^{-4}$). This ablated version still leverages language likelihood, which lets it outperform learning from experience alone ($\Delta(\text{nAUC}) = -0.030$, $p = 5.7 \times 10^{-3}$), see Method Section 3.1 and Appendix Figure 8.

Variability across games: Our results reveal interesting patterns in the success and failures of linguistic guidance (Figure 5b). Models learning from other models show consistent benefits across all games because models generate more comprehensive messages and process advice more reliably than humans on average. In contrast, human learners show minimal improvements in games requiring rapid reactions and precise motor control (portals, plaqueAttack, aliens, missileCommand), where linguistic advice cannot substitute for motor practice. Models struggle to learn from humans in relational, where humans often provide imprecise descriptions of complex rule interactions, and plaqueAttack, where humans frequently omit important mechanics they discovered. In contrast, all learners benefit from both human and model advice in games like avoidGeorge, where the critical strategy is non-obvious, and beesAndBirds, jaws, and pushBoulders, where key dangers are memorable and straightforward to describe. These patterns suggest that social learning strategies like selecting high-performing teachers or aggregating multiple sources of advice could mitigate some failure cases.

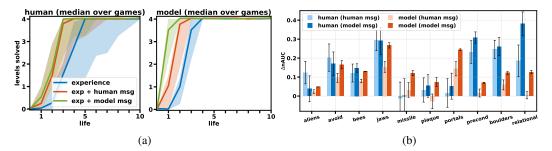


Figure 5: **Bidirectional human–model learning.** a) Median performance across games for humans and models learning from experience alone (blue), or experience combined with human-generated (orange) or model-generated (green) linguistic guidance (N=10, shaded: interquartitle range). b) Performance improvements over individual learning baselines for each game (error bar: sem).

5.4 GENERATIONAL LEARNING

While the previous experiments allowed social learners to benefit from fully explored game mechanics, real-world learning often relies on partial and imperfect knowledge transmission. To investigate whether our model could replicate this gradual accumulation of knowledge, we designed an iterated learning experiment inspired by Tessler et al. (2021). In this setting, each agent interacts with the game environment for only two lives before generating advice to the next agent. This cycle continues across 10 generations for each of the 10 games, with performance tracked generation by generation.

Our results show that learning reliably accumulates across generations in all games where models do not already achieve mastery from generation 1 (preconditions and aliens) (Figure 6). Some games reached near-complete mastery by Generation 2, while others showed more gradual improvements. A fixed-effect model controlling for game variability confirmed this trend, showing significant improvements over the first generation for all subsequent ones ($\Delta(\text{nAUC})_{i>1} \in [0.44, 0.57]$, all $p < 10^{-10}$). However, in plaqueAttack and relational, we observe occasional regressions where later generations underperform earlier ones, highlighting the brittleness of single-teacher transmission.

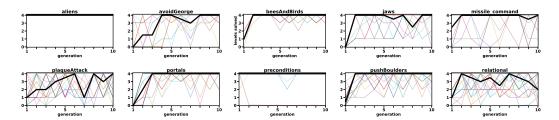


Figure 6: Generational learning via linguistic advice. Median performance of N=20 agents (black) across generations (2 lives each). Our models affords generational learning: performance increases across generations, beyond the performance of individual agents (generation 1).

This phenomenon could be mitigated through the integration of multiple teachers or teacher selection (Kendal et al., 2018; Schultner et al., 2024). Together, these findings indicate that agents can build upon fragmented experiences to gradually refine their world models over generations, mirroring learning dynamics observed in human populations (Tessler et al., 2021).

6 DISCUSSION

Experiments showed that our computational model reproduces key features of human social learning: advice reduces the attempts needed for success by shaping exploration, supports generational knowledge accumulation, and even allows model-generated guidance to help human learners—demonstrating bidirectional knowledge exchange.

While VGDL is only a coarse approximation of human game representations, our results show it captures enough structure to study social learning and afford bidirectional human–model knowledge transfer. Future work could leverage library-learning methods (Ellis et al., 2020; Wong et al., 2021) to model the emergence of shared representations through linguistic interaction, potentially driving representational convergence rather than requiring pre-aligned representations.

Our results show that linguistic guidance often speeds and safeguards exploration, but when advice is wrong it can restrict search—mirroring the "double-edged sword of pedagogy" observed in developmental psychology (Bonawitz et al., 2011). Humans counter this issue by evaluating testimony against prior causal theories before integrating it (Harris et al., 2018; Sobel and Kushnir, 2013). To match this sophistication, computational models will need meta-cognitive mechanisms to judge the reliability of advice and adjust their exploration accordingly.

This paper opens several avenues for future work. Human-generated messages often include game abstractions, high-level strategies, and planning heuristics that our model currently cannot leverage. Extending the framework to interpret and learn from these richer forms of guidance—e.g. through inference of auxiliary reward functions, planning abstractions and strategies—could unlock enhanced social learning capabilities. Beyond passive learning, our model could be further extended to make decisions about who to learn from based on perceived expertise or success—a capacity known as prestige-based social learning (Kendal et al., 2018; Schultner et al., 2024). Finally, one limitation is that our implementation is tailored to the VGDL game space. Scaling to more open-ended game spaces will require advances in program synthesis and probabilistic inference over richer state and action spaces. Leveraging recent progress in probabilistic programming (Cusumano-Towner et al., 2019; Lew et al., 2023) and direct inference over Python code (Tang et al., 2024) could provide the computational efficiency and expressiveness needed to model more complex, dynamic environments.

Finally, our results point to exciting possibilities for human-machine collaborative learning. Our model not only benefits from human-generated guidance but also contributes back through effective, pedagogical advice—closing the collaborative loop. This demonstrates a first step toward bidirectional learning systems capable of supporting human learners. The future directions outlined above—handling richer linguistic guidance, adaptive trust in information sources, and scaling to open-ended domains—would represent major steps toward AI systems that not only learn efficiently but also teach, collaborate, and adapt within complex social learning networks, augmenting collective intelligence in hybrid human—AI communities (Colas et al., 2022; Brinkmann et al., 2023; Collins et al., 2024).

REPRODUCIBILITY STATEMENT

Section 3 details the full model specification and inference procedure. Section 4 describes the human and model experimental design. The appendices provide additional details including: the description of all VGDL primitives, the inference pseudo-code, details about guided proposals and planning, instructions used for human data collection and full prompts used by the LM components of our model. The codebase will be released with the camera-ready version of the paper. Together these resources allow full replication of the experiments.

