Learning a natural-language to LTL executable semantic parser for grounded robotics

Christopher Wang, Candace Ross, Yen-Ling Kuo, Boris Katz, Andrei Barbu

Abstract

Children acquire their native language with apparent ease by observing how language is used in context and attempting to use it themselves. They do so without laborious annotations, negative examples, or even direct corrections. We take a step toward robots that can do the same by training a grounded semantic parser, which discovers latent linguistic representations that can be used for the execution of natural-language commands. In particular, we focus on the difficult domain of commands with a temporal aspect, whose semantics we capture with Linear Temporal Logic, LTL. Our parser is trained with pairs of sentences and executions as well as an executor. At training time, the parser hypothesizes a meaning representation for the input as a formula in LTL. Three competing pressures allow the parser to discover meaning from language. First, any hypothesized meaning for a sentence must be permissive enough to reflect all the annotated execution trajectories. Second, the executor — a pretrained end-to-end LTL planner — must find that the observed trajectories are likely executions of the meaning. Finally, a generator, which reconstructs the original input, encourages the model to find representations that conserve knowledge about the command. Together these ensure that the meaning is neither too general nor too specific. Our model generalizes well, being able to parse and execute both machine-generated and human-generated commands, with near-equal accuracy, despite the fact that the human-generated sentences are much more varied and complex with an open lexicon. The approach presented here is not specific to LTL: it can be applied to any domain where sentence meanings can be hypothesized and an executor can verify these meanings, thus opening the door to many applications for robotic agents.

Learning a natural-language to LTL
executable semantic parser for grounded robotics

Christopher Wang
MIT CSAIL & CBMM
czw@mit.edu

Candace Ross
MIT CSAIL & CBMM
ccross@mit.edu

Yen-Ling Kuo
MIT CSAIL & CBMM
ylkuo@mit.edu

Boris Katz
MIT CSAIL & CBMM
boris@mit.edu

Andrei Barbu
MIT CSAIL & CBMM
abarbu@mit.edu

Abstract: Children acquire their native language with apparent ease by observing how language is used in context and attempting to use it themselves. They do so without laborious annotations, negative examples, or even direct corrections. We take a step toward robots that can do the same by training a grounded semantic parser, which discovers latent linguistic representations that can be used for the execution of natural-language commands. In particular, we focus on the difficult domain of commands with a temporal aspect, whose semantics we capture with Linear Temporal Logic, LTL. Our parser is trained with pairs of sentences and executions as well as an executor. At training time, the parser hypothesizes a meaning representation for the input as a formula in LTL. Three competing pressures allow the parser to discover meaning from language. First, any hypothesized meaning for a sentence must be permissive enough to reflect all the annotated execution trajectories. Second, the executor — a pretrained end-to-end LTL planner — must find that the observed trajectories are likely executions of the meaning. Finally, a generator, which reconstructs the original input, encourages the model to find representations that conserve knowledge about the command. Together these ensure that the meaning is neither too general nor too specific. Our model generalizes well, being able to parse and execute both machine-generated and human-generated commands, with near-equal accuracy, despite the fact that the human-generated sentences are much more varied and complex with an open lexicon. The approach presented here is not specific to LTL: it can be applied to any domain where sentence meanings can be hypothesized and an executor can verify these meanings, thus opening the door to many applications for robotic agents.

Keywords: LTL, semantic parsing, weak supervision

1   Introduction

Natural language has the potential to be the most effective and convenient way to issue commands to a robot. However, machine acquisition of language is difficult due to the context- and speaker-specific variations that exist in natural language. For instance, English usage differs widely throughout the world: between children and adults, and in businesses vs. in homes. This does not pose a significant challenge to human listeners because we acquire language by observing how others use it and then attempting to use it ourselves. Upon observing the language use of other humans, we discover the latent structure of the language being spoken. We develop a similar approach for machines.

