



HAL
open science

Policy Learning with a Language Bottleneck

Megha Srivastava, Cédric Colas, Dorsa Sadigh, Jacob Andreas

► **To cite this version:**

Megha Srivastava, Cédric Colas, Dorsa Sadigh, Jacob Andreas. Policy Learning with a Language Bottleneck. 2024. hal-04810194

HAL Id: hal-04810194

<https://inria.hal.science/hal-04810194v1>

Preprint submitted on 28 Nov 2024

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

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



Distributed under a Creative Commons Attribution 4.0 International License

Policy Learning with a Language Bottleneck

Megha Srivastava[◇], Cédric Colas[‡], Dorsa Sadigh[◇] and Jacob Andreas[‡]

[◇] Stanford University
{megha, dorsa}@cs.stanford.edu

[‡]Massachusetts Institute of Technology
{ccolas, jda}@mit.edu

[†]Inria

Abstract

Modern AI systems such as self-driving cars and game-playing agents achieve superhuman performance, but often lack human-like features such as generalization, interpretability and human inter-operability. Inspired by the rich interactions between language and decision-making in humans, we introduce *Policy Learning with a Language Bottleneck* (PLLB), a framework enabling AI agents to generate linguistic rules that capture the strategies underlying their most rewarding behaviors. PLLB alternates between a *rule generation* step guided by language models, and an *update* step where agents learn new policies guided by rules. In a two-player communication game, a maze solving task, and two image reconstruction tasks, we show that PLLB agents are not only able to learn more interpretable and generalizable behaviors, but can also share the learned rules with human users, enabling more effective human-AI coordination.¹

1 Introduction

As AI systems play an increasingly central role in automation, there is a growing need for them to better model and emulate human behavior. Such systems must be interpretable, and both act and generalize in human-predictable ways so that users can effectively interact with them. Unfortunately, many of today’s AI systems do not meet these standards. Self-driving cars or game-playing agents may achieve super-human performance but lack interpretability (McIlroy-Young et al., 2020) and often act unpredictably, especially outside their training distribution (Wang et al., 2023).

Humans acquire most of their skills and knowledge from others, often using language—via instructions, advice, or explanations that improve

¹We provide source code for our experiments at <https://github.com/meghabyte/bottleneck>.

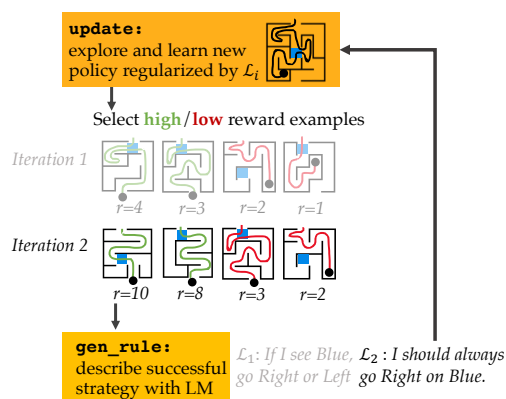


Figure 1: *Policy Learning with a Language Bottleneck* (PLLB) alternates between two steps: 1) **gen_rule** (bottom) generates a linguistic rule \mathcal{L}_i explaining the agent’s best behaviors by prompting a language model with contrastive (positive and negative) episodes (middle); 2) **update** (top) learns a new policy conditioned on that rule. PLLB facilitates stronger human-AI collaboration, while constraining policies to be more interpretable and generalize better.

their decision-making capabilities (Carruthers and Boucher, 1998; Mesoudi and Thornton, 2018). Language acts as a *communicative* medium, enabling us to teach, learn from, and coordinate with others to solve complex problems. It also supports other *cognitive* functions: even when not used for communication, it allows us to represent abstract concepts (Hesse, 1988; Lakoff and Johnson, 2008), and plan (Vygotsky, 1965; Clark, 1998); it guides our attention (Waxman, 1994; Yoshida and Smith, 2003), and prompts relational thinking (Gentner and Loewenstein, 2002).

Consider a driver learning to navigate novel social conventions (e.g., triangle-shaped stop signs). While adapting to the environment, they might verbalize strategies to themselves to avoid future mistakes (e.g., *If the sign is triangular, I should stop*, a cognitive use), or transmit this convention to others (e.g., telling a friend *In Japan, stop signs*

are triangles, a communicative use). Representing learned information in language helps humans solve problems and transmit knowledge by effectively capturing abstract problem structures that facilitate learning and generalization (Boutonnet and Lupyan, 2015; Chopra et al., 2019; Tessler et al., 2021).

The recent literature contains many examples of AI systems leveraging language-based representations or linguistic feedback, but these often rely on external supervision: humans in the loop or hard-coded feedback functions (see related work and reviews in Luketina et al., 2019; Colas et al., 2022). We argue that more human-like AI systems should not only *use* language-based supervision but also *generate their own language-based feedback* to leverage both the communicative and cognitive functions of language.

We introduce *Policy Learning with a Language Bottleneck* (PLLB), a framework providing artificial embodied agents the ability to generate linguistic rules that capture the strategies underlying their most rewarding behaviors. As shown in Figure 1, PLLB alternates between a **gen_rule** step that explains the agent’s most rewarding behaviors by prompting a language model (LM) with contrastive episodes, and an **update** step that learns a new policy conditioned on these rules. Unlike past work that solely leverages LMs for modeling agent behavior and multi-step reasoning (Park et al., 2023; Wei et al., 2023), PLLB is applicable even when aspects of the target policy cannot be expressed with language. The rules learned with PLLB can not only be shared with users to increase the agents’ interpretability and collaborative capabilities (communicative use), but also used internally by the agents to constrain their exploration, leading to more human-like and generalizable behaviors (cognitive use).

Policy Learning with a Language Bottleneck helps agents learn more human-like policies, which we examine across four diverse tasks. **They perform better:** in two image reconstruction tasks, PLLB agents generate instructions increasing the listeners’ performance compared to non-linguistic baselines (Section 7). **They generalize better:** in a maze task, agents uncover the abstract problem structure and generalize better to similar mazes (Section 6). **They are more interpretable:** in a coordination task, agents converge on the most human-interpretable policy when several optimal

ones exist (Section 5). **They are more interoperable:** in maze and image reconstruction tasks, humans achieve better rewards when interacting with PLLB agents than agents trained without a bottleneck (Sections 6 and 7).

2 Background & Related Work

PLLB is inspired by the coupling of language and decision-making in both humans (Carruthers and Boucher, 1998) and machines (Colas et al., 2022). **Language for communication.** Language facilitates cooperation and coordination between humans and machines via instructions (Hermann et al., 2017; Chevalier-Boisvert et al., 2019), advice (Watkins et al., 2021), explanations (Zhong et al., 2020; Lampinen et al., 2022), or the formation of conventions (Hawkins et al., 2020; Hu and Sadigh, 2023). Such communicative functions increase the fidelity and breadth of cultural transmission—a process of social learning that underlies human ecological success (Mesoudi and Thornton, 2018). PLLB agents not only learn from language, but also generate their own to be shared with others.

Language for cognition. Language also augments a learner’s cognitive abilities. Language-augmented RL agents represent more abstract goals (Jiang et al., 2019), generalize better (Hill et al., 2020; Colas et al., 2020; Wong et al., 2021), explore more efficiently (Colas et al., 2020; Tam et al., 2022; Klissarov et al., 2023) and can decompose complex goals into simpler ones (Chen et al., 2021; Ahn et al., 2022; Hu et al., 2022; Sharma et al., 2021; Hu and Clune, 2023). Our work extends these benefits to agents that learn from *self-generated* linguistic feedback.

Inner speech in humans. Generating linguistic rules for oneself is a form of *inner speech*. In the Vygotskian tradition, inner speech is seen as the internalization of the social speech generated by caretakers to help children solve problems (Vygotsky, 1965; Luria, 1959). As a result, it is thought to support our capacities for complex, long-term behaviors (Vygotsky, 1965; Luria, 1959; Hermer-Vazquez et al., 2001; Spelke, 2003).

“Inner speech” in machines. AI agents endowed with forms of inner speech (explanations, descriptions or subgoals) have been found to perform and generalize better than agents trained with purely neural representations (Wong et al., 2021; Lampinen et al., 2022; Roy et al., 2022; Hu and

Clune, 2023). Unlike these approaches, PLLB approach generates language in an unsupervised way and maximizes downstream performance as well as interpretability and inter-operability.

