

# *lil*Gym: Natural Language Visual Reasoning with Reinforcement Learning

Anne Wu, Kianté Brantley, Noriyuki Kojima, and Yoav Artzi

Department of Computer Science and Cornell Tech, Cornell University

{aw588, kdb82, nk654}@cornell.edu, {yoav}@cs.cornell.edu

## Abstract

We present *lil*Gym, a new benchmark for language-conditioned reinforcement learning in visual environments. *lil*Gym is based on 2,661 highly-compositional human-written natural language statements grounded in an interactive visual environment. We introduce a new approach for exact reward computation in every possible world state by annotating all statements with executable Python programs. Each statement is paired with multiple start states and reward functions to form thousands of distinct Markov Decision Processes of varying difficulty. We experiment with *lil*Gym with different models and learning regimes. Our results and analysis show that while existing methods are able to achieve non-trivial performance, *lil*Gym forms a challenging open problem. *lil*Gym is available at <https://lil.nlp.cornell.edu/lilgym/>.

## 1 Introduction

The ability to reason about natural language has the potential to transform how reinforcement learning (RL) is used to train grounded agents. Language provides an expressive and accessible conduit for task specification, enabling agents to address a broad set of tasks, rather than to learn single-behavior policies. At the same time, RL is a promising framework to study core grounded language learning problems within interactive environments.

Prerequisite to realizing this potential are expressive benchmark environments, as has been instrumental for progress in RL more broadly. However, natural language poses unique challenges to building such benchmarks. Beyond the design of the environment itself, which must support rich linguistic reasoning, accurate reward computation requires resolving language semantics. Existing approaches adopt different strategies to address this issue, most often by using synthetic language (e.g., Côté et al., 2018; Co-Reyes et al., 2019), or by removing the

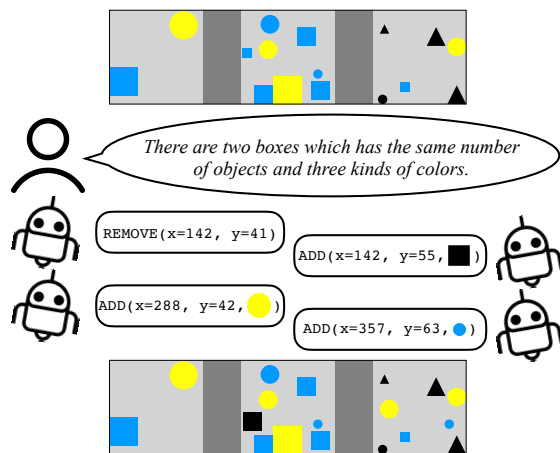


Figure 1: An illustrative example of *lil*Gym. Given a natural language statement with a target truth value, True in this example (omitted from the figure), and an initial image, the agent needs to manipulate objects in the environment so that the truth value of the statement with respect to the image is the target truth value.

language problem from reward computation by restricting to a single goal state (e.g., Anderson et al., 2018; Chen et al., 2019). While these approaches open new research avenues, both have significant drawbacks. The simplifications of synthetic language limit the relevance of methods and results to the complexities of human language. Single-goal formulations forgo or restrict language’s ability to efficiently abstract over many possible solutions, a core argument for its potential to RL.

We present *lil*Gym, an RL benchmark where an agent manipulates a visual environment by adding and removing objects conditioned on a natural language statement. The agent’s goal is to modify the environment so that a given statement will have a pre-specified truth-value with regard to the environment (i.e., the constraints specified in the language are either satisfied or violated). Figure 1 illustrates the scenario. *lil*Gym includes highly-compositional and semantically-diverse natural language statements and visual environments from the Natural Language for Visual Reasoning (NLVR)

corpus (Suhr et al., 2017), that combine with a configurable environment backbone to create thousands of Markov Decision Processes (MDP) of varying complexity.

A key challenge in constructing *lilGym* is accurate reward computation. Because of the flexibility of the environments and language, there are many possible equally correct termination states (i.e., goals) for each MDP. Correct reward computation at every possible state requires taking into account the semantic nuances of the highly-compositional language. We address this by annotating all statements with executable Python programs representing their meaning, effectively creating a supervised semantic parsing corpus (e.g., Zelle and Mooney, 1996; Zettlemoyer and Collins, 2005). The executable programs allow for exact and reliable reward computation at every possible state.

Our experiments with *lilGym* show that existing models demonstrate non-trivial performance given sufficient training time, with multi-modal pre-training providing further benefit. However, there remains significant room for improvement. For example, on the simplest configuration, our agent can solve 76.23% of the test environments, but performance drops significantly to only 16.81% on the most complex configuration. Our experiments also confirm the importance of modeling language meaning for reward computation in learning. The *lilGym* benchmark and trained models are available under the MIT license at <https://lil.nlp.cornell.edu/lilgym/>.

## 2 Related Work

RL research has benefited greatly from benchmarks such as Atari (Bellemare et al., 2013) and MuJoCo (Todorov et al., 2012). Although both are synthetic with limited potential to train policies directly applicable to realistic environments, their accessibility and focus on core RL problems have made them impactful drivers of algorithm development. For example, the first demonstration of an effective neural network policy (Mnih et al., 2013) and the development of proximal policy optimization (PPO; Schulman et al., 2017) both used these benchmarks, and both were later shown to generalize to more complex scenarios. *lilGym* is inspired by these benchmarks. Rather than aiming for training models that transfer to realistic domains, it aims to enable fundamental algorithmic development by emphasizing semantic diversity and compositional-

ity within an accessible benchmark.

There is significant interest in RL for training policies conditioned on natural language. This includes multiple efforts at developing RL environments with language-informed tasks, mainly via grounded language learning in 2D or 3D environments or using text games. Oftentimes, using synthetic language (Narasimhan et al., 2015; Johnson et al., 2017; Chevalier-Boisvert et al., 2019; Cao et al., 2020; Côté et al., 2018; Urbanek et al., 2019; Hermann et al., 2017; Co-Reyes et al., 2019; Hausknecht et al., 2020; Jiang et al., 2020). Although synthetic language allows studying the problem of learning high-level concepts, many of the complexities of natural language may be stripped away, and such approaches run the risk of reducing the language learning challenge to reverse engineering the hand-crafted generation process. In contrast, *lilGym* is based on semantically-diverse human-written language grounded in a visual environment, and requires both the ability of reasoning over highly-compositional language including sets and spatial relations, and precise alignment between statements and states.

Another approach is to simplify the task so only a few annotated termination states or trajectories are correct (Anderson et al., 2018; Chen et al., 2019; Ku et al., 2020). This forgoes much of the abstractive potential of natural language, where it can succinctly define extremely large set of states. Thereby reducing the utility of language, and not exposing learning algorithms to some of the core challenges it introduces to the RL problem. *lilGym* does not adopt such simplifications. Our experiments show the importance of considering language meaning for reward computation when there is a large set of valid goal states.

Alternatively, other benchmarks are created by first generating the target sequence of decisions (i.e., task demonstration), and then soliciting post-hoc instructional language (Shridhar et al., 2020, 2021; Hanjie et al., 2021). This process uses human-written language, but retains the regularities of the demonstration generation procedure. *lilGym* uses language from NLVR, which was crowdsourced via a contrastive task that was shown to elicit high semantic diversity.

## 3 Background: the NLVR Corpus

*lilGym* uses data from the NLVR corpus (Suhr et al., 2017). NLVR was initially created as a su-

pervised learning benchmark. We formalize an interactive task using the NLVR data and collect additional annotations for reward computation.

NLVR includes human-written natural language statements paired with synthetic images. Each pair is annotated with the boolean truth-value of the statement with regard to the image (i.e., True if the statement is true with regard to the image, or False otherwise). The images are designed to support complex reasoning, including about spatial and set relations. The original learning task posed by NLVR is to classify statement-image pairs as True to indicate the statement is true with regard to the image, or False otherwise. NLVR has been studied extensively (Suhr et al., 2017; Tan and Bansal, 2018; Goldman et al., 2018; Pavez et al., 2018; Yao et al., 2018; Hudson and Manning, 2018; Perez et al., 2018; Dasigi et al., 2019; Zheng et al., 2020; Gupta et al., 2021), and a separate version using photos was also released (Suhr et al., 2019).<sup>1</sup>

Qualitative analysis of the NLVR data (Table 2 in Suhr et al., 2017) showed it to provide diverse representation of semantic and compositional phenomena, including requiring joint visual-linguistic reasoning about spatial relations, quantities, and set relations. NLVR also provides an underlying structured representation for every image, which supports image manipulation. The combination of an interface for image manipulation with complex reasoning via natural language makes NLVR ideal to support an interactive benchmark environment.

## 4 The *lilGym* Benchmark

*lilGym* consists of a collection of environments that share a common backbone. The backbone is a 2D plane that is manipulated by placing and removing objects of different types. Each environment instance is a Markov Decision Process (MDP) created by pairing a natural language statement and a target boolean value with a configuration of the shared backbone. The goal of the agent is to manipulate the environment by adding and removing objects so that the truth-value of the statement with regard to the environment is the target boolean.