Our grounded semantic parser learns the latent structure of natural language utterances. This knowledge is executable and can be used by a planner to run commands. The parser only observes how language is used: what was said i.e., the command, and what was done in response i.e., a trajectory in the configuration space of an agent. We provide two additional resources to the parser, which children also have access to. First, we build in an inductive bias for a particular logic, in our
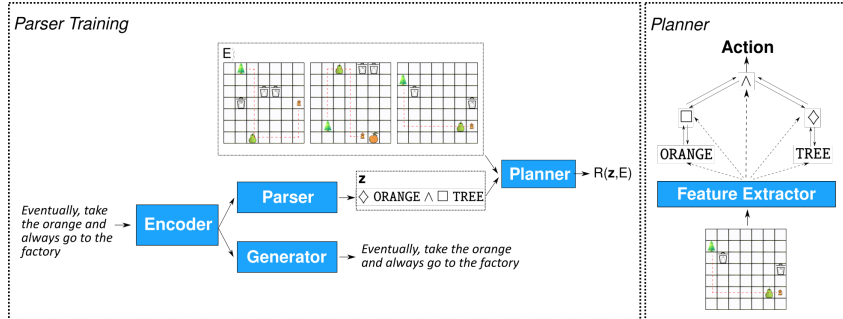
Figure 1: (left) The paradigm for training the parser is shown on the left. Given an input sentence, the parser proposes an LTL formula, $\mathbf{z}$, that could encode the meaning of that sentence. That formula is executed by a robotic agent to estimate the likelihood of the observed behavior, E, given the interpretation of the sentence, i.e., determining whether many more efficient approaches to executing this formula might exist. This likelihood is used to compute the reward $R(\mathbf{z}, E)$. At the same time, a generator attempts to reconstruct the sentence. (right) An example of the planner in action. We use the planner described by Kuo et al. [3] which learns to execute LTL formulas end-to-end from pixels to actions. Each predicate, ORANGE and TREE here, and each operator, are neural networks which together output an action that the robot should take given the current state of the world. Not shown are recurrent connections that enable each component to keep track of execution progress.

case Linear Temporal Logic (LTL) over finite sequences. We believe this is a reasonable starting place for our model, since it is widely assumed that priors over possible languages are built in by evolution and are critical to human language learning [1]. Second, we provide feedback from an executor: a planner trained end-to-end that is capable of executing formulas in whatever formalism we use in the prior, in our case LTL. With only this knowledge, our grounded semantic parser learns to turn sentences into LTL formulas and to execute those formulas.

The parser strives to find an explanation for what a sentence might mean by hypothesizing potential meanings and then updating its parameters depending on how suitable those meanings were. Four sources of information combine together to inform the semantic parser and are woven together into a single loss function. First, all outputs are verified to be syntactically, but not semantically valid LTL formulas, i.e., only valid LTL formulas are accepted. Second, the parser aims to create interpretations that are generic enough and whose LTL formulas actually admit the behavior that was observed. Third, given an interpretation of a sentence, the executor validates that the observed behavior is rational, i.e., has a high likelihood conditioned on that interpretation. Fourth, a generator attempts to reconstruct the input to maximize the knowledge conserved when translating sentences into some formalism. We do not provide any LTL-specific or domain-specific knowledge.

Our approach performs well on both machine and human generated sentences. In both cases, it is able to execute around 80% of commands successfully. A traditional fully-supervised approach on the machine-generated sentences outperforms ours with 95% accuracy, but requires the ground-truth LTL formulas. Where this approach of discovering latent structures shines is on real-world data. We asked human subjects, unconnected with this research and without knowledge of robotics or ML, to produce sentences that describe the behavior of robots. This behavior was produced according to an LTL formula that we randomly generated, but the humans were free to describe whatever they wanted with whatever words they desired as long as it was true. Despite the human sentences being much more varied and complex, including structures which our formalism cannot exactly represent, our method still finds whatever latent structure is present and required to execute the natural-language commands produced by humans with nearly the same accuracy as those produced by machines.