Multi-step reasoning. Recent work has proposed to guide LMs’ reasoning by prompting them to step through sequences of “thoughts” (Wei et al., 2023; Yao et al., 2022; Li et al., 2023b; Shinn et al., 2024). Similarly, our approach uses a language bottleneck to concisely express intermediate information useful for later behavior (here policy learning). Unlike LM reasoning approaches, our generated language is used both for increasing interpretability and inter-operability of initially uninterpretable AI systems (low level policies) and for enhancing the performance and generalization capacities of an RL agent.

3 The Language Bottleneck

PLLB builds on the standard RL framework to train agents to solve decision-making tasks in human-like ways. We formalize decision-making tasks (e.g., solving a maze) as Markov decision processes (S, A, f, R, T) with reward function $R : S \times A \rightarrow \mathbb{R}$ (e.g., solving speed) over states S (e.g., cell coordinates) and actions A (e.g., directions), finite time horizon T , and a deterministic transition function $f : S \times A \rightarrow S$ that maps state–action pairs (s, a) to next states s (e.g., moving between cells). Standard RL then involves training a policy π to maximize expected reward (e.g., by learning to solve the maze faster). This is usually done by alternating between two steps: (1) collecting data with the current policy: $D \leftarrow \pi_i$ and (2) updating the policy using the data: $\pi_{i+1} \leftarrow \text{update}(\pi_i, D)$, where `update` is an RL algorithm (Sutton and Barto, 2018). This procedure often fails to yield interpretable, inter-operable, or generalizable behaviors.

Our key idea is to introduce a *language bottleneck* between data collection and policy update. We extract linguistic rules explaining past rewarding behaviors and then use them to regularize the policy’s behavior in the next learning iteration (e.g., *I should go right in blue cells*, see Figure 1). Our experiments will show that this improves the interpretability, interoperability (communicative use), performance and generalization (cognitive use) of our agents over standard RL baselines in a variety of tasks (see Figure 2 and Sections 5 to 7). The resulting algorithm thus

alternates between three steps: (1) data collection, $D_i \leftarrow \pi_i$; (2) language bottleneck generation, $\mathcal{L}_i \leftarrow \text{gen_rule}(D_i)$ and (3) policy updating, $\pi_{i+1} \leftarrow \text{update}(\pi_i, D_i, \mathcal{L}_i)$, see Figure 1. In PLLB, the first step is implemented just as in ordinary policy learning approaches; below describe the last two steps in more detail.

3.1 Rule Generation (`gen_rule`)

Using all the experience D_i collected by the policy π_i in the current iteration, the `gen_rule` step aims at inferring an abstract rule \mathcal{L}_i that best explains the agents’ successful behaviors: $\mathcal{L}_i \leftarrow \text{gen_rule}(D)$. This is done by prompting an LM with contrastive episodes from D_i (top- N highest vs. top- N lowest total rewards) and asking it to *provide the rule that should be followed to obtain high rewards* (see first row of Figure 2 and Appendix Section A.3 for full prompts). We found this contrastive approach, inspired by Zhong et al. (2023) and Dunlap et al. (2023), to outperform simple summarization of high-reward samples.

3.2 Rule-Guided Policy Update (`update`)

Given a rule \mathcal{L}_i , the update step produces a new policy $\pi_{i+1} \leftarrow \text{update}(\pi_i, D_i, \mathcal{L}_i)$ that is better aligned with \mathcal{L}_i . One existing approach to solving this problem is the InstructRL approach of Hu and Sadigh (2023), which regularizes the learned policy with another policy induced by the linguistic rule $\pi_{\mathcal{L}}$. For instance, in the maze example shown in Figure 1, if the rule is *I should go right on every blue cell*, the induced policy $\pi_{\mathcal{L}}$ should assign probability 1 to the *right* action in blue cells, but equal probabilities to all actions in every other situations. In the Q-learning algorithm, this approach simply adds a *regularizing term* (orange) to the standard Q-learning update rule ²:

$$Q^\theta(s_t, a_t) \leftarrow r_{t+1} + \gamma Q^\theta(s_{t+1}, a_{t+1}), \text{ where} \\ a_{t+1} = \arg \max_a [Q(s_{t+1}, a) + \lambda \log \pi_{\mathcal{L}}(a | s_t)].$$

Here γ is a discount factor and λ controls the strength of the rule-induced regularization and could be made time-dependent with a pre-defined or learned schedule (e.g., stronger regularization at the beginning of learning).

But how do we induce $\pi_{\mathcal{L}}$ from \mathcal{L} ? Since \mathcal{L} is expressed in natural language, we emulate it with

²This regularizing term can also be applied in the function approximation setting (Mnih et al., 2013).

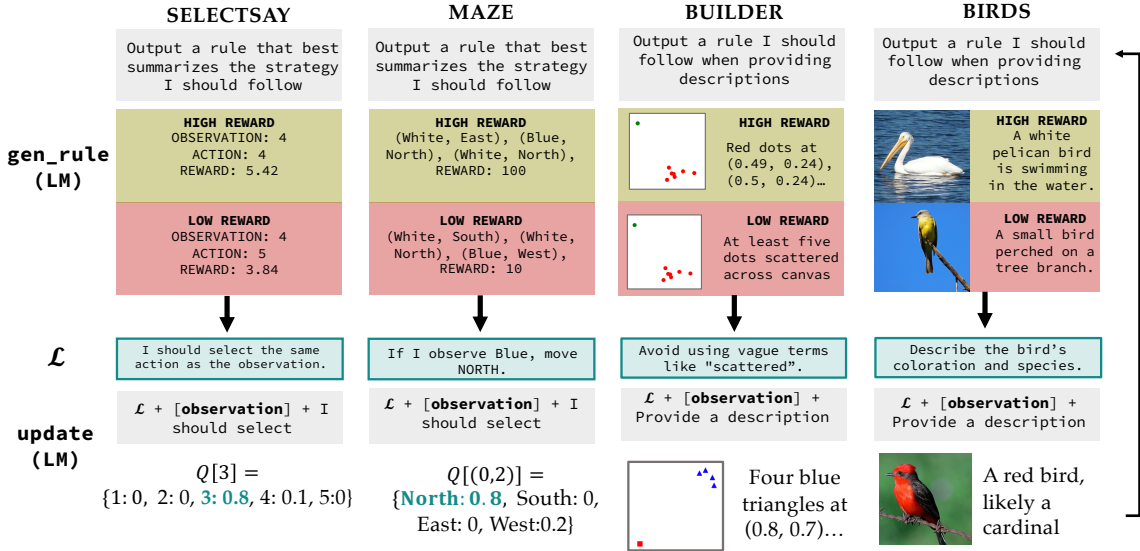


Figure 2: PLLB can be applied to a diversity of domains. In all domains, it iterates between `gen_rule`, a function that prompts an LM to extract a linguistic rule (\mathcal{L} , blue) by contrasting high and low reward episodes from past experience (green and red boxes), and `update`, a function for updating an agent’s policy through interaction with the environment conditioned on \mathcal{L} . PLLB can be applied to multi-step decision making tasks and visual tasks alike with minimal implementation variations (replacing the LM by a VLM in `gen_rule`, or replacing policy regularization by simple instruction conditioning in `update`). Exact prompts are in Appendix Section A.3.

an LM. We prompt it with both the rule and the current state of the agent s_t and instruct it to generate the next action to obtain a probability distribution over admissible actions (e.g directions). This can be done by normalizing the probability of generating each of the action tokens, or by counting the frequency of these tokens over multiple LM generations. Running Q-learning updates with the regularization term from experience data gives us a new Q-table Q_{i+1}^θ from which we can derive the new policy π_{i+1} by taking actions with maximum expected value in every state $\pi_{i+1}(s_t) = \max_a Q_{t+1}^\theta(s_t, a)$.

As an alternative, in certain domains, the policy can be directly implemented by an LM or VLM. In that case, simply conditioning the policy on the rule by adding it to the prompt $\pi_{\mathcal{L}}$ might be enough to steer the agent’s behavior to become more human-like: $\pi_{i+1} \leftarrow \pi_{\mathcal{L}_i}$. Given a new policy π_{i+1} , we can now close the loop and generate a new rule \mathcal{L}_{i+1} . We implement PLLB with policies from RL agent exploration (Sections 5 and 6) and LM/VLMs (Section 7) in our experiments.