The learning problem *lilGym* presents is to induce a policy that generalizes across MDPs. We split the MDPs to training, development, and held-out testing sets. The training environments are for parameter estimation, while the two other sets are

for testing during development and for final held-out testing to report approach performance.<sup>2</sup>

There are two dimensions of configuration: appearance and starting condition. The appearance determines the state space, transition function, and action space. The appearance of the environment can be (a) TOWER: the objects include squares only, and they can be stacked into towers in specific positions only; or (b) SCATTER: objects of different types can be freely distributed. The two leftmost examples in Figure 2 are from TOWER, and the two rightmost are from SCATTER. TOWER gives a more constrained problem with much smaller state and action spaces compared to SCATTER.

There are two starting conditions, which also determine the agent’s goal: (a) SCRATCH: the environment starts without any objects and the goal is to modify it so that the statement’s truth-value is True; or (b) FLIPIT: the environment starts with a set of objects and the agent’s goal is to flip the truth-value of the statement, by modifying the environment. The first row of images in Figure 2 shows start states in both conditions. SCRATCH generally only requires adding objects, except in cases of correcting for agent’s errors, while FLIPIT requires both adding and removing, because there are already objects present.

The four configurations are TOWER-SCRATCH, TOWER-FLIPIT, SCATTER-SCRATCH, and SCATTER-FLIPIT. In our experiments (Section 6), we observe the different configurations provide different levels of difficulty. For example, SCATTER configurations are generally harder than TOWER, due to the larger state and action spaces.

Formally, each configuration is a Contextual Markov Decision Process (CMDP; Hallak et al., 2015). CMDP is an abstraction over a set of MDPs to account for a context that remains constant throughout an interaction with an MDP. The context includes the statement and the target boolean the interaction is conditioned on. A CMDP is a tuple  $(\mathcal{C}, \mathcal{S}, \mathcal{A}, \mathcal{M}(c))$ , where  $\mathcal{C}$  is the context space,  $\mathcal{S}$  the state space,  $\mathcal{A}$  the action space, and  $\mathcal{M}$  a function mapping a context  $c \in \mathcal{C}$  to an MDP  $\mathcal{M}(c) = (\mathcal{S}, \mathcal{A}, T, R^c, \beta^c)$ . Here,  $T : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$  is a transition function,  $R^c : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$  a reward function, and  $\beta^c$  an initial state distribution. This means that a CMDP is a set of MDPs that share the same states and actions. The policy learning

<sup>1</sup>We do not use the photographic NLVR2 in this work.

<sup>2</sup>We recommend reporting both development and held-out test results in future work for easy comparison.

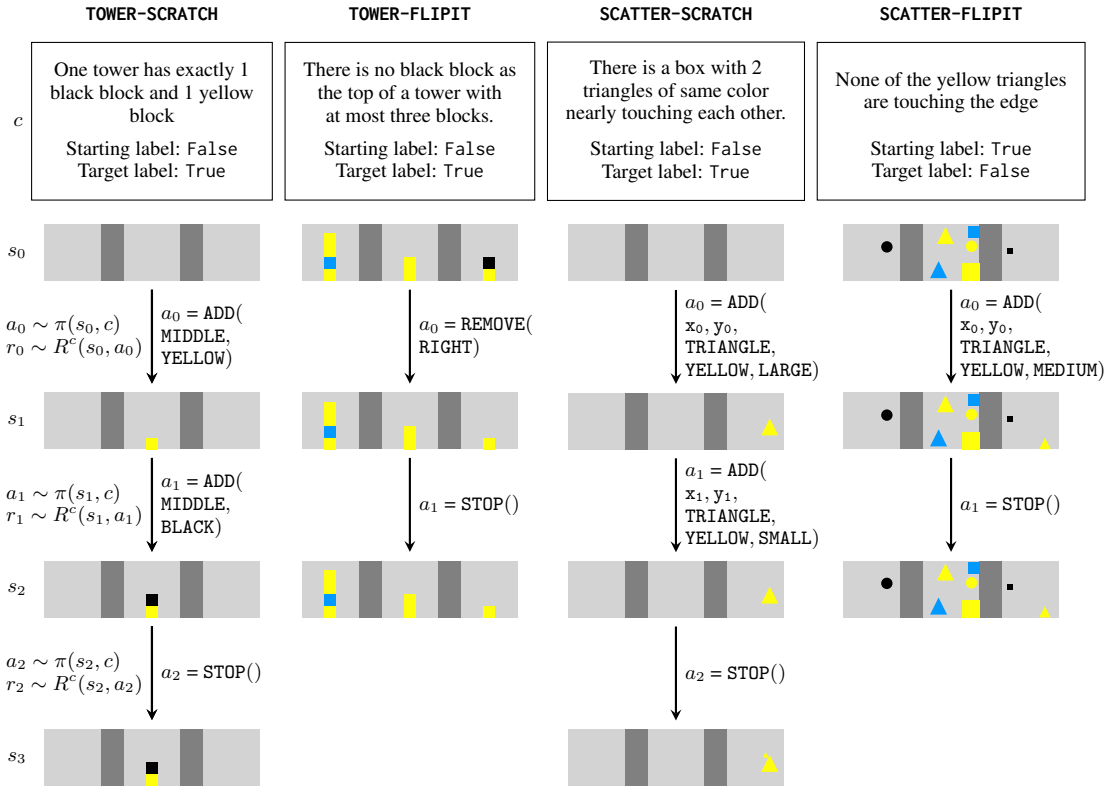


Figure 2: Examples from the four CMDP configurations. Each example is conditioned on a context  $c = (\bar{x}, b)$ , and starts with a state  $s_0$ , sampled from the initial state distribution  $\beta^c$ . For example, for TOWER-SCRATCH (left column), the context  $c$  pairs the statement *one tower has exactly 1 black block and 1 yellow block* with the target boolean True. The initial state  $s_0$  is an image with three empty grey box regions separated by darker grey separators. The agent  $\pi$  is given  $(s_0, c)$ , and samples an action  $a_0 \sim \pi(s_0, c)$ . The environment transitions to the next state  $s_1$ , while the context remains the same. This process continues until the agent selects the STOP action.

problem is to estimate parameters  $\theta$  of a policy  $\pi_\theta : \mathcal{S} \times \mathcal{C} \rightarrow \mathcal{A}$ , which maps the current state and the context underlying the MDP to an action. The policy must generalize across different contexts from  $\mathcal{C}$ . Figure 2 shows example action trajectories in MDPs for each of the four CMDPs. Table 1 shows the number of MDPs in each configuration.<sup>3</sup>

**Contexts  $\mathcal{C}$**  A context  $c \in \mathcal{C}$  is a pair  $c = (\bar{x}, b)$ , where  $\bar{x}$  is a natural language statement and  $b \in \{\text{True}, \text{False}\}$  is a target boolean value for the statement  $\bar{x}$  with respect to the state  $s$ . The set of statements is predefined for TOWER and SCATTER based on the NLVR data, but identical across the choice of SCRATCH and FLIPIT. The target boolean value in SCRATCH is always True. In FLIPIT, the target boolean value is either True or False. Depending on the context, different types of reasoning are required. For example, in the second column

of Figure 2, the statement *there is no black block as the top of a tower with at most three blocks* requires reasoning about negation, soft cardinality, color, and position, while the statement in the third column *there is a box with 2 triangles of same color nearly touching each other* requires a comparison and to reason about several object attributes (shape, color, position). Both require high-level relational reasoning about single objects or sets.

**States  $\mathcal{S}$**  A state  $s \in \mathcal{S}$  is an RGB image. Images in *lilGym* are divided into three box regions of identical dimensions by two dark gray separators (Figure 2). The objects in *lilGym* have three properties, each can take multiple values: shape (CIRCLE, SQUARE, or TRIANGLE), color (BLACK, BLUE, or YELLOW), and size (SMALL, MEDIUM, or LARGE). In TOWER, states are constrained to have stacks of up to four SQUAREs of MEDIUM size and any color at the center of each box. SCATTER states support all object shapes, sizes, and colors, and they may be positioned freely. In both conditions, objects cannot cross image boundaries or into the separators.

<sup>3</sup> NLVR includes 18,322 images. This allows further expanding the number of initial states to 92,179 initial states through box element permutations. We do not manipulate this property in this work, but future work could take advantage of it. Our reward computation is invariant to such permutations.

	TOWER-SCRATCH	TOWER-FLIPIT		SCATTER-SCRATCH	SCATTER-FLIPIT	
	MDPs	MDPs	Init.	MDPs	MDPs	Init.
Train	989	1,910	5,704	1,241	2,340	6,696
Dev	163	317	676	87	164	313
Test	324	619	1,383	155	285	591
Total	1,476	2,846	7,763	1,483	2,789	7,600

Table 1: Data statistics per CMDP configuration and data split. The number of MDPs corresponds to the number of contexts under each CMDP. For FLIPIT, “Init.” corresponds to the total number of initial states across all MDPs for this CMDP.<sup>3</sup> The number of initial states and MDPs is equal for SCRATCH CMDPs.

The choice of starting condition between SCRATCH or FLIPIT does not influence the state space.