Executing LTL commands and understanding the kinds of temporal relations that require LTL is particularly difficult due to the richness and openness of natural language [2]. But LTL is just a stepping stone. It remains an important open question in grounded robotics: what representation will be enough to capture the richness of how humans use language? For instance, notions such as modal operators to reason about hypothetical futures will likely be required, but what else is unclear. Having a general-purpose mechanism for learning to execute commands is extremely helpful under these conditions; we can experiment with different logics and domains with the same agents by changing the priors and planners while leaving the rest of the system intact.

Figure 2: Execution tracks for the command "Always take the pear and go to the tree and stay there."

The main contributions of this work are:

1. a semantic parser that maps natural language to LTL formulas trained without access to any annotated formulas — no annotations were even collected for human-generated commands,
2. a variant of Craft by Andreas et al. [4] suited for grounded semantic parsing experiments, and
3. a recipe for creating grounded semantic parsers for new domains that results in executable knowledge without annotations for those domains.

2 Related Work

Semantic parsing Early attempts at semantic parsing, such as the famous SHRDLU system [5], were rule-based language systems. Later, advances in statistical language modeling gave rise to grammar-based approaches that were trained on full supervision [6, 7]. Recent approaches have focused on training grounded semantic parsers using weak supervision. The work of [8, 9, 10, 11] trains a language model on a dataset of question-answer pairs to produce queries in a formal language. The approaches of [12, 13] present a planner and CCG-based parser to generate programs for a deterministic robotic simulator. However, in these approaches, it is not clear whether constraints on behavior across time are learned, because the tasks do not require this knowledge and the formalisms used cannot easily represent such concepts. Ross et al. [14] use videos to supervise a grounded semantic parser. Although the videos contain information about events across time, the predicate logic used in [14] does not contain the operators necessary to represent the constraints that we focus on in this work. Nor did that work result in executable knowledge, i.e., plans that could drive a robot, merely descriptions of videos and actions. The semantic parsing literature is relatively sparse on the topic of time and temporal relations; for example Lee et al. [15] parse natural-language expressions denoting relative times, in many ways a much easier task than parsing into LTL. Even among the parsing approaches that do use the LTL formalism, many require a repository of general-purpose templates [16, 17, 18, 19] and do not truly address unbounded natural language input.

Our work follows recent approaches that cast the problem of semantic parsing as a machine translation task. Instead of using chart parsers, as is typical for grammar-based approaches [12, 14, 13], we use an encoder-decoder sequence-to-sequence model where the input and output are sequences of tokens; in our case, natural language commands and LTL formulas, respectively. Attempts to train sequence models using weak supervision usually warm start the model by pre-training with full supervision [20, 21]. By contrast, we train our model from scratch using a relatively small set of command-execution pairs. This is difficult in the initial stages of training due to the large exploration space. To address these challenges, we follow Guu et al. [22] and Liang et al. [9] in using randomized search with a search space restricted to formulas that are syntactically valid in a target logic, LTL.

Planning Closely related to our task is the work that has been devoted to mapping LTL formulas to execution sequences [23, 24, 25]. Kuo et al. [3] presents a compositional neural network that learns to map LTL formulas to action sequences. It can also compute the likelihood of a sequence of actions conditioned on a formula. We adopt this agent as the executor in this work and train it for our domain; it never sees a single natural-language utterance, it merely learns how to execute LTL formulas.

Most relevant to our work, Patel et al. [26] trains a weakly supervised semantic parser for LTL formulas. This work uses an algorithm that requires significant knowledge about LTL and is entirely LTL-specific, while here we merely reject syntactically invalid formulas. Their approach requires access to the ground truth observations of the environment to execute LTL formulas step by step

while our approach takes as input images observed by the robot of its environment. Fundamentally, their approach requires reasoning about locations and paths, while our approach includes object interactions. No prior work or prior dataset includes a robot that can interact with objects, that perceives its environment, and must understand natural-language commands that have a temporal aspect to them and thus require LTL.