REFERENCES

- Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Hausman, et al. Do as i can, not as i say: Grounding language in robotic affordances. *arXiv preprint arXiv:2204.01691*, 2022.
- Kelsey Allen, Franziska Brändle, Matthew Botvinick, Judith E Fan, Samuel J Gershman, Alison Gopnik, Thomas L Griffiths, Joshua K Hartshorne, Tobias U Hauser, Mark K Ho, et al. Using games to understand the mind. *Nature Human Behaviour*, pages 1–9, 2024.
- Chris Baker, Rebecca Saxe, and Joshua Tenenbaum. Bayesian theory of mind: Modeling joint belief-desire attribution. In *Proceedings of the annual meeting of the cognitive science society*, volume 33, 2011.
- Chris L Baker, Julian Jara-Ettinger, Rebecca Saxe, and Joshua B Tenenbaum. Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. *Nature Human Behaviour*, 1(4): 0064, 2017.
- Elizabeth Bonawitz, Patrick Shafto, Hyowon Gweon, Noah D Goodman, Elizabeth Spelke, and Laura Schulz. The double-edged sword of pedagogy: Instruction limits spontaneous exploration and discovery. *Cognition*, 120(3):322–330, 2011.
- Robert Boyd, Peter J Richerson, and Joseph Henrich. The cultural niche: Why social learning is essential for human adaptation. *Proceedings of the National Academy of Sciences*, 108(supplement_2): 10918–10925, 2011.
- Levin Brinkmann, Fabian Baumann, Jean-François Bonnefon, Maxime Derex, Thomas F Müller, Anne-Marie Nussberger, Agnieszka Czaplicka, Alberto Acerbi, Thomas L Griffiths, Joseph Henrich, et al. Machine culture. *Nature Human Behaviour*, 7(11):1855–1868, 2023.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Valerie Chen, Abhinav Gupta, and Kenneth Marino. Ask Your Humans: Using Human Instructions to Improve Generalization in Reinforcement Learning. *Proc. of ICLR*, 2021.
- Nicolas Chopin. A sequential particle filter method for static models. *Biometrika*, 89(3):539–552, 2002.
- 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. *Advances in Neural Information Processing Systems*, 33:3761–3774, 2020.
- Cédric Colas, Tristan Karch, Clément Moulin-Frier, and Pierre-Yves Oudeyer. Language and culture internalization for human-like autotelic ai. *Nature Machine Intelligence*, 4(12):1068–1076, 2022.
- Katherine M Collins, Ilia Sucholutsky, Umang Bhatt, Kartik Chandra, Lionel Wong, Mina Lee, Cedegao E Zhang, Tan Zhi-Xuan, Mark Ho, Vikash Mansinghka, et al. Building machines that learn and think with people. *Nature Human Behaviour*, 8(10):1851–1863, 2024.
- Marco F Cusumano-Towner, Feras A Saad, Alexander K Lew, and Vikash K Mansinghka. Gen: a general-purpose probabilistic programming system with programmable inference. In *Proceedings* of the 40th acm sigplan conference on programming language design and implementation, pages 221–236, 2019.

- Pierre Del Moral, Arnaud Doucet, and Ajay Jasra. Sequential monte carlo samplers. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 68(3):411–436, 2006.
 - Kevin Ellis, Catherine Wong, Maxwell Nye, Mathias Sable-Meyer, Luc Cary, Lucas Morales, Luke Hewitt, Armando Solar-Lezama, and Joshua B. Tenenbaum. DreamCoder: Growing generalizable, interpretable knowledge with wake-sleep Bayesian program learning. *arXiv:2006.08381 [cs]*, June 2020. URL http://arxiv.org/abs/2006.08381. arXiv: 2006.08381.
 - Michael C Frank and Noah D Goodman. Predicting pragmatic reasoning in language games. *Science*, 336(6084):998–998, 2012.
 - Noah D Goodman and Michael C Frank. Pragmatic language interpretation as probabilistic inference. *Trends in cognitive sciences*, 20(11):818–829, 2016.
 - Prasoon Goyal, Scott Niekum, and Raymond J Mooney. Using natural language for reward shaping in reinforcement learning. *arXiv preprint arXiv:1903.02020*, 2019.
 - Thomas L Griffiths, Nick Chater, Charles Kemp, Amy Perfors, and Joshua B Tenenbaum. Probabilistic models of cognition: Exploring representations and inductive biases. *Trends in cognitive sciences*, 14(8):357–364, 2010.
 - Paul L Harris, Melissa A Koenig, Kathleen H Corriveau, and Vikram K Jaswal. Cognitive foundations of learning from testimony. *Annual Review of Psychology*, 69(1):251–273, 2018.
 - Felix Hill, Andrew K. Lampinen, Rosalia Schneider, Stephen Clark, Matthew Botvinick, James L. McClelland, and Adam Santoro. Emergent Systematic Generalization in a Situated Agent. *Proc. of ICLR*, 2020.
 - Hengyuan Hu, Denis Yarats, Qucheng Gong, Yuandong Tian, and Mike Lewis. Hierarchical Decision Making by Generating and Following Natural Language Instructions. *Proc. of NeurIPS*, 2019.
 - Wenlong Huang, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. Language models as zero-shot planners: Extracting actionable knowledge for embodied agents. *Proc. of ICML*, 2022.
 - Julian Jara-Ettinger, Hyowon Gweon, Laura E Schulz, and Joshua B Tenenbaum. The naïve utility calculus: Computational principles underlying commonsense psychology. *Trends in cognitive sciences*, 20(8):589–604, 2016.
 - Yiding Jiang, Shixiang Gu, Kevin Murphy, and Chelsea Finn. Language as an abstraction for hierarchical deep reinforcement learning. *Proc. of NeurIPS*, 2019.
 - Leslie Pack Kaelbling, Michael L Littman, and Anthony R Cassandra. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2):99–134, 1998.
 - Rachel L Kendal, Neeltje J Boogert, Luke Rendell, Kevin N Laland, Mike Webster, and Patricia L Jones. Social learning strategies: Bridge-building between fields. *Trends in cognitive sciences*, 22 (7):651–665, 2018.
 - Brenden M Lake, Tomer D Ullman, Joshua B Tenenbaum, and Samuel J Gershman. Building machines that learn and think like people. *Behavioral and brain sciences*, 40:e253, 2017.
 - Alexander K Lew, Tan Zhi-Xuan, Gabriel Grand, and Vikash K Mansinghka. Sequential monte carlo steering of large language models using probabilistic programs. *arXiv* preprint arXiv:2306.03081, 2023.
 - Jelena Luketina, Nantas Nardelli, Gregory Farquhar, Jakob Foerster, Jacob Andreas, Edward Grefenstette, Shimon Whiteson, and Tim Rocktäschel. A survey of reinforcement learning informed by natural language. *arXiv preprint arXiv:1906.03926*, 2019.
 - Nicholas Metropolis, Arianna W Rosenbluth, Marshall N Rosenbluth, Augusta H Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The journal of chemical physics*, 21(6):1087–1092, 1953.

- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.
- Karthik Narasimhan, Regina Barzilay, and Tommi Jaakkola. Grounding Language for Transfer in Deep Reinforcement Learning. *Journal of Artificial Intelligence Research*, 63:849–874, 2018.
- Kolby Nottingham, Prithviraj Ammanabrolu, Alane Suhr, Yejin Choi, Hannaneh Hajishirzi, Sameer Singh, and Roy Fox. Do embodied agents dream of pixelated sheep: Embodied decision making using language guided world modelling. In *International Conference on Machine Learning*, pages 26311–26325. PMLR, 2023.
- Maxwell Nye, Anders Johan Andreassen, Guy Gur-Ari, Henryk Michalewski, Jacob Austin, David Bieber, David Dohan, Aitor Lewkowycz, Maarten Bosma, David Luan, et al. Show your work: Scratchpads for intermediate computation with language models. *arXiv preprint arXiv:2112.00114*, 2021.
- Davide Paglieri, Bartłomiej Cupiał, Samuel Coward, Ulyana Piterbarg, Maciej Wolczyk, Akbir Khan, Eduardo Pignatelli, Łukasz Kuciński, Lerrel Pinto, Rob Fergus, et al. Balrog: Benchmarking agentic llm and vlm reasoning on games. *arXiv preprint arXiv:2411.13543*, 2024.
- Tom Schaul. A video game description language for model-based or interactive learning. In 2013 IEEE Conference on Computational Inteligence in Games (CIG), pages 1–8. IEEE, 2013.
- David Schultner, Lucas Molleman, and Björn Lindström. Feature-based reward learning shapes human social learning strategies. 2024.
- Patrick Shafto, Noah D Goodman, and Thomas L Griffiths. A rational account of pedagogical reasoning: Teaching by, and learning from, examples. *Cognitive psychology*, 71:55–89, 2014.
- Pratyusha Sharma, Antonio Torralba, and Jacob Andreas. Skill Induction and Planning with Latent Language. *Proc. of ACL*, 2021.
- Mohit Shridhar, Xingdi Yuan, Marc-Alexandre Cote, Yonatan Bisk, Adam Trischler, and Matthew Hausknecht. ALFWorld: Aligning Text and Embodied Environments for Interactive Learning. *Proc. of ICLR*, 2021.
- David M Sobel and Tamar Kushnir. Knowledge matters: how children evaluate the reliability of testimony as a process of rational inference. *Psychological Review*, 120(4):779, 2013.
- Richard S Sutton and Andrew G Barto. Reinforcement learning: An introduction. MIT press, 2018.
- Hao Tang, Darren Key, and Kevin Ellis. Worldcoder, a model-based llm agent: Building world models by writing code and interacting with the environment. *arXiv preprint arXiv:2402.12275*, 2024.
- Michael Henry Tessler, Jason Madeano, Pedro A Tsividis, Brin Harper, Noah D Goodman, and Joshua B Tenenbaum. Learning to solve complex tasks by growing knowledge culturally across generations. *arXiv preprint arXiv:2107.13377*, 2021.
- Michael Tomasello. The cultural origins of human cognition. Harvard university press, 2009.
- Momchil S Tomov, Pedro A Tsividis, Thomas Pouncy, Joshua B Tenenbaum, and Samuel J Gershman. The neural architecture of theory-based reinforcement learning. *Neuron*, 111(8):1331–1344, 2023.
- 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. *arXiv preprint arXiv:2107.12544*, 2021.
- Karthik Valmeekam, Matthew Marquez, and Subbarao Kambhampati. Can large language models really improve by self-critiquing their own plans? *arXiv preprint arXiv:2310.08118*, 2023.
- Hado Van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30, 2016.