4 Experiment Set-Up

We run experiments on a set of four diverse tasks to showcase the capacity of PLLB to train more human-like agents. Our domains include a sim-

pler two-player communication game called SELECTSAY (Section 5), a MAZE solving task (Section 6), and two collaborative image reconstruction tasks with either synthetic (BUILDER) or natural (BIRDS) images (Section 7). For all tasks, the exact prompts used in `gen_rule` and `update` and examples of rules generated across iterations can be found in Appendix A.3 and B.1 respectively.

To evaluate the *cognitive* function of PLLB, we first evaluate the policies learned by artificial agents in each environment based on performance metrics such as task reward, similarity to known human-like policies, or generalization. For all tasks, we compare our **Bottleneck** approach with baseline non-linguistic methods (e.g., standard RL, base LM) and methods leveraging invalid linguistic rules (e.g., obtained by labeling high reward examples as low reward) in `gen_rule` (**Adversarial**). These baselines allow us to study the impact of generating linguistic rules aligned with the underlying environment structure.

To evaluate the *communicative* function of PLLB, we conduct a variety of human subject studies showing that generated rules \mathcal{L} can help improve human task performance on a variety of metrics (e.g., reward, solving time). Human subject study participants were either recruited in-person (for MAZE) or on the Prolific crowdsourc-



Figure 3: Overview of SELECTSAY game, reproduced from Hu and Sadigh (2023).

ing platform³ (BUILDER, BIRDS), compensated at a rate of US\$12/hour. All studies were approved by our institution’s IRB.

5 SelectSay

As a proof-of-concept, we first consider SELECTSAY, a simple collaborative game introduced in Hu and Sadigh (2023). Here, a *speaker* can see a hidden set of five balls including two blue ones and must help a *listener* find the two blue balls by communicating with them only via numbers, see Figure 3. In principle, the speaker and listener could converge on any bijective convention between messages and balls, and this is what we find with RL agents (multiple optimal policies). Humans, on the other hand, would prefer the intuitive mapping $1 \rightarrow 1, 2 \rightarrow 2$, etc. that we call the *human-interpretable policy*. Hu and Sadigh (2023) showed that regularizing the listener’s policy with the instruction *I should select the same number as my partner*, successfully guided the two RL agents to the human-interpretable policy. We investigate whether PLLB could generate a similar rule and achieve a similar interpretability *on its own*.

Task Overview. Like in Hu and Sadigh (2023), the speaker and listener are RL agents trained with Q-Learning that receive positive rewards when the listener collects blue balls, negative rewards otherwise. After training both agents for 200 episodes, we prompt llama-2-70b-chat⁴ with the task description and a contrastive set of high- and low-reward episodes to generate a rule \mathcal{L} , e.g., *I should choose action 4 whenever the observation is 2, 3, 4, or 5*. (gen_rule, see Figure 2).

We generate new rules every 500 training episodes. In between, we regularize the listener’s policy with a another policy $\pi_{\mathcal{L}}$ induced from the current rule \mathcal{L} . This is done by applying a *softmax* on the LM’s logits (see method in Section 3) across all possible actions and obtaining a distribution. In the following experiments, we com-

³Prolific, <http://prolific.com>.

⁴<https://huggingface.co/meta-llama/Llama-2-70b-chat-hf>

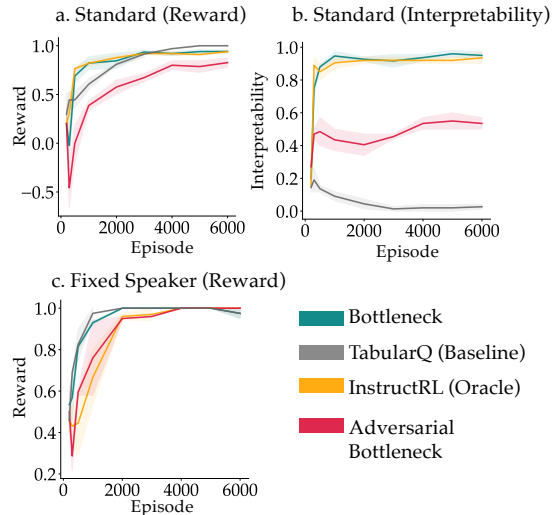


Figure 4: SELECTSAY. a) **Bottleneck** agents learn as fast as **TabularQ** and **InstructRL** and faster than agents using adversarial rules **Adversarial**. b) Unlike **TabularQ**, they learn human-interpretable policies; this without relying on an external instruction like **InstructRL**. c) When faced with speakers enforcing a non-human-interpretable policy, **Bottleneck** converges faster than **InstructRL**.

pare our **Bottleneck** method, with the **Adversarial** version (corrupted rule), a **TabularQ** baseline, as well as an **InstructRL** upper bound using the same predefined instruction as (Hu and Sadigh, 2023).

5.1 PLLB helps learn human-interpretable policies

We first compare the four methods along two dimensions: their reward and their *interpretability* measured as the proportional (0 to 1) similarity between the listener’s policy and the optimal human-interpretable policy (the intuitive message-ball described above). All methods but **Adversarial** quickly solve the task (Figure 4a), showing that corrupted rules hinder learning. **Bottleneck** and **InstructRL** both converge on the human-interpretable policy thanks to their linguistic rules while **TabularQ**—deprived of such rule—does not (Figure 4b). Interestingly, we note that **Adversarial** policies are more interpretable than **TabularQ** ones, suggesting that even corrupted language may provide inductive biases towards more human-like behavior. Qualitatively, we observe that generated rules converge towards describing the human-interpretable policy (e.g., *I should follow the strategy of choosing the same action as Agent 1*, see examples in Appendix Section B.1).

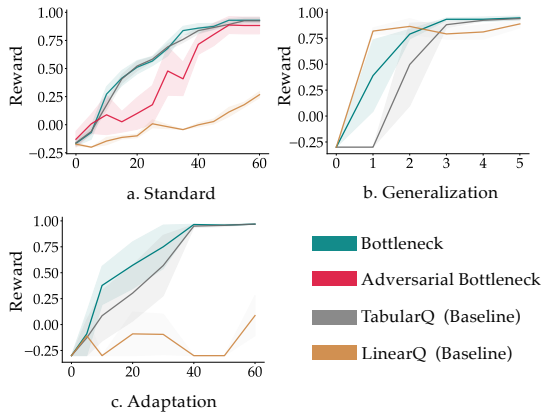


Figure 5: Results in MAZE. a) **Bottleneck** agents learn as fast as the non-linguistic **Baseline agents**, but faster than **Adversarial** and **LinearQ** agents. b) When faced with a new maze with similar structure, **Bottleneck** agents learn faster than **TabularQ** (which does not perceive color) and **LinearQ** (which does, but learns slower). c) When faced with a maze of a different structure, **Bottleneck** agents adapt swiftly while **LinearQ** cannot recover. We do not evaluate **Adversarial** for the Generalization or Adaptation experiments due to its poor performance in the standard setting.

5.2 PLLB can learn counter-intuitive policies

A potential confound in the preceding experiment is the possibility that PLLB helped converge to the human-interpretable policy because this is the *only* behavior an LM is able to describe. To test this hypothesis we fixed the speaker to use a counter-intuitive policy using a random mapping between balls and messages (e.g., $1 \rightarrow 3$, $5 \rightarrow 2$, etc.). Here, a listener following the human-interpretable policy would fail. In Figure 4c we observe that while **InstructRL** and **Adversarial**, both regularized by misaligned rules (the human-interpretable one for **InstructRL**, another random permutation for **Adversarial**), are able to eventually adapt to the fixed speaker, they converge at a slower rate than both **TabularQ** (no rule) and **Bottleneck**. Overall, **Bottleneck** is the only algorithm able to flexibly adapt to a fixed speaker policy while converging on the human-like policy when the speaker is also learning.

6 Maze

In the MAZE domain, agents must navigate a maze to find the goal using four directional actions (N/S/E/W). We study mazes in which environment cues provide hints about the optimal path. This setup helps us study the impact of the language

bottleneck: can PLLB agents infer the underlying structure and use it to generalize across mazes (cognitive use)? Can they transmit this knowledge to humans and help them perform better in novel settings (communicative use)?