3 Model

Figure 1 (left) provides an overview of our model. Following one-to-many sequence-to-sequence approaches for multi-task learning such as [27], our model consists of one encoder and two decoders. A natural language input command, $\mathbf{x}$, is first encoded to a high-dimensional feature vector then separately decoded into the predicted LTL formula, $\hat{\mathbf{z}}$, and reconstructed command, $\hat{\mathbf{x}}$. During training, the planner assigns a reward, $R(\hat{\mathbf{z}}, E)$, to each hypothesized formula, $\hat{\mathbf{z}}$. The training dataset, $D = \{(\mathbf{x}, E)\}$, contains pairs of natural language commands $\mathbf{x}$ and execution demonstrations E. Each demonstration, E, consists of k trajectory-environment pairs $E = \{\mathbf{y}_k, e_k\}_{i=1}^k$, see Figure 2. A trajectory is a sequence of actions e.g., $\mathbf{y} = [\texttt{left}, \texttt{left}, \texttt{up}, \texttt{grab}]$.

3.1 Parser and Generator

Our architecture is similar to the sequence-to-sequence models of Guu et al. [22] and Dong and Lapata [28]. Given input sentence $\mathbf{x}$, our model defines a probability for an output formula $\mathbf{z} = [z_1, ..., z_m]$:

$$p(\mathbf{z} \mid \mathbf{x}) = \prod_{i=1}^{|\mathbf{z}|} p(z_i \mid z_{<i}, \mathbf{x}; \theta_{\text{parse}}) \tag{1}$$

where θ_{parse} are the parser parameters.

Encoder The encoder is a stacked bi-directional LSTM that takes word embeddings as produced by the pretrained English GloVe model [29]. The encoder maps the input $[x_1, ..., x_n]$ to a feature representation $\mathbf{h} = [h_1, ..., h_n]$.

Parser-Decoder The parser decoder takes the feature vector from the encoder and generates a sequence of tokens: the LTL formula $\hat{\mathbf{z}} = [z_1, ..., z_m]$. The decoder is a stacked LSTM with an attention mechanism [30], as implemented by Bastings [31]. Dropout is applied before the final softmax.

At each step, the attention mechanism produces a context vector c_i from the decoder hidden state s_i and $\mathbf{h}$. The previous decoder input z_{i-1}, along with c_i and s_i, are used to produce a distribution over output tokens:

$$p(z_i \mid z_{<i}, \mathbf{x}; \theta_{\text{parse}}) = \text{Softmax}(W_o[z_{i-1}; c_i; h_i]) \tag{2}$$

We keep a stack of generated tokens and output LTL formulas in post-order, since all operators in LTL have a fixed arity. This allows us to avoid the problem of parentheses matching. We condition the decoder to sample syntactically-valid continuations of the formula being decoded. Note that this assumes nothing about the formula's meaning. We simply build in the fact that certain continuations are trivially guaranteed to never be syntactically valid; for example, the formula $\wedge\wedge$ is not a valid LTL formula. Practically, this property is trivially computable since LTL formulas, and virtually all logics in general, use notation that is context-free, which allows us to remove invalid options from the softmax output before sampling the next token. Following Guu et al. [22], we use ϵ-randomized sample decoding for better exploration during training. At each timestep i, with probability $p = \epsilon$, we draw the next token z_i according to $p(z_i \mid z_{<i}, \mathbf{x}; \theta_{\text{parse}})$. With probability $p = 1 - \epsilon$, we pick a valid continuation uniformly at random.