- Natalia Vélez and Hyowon Gweon. Learning from other minds: An optimistic critique of reinforcement learning models of social learning. *Current opinion in behavioral sciences*, 38:110–115, 2021.
- Catherine Wong, Kevin Ellis, Joshua B. Tenenbaum, and Jacob Andreas. Leveraging Language to Learn Program Abstractions and Search Heuristics. *Proc. of ICML*, 2021.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629*, 2022.
- Lance Ying, Tan Zhi-Xuan, Vikash Mansinghka, and Joshua B Tenenbaum. Inferring the goals of communicating agents from actions and instructions. In *Proceedings of the AAAI Symposium Series*, volume 2, pages 26–33, 2023.
- Lance Ying, Tan Zhi-Xuan, Lionel Wong, Vikash Mansinghka, and Joshua B Tenenbaum. Understanding epistemic language with a bayesian theory of mind. *arXiv preprint arXiv:2408.12022*, 2024.
- Tan Zhi-Xuan, Paul Stefan Lunis, Nathalie Fernandez Echeverri, Vikash Mansinghka, and Josh Tenenbaum. Language models as informative goal priors in a bayesian theory of mind. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 45, 2023.
- Tan Zhi-Xuan, Lance Ying, Vikash Mansinghka, and Joshua B Tenenbaum. Pragmatic instruction following and goal assistance via cooperative language-guided inverse planning. arXiv preprint arXiv:2402.17930, 2024.
- Victor Zhong, Tim Rocktäschel, and Edward Grefenstette. Rtfm: Generalising to new environment dynamics via reading. In *ICLR*, pages 1–17. ICLR, 2020.

Appendices

702

703704705

706 707

708

709

710

711

712

713

714

715

716

717

718

719

720

721

722

723

724

725

726

727

728

729

730

731

732

733

734

735

736

737

738

739

740

741

A EXTENDED RELATED WORK

Language and RL. Language has emerged as a powerful tool to enhance reinforcement learning (RL) by conveying state abstractions (Narasimhan et al., 2018), world dynamics (Zhong et al., 2020), auxiliary reward functions (Goyal et al., 2019), and task decompositions (Shridhar et al., 2021; Sharma et al., 2021). Its inherent compositionality and abstraction capabilities allow agents to pursue more abstract goals (Jiang et al., 2019), generalize effectively across diverse environments (Zhong et al., 2020) and goals (Hill et al., 2020; Colas et al., 2020), and structure long-term decision-making more effectively (Hu et al., 2019; Chen et al., 2021). Despite these advances, most language-based RL approaches rely on substantial amounts of paired data to ground linguistic information in agent experience, limiting their scalability and resemblance to human-like social learning. Language models (LMs) promise more flexible combination of language and decision-making (Ahn et al., 2022; Huang et al., 2022), yet they struggle to learn complex embodied skills that require lowlevel perception and temporally extended actions (Valmeekam et al., 2023; Paglieri et al., 2024). Nottingham et al. (2023) use LLMs to generate complete world model hypotheses that are then verified through experience. Our approach introduces a Bayesian framework that treats experience and language as two complementary sources of evidence in the inference of world models. This joint inference enables rapid adaptation to new linguistic inputs and tasks from the very first interaction, bypassing the need for extensive paired data.

Bayesian models of social cognition. Bayesian models of social cognition have proven to be powerful tools for modeling theory of mind (ToM)—the ability to infer the hidden goals, beliefs, and intentions of others from observable behavior or linguistic cues (Baker et al., 2011; 2017; Frank and Goodman, 2012; Goodman and Frank, 2016). These models formalize social inference as an inverse planning problem, where observers assume that agents act approximately rationally towards their goals and use this assumption to infer likely mental states (Baker et al., 2011; 2017). Extensions of these frameworks to language understanding have resulted in the Rational Speech Acts model, which interprets communication as recursive social reasoning: listeners reason about what a speaker intends to convey based on the assumption that speakers choose utterances optimally, given their own beliefs (Frank and Goodman, 2012; Goodman and Frank, 2016). More recent work integrates Bayesian reasoning with modern machine learning to enable richer social inferences. For instance, (Zhi-Xuan et al., 2023) leverage large language models (LLMs) as priors in a Bayesian goal inference system, allowing for the efficient suggestion and evaluation of likely goals in complex environments. Similarly, (Ying et al., 2024) introduce a language-augmented ToM model that translates natural language statements about beliefs into formal epistemic representations, enhancing agents' ability to reason about others' knowledge and intentions. By combining linguistic input with structured probabilistic models, these approaches extend traditional ToM beyond purely behavioral cues, opening new avenues for interactive and socially aware AI agents (Vélez and Gweon, 2021). Our work builds on these ideas by leveraging Bayesian inference to jointly interpret linguistic and experiential data, allowing agents to efficiently acquire world models from sparse social interactions.

B LEARNING FROM EXPERIENCE ALONE

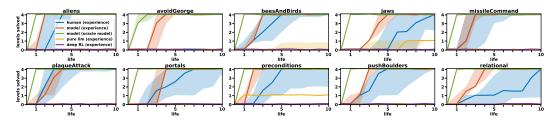


Figure 7: **Learning from experience.** Both humans and our model achieve high sample efficiency across most games. Systematic differences emerge in *relational*, which requires systematic testing of object combinations, and *avoidGeorge*, which demands rapid and precise planning. *Deep RL* and *pure LM* baselines fail to solve games efficiently, highlighting the value of structured representations (N=20, median \pm interquartile range).

C VGDL: GAME PRIMITIVES, STATE AND ACTION SPACES

We work with a subset of the VGDL domain. Here are the possible types of avatars:

- *MovingAvatar*: controllable player that can move in the four directions with a certain speed based on keyboard presses
- ShootAvatar: MovingAvatar that can also shoot objects stype when the player presses the space bar
- FlakAvatar: ShootAvatar that can only move sideways and always shoot upwards.