Task Overview. We randomly generate different 7×7 mazes with the `gym-maze` code base.⁵ Agents are rewarded when they achieve the goal, by the inverse of the number of steps they needed to reach it as an incentive to learn the optimal solving path. We introduce additional structure to all mazes by coloring cells for which the optimal action is *south* (red) or *north then east* (blue) with 50% probability (or they are left blank). Therefore, agents that leverage color information may better generalize to new mazes where the coloring structure is preserved, but the optimal action sequence differs.

The first iteration of PLLB correspond to training a standard RL algorithm to obtain π_1 , in our case a tabular Q-learning agent (**TabularQ**), since there are no rules yet. After the agent observes at least two solved mazes, we run `gen_rule` by prompting a llama-2-70b-chat LM with contrastive episodes and a task description. This gives us a linguistic rule \mathcal{L}_1 (e.g., *Upon observing RED, take SOUTH*) that we can use to update the policy. We first induce the regularizing policy $\pi_{\mathcal{L}_1}$ with rule \mathcal{L}_1 by obtaining a probability distribution over the 4 actions from the LM, and then run the RL algorithm for 5 episodes to obtain the new policy π_2 (procedure described in Section 3). We repeat these steps every 5 episodes of interactions with the environment.

Baselines include the base tabular Q-learning algorithm without language bottleneck (**TabularQ**) and a variant of **Bottleneck** generating rules from reward-randomized episode samples (**Adversarial**). We additionally evaluate **LinearQ**, an agent learning a linear model Q-function with an additional feature for cell color. We train **LinearQ** with a batch update of size 10 and learning rate 0.001, after performing a hyperparameter sweep, and note that we found it generally took significantly longer to converge than the other methods.

6.1 PLLB learns more generalizable policies

Figure 5a shows that learning a valid rule (**Bottleneck**) does not increase learning speed

⁵<https://github.com/MattChanTK/gym-maze>

over **TabularQ** (already quite fast), but using a corrupted rule (**Adversarial**) or learning from linear features (**LinearQ**) slows down learning. We next replace the maze by another 7×7 maze sharing the same underlying color semantics, but with a different optimal action sequences. For fair comparison, we start with the fully converged policy for each method (requiring 100 episodes for **LinearQ**). Figure 5b shows that **Bottleneck** leverages the learned rule and adapts to the new maze more effectively than both **TabularQ** and a fully converged **LinearQ** agent. While **TabularQ** cannot generalize because it does not perceive colors, **LinearQ** does generalize faster at first but converges slower.

We find that the generated rules indeed improve over time to better capture the underlying structure of the maze (e.g., *if I observe BLUE, then take the NORTH action*; see other examples in Appendix B.1). Across all trials, 100% of the final rules mention the red \rightarrow south rule and 60% uncovered the more complex blue \rightarrow north-then-east rule.

6.2 PLLB learns adaptable policies

In Figure 5c, we reproduce this same experiment on a maze with a different underlying structure (red now indicates *west* while blue indicates *east then south*). Although the rule **Bottleneck** learned does not apply anymore, it can still adapt faster than the baselines. The language bottleneck can indeed quickly capture the new structure of the task and regularize learning to steer adaptation. Here **LinearQ** seems to have overfit to the first maze and struggled to adapt. We find that all trials end with rules capturing the new mapping (100% red \rightarrow west, 50% blue \rightarrow east-then-south). Overall, these results show that the language bottleneck supports human-like cognitive functions by finding a tradeoff between the efficiency of **TabularQ** and the generalization capabilities **LinearQ** while remaining more adaptable.

6.3 PLLB is more interpretable and inter-operable

Can generated rules be useful to humans as well? We asked 50 university students to solve a 7×7 maze in the fewest steps possible using arrow keys. They could only observe the cells they had already visited and the walls they had already bumped into (see Figure 12 in Appendix). We split them into three groups: a **Control** group receiving no assistance, and two others receiving infor-

mation about a *similar* but different maze sharing the same underlying color semantics. The **Visual** group is shown a visual representation of the optimal policy (arrows in each cell, see Figure 11 in Appendix), while the **Bottleneck** group is provided a randomly sampled language rule \mathcal{L} learned by PLLB. This set-up lets us evaluate how generalizable the two different aids are to a new maze, as well as how quickly participants can account for incorrect information depending on modality.

Participants using PLLB rules solve the new maze with fewer steps than others and find this aid more useful than the visual one in average (statistically significant at level 0.05 with two-sided Mann-Whitney test after Bonferroni correction, see Figure 6). Participants indeed found it harder to extract the relevant information from the visual aid, which contained a great deal of non-transferable information (optimal actions in all non-colored cells). PLLB rules focus on the useful and transferable information learned by the previous agent, which is much more readily usable by humans.

Overall, the results in MAZE demonstrate the ability of PLLB to train agents that are more generalizable and adaptable (cognitive use) as well as more interpretable by, and inter-operable with humans (communicative use).

7 Collaborative Image Reconstruction

PLLB is not restricted to the training of RL agents. This section introduces two collaborative image reconstruction tasks inspired by the collaborative assembly task of McCarthy (2020). **BUILDER** and **BIRDS** both consider two agents: a *speaker*, who can see a hidden target image, and a *listener*, who must accurately reconstruct

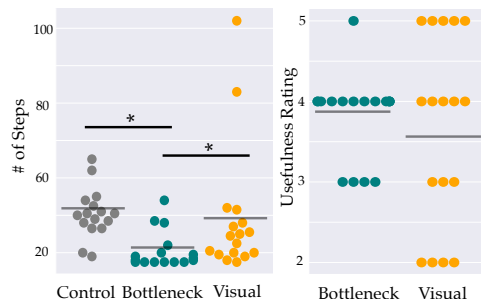


Figure 6: Participants provided rules from PLLB solve new mazes faster than those given either visual or no aid (control), left. They also self-reported linguistic feedback as more useful than visual (right).

the target image based on a description from the speaker. Both agents aim to converge to a description style leading to high reconstruction accuracy. We consider both synthetic images built from sequences of discrete actions by the listener (BUILDER), and natural images of birds that the listener reconstructs using a text-to-image generation model (BIRDS). Here, PLLB combines with state-of-the-art multi-modal models to let agents perceive semantic visual features.

7.1 Task Overview

Dataset. In BUILDER, we construct a dataset of synthetic images with a variable number of shapes (triangle, square, circle) in different colors (magenta, blue, red, green) and locations on a 2D grid. Each image is created by a sequence of discrete actions, each representing a particular combination of shape, color, and area of the grid (e.g., [ACT12 ACT2 ACT4], where ACT2 adds a magenta circle in the lower right area of the grid). In BIRDS, we construct a dataset of natural bird images by selecting 5 images for each of 10 bird species from the CUB-200-2011 dataset provided by Wah et al. (2011) (see Appendix Section A.2 for species list).

Reward. In BUILDER, because target and reconstructed images are completely defined by the action sequences that generate them, we measure task success with the Levenshtein similarity between the (sorted) corresponding action sequences. However, the listener in the BIRDS environment outputs images using a text-to-image model that exhibit a variety of properties. We therefore introduce reward functions corresponding to three different properties for BIRDS: **Color**, **Background**, and **Species**. To evaluate a listener’s reconstructed image x_r with respect to a particular target image x_t and reward function, we create a *contrast* set C by selecting the three images in CUB-200-2011 dataset that have a different value for the target reward property (e.g., color), but are most similar with respect to values for all other annotated properties. We then define reward r as $r = \sum_{c \in C} d(\text{CLIP}(x_r), \text{CLIP}(x_c)) - d(\text{CLIP}(x_r), \text{CLIP}(x_t))$, where CLIP refers to the CLIP image embedding (Radford et al., 2021). For both BUILDER and BIRDS tasks, reward is reported over a separate held-out validation set.

Models. As open-source VLMs struggle with describing synthetic images, we use the llama-2-70b-chat LM as our speaker for BUILDER, rep-

resenting images in raw text that list all shapes’ type, color, and exact coordinates individually in sequence. On the other hand, we use the open source Llava VLM as our speaker for BIRDS (Liu et al., 2023). We prompt both speaker models using the target image and a general task description, as well generated rules after the first iteration.

We implement the listener for BUILDER with BART, a neural sequence-to-sequence model pre-trained on English text by Lewis et al. (2019), which we fine-tune at each episode on a training set of (description, action sequence) pairs using descriptions provided by the speaker. Meanwhile, we use the Stable Diffusion text-to-image diffusion model as the listener for BIRDS.