Generator-Decoder The parser-decoder produces formulas which are scored by the planner, but this does not ensure that the full content of the utterance is reflected in the parse. To encourage this, we include a second, separate decoder, which is trained to reconstruct the original natural language command, $\mathbf{x}$, from the feature representation $[h_1, ..., h_n]$. This is a standard multi-task learning approach that is often used in the literature to improve generalization. The architecture of this component is identical to the parser-decoder.

	Constituents	Description
Logical operators	$\wedge$, $\vee$	And, or
Temporal operators	$\square$, $\diamond$, U	Eventually, always, until
Objects	APPLE, ORANGE, PEAR	Objects that can be held
Relations	CLOSER_APPLE, CLOSER_ORANGE, CLOSER_PEAR	Spatial relations
Destinations	FLAG, HOUSE, TREE	Destinations

Table 1: The various constituents of our target formalism are shown above. We ground the meaning of natural-language sentences produced by humans into LTL$_f$ ([32]) without negation. Predicates from the Craft domain are renamed as users found these labels easier to understand. Note that this is the size of the target formalism; it is unrelated to the complexity of the input. Humans produced sentences that contained 266 words across them which had to be grounded to these semantics.

Planner Formulas are scored using a planner. We adopt the one described by Kuo et al. [3] because it learns to execute LTL formulas end to end and is pretrained for the CRAFT environment that we evaluate on. Given a formula and an environment, the planner is trained to create an execution sequence; it never has access to the natural language utterance. It learns to extract features from images of the environment around the robot and acquires knowledge about LTL predicates and operators in order to execute novel formulas in novel environments. Figure 1 (right) shows an example of the planner configured to execute an LTL formula. Given an LTL formula, the planner is configured by assembling a compositional recurrent network specific to that formula; it then guides the robot to execute the formula. This compositionality enables zero-shot generalization to new formulas. Any planner could in principle be used as long as it could learn to execute formulas for the target domain and if it could score an arbitrary trajectory against a formula.

The planner does not have access to the ground truth environment. Instead, it observes an image of a 5×5 patch of the world around it, which is passed through a learned feature extractor CNN. At every time step, each module within the planner takes as input the features extracted from the surroundings of the robot, the previous state of that module, and the previous state of the parent. The state of the root of the LTL formula, according to an arbitrary but consistent parse of the formula, is decoded to predict a distribution over actions. The model is pretrained on randomly generated environments and LTL formulas using A2C, an Advantage Actor-Critic algorithm [33, 34].

4 Training

The model described above produces a candidate LTL formula $\hat{\mathbf{z}}$, along with a reconstruction of the input, $\hat{\mathbf{x}}$. Each candidate formula is used to compute a reward that incorporates the likelihood, as computed by the planner, of the observed trajectories given the hypothesized LTL formula. This plays two roles: first it ensures that the observed trajectories are actually feasible given the hypothesized LTL formula; otherwise they will have zero likelihood. Secondly, it provides a score for how rational the planner judges the behavior to be. Not all feasible paths are equally rational, and so by extension not equally likely. For example, a complex observed behavior is unlikely to be the consequence of a simple parse: it is more likely that the parser is producing an overly broad interpretation rather than the observed trajectory going out of its way to do something unnecessary. The reward is then

$$R(\hat{\mathbf{z}}, E) = \begin{cases} \frac{1}{k} \sum_{(e_i, \mathbf{y}_i)} \frac{1}{|\mathbf{y}_i|} \sum_j p(\mathbf{y}_{i,j} \mid \hat{\mathbf{z}}, \mathbf{y}_{i,<j}, e) & \forall \mathbf{y} \in Y.\mathbf{y} \in \hat{\mathbf{a}} \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

where $\hat{\mathbf{a}}$ is an NFA representation of the formula $\hat{\mathbf{z}}$, so that $\mathbf{y} \in \hat{\mathbf{a}}$ indicates that $\mathbf{y}$ is feasible. We optimize the reward of the output with either REINFORCE [35] or Iterative Maximum Likelihood (IML). This reward computes an average of the likelihood over the k execution traces in E, conditioned on the candidate formula $\hat{\mathbf{z}}$ (recall that each sentence is paired with k demonstration trajectories, each in a different randomly-generated environment). Since the size of the search space for LTL formulas grows exponentially in the length of the formula, we employ curriculum learning as in Liang et al. [9]. Every 10 epochs, we increase the maximum length of the predicted formulas by 3.