Here are the possible object types and their parameters:

- Immovable: object that cannot move
- Flicker: object that disappears after total steps
- SpawnPoint: object that spawn objects stype with probability p
- ResourcePack: object that can be collected (see interaction addResource and removeResource)
- Passive: object that can be pushed (see interaction bounceForward)
- *Missile*: object that moves in one direction with a certain speed and an original orientation. They can change direction (see interactions *turnAround* and *reverseDirection*
- Bomber: the combination of a missile and a spawner (with their combined parameters)
- Chaser: object that moves in the direction of the nearest target object stype, with a certain speed
- RandomNPC: object that moves in a random direction with a certain speed
- *Portal*: object that can teleport another object contacting it to another exit object (see interaction *teleportTo*).

All moving objects move every cooldown environment steps.

Interactions describe what happens when two objects contact. Here are the possible interaction types and their parameters:

- noInteraction: nothing happens
- killSprite: the second object kills the first object
- transformTo: the second object transforms the first object into a third object stype
- removeResource: the second object decreases the count of resource from the first object
- addResource: the second object increases the count of resource of the first object
- killfHasLess: the second object kills the first if it has less than 1 resource stype
- stepBack: objects step back (second steps back first, if not possible the second does)
- bounceForward: the second object pushes the first if possible (e.g., unless it is blocked)

- turnAround: the first object (a missile) does one block down and switches direction when encountering the second object
- reverseDirection: the first object (a missile) reverses direction when encountering the second object

killSprite and transformTo interactions can further lead to a positive (+1), negative (-1) or null (0) reward.

Lose conditions can be:

- Timeout: the player loses if it runs out of time before solving the task
- *CountIsZero*(objs): the player loses if at least one of the objects objs has no remaining instances in the game.

Win conditions can be:

- Survive: the player wins if it survives long enough
- CountIsZero(objs): the player wins if the count of all objects objs goes to zero (e.g., they were all killed or disappeared)

Perception and action. Players have five possible actions: four directional moves (arrow keys) and a shoot action (space bar) when the game includes a *ShootAvatar* or a *FlakAvatar*. Computational agents also have access to a NOOP action, while human players can simply choose not to act. After taking a non-NOOP action, computational agents must wait for 4 environment steps before being able to take a new one, a way of capturing human reactive time (5 fps). The game runs at 20 frames per second. Humans perceive the game visually, with the current reward displayed at the top (see Figure 2), whereas computational agents perceive symbolic states consisting of object properties (object id, color, position), rewards, and win/loss events.

D INFERENCE PSEUDO-CODE

Algorithm 1 shows the pseudo-code of the inference algorithm—a particle filter with Metropolis rejuvenations—used to update the agent's beliefs about possible games it might be playing, supported by experiential and linguistic evidence (Del Moral et al., 2006; Metropolis et al., 1953). The *Perturb* operator replaces one rule of the VGDL game description at a time, according to guided proposals (see Appendix Section E below).

Algorithm 1 Particle filter with Metropolis rejuvenation

Params: N number of particles/theories, p(T) prior over theories, S, R number of steps and rejuvenation steps.

Inputs: inference step s, new experiential data D_s , particles θ_s (for s > 0).

```
\triangleright bold indicates vectors of size N
if s == 0 then

    Sample initial particles from prior

    T_0 \sim p(T)
for s \in [1, S] do

    inference step

    T_s \leftarrow T_{	ext{s-1}}
    for r \in [1, R] do
                                                                                            T_s' \leftarrow \operatorname{Perturb}(T_s)
                                                                                                          ▶ Propose move
        lpha \leftarrow rac{p(D_s, T_s')}{p(D_s, T_s)}
                                                                                                for i \in [1, N] if u \sim U(0, 1) < \alpha^i do
             T_s^i \leftarrow T_s^{'i}
                                                                                                           w \leftarrow \text{Norm}(p(D_s, T_s))

    □ Update weights

    T_s \leftarrow \text{Resample}(T_s, w)

⊳ SMC resampling step

Return: arg max_{T_{t+S}^i \in T_{t+S}} p(D_s, T_{t+S}^i)
```

E GUIDED PROPOSALS

We use biased proposals to initialize the set of 20 candidate theories and to generate rejuvenation moves (*Perturb* operator in Algorithm 1). These proposals are guided by both experience and linguistic evidence, accelerating convergence towards theories that better explain the agent's observations.

Experience-driven proposals bias the sampling process in the following ways:

- objects observed to have moved cannot be assigned an object type incompatible with movements (e.g., *Immovable*, or *Flicker*),
- objects moving in one direction are more likely to be assigned type object types that move linearly (*Missile* and *Bomber* than object types that allow movements in all directions (*RandomNPC*,, *Chaser*).
- objects pairs involved in collisions preceding observed rewards are more likely to be assigned reward-generating interactions.

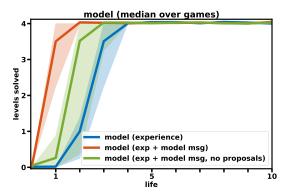


Figure 8: **Ablation of language-guided proposals.** a) Median performance across games for models learning from experience alone (blue), or experience and language, using language likelihood and language-guided proposals (orange), or ablating language-guided proposals (green). (N=20, error bar: interquartile range).

F DETAILS ABOUT PLANNING

Our model strategically balances curiosity-driven exploration with goal-directed exploitation, mirroring human problem-solving strategies (Tsividis et al., 2021). From the current best theory $T_{\rm MAP}$, we generate candidate subgoals: specific collisions between pairs of objects that the agent can cause to occur by either touching another object itself, or by pushing or spawning an object onto it. Candidate subgoals are assigned values based on their exploration and exploitation potentials:

$$Value(g) = ExplorationValue(g) + ExploitationValue(g).$$

The exploration value measures disagreement between theories in the current population:

ExplorationValue
$$(g) = 1 - \frac{\max_{i} \operatorname{count}(i \mid g)}{\sum_{i} \operatorname{count}(i \mid g)}$$

where count (i | g) is the count of theories assigning interaction type i to subgoal g.

The exploitation value rewards key game mechanics:

$$\text{ExploitationValue}(g) = \begin{cases} 10 & \text{if } g \text{ contributes to win condition} \\ 8 & \text{if } g \text{ protects essential resources} \\ 6 & \text{if } g \text{ creates useful tools} \\ 2 & \text{if } g \text{ collects resources or rewards} \\ 0 & \text{else} \end{cases}$$

To achieve these goals, the model evolves short 10-step action sequences (e.g. move left, shoot, move up) through stochastic search, using $T_{\rm MAP}$ to simulate outcomes. Action plans are iteratively mutated and refined over three generations using a simple genetic algorithm, where mutations crop and regrow action sequences from a uniformly sampled mid-point. Each sequence a is evaluated according to:

$$V(a) = R_{\text{game}}(a) + R_{\text{goal}}(a) + R_{\text{win/loss}}(a),$$

where R_{game} is the game's reward function under T_{MAP} , R_{goal} rewards progress toward the selected subgoal (between 0 and 1), and $R_{\text{win/loss}}$ provides +100 for winning and -100 for losing.

Additionally, the model performs 10 3-step lookaheads to avoid catastrophic errors, triggering replanning when the originally predicted value deviates significantly from the distribution of values for newly simulated trajectories:

$$V(a_{\text{original}}) > \max_{i \in 1...10} V(a_{\text{lookahead}_i}).$$

This safety mechanism prevents the agent from executing plans that appeared promising under limited simulation but fail under more extensive testing, or when unexpected changes in the environment render the original plan ineffective.

G INSTRUCTIONS FOR HUMAN DATA COLLECTION

Before playing the games, participants read the following instructions:

Instructions for human data collection

In this experiment, you will play different games before describing each game to another participant who has not yet played. You will play 5 different games. Each game will be played using the arrow keys and spacebar on your keyboard. You will start each game with 15 lives. Each game has 4 levels. Your goal is to win all 4 levels of the game in as few lives as possible.