For both tasks, we consider two settings: **From Scratch**, where the listener is initialized using the original pre-trained model weights at every episode, and **Continual Training**, where the listener is continuously updated over time using descriptions by the speaker. We therefore split our datasets into separate held-out sets for training, early stopping, selecting samples for gen_rule, and evaluation.

Implementation. We first generate base image descriptions for both tasks by providing the speaker a prompt containing a general task description and target image, but without any rule \mathcal{L} . For the BUILDER task, an original image description is: *At least five dots in total - four red and one green - scattered across the canvas, with two sets of matching locations ($x = 0.49 / y = 0.24$ and $x = 0.5 / y = 0.24$), another set at $x = 0.52 / y = 0.3$, and finally, one dot located near $x = 0.76 / y = 0.71$. This description is overly complex, mixes different coordinate formats, and difficult for the listener to use to reconstruct the target image, leading to low reward. Meanwhile, an example original image description for the BIRD task is: “a bird on a branch”. While simple, this description does not contain sufficient information about the **Color**, **Background**, and **Species**, causing the listener to generate an image reconstruction with low reward.*

We next evaluate these base descriptions using each task’s respective reward functions, and select the (image, description) examples with the 5 highest and lowest rewards. Using the prompts shown in Section A.3, we implement gen_rule using the llama-2-70b-chat and llava models for BUILDER and BIRDS, respectively. For BIRDS, we observe that output rules \mathcal{L} are specific to the reward func-

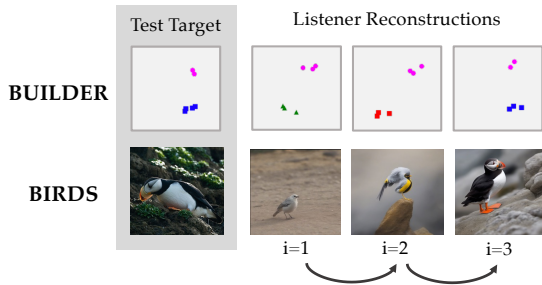


Figure 7: Both the BUILDER and BIRDS tasks consist of speaker and listener agents. At test time, the speaker needs to provide a language description to the listener that helps them recreate the image accurately. For both tasks, PLLB helps improve listener accuracy over time.

tion: an example rule we get for the **Species** reward is: “Identify the bird’s species if possible, and include any distinctive characteristics that set it apart from other birds.”, while an example rule we get for the **Background** reward is: “Describe the setting or background of the image, such as the presence of snow, water, or other elements”.

We implement update for both tasks by simply appending the output rule \mathcal{L} to the original prompt to the speaker, as shown in Figure 2, using the speaker’s output as the new description for a given target image. We repeat this cycle of eliciting rule \mathcal{L} , appending \mathcal{L} to the same prompt used to generate the base instructions in order to re-label our data with new instructions, and training and evaluating the listener model for a total of 5 iterations for BUILDER and 3 iterations for BIRDS. We qualitatively show this leads to improved image reconstruction from the listener in Figure 7, and next discuss our experimental results.

7.2 PLLB helps speakers provide more usable instructions

For both tasks, listeners following rules generated by PLLB (**Bottleneck**) outperform listeners trained from uninformed (**Baseline**) or misinformed (**Adversarial**) speakers (Figures 8 and 10, top row). This holds true for all three different reward functions in BIRDS, as well as both training settings, with **Continual Learning** learning enabling **Bottleneck** to have an even stronger improvement over **Baseline** image descriptions by leveraging learned task experience. **Continual Learning** does not significantly help the **Adversarial** method, emphasizing that the linguistic rules \mathcal{L} in **Bottleneck** do capture successful

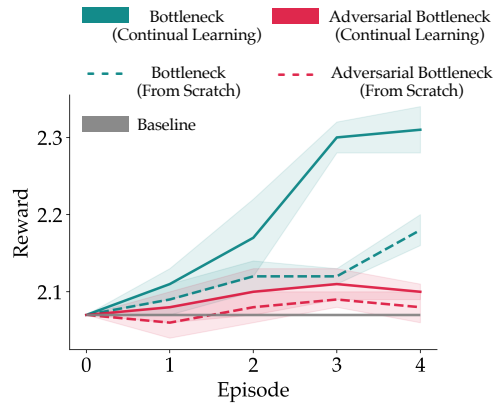


Figure 8: In BUILDER, performance of the **Bottleneck** listener improves upon **Baseline** descriptions over time with PLLB for both **From Scratch** and **Continual Learning** training settings, unlike an (**Adversarial**) listener that inaccurately represents rewards.

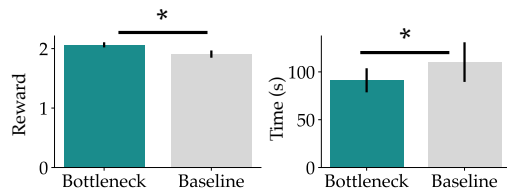


Figure 9: In BUILDER, study participants reconstructed target images with higher accuracy (left), and less time (right), when provided instructions that were generated using **Bottleneck** vs. original **Baseline** descriptions.

task strategies. However, **Continual Learning** is not required for **Bottleneck** to outperform **Baseline** image descriptions. The success of the **From Scratch** training setting can be interpreted as the speaker generating abstract rules to guide its own learning, leading to improved performance and showing evidence for a cognitive use of language.

Finally, we observe that the generated rules \mathcal{L} encourage the speaker to reduce vague and general language, and describe properties relevant to the corresponding underlying reward function in BIRDS, although this might take several cycles to fully converge (see example rules in Appendix Section B.1). Furthermore, rules \mathcal{L} become more complex over time, and across different trials we observe path dependency where different trials converge to different rules (e.g., (x, y) vs. $x = 0.xx, y = 0.yy$ for coordinates in BUILDER). Overall, these rules help guide the speaker and listener agents in both tasks to iteratively improve image reconstruction over time, which can also be seen qualitatively in Figure 7.

7.3 PLLB can collaborate with human listeners

Do human listeners also benefit from image descriptions provided by speakers following **Bottleneck** rules? We next conduct a human subject study for the BUILDER task, where we replace the listener with study participants at the end of one iteration of PLLB, and evaluate how accurately and quickly they can reconstruct target images from a held-out test set. We recruit 20 crowdworkers on Prolific to reconstruct 5 images using descriptions from a speaker following a rule \mathcal{L} generated with (**Bottleneck**), and 5 images using the original (**Baseline**) descriptions. We provided participants an interface that included a drawing canvas and buttons controlling shape types and colors (Section A.1), allowing us to use our reward function of Levenshtein similarity between user actions and the underlying action sequences for a target image.

Participants using **Bottleneck** descriptions built target images faster and more accurately than participants using control descriptions (Bonferroni-corrected Wilcoxon signed-rank test, $p < 0.05$, Figure 9). 55% of them preferred **Bottleneck** descriptions when asked in a post-study survey, describing them as “*more direct and less ambiguous*”. 10% answered that both descriptions types were equally easy to follow while remaining participants preferred **Baseline** descriptions because they *they gave more flexibility*, showing a preference not captured in our reward function.

7.4 PLLB collaborates with human speakers

Can humans benefit from rules generated by PLLB? We asked 12 Prolific crowdworkers to act as speakers and provide descriptions for our dataset of BIRDS images. Half of the participants were provided rules corresponding to one of the three reward functions (\mathcal{L}_3) while the other half were not. Figure 10 (bottom) shows that listeners instructed by human speakers provided with PLLB rules outperform those instructed by uninformed human speakers (**Baseline**) (Bonferroni-corrected t-tests $p < 0.05$ for **species** and **colors**).

The descriptions of workers not provided a rule are more diverse and less focused e.g., *A barbed wire is an uncomfortable stop for a bird*. With PLLB rules on the other hand (e.g., *Describe the bird’s coloration accurately*), they generate more specific descriptions: *Red and black bird on barbed wire. Bright red chest, red at top of*

head, black wings and beak for the same target image (see other example rules in Appendix Table 4). PLLB agents can easily transmit what they learned from experience (the rule) to humans who can then benefit from that knowledge (cultural transmission, communicative use).