**Actions  $\mathcal{A}$  and Transitions  $T$**  There are three action types STOP, ADD, and REMOVE. STOP terminates the episode. The truth-value of the statement is only evaluated and compared to the target boolean after the STOP action is taken. ADD adds objects to the environment, and REMOVE removes objects. ADD and REMOVE take arguments that differ between TOWER and SCATTER:

**TOWER:** Both ADD and REMOVE take a position argument, which has three possible values corresponding to the three box regions. Objects are added or removed at the top of the stack. Adding an object on top of a stack of four objects or removing an object from an empty box are both invalid actions. ADD also takes a color argument. For example, the first action on the left trajectory in Figure 2 is adding a yellow square in an empty box. Including STOP, there are  $1 + (3 + 1) \times 3 = 13$  actions.

**SCATTER:** Unlike TOWER, objects of any type can be placed freely in the box regions. Both ADD and REMOVE take 2D coordinates that specify pixel location. Adding an object places it so that its top-left coordinates are the given coordinates. Removing an object will remove the object at the given coordinates. Adding also requires specifying the shape, color, and size. The action is invalid if adding results in objects’ overlap or boundary crossing with the separators or image boundaries. Removing from a position that does not include an object is also an invalid action. The native resolution of images in *lilGym* is  $380 \times 100$  pixels. Including STOP, there are  $1 + (380 \times 100) \times ((3 \times 3 \times 3) + 1) = 1,064,001$  actions. Because of the extremely large action space, *lilGym* also supports a simplification through a coarser grid system for SCATTER that is automatically mapped to the original resolution (Appendix A). The grid simplification includes heuristics that assist in identifying locations

to place objects in the original pixel space or objects to remove once a grid cell is selected. In our experiments (Section 6), we use a grid simplification of  $19 \times 5$ , giving a total of 2,661 actions. The difficulty of SCATTER can be adjusted by modifying the grid size, or acting at the original resolution.

The transition function  $T : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$  depends on the choice between TOWER and SCATTER configurations, because this choice determines the action space. Similar to the action space, the transitions in TOWER are more constrained compared to SCATTER. The transition function does not modify the context, which is fixed for a given MDP.

**Reward Function  $R^c$**  The reward function  $R^c$  is computed with respect to the context  $c = (\bar{x}, b)$ , and is based on evaluating the truth-value of the natural language statement  $\bar{x}$  with respect to a state  $s$ , and comparing it to the target boolean  $b$ . *lilGym* includes an evaluation function  $\mathcal{E}^{\bar{x}} : \mathcal{S} \times \mathcal{A} \rightarrow \{\text{True}, \text{False}\}$  for every statement  $\bar{x}$ . Section 5 describes how we create the evaluation functions.

The agent receives a positive reward for terminating the episode using the STOP action if the evaluation  $\mathcal{E}^{\bar{x}}(s)$  is equal to the target boolean  $b$ . If  $\mathcal{E}^{\bar{x}}(s)$  does not equal  $b$  when taking the STOP action, the agent receives a negative reward. If the episode terminates because the current time step  $t$  reached the action horizon  $H$  or because of an invalid action, the agent also receives a negative reward. Action validity depends on the current state  $s$  and on the configuration, because TOWER and SCATTER have different action spaces. For example, in TOWER, adding an object to a box (e.g., `ADD(MIDDLE, BLUE)`) is only valid if the box has less than four objects, because towers have a maximum height of four. There is also a verbosity penalty of  $\delta$ . Formally, the reward is:

$$R^c(s, a) = \begin{cases} 1.0 & a = \text{STOP} \wedge \mathcal{E}^{\bar{x}}(s) = b \\ -1.0 & a = \text{STOP} \wedge \mathcal{E}^{\bar{x}}(s) \neq b \\ -1.0 & (a \text{ is invalid in } s) \vee (t = H) \\ -\delta & \text{otherwise} \end{cases} \quad (1)$$

**Initial State Distribution**  $\beta^c$  The initial state distribution  $\beta^c$  is parameterized by the context  $c \in \mathcal{C}$ , and differs between SCRATCH and FLIPIT. In SCRATCH, the agent modifies an empty environment to satisfy the truth-condition of the statement  $\bar{x}$  in the context  $c$ , so the initial state  $s_0$  is always an empty image. The set of initial states  $\beta^c$  for every context  $c \in \mathcal{C}$  is the set of images associated with the statement  $\bar{x}$  in the NLVR data. This set includes between 1 to 43 images. Table 1 shows the total number of initial states in each configuration.

## 5 The *lilGym* Data and Annotation

We use the NLVR data to create each of the CMDPs (Table 1). SCRATCH CMDPs include contexts for all natural language statements from NLVR, each paired with the empty initial state containing no shapes (Figure 2, left and center-right columns). FLIPIT CMDPs include the natural language statements with their corresponding images, both from NLVR. The images are used as initial states. The target boolean is set so that the initial state does not fulfil it. The split between TOWER and SCATTER also follows from NLVR. Statements corresponding to TOWER images in NLVR are included in our TOWER CMDPs, and the same for SCATTER sentences.

NLVR has four splits for training, development, public testing, and hidden testing. We adopt the original training and development sets splits. Following the recent public release of the hidden testing set, we merge the public and hidden testing sets into a single public test split.

The NLVR annotations include the truth-value of each statement with regard to the images paired with it in the data. Once we manipulate an image (i.e., change the state in our environment), the truth-value annotation does not necessarily hold. A key challenge for creating an interactive environment using this data is an accurate evaluation of the natural language statement for *every* possible state (i.e., image) for reward computation (Section 4). We address this by annotating each statement  $\bar{x}$  with an executable boolean Python program representing its meaning,  $\mathcal{E}^{\bar{x}}$  in Section 4. This process is inspired by data annotation for supervised semantic parsing (Zelle and Mooney, 1996; Zettlemoyer and

Collins, 2005; Suhr et al., 2018), where sentences are annotated with formal meaning representations.

The Python programs operate on the underlying structured representation. Each program returns True for every image that satisfies the constraints specified in the corresponding statement, and False otherwise. In general, there are many states that satisfy any given statement, many more than provided with the original NLVR images.

The programs are written using an API defined over the structured representations. We base the API design on the ontology designed for NLVR’s structured representations by Goldman et al. (2018), which we extend to include 66 functions. Figure 7 in Appendix B shows two example programs with their corresponding statements.

We use the freelancing platform Upwork<sup>4</sup> for annotation. We recruit three programmers based on preliminary screening of their fluency in English and competency in Python. We de-duplicate the naturally occurring sentences in the data, and distribute sentences to annotators randomly, each with a single example NLVR image. Each program is evaluated against a corresponding hidden validation set made of all remaining NLVR images paired with the sentence, and must pass all the tests. Appendix B provides a screenshot of the interface and more details. We collect 2,666 annotations at a total cost of \$3,756, and keep 2,661 valid annotations.

## 6 Experiments

### 6.1 Methods

We experiment with each of the four CMDPs separately, training on the training split and testing on the development and test splits. We sample a validation set from the training split for model selection. For SCATTER we use a simplified grid action space of  $19 \times 5$  (Section 4). Each grid cell is  $20 \times 20$  pixels. We set the action horizon  $H = 12$ . Appendix C provides implementation details.

We use PPO (Schulman et al., 2017) for parameter estimation,<sup>5</sup> with a separate network as a critic. The critic network is identical to the policy, except that we add a tanh activation for the value output. Because of the large action space, especially for SCATTER, the agent rarely observes positive reward, which requires taking a STOP action at an appropriate state. We design a simple variant of PPO called

<sup>4</sup><https://www.upwork.com>

<sup>5</sup>We use the PPO implementation of (Kostrikov, 2018).

PPO+SF (PPO with stop forcing) to study this issue. PPO+SF is identical to PPO, except that during training, we mask all actions except STOP when the agent reaches a state where selecting STOP will give a positive reward. This modification is present only during training. All testing is done under the same conditions, without stop forcing.

We also study the importance of our annotation for reward computation using a reward function for SCRATCH that does not require any annotation beyond what is already in NLVR. The reward uses NLVR images that are associated with each statement (between 1–43) and are labeled with the target boolean.<sup>6</sup> Instead of testing the state using a program, it compares the state to the available NLVR images, and only if it equals one of them, the learner receives a positive task completion reward.

We experiment with three models: C3+BERT, C10+BERT and ViLT.<sup>7</sup> In C3+BERT and C10+BERT, we process the statement  $\bar{x}$  using BERT (Devlin et al., 2019), and do mean pooling across all layers and tokens to get the statement representation. We use a three-layer CNN (Fukushima and Miyake, 1982) in C3+BERT to embed the image of the current state  $s$ , and a ten-layer CNN in C10+BERT. We concatenate the statement and image representations with an embedding for the target boolean  $b$ , and use a multi-layer perceptron to compute the action distribution. ViLT is a pre-trained multi-modal Transformer that jointly processes text and image inputs (Kim et al., 2021). We create a sequence of tokens by concatenating the statement, a token for the target boolean, and image patches, separated by special tokens. The image patches are the same size as the  $19 \times 5$  grid cells, including in TOWER, where the action space does not use a grid.

## 6.2 Results and Analysis

Table 2 shows task-completion accuracies for all CMDPs, and Figure 3 shows reward statistics. We observe only minor differences between C3+BERT and C10+BERT, so conduct the bulk of our analysis on C3+BERT and ViLT. Figure 4 plots training curves for SCRATCH.<sup>8</sup> Figure 5 breaks down development set accuracies for FLIPIT CMDPs by the

<sup>6</sup>We dropped 1% of MDPs without a single target state.