You lose a life when you lose a level. There are different ways of winning and losing. Possible ways to win a level include removing all objects of particular colors, reaching a particular object, or surviving for long enough.

Possible ways to lose a level include allowing all objects of particular colors to die or disappear, or by not solving the level fast enough.

You move on to the next level of a game only after winning the current level. You will finish a game when you win all 4 levels, or when you have lost $15 \ \mathrm{lives}$.

After finishing a game, you will be asked to describe the game to another participant who has not yet played. Your goal is to help them solve the game in as few lives as possible. You will be required to spend at least 30 seconds describing each game.

\textbf{Bonus opportunity: Solve each game in as few lives as
possible!}

On each game, you can win up to \$\\$0.20\$ based on your performance. Your rank on each game will be computed relative to other participants, based on the number of levels you complete, and how many lives you take to complete them. For each game you will earn the max bonus multiplied by the proportion of other participants you outrank!

\textbf{Bonus opportunity: Write useful descriptions!}
On each game, you can win up to an additional \$\\$0.20\$ based on your
description. Another group of participants will read descriptions

before playing each game, and receive bonuses up to \$\\$0.20\$ based on their relative rank. On each game, you will additionally receive the bonus earned by the player who reads your description, so help them win!

// [Optionally, for participants in condition 2 and 3 (social learning)]

When starting each game, you will be shown a message from another player who has already played the game. You should use this message to help you play the game, but keep in mind that it's possible that it contains mistakes.

When starting each game, participants in Condition 1 received no game specific advice, as shown in Figure 9a, while participants in Conditions 2 and 3 saw a message from a previous player, as shown in Figure 9b. Players in Conditions 2 and 3 can still read the message as they play. After completing each game, participants in all conditions were asked to describe the game to a future player, as shown in Figure 9c. In the generational learning study, players could see the message from the previous player as they are asked to write their own, but could not copy-paste it. This was aimed at getting participants to engage and write a message of their own, as we found some participants to simply copy and paste messages containing false information without updating them in our pilot study.

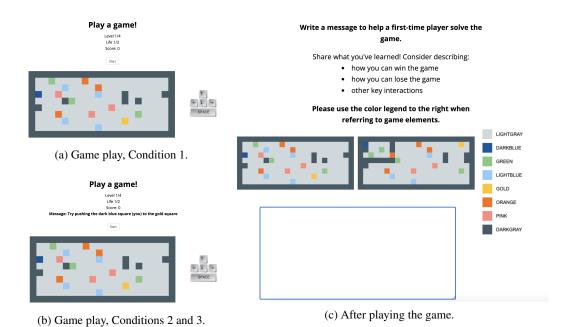


Figure 9: Screenshots of the game screens for the different experimental conditions. (a) Condition 1: individual gameplay instructions. (b) Conditions 2 and 3: social gameplay instructions. (c) Message interface shown to participants for writing advice to future players. Note that players in the social conditions can see the message during gameplay.

H COMPUTATIONAL RESOURCES

The model requires a GPU to run the LM, we use one NVIDIA A100 (80Gb). We use prompt caching using the vLLM library to speed up the generation of proposals. Simulation runs vary in function of game complexity (slower with more objects) and the learning speed of the agent (runs end early when all levels are solved), taking anywhere between an hour and a day, depending on games and conditions.

I EXAMPLE MESSAGES

Below are example messages generated by humans and model players in the social learning experiments.

Example of human messages in social learning experiment

[beesAndBirds] "Move the darkblue square by pressing the arrows keys. You win the game the touching the green square, you can also touch red but not needed. Avoid touching the moving yellow square because they can kill you. The green squares can destroy the yellow and orange squares when you free it from the purple squares. You can destroy purple squares by touching it."

[aliens] "You must use space bar to shoot and hit the pink moving squares. You are also needed to dodge the red squares that are being shot by the pink squares. You will dodge by using the left, right, up and down keys. You must destroy all the pink squares before they reach the bottom."

[jaws] "I think if you move the block over the purple and reds you will win the game. You will lose the game if you go over other color blocks." $\,$

[preconditions] "eat the white tiles to enable you knock down the green tiles toget to the yellow tiles. your objective is to get to the yellow tiles"

[missile_command] "So the main goal is to protect the GREEN tiles by making sure the RED tiles and the YELLOW tiles don't touch them and you can win even if you only have ONE GREEN tile left. As soon as you click START, they will start moving and you need to act fast by trying to interact with the RED and YELLOW tiles to make them disappear using the SPACEBAR while next to them. The YELLOW tiles move faster than the red so try to prioritize them if possible but there will be times that you would need to sacrifice one GREEN tile to make one YELLOW tile disappear because at that same time, you can possibly eliminate 2-3 RED tiles because they are slower. You win the level by keeping at least one GREEN tile intact and you lose when you lose all five of them."

[pushBoulders] "Move darkblue square by using the arrow keys. Avoid touching the orange squares because they can kill you, push the green into orange to make orange disappear. Avoid pink square because it can kill, move pink by pushing green into it. You can touch lightblue square to make it disappear. Also you can touch yellow square to make it disappear and win the game."

[avoidGeorge] "Game 4: Move your dark blue square quickly to avoid the light blue square. Use the brown spacebar to turn the green squares back to purple. Shoot the squares next to you with the brown square to eliminate them."

[plaqueAttack] "Use your arrows to shoot from your square to protect the red and yellow squares. If the other colors get to red and yellow and turn them before you can eliminate them, you lose."

[relational] "Use the arrow keys to move the orange cubes to merge the darkblue and white cubes by pressing the space bar key. Avoid colliding with the rest of the cubes in the game"

[portals] "You need to move your character over the gold and light blue squares. To win the game you need to get to the light green square. In the levels coming up you needs to move up from light blue to dark blue to get to different areas. Do not let the other colors touch you because you will lose a life."

Example of model messages in social learning experiment

[beesAndBirds] "Use arrow keys to move your darkblue. To win, eliminate all green by touching them, earning points for each kill. Be careful: you'll die if you touch yellow or orange, but you can use lightgreen to your advantage — it can kill both yellow and orange. Touch red and purple to kill them, but they're not your main objective."

[aliens] "Use left and right arrow keys to move your darkblue square and spacebar to shoot blue squares. To win, eliminate all orange, pink, and brown - shooting blue at pink will kill them and earn points, but watch out as pink can kill you on contact. Protect yourself from pink and red, which can also kill you, and avoid touching them. Blue objects can help by killing pink, but be aware that blue and white objects kill each other, and pink can generate more red objects."

[jaws] "Control your darkblue square with arrow keys. Your goal is to survive for as long as possible. Be careful, as orange, brown, and pink objects will kill you on contact - orange objects will even chase you. Stay away from them and try to navigate around the safe darkgray, purple, and green objects to stay alive."

[preconditions] "Use arrow keys to move your darkblue square. Your goal is to eliminate all gold objects by touching them to earn points. Be careful, as green objects will kill you on contact unless you have a white resource to protect yourself - collect white resources to stay safe. You can also kill green objects by touching them, but they'll take a resource from you in the process."

[missile_command] "Use arrow keys to move the darkblue square and spacebar to shoot pink. To win, eliminate all red and gold - shoot pink at them to kill and earn points. Protect green objects, as you'll lose if they all die, and try to avoid letting red and gold kill them as this will also cost you points. Red and gold objects will chase green, so use this to your advantage and take them out with pink squares."

[pushBoulders] "Use arrow keys to move your darkblue square. To win, eliminate all gold objects by touching them or pushing green onto them - you'll earn points for each gold destroyed. Be careful: touching orange or pink will kill you. Green is a valuable ally, as it can also kill gold and earn you points, so use it to your advantage by pushing it onto gold. Protect yourself and keep pushing green to clear the way to victory, as dying means losing the game."