5

4.1 Reinforcement Learning

In the reinforcement learning setting, our objective is to maximize the expected reward, marginalizing over the space of possible formulas: $J_{RL} = \sum_{\mathbf{x}} \sum_{\mathbf{z}} R(\mathbf{z}) p_\theta(\hat{\mathbf{z}} \mid \mathbf{x})$. We use the REINFORCE algorithm [35] to learn the policy parameters with Monte-Carlo sampling. For better exploration, we use ϵ-dithering when sampling as described in 3.1. To incorporate the generator, we optimize a linear combination of this reward and the reconstruction loss, $J_{GEN} = -\sum_{\mathbf{x}} \log p(\mathbf{x} \mid \mathbf{x}; \theta_{GEN})$, so that $J = J_{RL} + \alpha J_{GEN}$. We adjust α at training time to balance the two components. In particular, it is important to start with a small α initially, since J_{RL} is small when θ is untrained and few candidate formulas have non-zero reward.

4.2 Iterative Maximum Likelihood

Iterative Maximum Likelihood, IML has proven itself to be as efficient if not more efficient when acquiring semantic parsers compared to RL. We adopt a method similar to that of [9] and [11]. First, we explore the output space by sampling K formulas from the parser. We keep the highest reward formulas $\hat{z}^*$ and use them as a pseudo-gold. We then maximize the likelihood of the pseudo-gold formulas over the course of 10 epochs: $J_{IML} = \sum_{\mathbf{x}} \sum_i \log p(\hat{\mathbf{z}}_i^* \mid \mathbf{x}; \theta_{\text{parse}})$. To incorporate the generator, we again combine the two objective functions, but this time no scaling parameter is required: $J = J_{IML} + J_{GEN}$. Sampling and MLE steps are then iterated.

5 Experiments

We test the parser in two experiments. The first verifies that machine-generated natural-language commands derived from LTL formulas can be understood and followed by a trained agent in a way that reflects the formula correctly. The second verifies that our model can 1) understand sentences produced by humans which describe a given behavior and 2) express a plan in LTL that will result in this behavior. Note that the humans never see the LTL formulas; they produce natural language descriptions for the behavior of robots.

In all cases, our model has the same hyperparameters. The stacked LSTMs all have two layers with hidden dimensions of size 1000 and dropout probability 0.2. We use Adam with learning rate $1e - 3$. REINFORCE and IML both sample 128 formulas to compute the expectation and generate sentences for the next iteration with $\epsilon = 0.15$ when exploring. We train using $k = 3$ trajectories for 50 epochs. Results are reported for the model with highest validation set performance as measured by the **Exec** metric, see section 5.1. At test time, we decode using a beam search with width 10.

Temporal Phenomena Most randomly generated LTL formulas are uninteresting, similar to the way in which most instances of the boolean satisfiability problem SAT are uninteresting [36]. To avoid such issues, we adopt the standard classification of LTL formulas produced by Manna and Pnueli [37] and generate formulas in their six partially-overlapping categories: *safety*, *guarantee*, *persistence*, *recurrence*, *obligation*, and *reactivity*. Respectively, *safety*, *guarantee*, *persistence*, and *recurrence*, ensure that a property will always hold, will hold at least once, will always hold after a certain point, or will hold at repeated points in time. *Obligation* and *reactivity* are compound classes formed by unrestricted boolean combinations of the safety and recurrence classes respectively. While we adopt the LTL over finite sequences, LTL_f as described by De Giacomo and Vardi [32], the target formalism uses the $\diamondsuit$ *eventually* and $\square$ *always* temporal operators rather than $\bigcirc$ *next* to readily generate instances of the Manna and Pnueli [37] classes. Since we found humans to be very unlikely to spontaneously generate sentences that required negation, it was not included. The components of the target formalism are shown in section 3.1.