<sup>7</sup>Appendix E describes preliminary experiments with one more model.

<sup>8</sup>We plot the training curves with no regard to the patience stopping criteria we use for model selection. This better reflects long term trends.

target boolean, and Figure 6 shows development rollout statistics for PPO.<sup>9</sup> We sample 50 development examples for each CMDP, and annotate them with expert<sup>10</sup> trajectories to estimate the expert reward and rollout statistics. All expert rollouts are successful.

Overall, we observe stronger task-completion performance (Table 2) on TOWER CMDPs compared to SCATTER, especially with ViLT, which shows stronger performance than C3+BERT and C10+BERT in most cases. The development rewards (Figure 3) and training curves (Figure 4) show similar trends. The training curve comparison to the alternative reward that uses NLVR images instead of our program annotations shows no effective learning. This illustrates the importance of exact reward computation, such as possible with our program annotations. The comparison to the estimate of expert rewards shows there remains significant room for improvement across the board. Even when the learned policies are able to complete the task, they are doing it inefficiently, so much so that the mean rewards are often negative. The mean rewards of the random baseline policy illustrate the task is far from trivial. Both task accuracies and reward statistics indicate TOWER-SCRATCH is the easiest of the CMDPs, and SCATTER-FLIPIT is the hardest.

The additional guidance of PPO+SF compared to PPO helps with exploration, especially on SCATTER CMDPs. On SCATTER-FLIPIT, PPO+SF improves performance by 13.25% compared to PPO. This illustrates the exploration challenges SCATTER CMDPs pose.

ViLT generally outperforms C3+BERT and C10+BERT, except on SCATTER-SCRATCH, where the ViLT policy more often selects invalid actions. ViLT general advantage is expected given the joint reasoning architecture and multi-modal pre-training of ViLT. FLIPIT policies generally do better on examples with a False target boolean, except when learning fails (Figure 5). The other direction is harder, because the set of states that invalidates a statement is usually larger than the set that validates it, and it generally requires fewer actions to invalidate a statement.

We observe more rollouts that are terminated either by reaching the action horizon  $H$  or by taking

<sup>9</sup>Appendix D provides further analysis, including error analysis, performance breakdown by semantic phenomena, and rollout statistics for PPO+SF.

<sup>10</sup>The expert is an author of this paper.

		TOWER-SCRATCH		TOWER-FLIPIT		SCATTER-SCRATCH		SCATTER-FLIPIT	
		Dev	Test	Dev	Test	Dev	Test	Dev	Test
PPO	C3+BERT	72.80±9.92	68.72±6.60	28.11±6.03	27.84±4.50	<b>59.00±5.42</b>	<b>68.17±2.89</b>	0.00±0.00	0.06±0.10
	C10+BERT	80.16±1.77	73.66±3.04	30.13±8.22	28.75±7.42	43.68±7.18	50.97±12.95	0.00±0.00	0.00±0.00
	ViLT	<b>81.19±2.90</b>	<b>76.23±5.58</b>	<b>53.25±4.28</b>	<b>55.19±4.92</b>	40.23±6.99	47.74±11.47	<b>13.31±6.88</b>	<b>16.81±7.43</b>
PPO+SF	C3+BERT	81.80±0.94	76.54±1.85	32.59±3.14	29.55±4.74	72.03±3.32	74.41±3.25	17.04±1.82	18.16±3.07
	C10+BERT	77.91±2.21	75.62±1.41	37.38±2.02	35.70±2.08	<b>73.18±2.65</b>	<b>77.85±1.62</b>	8.52±2.72	10.55±2.45
	ViLT	<b>84.05±3.25</b>	<b>81.48±1.93</b>	<b>65.68±9.17</b>	<b>65.51±8.43</b>	67.43±1.35	73.98±0.30	<b>28.01±5.32</b>	<b>30.06±4.51</b>

Table 2: Mean task-completion accuracy and standard deviation computed over three runs for all four CMDPs.

an invalid action (i.e., without STOP) on SCATTER CMDPs compared to TOWER (Figure 6, upper left). This difference is partially explained by a higher rate of invalid actions in SCATTER (Figure 6, upper right), which cause immediate rollout termination. On SCATTER CMDPs, where we have a higher rate of invalid actions, the type of invalid actions we most often observe for both models are actions hitting one of the separators, except when learning fails completely. This is related to the action selection bias of the models, which tend to select some coordinates more often than others. Section D.2 provides further error analysis for both models for SCATTER CMDPs trained with PPO, and Section D.3 illustrates the action selection bias.

There is no consistent difference in the length of rollouts between the two models (Figure 6, bottom left). Expert trajectory length is similar on TOWER, where models perform fairly well. However, on SCATTER, where our models are weaker, expert trajectories are significantly longer. This is partially explained by the models not learning to effectively avoid invalid actions, which terminate the execution immediately. Using REMOVE actions is generally difficult for the learned policies. TOWER-FLIPIT is an exception with REMOVE dominating the rollouts (Figure 6, bottom right), potentially because removing objects generally provides a more efficient path to flip the boolean value. While PPO policies generate REMOVE actions for SCATTER-FLIPIT, the extremely low performance indicates that these actions are not used effectively. Expert statistics indicate that REMOVE actions are beneficial for FLIPIT CMDPs.

We also performed semantic and syntactic analyses using the 200 development examples manually annotated by Suhr et al. (2017). Table 3 shows the performance on this data of policies trained with PPO. We only include categories with more than 10 instances across all CMDPs. Appendix D.4 pro-

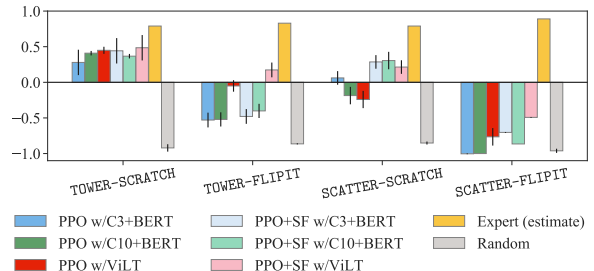


Figure 3: Mean development set rewards, averaged over three runs.

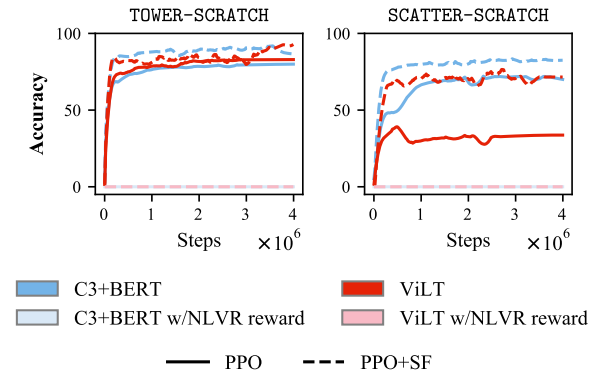


Figure 4: Training curves for SCRATCH, averaged over three runs. All NLVR reward curves are superposed because accuracy remains zero throughout training.<sup>8</sup>

vides the complete tables with examples, including for PPO+SF. The two models mostly follow similar trends with respect to the categories on which they perform above and below overall performance. Both models perform better than they do overall on hard cardinality (e.g., ... *exactly four objects* ...) for TOWER CMDPs, and on presupposition for SCRATCH CMDPs. However, on spatial relations, both models perform below overall performance for all CMDPs except TOWER-SCRATCH.

## 7 Conclusion

We introduce *lilGym*, an RL benchmark that focuses on natural language visual reasoning. *lilGym* is designed to be accessible for researchers, while



	TOWER-SCRATCH		TOWER-FLIPIT		SCATTER-SCRATCH		SCATTER-FLIPIT	
	Total	Correct %	Total	Correct %	Total	Correct %	Total	Correct %
Cardinality (hard)	98	<b>76.5</b> <b>83.7</b>	480	<b>28.9</b> <b>56.2</b>	35	49.5   40.0	119	0.0   <b>13.7</b>
Cardinality (soft)	21	68.2   81.0	82	<b>30.1</b> <b>56.5</b>	11	<b>60.6</b>   24.3	42	0.0   10.3
Existential	122	<b>75.1</b> <b>81.7</b>	577	<b>29.0</b>   55.1	55	<b>55.7</b>   35.1	192	0.0   12.3
Coordination	19	<b>85.9</b> <b>84.2</b>	86	27.1   52.7	15	<b>64.5</b>   40.0	55	0.0   9.7
Spatial Relations	94	<b>74.8</b> <b>81.6</b>	438	26.6   53.1	39	53.8   32.5	128	0.0   10.1
Presupposition	17	<b>74.5</b> <b>90.2</b>	74	27.0   <b>54.1</b>	22	<b>66.7</b> <b>51.5</b>	78	0.0   12.4
<b>Overall</b>		72.80   81.19		28.11   53.25		59.00   40.23		0.00   13.31

Table 3: Performance on a set of development examples annotated for semantic categories by Suhr et al. (2017) for both models (C3+BERT | ViLT) when trained with PPO. Development performance refers to mean performance on the respective full development set. Results outperforming dev performance are in bold.




Figure 5: Mean development set accuracies for FLIPIT CMDPs, averaged over three runs, reported according to the value of the context target boolean (Red for True, Gray for False). Dashed gray line: full development set accuracies.