[avoidGeorge] "Use arrow keys to control your darkblue square and spacebar to shoot brown. To win, survive long enough while keeping green objects alive - they can be created by shooting brown at purple or letting orange touch purple. Watch out: lightblue objects will kill you on contact and can also turn green back into purple, so protect your green! You'll earn points when brown or orange transform purple into green, but lose points if lightblue kills you."

[plaqueAttack] "Use arrow keys to move your darkblue and spacebar to shoot lightred. Your goal is to eliminate all purple, orange, brown, and lightblue while keeping yellow and red alive. To help you, lightred objects will kill orange and brown on contact, so position them to clash. You can also shoot lightred at orange and brown to kill them directly. Touching green objects transforms them into yellow, earning you points, and touching blue objects transforms them into red, also earning points. Be careful, as orange and brown can transform yellow into green, losing you points, and blue objects can"

[relational] "Use arrow keys to move your darkblue square. To win, eliminate all blue objects, which can be done by pushing them onto yellow objects - when they kill each other, you gain points. Be careful not to touch white objects, as they'll kill you. You can also transform objects to help your strategy: touch red to turn them into yellow, touch pink to turn them into orange, and push orange onto purple to turn them into yellow."

[portals] "Use arrow keys to move your darkblue square. Your goal is to eliminate all lightgreen objects by touching them to earn points and win the game. Be careful - darkgreen, pink, and lightred objects will kill you on contact, so avoid them. You can teleport using lightblue to gold and blue to lightorange to help navigate. Brown objects move randomly, but won't harm you, and other objects are stationary, so use them to your advantage."

J PROMPTS

Prompt for language likelihood estimation and language generation

Game Message Generation Task

You are writing standardized messages to describe a 2D video game based on the game's mechanics.

Instructions

Each message should explicitly state:

- How to control the game
- How to win
- How you could lose/die
- How to get/lose points
- Key game mechanics that aren't obvious while playing

Context

- The player controls the darkblue object
- In some games, it can shoot other objects of a particular color
- Each object color is associated to a particular object type conditioning how they move
- Effects can apply when two colors collide, eg one can kill the other, one can transform the other into a third color, one can push the other, etc $\,$
- Possible wining conditions include: surviving long enough, or killing all objects of one or several target colors
- Possible losing conditions include: reaching a timeout before wining, or seeing all objects in a list of colors dying/disappearing

Writing Rules

- Mention all important interactions in the game
- State each interaction in a single clear sentence
- List all killing and transformation interactions

```
1188
         - Always specify win and loss conditions
1189
1190
         ### Example
1191
        Who you are and how you move:
1192
        - You control the darkblue square with arrow keys.
1193
        How you win and lose:
1194
         - You get points when pink objects transform blue objects into purple
1195
         objects.
1196
         - You lose points when darkred objects kill purple objects.
         - You win the game when all blue and orange objects are dead.
1197
          You lose if all purple and lightred objects die.
1198
        What you can do:
1199
         - You can shoot pink squares by pressing space bar.
1200
        - You can kill orange objects by shooting pink objects at them.
1201
        - You can transform blue objects into purple objects by shooting pink
1202
         objects at them.
        What can kill you:
1203
        - Orange objects will kill you if you touch them.
1204
        Other possible interactions:
1205
         - You can push lightred objects.
1206
        - Darkred objects can kill purple objects by touching them.
        - Darkred objects can kill lightred objects by touching them.
1207
        Other objects:
1208
        - Pink objects move along one direction.
1209
        - Darkred objects chase the nearest purple object.
1210
        - Orange objects move along one direction.
1211
        - Purple and blue objects do not move.
1212
        Message:
1213
        You control a darkblue square with arrow keys. You can shoot pink
1214
        objects by pressing space bar. Orange objects kill you when they
1215
        touch you. Pink objects kill orange objects when they touch them.
1216
        Pink objects transform blue objects into purple objects when they
1217
        touch them. You can push lightred objects. Darkred objects kill
        purple and lightred objects when they touch them. Pink objects move
1218
        in one direction. Darkred objects chase the nearest purple object.
1219
        Orange objects move in one direction. Purple and blue objects don't
1220
        move. You get points when pink objects transform blue objects into
1221
        purple objects. You lose points when darkred objects kill purple
1222
        objects. You win when all blue and orange objects are dead. You lose
        if all purple and lightred objects die.
1223
1224
         ### Task
1225
1226
         {THEORY DESCRIPTION}
1227
        Message:
1228
         {MESSAGE} # we compute the likelihood of these tokens
1229
```

Prompt for making rules proposals

Game Rules Analysis Task

You are analyzing messages written by video game players to help you play a game for the first time. These games feature colored objects interacting in a 2D space, where players control a darkblue object. The messages describe controls, scoring systems, win/loss conditions, hazards, core mechanics, and give you tips to help you win the game.