8 Limitations & Discussion

By inferring the rule explaining its best past behaviors, PLLB becomes more interpretable by and inter-operable with humans. When it uses a rule to condition its policy updates, it steers its own behavior to become more interpretable and generalizable. In our current implementation, rule generation remains limited: we require converting agent histories into representations the LM can handle (e.g., text or images), which may prevent the use of long-horizon sensorimotor trajectories (e.g., robotic systems). Rule generation can also fail when the LM overfits or abstracts too much (e.g., in BIRDS we found the rule: “*Use descriptive language that conveys the mood, such as the serenity of a snowy day or the freedom of flight*”). Advances in multi-modal models and better approaches for selecting contrastive examples for `gen_rule` might help mitigate these issues.

An interesting direction for future work is to apply PLLB to tasks with complex reward functions that include hard-to-articulate human preferences. In applications like image captioning for accessibility (Nie et al., 2020) and personalization of language models (Li et al., 2023a), linguistic rules might improve transparency and predictability even in cases where users find it difficult to describe their own goals.

Finally, we could extend PLLB to generate other kinds of linguistic information: not only abstract rules, but also explanations, goals, or learning strategies. Such information could then be shared with others (cultural transmission) or used by update functions in new ways: mapping generated language to state or action abstractions, goals, reward functions, or even learning strategies. This opens the door for more human-like AI agents leveraging the deep interactions between language and decision-making (Colas et al., 2022).

Acknowledgements

We would like to acknowledge support from NSF Awards #2006388, #2125511, #2238240 and #2212310, AFOSR YIP, ONR #N000142112298,

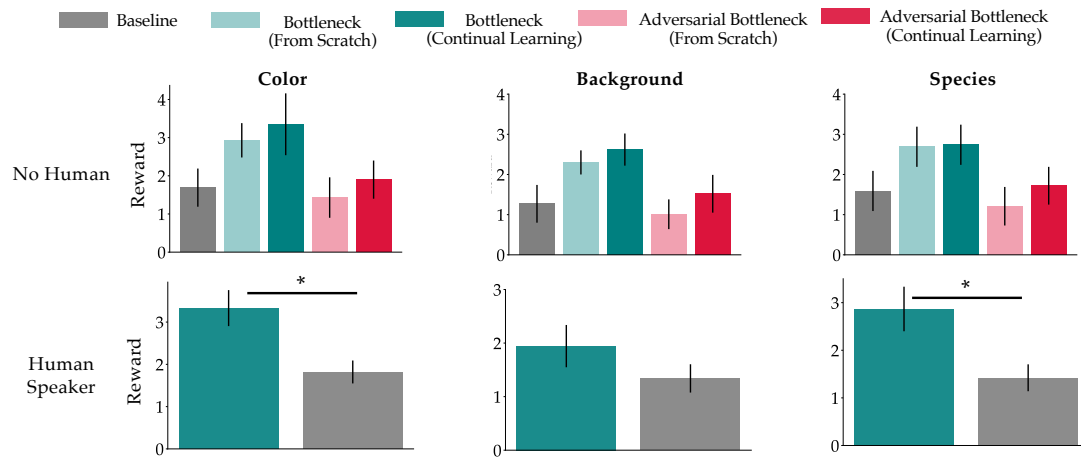


Figure 10: For the BIRDS task, rules generated after 3 iterations of LANGUAGE BOTTLENECK help automated (top) and human (bottom) speakers generate better descriptions for naturalistic image reconstruction for three different underlying reward functions (columns). **Continual learning** (dark teal, red in top graphs) further improves reward over listener agents learning from scratch (dark teal, red in top graphs), but not sufficiently enough to help **Adversarial** outperform baseline descriptions, showing that **Bottleneck** rules describe rewarding behavior.

and DARPA. Megha Srivastava was additionally supported by an IBM PhD Fellowship. Cédric Colas received funding from the European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No. 101065949.

References

- Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Hausman, et al. 2022. [Do as i can, not as i say: Grounding language in robotic affordances.](#)
- Bastien Boutonnet and Gary Lupyan. 2015. [Words jump-start vision: A label advantage in object recognition.](#) *Journal of Neuroscience*, 35(25):9329–9335.
- Peter Carruthers and Jill Boucher. 1998. *Language and Thought*. Cambridge University Press.
- Valerie Chen, Abhinav Gupta, and Kenneth Marino. 2021. Ask your human: Using human instructions to improve generalization in reinforcement learning. *Proc. of ICLR*.
- Maxime Chevalier-Boisvert, Dzmitry Bahdanau, Salem Lahlou, Lucas Willems, Chitwan Saharia, Thien Huu Nguyen, and Yoshua Bengio. 2019. Babyai: A platform to study the sample efficiency of grounded language learning. *Proc. of ICLR*.
- Sahil Chopra, Michael Henry Tessler, and Noah D Goodman. 2019. The first crank of the cultural ratchet: Learning and transmitting concepts through language. *Proc. of CogSci*.
- Andy Clark. 1998. Magic Words: How Language Augments Human Computation. In Peter Carruthers and Jill Boucher, editors, *Language and Thought*, 1 edition. Cambridge University Press.
- Cédric Colas, Tristan Karch, Nicolas Lair, Jean-Michel Dussoux, Clément Moulin-Frier, Peter F. Dominey, and Pierre-Yves Oudeyer. 2020. Language as a cognitive tool to imagine goals in curiosity driven exploration. *Proc. of NeurIPS*.
- Cédric Colas, Tristan Karch, Clément Moulin-Frier, and Pierre-Yves Oudeyer. 2022. Language and culture internalization for human-like autotelic ai. *Nature Machine Intelligence*, 4(12):1068–1076.
- Lisa Dunlap, Yuhui Zhang, Xiaohan Wang, Ruiqi Zhong, Trevor Darrell, Jacob Steinhardt, Joseph E. Gonzalez, and Serena Yeung-Levy. 2023. [Describing differences in image sets with natural language.](#)
- Dedre Gentner and Jeffrey Loewenstein. 2002.

- Relational Language and Relational Thought*. Erlbaum.
- Robert D Hawkins, Minae Kwon, Dorsa Sadigh, and Noah D Goodman. 2020. Continual adaptation for efficient machine communication. *Proc. of ACL*.
- Karl Moritz Hermann, Felix Hill, Simon Green, Fumin Wang, Ryan Faulkner, Hubert Soyer, David Szepesvari, Wojciech Marian Czarnecki, Max Jaderberg, Denis Teplyashin, Marcus Wainwright, Chris Apps, Demis Hassabis, and Phil Blunsom. 2017. [Grounded language learning in a simulated 3d world](#).
- Linda Hermer-Vazquez, Anne Moffet, and Paul Munkholm. 2001. Language, space, and the development of cognitive flexibility in humans: The case of two spatial memory tasks. *Cognition*, 79(3):263–299.
- Mary Hesse. 1988. The Cognitive Claims of Metaphor. *The Journal of Speculative Philosophy*.
- Felix Hill, Andrew K. Lampinen, Rosalia Schneider, Stephen Clark, Matthew Botvinick, James L. McClelland, and Adam Santoro. 2020. Emergent systematic generalization in a situated agent. *Proc. of ICLR*.
- Hengyuan Hu and Dorsa Sadigh. 2023. [Language instructed reinforcement learning for human-AI coordination](#).
- Hengyuan Hu, David J Wu, Adam Lerer, Jakob Foerster, and Noam Brown. 2022. [Human-AI coordination via human-regularized search and learning](#).
- Shengran Hu and Jeff Clune. 2023. Thought cloning: Learning to think while acting by imitating human thinking. *Proc. of NeurIPS*, 36.
- Yiding Jiang, Shixiang Gu, Kevin Murphy, and Chelsea Finn. 2019. Language as an abstraction for hierarchical deep reinforcement learning. *Proc. of NeurIPS*.
- Martin Klissarov, Pierluca D’Oro, Shagun Sodhani, Roberta Raileanu, Pierre-Luc Bacon, Pascal Vincent, Amy Zhang, and Mikael Henaff. 2023. [Motif: Intrinsic motivation from artificial intelligence feedback](#).
- George Lakoff and Mark Johnson. 2008. *Metaphors We Live By*. University of Chicago press.
- Andrew K. Lampinen, Nicholas A. Roy, Ishita Dasgupta, Stephanie C. Y. Chan, Allison C. Tam, James L. McClelland, Chen Yan, Adam Santoro, Neil C. Rabinowitz, Jane X. Wang, and Felix Hill. 2022. Tell me why! – explanations support learning of relational and causal structure. *Proc. of ICML*.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#).
- Belinda Z Li, Alex Tamkin, Noah Goodman, and Jacob Andreas. 2023a. [Eliciting human preferences with language models](#).
- Chengshu Li, Jacky Liang, Andy Zeng, Xinyun Chen, Karol Hausman, Dorsa Sadigh, Sergey Levine, Li Fei-Fei, Fei Xia, and Brian Ichter. 2023b. [Chain of code: Reasoning with a language model-augmented code emulator](#).
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2023. [Improved baselines with visual instruction tuning](#).
- Jelena Luketina, Nantas Nardelli, Gregory Farquhar, Jakob Foerster, Jacob Andreas, Edward Grefenstette, Shimon Whiteson, and Tim Rocktäschel. 2019. [A survey of reinforcement learning informed by natural language](#).
- Aleksander R Luria. 1959. The directive function of speech in development and dissolution. *Word*, 15(2).
- William McCarthy. 2020. [Emergence of compositional abstractions in human collaborative assembly](#).
- Reid McIlroy-Young, Siddhartha Sen, Jon Kleinberg, and Ashton Anderson. 2020. [Aligning superhuman ai with human behavior: Chess as a model system](#). In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery Data Mining*, KDD ’20. ACM.