Figure 6: Mean development statistics for PPO, averaged over three runs. Clockwise from top left: rollouts without a STOP action; rollouts with an invalid action; mean actions per rollout; relative rate of ADD/REMOVE actions.

still displaying the reasoning richness of natural language. It is relatively easy to deploy using the standard Gymnasium API (Brockman et al., 2016),

and has light compute requirements. Our data annotation approach allows including expressive and diverse natural language, while still providing accurate and automatic reward computation. It also exposes the potential connection between semantic parsing and reward evaluation in RL, thereby outlining how strong semantic parsers can benefit RL benchmarking. Our strong baselines illustrate the range of challenges *lilGym* presents, showing that existing methods can achieve non-trivial performance, but that there remain significant progress to be made. Our analysis lays out the framework for studying and reporting these future results.

*lilGym* has significant potential beyond the tasks we study. It can be used without the language, to create thousands of micro RL tasks requiring set and relational visual reasoning. Our annotations form a new semantic parsing corpus with annotated executable meaning representations. The semantic diversity of the data, its executability, and the focus on visual reasoning make it a unique asset in the landscape of corpora for semantic parsing. *lilGym* is also promising for program synthesis guided by natural language (Wong et al., 2021).

## 8 Limitations

*lilGym* uses synthetic visual stimuli, which does not reflect the complexity or characteristics of realistic visual observations. This is critical for our ability to control the environment and provide a lightweight and accessible RL benchmark. Our goal is not to provide a resource for the development of methods that aim to handle realistic visual input, and *lilGym* is not suitable for this purpose. The limited number of colors, shapes, and sizes used limits the visual and lexical complexity of the data. The synthetic nature of the data and the

modular library of functions we use allow to relatively easily extend the environment (e.g., with new colors). This will require collecting additional natural language data. In this work, we opted to rely on the NLVR data without further expanding it. Some annotators of the original NLVR data adopted annotation strategies that led to repetition of some common phrases (e.g., starting statements with *there is*). While this creates some implicit patterns in the data, [Suhr et al. \(2017\)](#) showed that NLVR demonstrates high semantic diversity and compositionality. Finally, *lilGym* includes English data only. Expanding this data to other language is an important direction for future work. Translating the data is a feasible low-cost solution, because the program annotations will not require updating.

## Ethics Statement

We paid U.S. standard market wage to our programmers ([Appendix B](#)). The rate was determined by the workers. The *lilGym* environment and data as is are intended to be used for research, including algorithm development and evaluation, and not for development of models to be deployed.

Commented for anonymous submission

## Acknowledgements

This research was supported by ARO W911NF21-1-0106, NSF under grant No. 1750499, and a gift from Open Philanthropy. KB is supported by NSF under grant No. 2127309 to the Computing Research Association for the CIFellows Project. Results presented in this paper were obtained using CloudBank ([Norman et al., 2021](#)), which is supported by the National Science Foundation under award No. 1925001. We thank Alane Suhr, Ge Gao, Justin Chiu, Woojeong Kim, Jack Morris, Jacob Sharf and the Cornell NLP Group for support, comments, and helpful discussions.

## References

Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, and Anton van den Hengel. 2018. Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In [The IEEE Conference on Computer Vision and Pattern Recognition](#), pages 3674–3683.

Marc G Bellemare, Yavar Naddaf, Joel Veness, and Michael Bowling. 2013. The arcade learning environment: An evaluation platform for general agents.

[Journal of Artificial Intelligence Research](#), 47:253–279.

Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. 2016. Openai gym. [arXiv preprint arXiv:1606.01540](#).

Tianshi Cao, Jingkang Wang, Yining Zhang, and Sivabalan Manivasagam. 2020. Babyai++: Towards grounded-language learning beyond memorization. [Beyond tabula rasa in RL \(BeTR-RL\) Workshop held in conjunction with the 8th International Conference on Learning Representations, ICLR](#).

Howard Chen, Alane Suhr, Dipendra Misra, Noah Snaveley, and Yoav Artzi. 2019. Touchdown: Natural language navigation and spatial reasoning in visual street environments. In [Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition](#), pages 12538–12547.

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. In [7th International Conference on Learning Representations, ICLR](#).

John D. Co-Reyes, Abhishek Gupta, Suvansh Sanjeev, Nick Altieri, John DeNero, P. Abbeel, and Sergey Levine. 2019. Guiding policies with language via meta-learning. In [7th International Conference on Learning Representations, ICLR](#).

Marc-Alexandre Côté, Ákos Kádár, Xingdi Yuan, Ben Kybartas, Tavian Barnes, Emery Fine, James Moore, Matthew J. Hausknecht, Layla El Asri, Mahmoud Adada, Wendy Tay, and Adam Trischler. 2018. Textworld: A learning environment for text-based games. In [Computer Games Workshop \(CGW\) held in conjunction with the 27th International Conference on Artificial Intelligence, IJCAI](#).

Pradeep Dasigi, Matt Gardner, Shikhar Murty, Luke Zettlemoyer, and Eduard Hovy. 2019. Iterative search for weakly supervised semantic parsing. In [Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 \(Long and Short Papers\)](#), pages 2669–2680.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In [Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 \(Long and Short Papers\)](#), pages 4171–4186.

Kunihiko Fukushima and Sei Miyake. 1982. Neocognitron: A self-organizing neural network model for a mechanism of visual pattern recognition. In

- Competition and Cooperation in Neural Nets, pages 267–285. Springer.
- Omer Goldman, Veronica Latcinnik, Ehud Nave, Amir Globerson, and Jonathan Berant. 2018. Weakly supervised semantic parsing with abstract examples. In Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 1809–1819.
- Nitish Gupta, Sameer Singh, and Matt Gardner. 2021. Enforcing consistency in weakly supervised semantic parsing. In Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers), pages 168–174.
- Assaf Hallak, Dotan Di Castro, and Shie Mannor. 2015. Contextual markov decision processes. arXiv preprint arXiv:1502.02259.
- Austin W Hanjie, Victor Y Zhong, and Karthik Narasimhan. 2021. Grounding language to entities and dynamics for generalization in reinforcement learning. In International Conference on Machine Learning, pages 4051–4062.
- Matthew J. Hausknecht, Prithviraj Ammanabrolu, Marc-Alexandre Côté, and Xingdi Yuan. 2020. Interactive fiction games: A colossal adventure. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 34.
- 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. arXiv preprint arXiv:1706.06551.
- Drew A. Hudson and Christopher D. Manning. 2018. Compositional attention networks for machine reasoning. In 6th International Conference on Learning Representations, ICLR.
- Minqi Jiang, Jelena Luketina, Nantas Nardelli, Pasquale Minervini, Philip HS Torr, Shimon Whiteson, and Tim Rocktäschel. 2020. Wordcraft: An environment for benchmarking commonsense agents. arXiv preprint arXiv:2007.09185.
- Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Judy Hoffman, Li Fei-Fei, C. Lawrence Zitnick, and Ross B. Girshick. 2017. Inferring and executing programs for visual reasoning. 2017 IEEE International Conference on Computer Vision (ICCV), pages 3008–3017.
- Wonjae Kim, Bokyung Son, and Ildoo Kim. 2021. Vilt: Vision-and-language transformer without convolution or region supervision. In Proceedings of the 38th International Conference on Machine Learning, pages 5583–5594.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In 3rd International Conference on Learning Representations, ICLR.
- Ilya Kostrikov. 2018. Pytorch implementations of reinforcement learning algorithms.
- Alexander Ku, Peter Anderson, Roma Patel, Eugene Ie, and Jason Baldridge. 2020. Room-across-room: Multilingual vision-and-language navigation with dense spatiotemporal grounding. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), pages 4392–4412.
- Ilya Loshchilov and Frank Hutter. 2019. Decoupled weight decay regularization. In 7th International Conference on Learning Representations, ICLR.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. 2013. Playing atari with deep reinforcement learning. arXiv preprint arXiv:1312.5602.
- Karthik Narasimhan, Tejas D. Kulkarni, and Regina Barzilay. 2015. Language understanding for text-based games using deep reinforcement learning. In Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP, pages 1–11.
- Michael Norman, Vince Kellen, Shava Smallen, Brian DeMeulle, Shawn Strande, Ed Lazowska, Naomi Alterman, Rob Fatland, Sarah Stone, Amanda Tan, Katherine Yelick, Eric Van Dusen, and James Mitchell. 2021. Cloudbank: Managed services to simplify cloud access for computer science research and education. In Practice and Experience in Advanced Research Computing, PEARC '21. Association for Computing Machinery.
- Juan Pavez, Héctor Allende, and Héctor Allende-Cid. 2018. Working memory networks: Augmenting memory networks with a relational reasoning module. In Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 1000–1009.
- Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. 2018. Film: Visual reasoning with a general conditioning layer. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 32.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In International Conference on Machine Learning, pages 8748–8763.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347.