Your Task

1230 1231

123212331234

1235 1236

1237

1238

1239

```
1242
        You will be asked to answer four types of questions:
1243
        - What is the type of objects of [color1]?
1244
        - What happens to objects of [color1] when they collide with objects
1245
        of [color2]?
        - Does the player need to kill all objects of [color1] to win?
1246
        - Does the player lose when all objects of [color1] die?
1247
1248
        Answer questions by looking for explicit statements in the message
1249
1250
        - How objects move and how they behave, eg can you use them as a
        portal, can you push them, do they chase other objects, etc.
1251
         Objects killing other objects
1252
        - Objects transforming into other objects
1253
        - Win conditions involving killing objects
1254
        - Loss conditions involving object death
1255
        ### Rules for Answering
1256
        - Only use information explicitly stated in the message
1257
        - You are the darkblue object.
1258
        - Do not make assumptions about interactions not mentioned
1259
          Pick the "I don't know" answer if the corresponding information is
1260
        ambiguous, contradicting, or absent
        - Pay attention to interaction direction - if A kills B, it doesn't
1261
        mean B kills A
1262
        - Each question comes with a fixed set of possible answer. You must
1263
        pick one of the answer. If the correct answer is not in the list,
1264
        pick the "I don't know" answer.
1265
         - If the way to win is to kill some object, then having all these
        objects disappear/die does NOT cause you to lose
1266
         - If you need to "protect [color1] objects" or "kill [color2] objects
         before they kill / transform [color1] objects", you must answer:
1268
          - "You lose if all [color1] objects die or disappear", note that
1269
          transformations kill the transformed objects.
1270
        - If you need to reach / touch / kill by touching [color1] objects to
1271
         win, this means you can kill color1 objects. You MUST answer:
          - "[color1] objects die when they collide with darkblue objects".
1272
        - If the message says you can "collect [color1]":
1273
          - either [color1] is a resource, and you must answer that "you can
1274
          collect [color1] resources"
1275
          - OR [color1] is not a resource, and you must answer that "you can
1276
          kill [color1] objects"
        - If the message says you can destroy/kill/eliminate [color1] objects
1277
         if you have enough [color2] resources, the two following statements
1278
        are true:
1279
          - "[color1] objects die when they collide with darkblue objects but
1280
          also take a [color2] resource from them'
          - AND "darkblue objects die if they don't have enough resources when
1281
           they collide with [color1] objects"
1282
        - Objects that kill you when you lack resources also die when they
1283
        touch you, and take one resource from you.
1284
          - "[color1] objects die when they collide with darkblue objects but
1285
          also take a [color2] resource from them"
        - When [color1] objects take away some resource / your health when
1286
        you touch them, they also die:
1287
          - "[color1] objects die when they collide with darkblue objects but
1288
          also take a [color2] resource from them'
1289
        - Use the ... after "Reasoning:" to reason about the game before
1290
        generating your answer.
1291
1292
        ### Example Game 1
1293
1294
        Message from the player:
```

```
1296
        "Orange objects kill you when they touch you. You kill purple objects
1297
         when you touch them. Lightgreen objects chase orange objects and can
1298
         kill them when they touch them. You win when you kill all purple
1299
        objects. You lose if orange kills you."
1300
1301
1302
        What happens to darkblue objects when they collide with orange
1303
        objects?
1304
        1) nothing happens to darkblue objects when they collide with orange
        objects
1305
        2) darkblue objects die when they collide with orange objects
1306
        3) darkblue objects get transformed when they collide with orange
1307
1308
        4) darkblue objects steal resource from orange objects when both
1309
        collide
        5) darkblue objects die if they don't have enough resources when they
1310
         collide with orange objects
1311
        6) I don't know / something else
1312
1313
        Reasoning: .....
1314
1315
        Answer (pick one from the list):
        2) darkblue objects die when they collide with orange objects
1316
1317
1318
1319
        What happens to orange objects when they collide with darkblue
        objects?
1320
        1) nothing happens to orange objects when they collide with darkblue
        objects
1322
        2) orange objects die when they collide with darkblue objects
1323
        3) orange objects get transformed when they collide with darkblue
1324
        objects
1325
        4) orange objects steal resource from darkblue objects when both
        collide
1326
        5) orange objects die if they don't have enough resources when they
1327
        collide with darkblue objects
1328
        6) I don't know / something else
1329
1330
        Reasoning: .....
1331
        Answer (pick one from the list):
1332
        6) I don't know / something else
1333
1334
        Note: orange kill darkblue (the player), but we don't know what
        happens to orange when they touch (maybe nothing)
1335
1336
1337
1338
        How do orange objects behave?
1339
        1) orange objects cannot move
        2) orange objects do not move and disappear after a certain time
1340
        3) orange objects regularly spawn/generate other objects
1341
        4) orange objects can be pushed
1342
        5) orange objects move along one axis
1343
        6) orange objects move along one axis and regularly spawn/generate
1344
        objects
        7) orange objects chase or flee another object
1345
        8) orange objects move randomly
1346
        9) orange objects are portals
1347
        10) I don't know
1348
1349
        Reasoning: .....
```

```
1350
1351
        Answer (pick one from the list):
1352
        10) I don't know
1353
1354
1355
        How do lightgreen objects behave?
1356
        1) lightgreen objects cannot move
1357
        2) lightgreen objects do not move and disappear after a certain time
1358
        3) lightgreen objects regularly spawn/generate other objects
        4) lightgreen objects can be pushed
1359
        5) lightgreen objects move along one axis
1360
        6) lightgreen objects move along one axis and regularly spawn/
1361
        generate objects
1362
        7) lightgreen objects chase or flee another object
1363
        8) lightgreen objects move randomly
1364
        9) lightgreen objects are portals
        10) I don't know
1365
1366
        Reasoning: .....
1367
1368
        Answer (pick one from the list):
1369
        7) lightgreen objects chase or flee another object
1370
1371
1372
        Do you need to kill purple objects to win?
1373
        1) To win you need to reach/touch purple objects
        2) To win you need to kill all purple objects
1374
        3) To win you don't need to kill all purple objects
        4) To win I don't know if you need to kill or reach/touch all purple
1376
        objects
1377
1378
        Reasoning: .....
1379
        Answer (pick one from the list):
1380
        2) To win you need to kill all purple objects
1381
1382
1383
1384
        Do you need to kill orange objects to win?
        1) To win you need to reach/touch orange objects
1385
        2) To win you need to kill all orange objects
1386
        3) To win you don't need to kill all orange objects
1387
        4) To win I don't know if you need to kill or reach/touch all orange
1388
        objects
1389
        Reasoning: .....
1390
1391
        Answer (pick one from the list):
1392
        4) To win I don't know if you need to kill or reach/touch all orange
1393
        objects
1394
1395
1396
        Do you lose if orange objects die?
1397
        1) You lose if all orange objects die or disappear
1398
        2) You don't lose if all orange objects die or disappear
        3) I don't know if you would lose if all orange objects die or
1399
        disappear
1400
1401
        Reasoning: .....
1402
1403
        Answer (pick one from the list):
```

```
1404
        2) You don't lose if all orange objects die or disappear
1405
1406
1407
        ### Example Game 2
1408
        Message from the player:
1409
        "You can shoot brown objects. Brown objects kill gold objects when
1410
        they touch them. Gold objects transform purple objects into lightred
1411
        objects when they touch them. You transform lightred objects into
1412
        purple objects when you touch them. You lose if all purple objects
        are transformed into lightred objects. You win if you survive long
1413
        enough."
1415
1416
1417
        What happens to gold objects when they collide with brown objects?
        1) nothing happens to gold objects when they collide with brown
1418
        objects
1419
        2) gold objects die when they collide with brown objects
1420
        3) gold objects get transformed when they collide with brown objects
1421
        4) gold objects steal resource from brown objects when both collide
1422
        5) gold objects die if they don't have enough resources when they
1423
        collide with brown objects
        6) I don't know / something else
1424
1425
        Reasoning: .....
1426
        Answer (pick one from the list):
        2) gold objects die when they collide with brown objects
1428
1429
1430
1431
        What happens to purple objects when they collide with gold objects?
1432
        1) nothing happens to purple objects when they collide with gold
1433
        objects
        2) purple objects die when they collide with gold objects
1434
        3) purple objects get transformed when they collide with gold objects
1435
        4) purple objects steal resource from gold objects when both collide
1436
        5) purple objects die if they don't have enough resources when they
1437
        collide with gold objects
1438
        6) I don't know / something else
1439
        Reasoning: .....
1440
1441
        Answer (pick one from the list):
1442
        3) purple objects get transformed when they collide with gold objects
1443
1445
        What happens to lightred objects when they collide with gold objects?
1446
        1) nothing happens to lightred objects when they collide with gold
1447
        objects
        2) lightred objects die when they collide with gold objects
1448
        3) lightred objects get transformed when they collide with gold
1449
        objects
1450
        4) lightred objects steal resource from gold objects when both
1451
1452
        5) lightred objects die if they don't have enough resources when they
1453
         collide with gold objects
        6) I don't know / something else
1454
1455
        Reasoning: .....
1456
1457
        Answer (pick one from the list):
```

```
1458
        6) I don't know / something else
1459
1460
1461
        How do lightred objects behave?
1462
        1) lightred objects cannot move
1463
        2) lightred objects do not move and disappear after a certain time
1464
        3) lightred objects regularly spawn/generate other objects
1465
        4) lightred objects can be pushed
1466
        5) lightred objects move along one axis
        6) lightred objects move along one axis and regularly spawn/generate
1467
        objects
1468
        7) lightred objects chase or flee another object
1469
        8) lightred objects move randomly
1470
        9) lightred objects are portals
1471
        10) I don't know
1472
        Reasoning: .....
1473
1474
        Answer (pick one from the list):
1475
        10) I don't know
1476
1477
1478
        Do you need to kill brown objects to win?
1479
        1) To win you need to reach/touch brown objects
1480
        2) To win you need to kill all brown objects
1481
        3) To win you don't need to kill all brown objects
1482
        4) To win I don't know if you need to kill or reach/touch all brown
        objects
1483
1484
        Reasoning: .....
1485
1486
        Answer (pick one from the list):
1487
        3) To win you don't need to kill all brown objects
1488
1489
1490
        Do you lose if purple objects die?
1491
        1) You lose if all purple objects die or disappear
        2) You don't lose if all purple objects die or disappear
1492
        3) I don't know if you would lose if all purple objects die or
1493
        disappear
1494
1495
        Reasoning: .....
1496
        Answer (pick one from the list):
1497
        1) You lose if all purple objects die or disappear
1498
1499
1501
        Do you lose if lightred objects die?
        1) You lose if all lightred objects die or disappear
1502
        2) You don't lose if all lightred objects die or disappear
1503
        3) I don't know if you would lose if all lightred objects die or
1504
        disappear
1505
1506
        Reasoning: .....
1507
        Answer (pick one from the list):
1508
        2) You don't lose if all lightred objects die or disappear
1509
1510
1511
        ### Task
```

```
1512
1513
        Message from the player:
1514
        {MESSAGE}
1515
        [QUESTION]
1516
        [CHOICES]
1517
1518
        Reasoning: .....
1519
1520
        Answer (pick one from the list):
        [CANDIDATE ANSWER] # we compute the likelihood of all possible
1521
        candidate answer tokens
1522
1523
```