- Alex Mesoudi and Alex Thornton. 2018. What is cumulative cultural evolution? *Proceedings of the Royal Society B*, 285(1880):20180712.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. 2013. [Playing atari with deep reinforcement learning](#).
- Allen Nie, Reuben Cohn-Gordon, and Christopher Potts. 2020. [Pragmatic issue-sensitive image captioning](#).
- Joon Sung Park, Joseph C. O'Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. [Generative agents: Interactive simulacra of human behavior](#).
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. [Learning transferable visual models from natural language supervision](#).
- Nicholas A Roy, Junkyung Kim, and Neil Rabinowitz. 2022. Explainability via causal self-talk. *Advances in Neural Information Processing Systems*, 35:7655–7670.
- Pratyusha Sharma, Antonio Torralba, and Jacob Andreas. 2021. Skill induction and planning with latent language. *Proc. of ACL*.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2024. Reflexion: Language agents with verbal reinforcement learning. *Proc. of NeurIPS*.
- Elizabeth S Spelke. 2003. What makes us smart? core knowledge and natural language. *Language in mind: Advances in the study of language and thought*, pages 277–311.
- Richard S Sutton and Andrew G Barto. 2018. *Reinforcement Learning: An Introduction*. MIT press.
- Allison Tam, Neil Rabinowitz, Andrew Lampinen, Nicholas A Roy, Stephanie Chan, DJ Strouse, Jane Wang, Andrea Banino, and Felix Hill. 2022. Semantic exploration from language abstractions and pretrained representations. *Proc. of NeurIPS*.
- Michael Henry Tessler, Jason Madeano, Pedro A. Tsividis, Brin Harper, Noah D. Goodman, and Joshua B. Tenenbaum. 2021. [Learning to solve complex tasks by growing knowledge culturally across generations](#).
- L.S Vygotsky. 1965. *Thought and Language*. MIT Press.
- C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. 2011. The caltech-ucsd birds-200-2011 dataset. Technical report.
- Tony T. Wang, Adam Gleave, Tom Tseng, Kellin Pelrine, Nora Belrose, Joseph Miller, Michael D. Dennis, Yawen Duan, Viktor Pogrebnik, Sergey Levine, and Stuart Russell. 2023. [Adversarial policies beat superhuman go ais](#).
- Olivia Watkins, Abhishek Gupta, Trevor Darrell, Pieter Abbeel, and Jacob Andreas. 2021. Teachable reinforcement learning via advice distillation. In *Proc. of NeurIPS*.
- Sandra R Waxman. 1994. The development of an appreciation of specific linkages between linguistic and conceptual organization. *Lingua*, 92:229–257.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023. [Chain-of-thought prompting elicits reasoning in large language models](#).
- Catherine Wong, Kevin Ellis, Joshua B. Tenenbaum, and Jacob Andreas. 2021. Leveraging language to learn program abstractions and search heuristics. *Proc. of ICML*.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2022. [React: Synergizing reasoning and acting in language models](#).
- Hanako Yoshida and Linda B Smith. 2003. Sound symbolism and early word learning in two languages. In *Proc. of CogSci*.
- Ruiqi Zhong, Peter Zhang, Steve Li, Jinwoo Ahn, Dan Klein, and Jacob Steinhardt. 2023. [Goal driven discovery of distributional differences via language descriptions](#).

Victor Zhong, Tim Rocktäschel, and Edward Grefenstette. 2020. RTFM: Generalising to new environment dynamics via reading. In *Proc. of ICLR*.

A Additional Methods

A.1 Maze

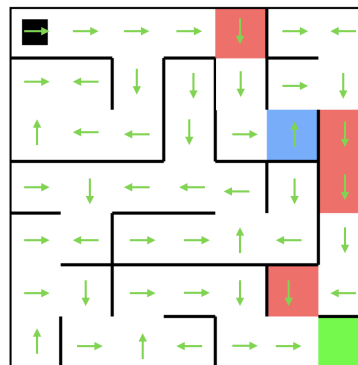


Figure 11: Visual aid provided to participants in the human subject study for the MAZE task.

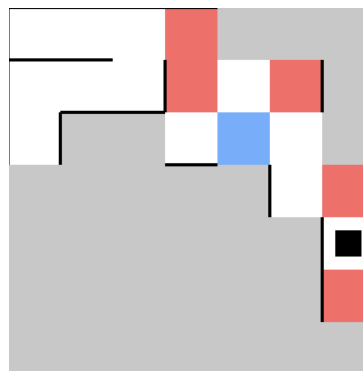


Figure 12: Interface used in the human subject study for the MAZE task. Cells not visited by the participants are hidden in gray.

A.2 Dataset Details for BIRDS.

We consider images of the following bird species in the CUB-200-2011 dataset from [Wah et al. \(2011\)](#) for the BIRDS task: [Indigo Bunting, Cardinal, Yellow Breasted Chat, American Crow, Vermillion Flycatcher, California Gull, Blue Jay, Tropical Kingbird, White Pelican, Horned Puffin].

A.3 Full Prompts

We first provide the full prompts used for `gen_rule`. For space constraints, we do not include example samples, but note they follow the format shown in [Figure 2](#).

SELECTSAY: You will be given a list of (OBSERVATION, ACTION, REWARD) examples collected from two agents learning to solve a task.

Welcome!

Your goal is to follow the instruction on the top right to create an image using the buttons on the right. Each instruction belongs to either Type A or Type B - you should follow all instructions the same way, but we will ask you after the study to choose whether you found Type A instructions or Type B instructions easier to follow.

When you are done with one image, click the **Finish** button to proceed to the next one. You will create 10 images in total. At the end, you will receive a link to a post-study survey to fill out.

Please complete the task in one sitting. Do your best to follow the instructions. We reserve the right to reject your submission if you create random or empty images.

8/10**Instruction #8**

(Type B) Three red shapes in the top left quadrant, including two triangles and a square. One triangle located at (0.2, 0.73), another at (0.27, 0.72), and a third at (0.31, 0.84). Additionally, there is a red square positioned at (0.24, 0.25) and two more triangles, one at (0.23, 0.68) and the other at (0.21, 0.73).

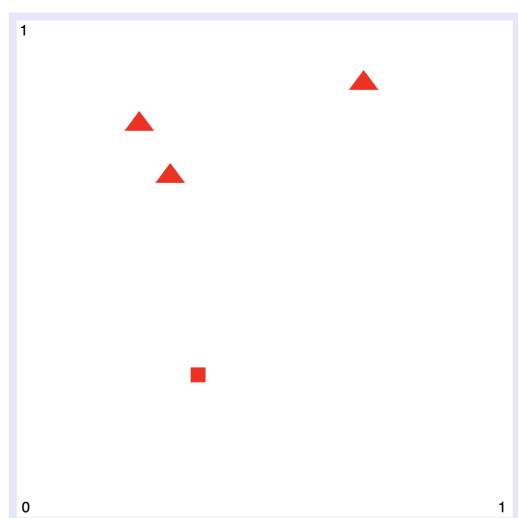


Figure 13: Interface used in the human subject study for the BUILDER task.