- Mohit Shridhar, Jesse Thomason, Daniel Gordon, Yonatan Bisk, Winson Han, Roozbeh Mottaghi, Luke Zettlemoyer, and Dieter Fox. 2020. ALFRED: A benchmark for interpreting grounded instructions for everyday tasks. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR, pages 10740–10749.
- Mohit Shridhar, Xingdi Yuan, Marc-Alexandre Côté, Yonatan Bisk, Adam Trischler, and Matthew J. Hausknecht. 2021. Alfworld: Aligning text and embodied environments for interactive learning. In 9th International Conference on Learning Representations, ICLR.
- Amanpreet Singh, Ronghang Hu, Vedanuj Goswami, Guillaume Couairon, Wojciech Galuba, Marcus Rohrbach, and Douwe Kiela. 2022. Flava: A foundational language and vision alignment model. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 15638–15650.
- Alane Suhr, Srinivasan Iyer, and Yoav Artzi. 2018. Learning to map context-dependent sentences to executable formal queries. In Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pages 2238–2249.
- Alane Suhr, Mike Lewis, James Yeh, and Yoav Artzi. 2017. A corpus of natural language for visual reasoning. In Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), pages 217–223.
- Alane Suhr, Stephanie Zhou, Ally Zhang, Iris Zhang, Huajun Bai, and Yoav Artzi. 2019. A corpus for reasoning about natural language grounded in photographs. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, pages 6418–6428.
- Hao Tan and Mohit Bansal. 2018. Object ordering with bidirectional matchings for visual reasoning. In Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers), pages 444–451.
- Emanuel Todorov, Tom Erez, and Yuval Tassa. 2012. Mujoco: A physics engine for model-based control. In 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 5026–5033.
- Jack Urbanek, Angela Fan, Siddharth Karamcheti, Saachi Jain, Samuel Humeau, Emily Dinan, Tim Rocktäschel, Douwe Kiela, Arthur Szlam, and Jason Weston. 2019. Learning to speak and act in a fantasy text adventure game. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, pages 673–683.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. Transformers: State-of-the-art natural language processing. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, pages 38–45.
- Catherine Wong, Kevin Ellis, Joshua B. Tenenbaum, and Jacob Andreas. 2021. Leveraging language to learn program abstractions and search heuristics. In Proceedings of the 38th International Conference on Machine Learning, ICML, pages 11193–11204.
- Yiqun Yao, Jiaming Xu, Feng Wang, and Bo Xu. 2018. Cascaded mutual modulation for visual reasoning. In Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, pages 975–980.
- J.M. Zelle and Raymond J. Mooney. 1996. Learning to parse database queries using inductive logic programming. In Proceedings of the National Conference on Artificial Intelligence.
- Luke S. Zettlemoyer and Michael Collins. 2005. Learning to map sentences to logical form: Structured classification with probabilistic categorial grammars. In Proceedings of the Conference on Uncertainty in Artificial Intelligence.
- Wenbo Zheng, Lan Yan, Chao Gou, and Fei-Yue Wang. 2020. Webly supervised knowledge embedding model for visual reasoning. 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 12442–12451.

## A SCATTER Grid Simplification

To reduce the large action space of SCATTER, *lilGym* allows to simplify the pixel-based action space with a grid that is coarser than the image resolution of  $380 \times 100$ . The actions applied in the environment remain in the original resolution, and the translation between the grid system to pixels is done heuristically. Without the heuristics, the transition to a grid coarser than the original image resolution would render many of the MPDs unsolvable.


The heuristics simplify two translation problems: in what pixel exactly to place an object and which object to remove from a grid cell. Depending on the grid size, it is possible to add multiple objects in a cell. To find the exact pixel within a cell to add an object, we search for a pixel in the grid box where we can add the object starting from the upper left corner. We can add an object in a pixel if the object fits there without overlapping with other objects, the image boundaries, or the columns. We also snap objects to touch each other if the distance between them is below a threshold. This is to allow adding objects that *touch* each other, a common constraint in *lilGym* statements. When removing an object from a grid cell, we remove the object with largest overlap with the cell.

## B Natural Language Annotation Details


We annotate each natural language statement in the NLVR corpus with a Python program representing its meaning. The programs return a boolean value, and are executable given the structured representation underlying each image. Figure 7 shows two examples of text statements with their annotated Python programs.

We provide the annotators with a web-based annotation interface (Figure 8), a tutorial, and an application programming interface (API) presenting a set of functions, classes and objects that they can use for annotation. We ask the annotators to prioritize the faithfulness of the program to the natural language sentence and to prefer shorter annotations. We also provide them with examples of spurious logical forms and ask them to avoid such expressions. Annotators can raise questions.

Figure 8 shows the annotation interface for a single sentence. For every sentence, annotators are provided with a single example image from NLVR and an associated boolean value. Other images for the same statement from NLVR are used as



```
exist(filter_obj(
  all_boxes, lambda x: x.is_tower() and
  exist(filter_obj(
    all_boxes, lambda y:
      y.is_tower() and
      count(x.all_items_in_box()) ==
      count(y.all_items_in_box()) and
      get_set_colors(filter_obj(
        y.all_items_in_box(),
        is_bottom)) !=
      get_set_colors(filter_obj(
        x.all_items_in_box(),
        is_bottom))))))
```



```
exist(filter_obj(
  all_boxes, lambda x:
    count(get_set_colors(
      x.all_items_in_box())) == 3 and
  exist(filter_obj(
    x.all_items_in_box(), lambda y:
      is_black(y) and
      is_triangle(y) and
      is_touching_wall(y, Side.TOP))))))
```

Figure 7: Example sentences with the example images displayed alongside them during annotation (left), and their annotated Python programs (right). Both sentences and programs are True for the corresponding image.

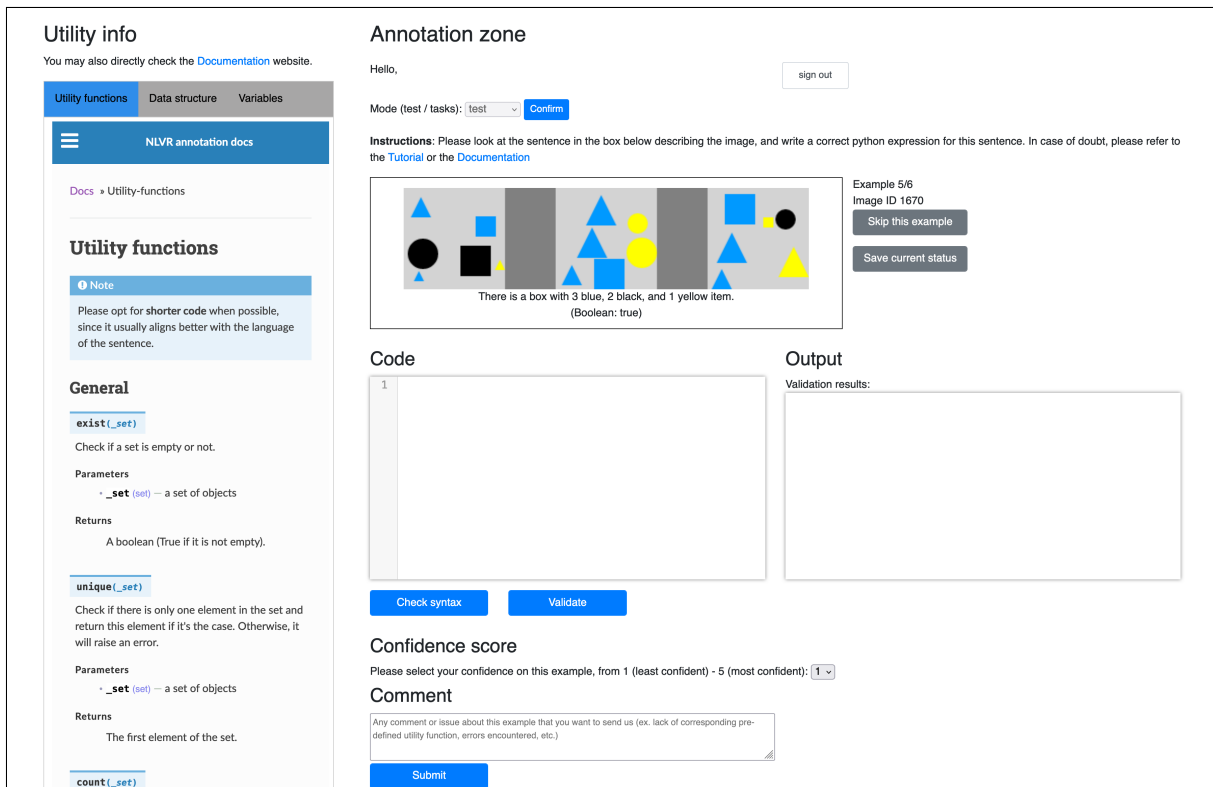


Figure 8: The annotation interface for collecting the Python program annotations in *lilGym*.

hidden validation examples. The annotator never sees these images.

The annotator can validate the program syntax and validate it within the browser. The validation executes the program against the given image and all hidden images. Validation passes only once the program returns the expected boolean value for all examples, including the visible and the hidden ones. The annotator can only submit their annotation after passing the syntax check and validation. They can assign a confidence score to their annotation and provide a comment.

Annotators can skip examples in case of doubt. When skipping, they need to explicitly provide the reason. We assess the annotations by batch, then randomly redistribute the skipped examples or examples with problematic annotations to the annotators after the questions have been solved. We iteratively communicate with the workers throughout the entire annotation process.

The annotation was done by four workers, one each from Croatia, India, Ukraine and United States. The hourly rate was roughly \$23.25 per hour. We communicated to the workers the purpose of the data collection and how data will be used at recruiting time.