Mechanical dataset We collect two sets of data. The first, a mechanically-generated dataset, consists of 1,000 natural language sentences paired with 3,000 execution traces, 3 per example, with a 70/15/15 training/val/test split. All commands and environments are given in the context of the Minecraft-like CRAFT environment, which we adapt from Andreas et al. [4] and Kuo et al. [3].

Following Jia and Liang [38] and Goldman et al. [20], we generate sentences and formulas by randomly and uniformly sampling productions and terminals from a synchronous context-free grammar (Appendix A).

For each command-formula pair, we populate three 7x7 grid environments with objects and landmarks. Each environment includes all the items and landmarks in the corresponding command in random locations. Other objects and landmarks that are not in the command are each included with probability 0.3, resulting in somewhat densely populated environments. Of course, this does not guarantee that a command can be executed on a particular map.

Given the formula, we generate a non-deterministic finite state automaton using Spot [39]. An oracle brute-force searches the action space and generates an action sequence which the automaton accepts; this can be quite slow. Rejection sampling over this process results in three environments for each command-formula pair. LTL with finite semantics [40, 32] requires a time horizon: we set it at 20 steps, by which point the command must be satisfied, or equivalently, the automaton must be in an accepting state.

# Total sentences	2,000
# **Machine sentences**	1,000
# Guarantee	204
# Safety	264
# Recurrence	243
# Persistence	214
# Obligation	52
# Reactivity	23
Avg. words/sent.	17.7 ± 8.4
# Lexicon size	44
# **Human sentences**	1,000
Avg. words/formula	5.2 ± 2.9
Avg. words/sent.	8.3 ± 3.3
# Lexicon size	266

Table 2: Dataset statistics. Note that the human-generated data is far more varied with a much larger lexicon.

Some commands mandate that a condition hold globally, e.g., "Always hold the gem". However, unless the agent's initial state satisfies this condition, e.g., the robot happens to start with the gem in hand, no satisfying action sequence is possible. To address this, we allow the robot time to approach and grab the object, which is surely the intent of any human speaker, replacing each predicate p with $\texttt{closer}(p) \cup p$, where $\texttt{closer}(p)$ means "closer to p". That is, instead of requiring that p be satisfied immediately and always, we mandate that the robot move closer to p until p is satisfied.

Human-generated dataset   We take all sampled environments from the mechanical dataset and present them to humans. Note that humans only see what the robots do, not why they did it. They do not see the LTL formulas or the machine-generated utterances. We asked six human annotators, who were working for pay, unconnected to this research, and unfamiliar with NLP or robotics, to describe what the robots are doing. Of course, this leads to different sentences than those that originally generated the behavior of the robots.

This process ensures that even though our mechanical dataset was generated from a context free grammar, no trace of that grammar remains; humans generate the sentences they are comfortable with. The distractor objects and landmarks were not removed for this experiment, giving annotators the opportunity to refer to them. The target object and the intended actions need never appear in the final human descriptions. We did not collect LTL formulas from the human annotators. Note that the size of the lexicon, 266 words, that the humans used is far larger than both what our formalism supports and what was produced by the machines.

5.1   Results

Results are shown in table 3. The machine-generated data is annotated with ground truth LTL formulas, but no equivalent concept exists for the human-generated dataset, since the humans only had to explain what they thought the robots were doing; it might not even be possible to fully capture the semantics of their sentences by LTL. **Exec** measures the fraction of formulas that accept all k execution traces. This is an overestimate of the performance of the grounded parser; merely accepting formulas does not guarantee any understanding. **Plan** measures the fraction of the environments that the planner executes correctly on, given the predicted formula as input. This is an underestimate of the performance of the grounded parser; even a human controlling a robot in such environments may not exactly carry out the expected actions, since many formulas are hard to interpret and there is much opportunity for error. As is common in linguistics, no single metric perfectly captures performance, but these two metrics do bracket the performance of the approach.