Prompt for the pure-LLM baseline

```
{basicstyle=\ttfamily\small}
You are trying to solve a video game you have never player before.
### Information about the game:
```

- The game contains colored square objects of size 1
- It is defined by:

1524 1525

1526 1527

1528

1529 1530

1531

1532

1533

1534

1535

1536

1537

1538

1539

1540

1541

1542

1543

1544

1545

1546

1547

1548

1549

1550

1551

1552

1553

1555

1556

1557

1558

1559

1560

1561

1562

1563 1564

1565

- one object type for each object color (eg a missile, a random NPC, an immovable object, a pushable object, a portal, etc.)
- interactions for each possible pair of object color (eg no interaction, one dies, one is transformed into the other, one pushes the other, one steps back, etc.). These define what happens when objects of these two colors collide.
- losing conditions: you may lose if you don't win fast enough, or if all objects of a specific color die (eg you lose if either all red die or all blue die)
- Each object color corresponds to a different type of object, but all objects of one color have the same type and behave the same way
- Each pair of object colors can have at most 1 type of interaction
- Some objects (color1) can transform certain objects (color2) into yet other objects (color3).
- Possible ways to act in the game are to touch objects, push objects or shoot. Yourself, the objects you push, or shoot can interact with other objects. Sometimes they will push them, kill them, transform them. Sometimes nothing happens.

State description:

- ${\hspace{0.25cm}\text{-}}$ You will be shown your recent actions and the corresponding history of game states, as well as object movements between these states
- A game state is composed of lists of object positions grouped by object color
 - Positions are indicated as (x, y) where x is between x=0 (LEFT) and x=max_x (RIGHT), and y is between y=0 (TOP) and y=max_y (BOTTOM).
 - When x increases, objects move to the right, when x decreases, objects move to the left. When y increases, objects move downwards, when y decreases, objects move upwards. Moving UP decreases y, movement = (0, -1).
 - Objects have size 1. They collide when delta_x or delta_y is below 1.
 - Look for wall positions, you cannot move through walls and shouldn't try to move in the direction of a wall.

Analysis:

- Observe the recent history of states and actions in the environment

```
1566
        - Observe your scratchpad where you store your analysis of the game
1567
        and strategy to solve it
1568
        - If your previous action was unsuccessful, try to understand why.
1569
        - If you won or lost in the previous episode, try to understand why.
        - Be careful about picking the right action to go where you want to
1570
1571
        - Try to predict the movements of important moving objects so you can
1572
         either catch or avoid them with your next action.
1573
        - Analyze your scratchpad, should you rephrase it to be more compact
1574
        and up to date?
        - Analyze both to update your understanding of the game, your
1575
        strategy, and prepare for action.
1576
        - Don't forget to plan for your next action. You might want to
1577
        explore the world and touch, push or shoot objects, or you might want
1578
         to plan towards the goal.
1579
        ### Scratchpad update:
1580
        - Decide whether to append information to the scratchpad or to
1581
        replace the text in the scratchpad by new text.
1582
        - The scratchpad should contain everything you know about the
1583
        environment: possible interactions, what kills you, what you think
1584
        the goal of the game is, what makes you lose, etc.
1585
        - The scratchpad should contain your current plan to explore or solve
         the game.
1586
        - The scratchpad is the only memory that will be transmitted to the
1587
        next decision making step, write down everything you need.
1588
        - To append text, use <append> text to append goes here </append>
1589
        - To replace text, use <replace> new text goes here </replace>
        - When replacing the scratchpad, all replaced text will be forgotten,
1590
         be careful.
        - No need to write down observations in the scratchpad, you will be
1592
        provided with the last observations every time you are asked to
1593
        generate a new action.
1594
        - Write your goal in the scratchpad so you can keep it in mind at the
1595
         next step.
        - The scratchpad has a maximum size of 400 words, you should
1596
        sometimes use the <replace> command to summarize its content and keep
1597
         it concise.
1598
        - You start each new episode with the scratchpad from the last
1599
        episode. It will contain useful information about the game you're
1600
        playing (eg how you can die, win, or important game dynamics), but
        may also contain outdated information (eg object positions, or
1601
        previous strategy from the last episode).
1602
        - When reaching the end of the episode, rewrite the scratchpad to
1603
        transmit information about what you learn so you can do better in the
1604
         next episode.
1605
        ### Actions:
1606
        - Possible actions are: UP, DOWN, LEFT, RIGHT, SPACE_BAR and NOOP.
1607
        - UP decreases y (0, -1), DOWN increases y (0, +1), LEFT decreases x
        (-1,\ 0), RIGHT increases x (+1,\ 0), SPACE_BAR and NOOP do not move
1609
        you.
        - SPACE_BAR let's you shoot / spawn objects in some games (not all).
1610
        - Try SPACE_BAR sometimes to see if anything happens, and write to
1611
        the scratchpad what happens.
1612
        - Make sure your actions get you closer to your goal. It's easy to
1613
        pick the wrong action and get further away, pay attention!
```

Answer format:

1614

1615

1616 1617

1618

1619

- Analyze what you know about the game and the history of recent game states

- Every time you pick an action that is not NOOP, you will have to wait 4 environment steps before picking a new action. If you select

NOOP you can select another action at the next environment step.

```
1620
        - Update the scratchpad
1621
        - Select the action to execute next
1622
1623
        Your response should be formatted in this way:
1624
        % first the analysis
1625
        <analysis>
1626
        write here your reflection and analysis of available information,
1627
        what you understand about the game, your plan to update the
1628
        scratchpad and your strategy for picking the next action.
        </analysis>
1629
1630
        % then update the scratchpad with either
1631
         <append>
1632
        text to append to the scratchpad goes here
1633
        </append>
        % or
1634
        <replace>
1635
        text to replace the scratchpad with goes here
1636
        </replace>
1637
1638
        % finally pick an action among UP, DOWN, LEFT, RIGHT, SPACE_BAR, NOOP
1639
        <act>
        action goes here
1640
        </act>
1641
1642
         ### Example
1643
        #### Recent history
1644
         % game history appears here
1645
1646
         #### Old scratchpad
1647
        Orange objects can kill you so you should watch out. Shooting red
1648
        objects seem to kill them. I'm not sure what the goal is yet. I can
1649
        push blue objects but not purple ones.
        Strategy: I should try to touch the yellow object to see what happens,
1650
         maybe I need to reach it?
1651
1652
         #### Example answer
1653
1654
        <analysis>
        It seems orange just killed the blue object? I should add this to the
1655
         scratchpad. I'm still relatively far from the yellow object I should
1656
         keep moving in that direction.
1657
        My next move should be UP.
1658
        </analysis>
1659
         <append>
1660
        Orange kill blue when they touch.
1661
        </append>
1662
1663
        <act>
        IJΡ
1664
         </act>
1665
1666
```