## C Experimental Setup Details

### C.1 Learning Details

**Model Parameters and Computational Resources** C3+BERT and C10+BERT use a BERT-base model with 110M parameters. For ViLT, we use a ViLT-B/32 model with 87.4M parameters (Table 6 in Kim et al. (2021)). We use 6 NVIDIA RTX A6000, 3 Titan RTX, and 8 GeForce GTX 2080 Ti for our computations. The total computational budget is 950 GPU hours.

**Tokenization** C3+BERT and C10+BERT use an uncased BERT WordPiece tokenizer with the default parameters. ViLT uses the default ViLT feature extractor and BERT tokenizer, based on the Hugging Face implementation (Wolf et al., 2020).

**Hyperparameters** For C3+BERT and C10+BERT, we optimize using Adam (Kingma and Ba, 2015) with a learning rate of  $3e-4$ , except on TOWER-FLIPIT and on SCATTER-FLIPIT, where we use  $3e-5$ . For ViLT, we use AdamW (Loshchilov and Hutter, 2019) with a cosine scheduler and a base learning rate of  $3e-5$  for all experiments. The learning rate is warmed up for 1% of the maximal total training steps of 4M. We use patience for early stopping. We set

entropy to 0.1 for all our TOWER experiments and to 0.3 for all our SCATTER experiments. We use a mini-batch of 64 actions for gradient updates. At each PPO iteration we sample 2,048 actions (i.e., for the internal update loop).

**PPO+SF Details** PPO+SF is a simple variant of PPO that applies masking to all the actions except for STOP when the agent reaches a state in which it will receive a positive reward if it would select STOP. PPO+SF allows the learner to observe STOP with positive reward with higher probability than with conventional PPO. A side effect of this masking is that the learner often samples action with very low probability, which can lead to exploding gradients. We clip the PPO ratio to address this. Formally, the original PPO objective is:

$$L(\theta) = \mathbb{E}_t \left[ \min(r_t(\theta)\hat{A}_t, \text{clip}(r_t(\theta), 1 - \epsilon, 1 + \epsilon)\hat{A}_t) \right], \quad (2)$$

where  $r_t(\theta) = \frac{\pi_\theta(a_t|s_t)}{\pi_{old}(a_t|s_t)}$ ,  $\hat{A}$  is the advantage function, and  $\epsilon$  is a hyperparameter (Schulman et al., 2017). In PPO+SF, we clip the ratio term  $r_t(\theta)$  to avoid very large value due to “force” sampling of actions with very low probability:

$$\hat{r}_t(\theta) = \min(r_t(\theta), M), \quad (3)$$

where  $M$  is a threshold bounding the ratio. We use  $\hat{r}_t(\theta)$  in place of  $r_t(\theta)$  for our experiments.

## C.2 Inference Details

There are three action types STOP, ADD, and REMOVE. Each type take a different number of arguments: STOP takes no arguments, ADD takes two arguments in TOWER and five in SCATTER, and REMOVE takes one argument in TOWER and two in SCATTER. During inference, actions  $a$  are sampled from the agent policy  $a \sim \pi(\cdot|s, c)$ , where  $s$  is a state and  $c$  is a context. We decompose the probability of an action to be a product of its type and arguments. This risks assigning generally lower probability to actions with more arguments, because of the multiplicative decomposition. We avoid this by sampling the required arguments as needed. We first sample an action type. Depending on the action type, we sample the required arguments. In practice, this means that when an argument slot is not used, the probability of that action marginalizes over all possible assignments to that argument slot.

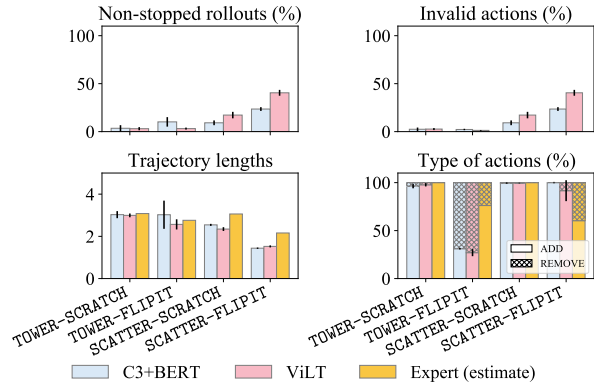


Figure 9: Mean development statistics for PPO+SF, averaged over three runs. Clockwise from top left: rollouts without a STOP action; rollouts with an invalid action; mean actions per rollout; relative rate of ADD/REMOVE actions.

## D Additional Results and Analysis

### D.1 Development Rollout Statistics

Figure 9 shows development rollout statistics for PPO+SF. The statistics follow similar trends for the ones we show for PPO in Figure 6. Compared to PPO, we observe more non-stopped rollouts for TOWER-FLIPIT when training with PPO+SF, and less for SCATTER. These non-stopped TOWER-FLIPIT rollouts often correspond to the model getting stuck in add-remove loops.

### D.2 Error Analysis

We analyze model errors by sampling 50 erroneous development examples,<sup>11</sup> for the two SCATTER CMDPs trained with PPO, over one run:

**SCATTER-SCRATCH with C3+BERT** 58% of the errors are due to invalid actions, and 42% due to direct or early termination. Among the invalid actions, all are due to trying to perform an action on a separator. Among the termination errors, 18% are due to direct termination, and 82% are due to early termination.

**SCATTER-SCRATCH with ViLT** 82% of the errors are due to invalid actions, and 18% due to direct or early termination. Among the invalid actions, 78% are due to trying to perform an action on a separator, 14% due to trying to remove an object from a position that does not include an object, 5% due to trying to put an item that cannot fit in

<sup>11</sup>If there are less than 50 errors in the development set, we analyze the entire set. This occurs only in SCATTER-SCRATCH with C3+BERT.

the box, and 3% due to trying to add an object on top of an existing one. Among the termination errors, 50% are due to direct termination and 50% due to erroneous termination.

**SCATTER-FLIPIT with C3+BERT** 58% of the errors are due to invalid actions, and 42% are due to direct or early termination. Among the invalid actions, 63% are due to trying to remove an object from a position that does not include an object, 24% are due to trying to perform an action on a separator, 10% due to trying to put an item that cannot fit in the box, and 3% due to trying to add an object on top of an existing one. Among the termination errors, 90% are due to direct termination and 10% due to erroneous termination.

**SCATTER-FLIPIT with ViLT** 64% of the mistakes are due to invalid actions, and 36% due to early termination. Among the invalid actions, 75% are due to trying to perform an action on a separator, 19% due to trying to remove an object from a position that does not include an object, and 6% due to trying to add an object on top of an existing one.

### D.3 Analysis of Action Selection Bias

We observe that the trained models often exhibit bias towards specific action arguments, which are sampled much more often than others during inference. Figure 10 illustrates this by visualizing coordinate selection frequencies on the development set for SCATTER CMDPs, for one of the runs. While the presence of bias is relatively persistent, the exact argument the models are biased towards vary. This indicates generalization limitations of our learned policies, which potentially converge to specific argument prematurely, and do not fully utilize the entire action space. We observe that this bias leads to selecting invalid actions, for example when attempting to place a large object on the edge so it crosses image boundaries.

### D.4 Performance Analysis by Semantic and Syntactic Phenomena

Suhr et al. (2017) manually annotated 200 development examples for semantic phenomena. Table 5 and Table 6 show the performance on this data of policies trained with PPO and PPO+SF. We provide an example sentence for each category. The two models mostly follow similar trends with respect to the categories on which they perform




Figure 10: Frequency of selecting the grid cell  $(x, y)$  in SCATTER CMDPs for rollouts sampled on the development set, over one run.  $x$ -axis and  $y$ -axis show the  $x$  and  $y$  position of the cells in the SCATTER  $19 \times 5$  grid approximation.

above and below overall performance. The two models mostly follow similar trends with respect to the categories on which they perform above and below overall performance. When trained with PPO, both models outperform overall performance on hard cardinality (e.g., ... *exactly four objects* ...) for TOWER CMDPs, and on presupposition for SCRATCH CMDPs. On spatial relations, both models perform above overall performance only for TOWER-SCRATCH, and below for all the other three CMDPs. We observe that PPO+SF is especially helpful for this category, bringing the performance of ViLT above average performance on all CMDPs.

## E Experiments with FLAVA

We conduct preliminary experiments with the base FLAVA model (350M parameters) (Singh et al., 2022).<sup>12</sup> Table 4 shows the results. On TOWER-FLIPIT, the results with PPO are outperforming ViLT in Table 2. On TOWER-SCRATCH, with PPO+SF, FLAVA’s results are on par with ViLT, and with PPO, below ViLT. On SCATTER environments, FLAVA’s performance is significantly lower than C3+BERT, C10+BERT and ViLT in Table 2. We tested different hyperparameters, using learning rates from  $1e-3$  to  $3e-6$ , but did not find a combination that significantly improves the learning behaviour. Due to the computational resources required in training FLAVA, and the results on TOWER environments that are comparable but not always

<sup>12</sup>We also experimented with CLIP (ViT-B/32) (Radford et al., 2021), but it performed poorly on the simplest TOWER-SCRATCH CMDP, so was discarded relatively early.