Possible ACTIONS an agent can take are: 1, 2, 3, 4, 5, and quit. Each OBSERVATION describes the ordered sequence of actions that AGENT 1 picks, and each ACTION describes the ACTION that AGENT 2 picks based on the given OBSERVATION. The examples are separated into HIGH REWARD and LOW REWARD examples. +[samples]+Output a language rule that best summarizes the strategy AGENT 2 should follow to receive HIGH REWARD, not LOW REWARD, based on the examples. Start the instruction with the prefix 'I should'.

MAZE: You will be given a list of example (OBSERVATION, ACTION) trajectories collected from an AGENT learning to solve a maze. Each trajectory receives a REWARD. Possible OBSERVATIONS an agent see are: WHITE, RED, BLUE Possible ACTIONS an agent can take are: NORTH, SOUTH, EAST, WEST. The examples are separated into HIGH REWARD and LOW REWARD examples + [samples] + Output a language rule that best summarizes the strategy the AGENT should follow when picking a sequence of ACTIONS to solve the maze and receive HIGH REWARD, not LOW REWARD, based on the examples. Start the instruction with the prefix 'I should'.

BUILDER: There are two agents. The goal of Agent 1 is to provide instructions to Agent 2 that helps Agent 2 to successfully recreate the image. You will be given a list of (ORIGINAL, AGENT 1 INSTRUCTION, REWARD) values where ORIGINAL is the original description of an image, INSTRUCTION is the instruction provided by Agent 1 to Agent 2, and REWARD is the reward Agent 2 receives when trying to re-create the image (higher is better). The examples are separated into HIGH REWARD and LOW REWARD examples. + [samples]+

Based on the examples above, output a list of 2 RULES for Agent 1 to follow when generating INSTRUCTION in order to receive HIGH REWARD, instead of LOW REWARD. Write the rules after the prefix RULES:

BIRDS: The top row of three images have the following HIGH REWARD descriptions:+high reward samples+The bottom row of three images have the following LOW REWARD descriptions:+low reward samples+Provide a rule I should follow in order to provide image descriptions with HIGH REWARD, not LOW REWARD. Provide the rule after the prefix RULE:"

We next provide the full prompts used in update for each task.

1. **SELECTSAY:** $[\mathcal{L}] + \text{Agent } 1 \text{ selected } [\text{observation}]$. So I should select

2. **MAZE:** You are an agent solving a maze following a provided **RULE**. You will be given a list of **PREVIOUS ACTIONS** and the **CURRENT OBSERVATION**. Follow the **RULE** to select your **NEXT ACTION** (East, West, South, North):

RULE: + $[\mathcal{L}] +$ **PREVIOUS ACTIONS:** + $[\tau_{1...t-1}] +$ **CURRENT OBSERVATION:** + $[\text{observation}] +$

What is the **NEXT ACTION** you should take? Output one of (East, West, South, North) after the prefix **NEXT ACTION:**.

3. **BUILDER:** You will be given a **DESCRIPTION** of an image. Your goal is to use this description to provide a short **INSTRUCTION** to help someone else, who cannot see the image, accurately re-construct it. You will also be given a list of **RULES** you must follow when providing the instruction.

DESCRIPTION: + observation +

RULES: + $\mathcal{L} +$

Please provide a short instruction following the prefix **INSTRUCTION:**

4. **BIRDS:** Provide a one-sentence description of this image, using the following **RULES:** + \mathcal{L}

B Additional Results

B.1 Examples of Generated Rules for All Environments

Setting	\mathcal{L}_1	\mathcal{L}_2	\mathcal{L}_5
Standard	I should follow the strategy of choosing action 4 whenever the observation is 2, 3, 4, or 5.	I should follow the strategy of choosing actions based on the observation provided by AGENT 1. If the observation is 1, I should choose action 1. If the observation is 2, I should choose action 3. If the observation is 3, I should choose action 3. If the observation is 4, I should choose action 4. If the observation is 5, I should choose action 5.	I should follow the strategy of choosing the same action as AGENT 1 for observations 1, 2, 3, 4, and 5.
Fixed Speaker	I should choose action 1 when observation is 1 or 2 or 4 or 5. I should choose action 2 when observation is 3.	I should choose action 1 when observation is 1, 2, 3 or 5. I should choose action 2 when observation is 4.	I should choose action 1 when observation is 1. I should choose action 4 when observation is 2. I should choose action 5 when observation is 3. I should choose action 2 when observation is 4. I should choose action 3 when observation is 5.

Table 1: Example \mathcal{L} rules generated for the SELECTSAY environment, for the Standard setting (both Listener and Speaker agents are RL agents trained from random initialization) and a Fixed Speaker agent.

Size	\mathcal{L}_1	\mathcal{L}_2	\mathcal{L}_5	\mathcal{L}_8
7x7	I should start by moving north, then continue moving north until I reach a wall, then turn east and move until I reach a wall, then turn south and move until I reach a wall, then turn west and move until I reach the end of the maze.	I should prioritize moving North, East, and South over moving West, as the HIGH REWARD EXAMPLES show a higher frequency of these actions. 2. I should avoid moving West unnecessarily, as it is rarely seen in the HIGH REWARD EXAMPLES and often leads to LOW REWARD EXAMPLES	I should take the EAST action when I see WHITE, and then repeat the EAST action until I see RED, then take the SOUTH action, and then repeat the EAST action until I see BLUE, then take the NORTH action, then take the EAST action until I see RED again, then take the SOUTH action, and finally take the WEST action.	I should take the EAST action when I observe WHITE, and keep taking EAST actions until I observe RED, then take the SOUTH action. After taking the SOUTH action, I should take the EAST action again until I observe BLUE, then take the NORTH action. After taking the NORTH action, I should take the EAST action until I observe RED again, then take the SOUTH action.

Table 2: Example \mathcal{L} rules generated for the MAZE environment for two different maze sizes.

Setting	\mathcal{L}_1	\mathcal{L}_2	\mathcal{L}_3
Re-Initialization	1. Be specific with location details: Agent 1 should provide detailed location information for each element in the image, such as x and y coordinates. 2. Use descriptive language for elements, such as "red dot" or "green triangle".	1. Use specific coordinates when instructing Agent 2 to draw shapes. 2. Use descriptive language to specify the color and shape of each element. For example, "a green triangle" instead of "a green thing".	1. Be specific with location coordinates: provide specific coordinates for the location of each shape, using the format $x=0.XX, y=0.YY$. 2. Use descriptive shape names: Instead of using generic terms like "dot" or "square," use more descriptive names that indicate the shape's color and size, such as "green triangle" or "red square."
Continual Training	1. Be specific and detailed in your instructions. High reward examples have specific coordinates and shapes, while low reward examples have more general descriptions. 2. Use a consistent format for your instructions. High reward examples have a consistent format for listing coordinates and shapes, while low reward examples have a more free-form format.	1. Provide explicit coordinates for each element in the image, using the format (x, y). 2. Use specific colors when referring to elements in the image, such as "red", "green", or "blue". Avoid using vague terms like "colored" or "shaded".	1. Use a consistent format for describing shapes, such as always listing the x-coordinate first, followed by the y-coordinate. For example, instead of "one green square at the point $x=0.53, y=0.24$ ", use "one green square at (0.53, 0.24)". 2. Avoid using vague terms like "various shades of green". Instead, use specific colors, such as "green" or "blue". Additionally, use specific shapes, such as "square" or "triangle", rather than vague terms like "rectangle".

Table 3: Example \mathcal{L} rules generated for the BUILDER environment.

Reward	\mathcal{L}_1	\mathcal{L}_2	\mathcal{L}_3
color	Describe the bird's color, species, and any distinctive markings or patterns.	Describe the bird's coloration accurately.	Describe the bird's coloration accurately.
background	Include details about the bird's surroundings, such as the type of branch or post it is on, and any additional elements in the background.	Include the bird's action (perched, flying, standing) and its location (on a branch, railing, pole, etc.)	Describe the bird's action (flying, perching, standing) and the environment it is in (sky, tree, water).
species	Describe the bird's color, markings, and any distinctive features.	Describe the subject's unique features, such as coloration, beak shape, or other distinguishing characteristics.	Include specific details about the bird's appearance, such as the color of its feathers, beak, or eyes, and any distinctive markings or patterns.

Table 4: Example \mathcal{L} rules generated for the BIRDS environment demonstrate reward-specificity over time.