	TOWER-SCRATCH		TOWER-FLIPIT		SCATTER-SCRATCH		SCATTER-FLIPIT	
	Dev	Test	Dev	Test	Dev	Test	Dev	Test
PPO	75.46	65.43	<b>62.72</b>	<b>58.79</b>	12.64	15.48	1.28	0.68
PPO+SF	<b>84.05</b>	76.24	58.88	58.21	17.24	27.10	7.35	6.94

Table 4: Mean task-completion accuracies for TOWER CMDPs using FLAVA, with seed 1. We optimize using AdamW, and use a learning rate of  $3e-5$ . Bold results are outperforming or on par with ViLT in Table 2.

outperforming ViLT, we choose to not perform further hyperparameter search on SCATTER.

## F Third-party Code

Whenever the intended use is provided, the use of existing artifacts comply with their intended use. Suhr et al. (2017) is under CC-BY-4.0, and Kostrikov (2018) is under MIT. The use of code from Goldman et al. (2018) was done with explicit approval from the authors, because no license was provided with the code.

	TOWER-SCRATCH		TOWER-FLIPIT		SCATTER-SCRATCH		SCATTER-FLIPIT		
	Total	Correct %	Total	Correct %	Total	Correct %	Total	Correct %	
<b>Semantics</b>									
Cardinality (hard)	98	<b>76.5</b> <b>83.7</b>	480	<b>28.9</b> <b>56.2</b>	35	49.5   40.0	119	0.0   <b>13.7</b>	<i>There are exactly four objects not touching any edge</i>
Cardinality (soft)	21	68.2   81.0	82	<b>30.1</b> <b>56.5</b>	11	<b>60.6</b>   24.3	42	0.0   10.3	<i>There is a box with at least one square and at least three triangles.</i>
Existential	122	<b>75.1</b> <b>81.7</b>	577	<b>29.0</b>   55.1	55	<b>55.7</b>   35.1	192	0.0   12.3	<b>There is a tower with yellow base.</b>
Universal	7	<b>85.7</b> <b>95.2</b>	28	<b>29.7</b>   46.4	9	<b>81.5</b> <b>59.3</b>	36	0.0   9.3	<i>There is a black item in every box.</i>
Coordination	19	<b>85.9</b> <b>84.2</b>	86	27.1   52.7	15	<b>64.5</b>   40.0	55	0.0   9.7	<i>There are 2 blue circles and 1 blue triangle</i>
Coreference	3	<b>100.0</b>   77.8	10	13.3   10.0	3	<b>44.4</b>   33.3	9	0.0   7.4	<i>There is a blue triangle touching the wall with its side.</i>
Spatial Relations	94	<b>74.8</b> <b>81.6</b>	438	26.6   53.1	39	53.8   32.5	128	0.0   10.1	<i>there is one tower with a yellow block above a yellow block</i>
Comparative	5	66.7   73.3	20	11.7   21.7	1	<b>100.0</b>   <b>100.0</b>	4	0.0   <b>16.7</b>	<i>There is a box with multiple items and only one item has a different color.</i>
Presupposition	17	<b>74.5</b> <b>90.2</b>	74	27.0   <b>54.1</b>	22	<b>66.7</b> <b>51.5</b>	78	0.0   12.4	<i>There is a box with seven items and the three black items are the same in shape.</i>
Negation	4	<b>75.0</b>   66.7	15	13.3   37.8	14	54.8   33.3	52	0.0   7.0	<i>there is exactly one black triangle not touching the edge</i>
<b>Syntax</b>									
Coordination	4	<b>83.3</b>   75.0	14	11.9   <b>59.5</b>	5	53.3   26.7	20	0.0   8.3	<i>There is a box with at least one square and at least three triangles.</i>
PP Attachment	44	<b>76.5</b> <b>81.8</b>	215	26.2   <b>54.3</b>	3	33.3   33.3	8	0.0   8.3	<i>There is a black block on a black block as the base of a tower with three blocks.</i>
<b>Overall</b>		72.80   81.19		28.11   53.25		59.00   40.23		0.00   13.31	

Table 5: Performance on a set of development examples annotated for semantic and syntactic categories by Suhr et al. (2017) for both models (C3+BERT | ViLT) when trained with PPO. Dev performance refers to mean performance on the respective full development set. Results outperforming dev performance are in bold.

	TOWER-SCRATCH		TOWER-FLIPIT		SCATTER-SCRATCH		SCATTER-FLIPIT						
	Total	Correct %	Total	Correct %	Total	Correct %	Total	Correct %					
<b>Semantics</b>													
Cardinality (hard)	98	<b>84.3</b>	85.1	480	30.3	<b>68.9</b>	35	60.9	59.1	119	<b>17.1</b>	27.7	<i>There are exactly four objects not touching any edge</i>
Cardinality (soft)	21	76.2	<b>87.3</b>	82	<b>33.7</b>	<b>66.7</b>	11	<b>72.7</b>	<b>78.8</b>	42	10.3	18.3	<i>There is a box with at least one square and at least three triangles.</i>
Existential	122	<b>84.1</b>	83.6	577	31.8	<b>69.5</b>	55	70.3	67.3	192	15.8	27.8	<b>There is a tower with yellow base.</b>
Universal	7	80.9	<b>95.2</b>	28	26.2	52.4	9	<b>88.9</b>	<b>88.9</b>	36	7.4	16.6	<i>There is a black item in every box.</i>
Coordination	19	<b>91.2</b>	<b>84.2</b>	86	<b>39.9</b>	65.1	15	<b>73.3</b>	64.5	55	6.1	17.0	<i>There are 2 blue circles and 1 blue triangle</i>
Coreference	3	<b>88.9</b>	<b>100.0</b>	10	23.3	30.0	3	44.4	44.4	9	3.7	<b>33.3</b>	<i>There is a blue triangle touching the wall with its side.</i>
Spatial Relations	94	<b>81.9</b>	<b>84.4</b>	438	26.9	<b>69.0</b>	39	68.4	<b>68.4</b>	128	16.7	<b>28.1</b>	<i>there is one tower with a yellow block above a yellow block</i>
Comparative	5	73.3	80.0	20	21.7	31.7	1	<b>100.0</b>	<b>100.0</b>	4	16.7	<b>41.7</b>	<i>There is a box with multiple items and only one item has a different color.</i>
Presupposition	17	<b>82.4</b>	<b>94.1</b>	74	29.3	65.3	22	<b>72.7</b>	<b>69.7</b>	78	<b>14.1</b>	<b>29.9</b>	<i>There is a box with seven items and the three black items are the same in shape.</i>
Negation	4	75.0	58.3	15	15.6	64.4	14	66.7	<b>71.4</b>	52	<b>17.3</b>	23.7	<i>there is exactly one black triangle not touching the edge</i>
<b>Syntax</b>													
Coordination	4	<b>91.7</b>	75.0	14	19.0	<b>69.0</b>	5	60.0	46.7	20	5.0	20.0	<i>There is a box with at least one square and at least three triangles.</i>
PP Attachment	44	<b>84.8</b>	<b>85.6</b>	215	27.0	<b>70.2</b>	3	<b>77.8</b>	66.7	8	8.3	20.8	<i>There is a black block on a black block as the base of a tower with three blocks.</i>
<b>Overall</b>		81.80	84.05		32.59	65.68		72.03	67.43		17.04	28.01	

Table 6: Performance on a set of development examples annotated for semantic and syntactic categories by Suhr et al. (2017) for both models (C3+BERT | ViLT) when trained with PPO+SF. Dev performance refers to mean performance on the respective full development set. Results outperforming dev performance are in bold.

## ACL 2023 Responsible NLP Checklist

---

### A For every submission:

- A1. Did you describe the limitations of your work?  
*See Section 8.*
- A2. Did you discuss any potential risks of your work?  
*Section 8, indicating that the work is solely based on a mainstream language (English).*
- A3. Do the abstract and introduction summarize the paper’s main claims?  
*See the abstract and Section 1.*
- A4. Have you used AI writing assistants when working on this paper?  
*Left blank.*

### B Did you use or create scientific artifacts?

*See Section 3 and Section 5. The benchmark, annotations and trained models are or will be available.*

- B1. Did you cite the creators of artifacts you used?  
*Yes, see Section 3, Section 5 and Section 6.1.*
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?  
*See Section 1.*
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?  
*See Appendix F.*
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?  
*Not applicable. The data used does not contain information that could identify individuals.*
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?  
*Section 3 and Section 5 include information about the languages and linguistic phenomena, and Section 6 and Section D.4 provide more information about the linguistic phenomena in a subset of development data.*
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.  
*Table 1 provides the information.*

### C Did you run computational experiments?

*See Section 6.*

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?  
*See Section 6.1 and Appendix C.1.*

*The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.*

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?  
*See Section 6.1 and Appendix C.1.*
  - C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?  
*See Section 6.*
  - C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?  
*See Section C.1.*
- D**  **Did you use human annotators (e.g., crowdworkers) or research with human participants?**  
*See Section 5 and Appendix B.*
- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?  
*See Appendix B.*
  - D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?  
*See Section 5 and Appendix B.*
  - D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?  
*See Appendix B.*
  - D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?  
*The annotation of the Python programs did not necessitate an ethics review board, and was done by contracting programmers through Upwork.*
  - D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?  
*See Appendix B.*