A more stringent metric would be to investigate how the annotated and predicted LTL formulas compare, which is possible on the machine-generated dataset. **Exact** measures the fraction of predicted formula which are equivalent to the ground truth. **Seq** is the F1 token overlap score between the predicted formula and the ground truth. These are extremely stringent metrics that even humans would perform poorly on, as the same actions can be carried out for many different reasons and many

7

Machine-generated dataset					Human-generated dataset		
	Exec	**Plan**	**Seq**	**Exact**		**Exec**	**Plan**
Supervised	94.7	36.7	94.9	91.3	Random	16.3	17.5
RL	82.0	41.3	22.9	8.7	RL	78.7	40.7
RL + generator	83.3	41.3	23.9	8.7	RL + generator	79.3	43.3
IML	81.3	32.2	14.0	2.0	IML	80.0	28.7
IML + generator	85.3	34.9	15.0	4.0	IML + generator	83.3	31.8

Table 3: Results on the machine-generated (left) and the human-generated (right) datasets. **Exec** measures the likelihood of formulas that recognize the ground truth trajectories, an overestimate of the real performance of the parser. **Plan** measures how often the planner produced a correct trajectory given the predicted formula, an underestimate of the real performance. **Seq** and **Exact** measure the overlap between predicted and ground-truth LTL formulas; note that many formulas have identical semantics, even humans may do poorly on this metric. The fully-supervised method outperforms our approach, as expected, but it is only relevant for the machine dataset where ground-truth annotations exist. Note that the drop between the machine- and human-generated datasets is small, despite the human sentences being more diverse.

LTL formulas are equivalent in context. Typical mistakes made when predicting the exact formula on the machine-generated dataset are shown on the right in table 4.

Overall, the supervised method outperforms our method, but when considering the percentage of correctly executed formulas, it only outperforms the weakly-supervised approach by 5-10%. Both RL and IML performed well, with the generator increasing performance by 1-4%. Overall, IML with the generator was the highest performing approach and recovers almost all the performance of the supervised approach.

We investigated how performance varied as we provided more examples per sentence using the machine-generated data. While the fraction of correctly executed sentences stays roughly the same, the exact match goes up significantly from 2% at $k = 3$ to 14% at $k = 7$. The ambiguity in this domain prevents exact matches, but allows for good executions.

Input	*Either grab the apple or the pear and hold them forever.*
Target	$\square\Diamond$FLAG $\wedge$ $\square\Diamond$ORANGE
RL+gen	$\square\Diamond(\square\Diamond$FLAG $\vee$ $\Diamond$ORANGE$)$
IML+gen	$\square\Diamond($FLAG $\vee$ ORANGE$)$

Table 4: Predicted output for the machine-generated test set showing typical mistakes. These formulas are hard to tell apart from observations of the robot's behavior. This makes it harder to learn the correct form while at the same time encouraging correct executions (See Appendices B and C).

On the human-generated dataset, despite the fact that the formalism we use is small compared to the 266 words that humans used, the grounded parser is able to understand most commands. It correctly executes in about 43% of the environments and accepts about 80% of commands.

6   Conclusion

We created a grounded semantic parser that, given only minimal knowledge about its environment and formalism, was able to discover the structure of an input language and produce executable formulas to command a robot. Its performance is competitive with a state of the art supervised approach, even though we provide no direct supervision. We were able to get similar performance on a challenging dataset produced by humans that could use any word and sentence construction to describe the actions of robots, even those that our formalism cannot completely capture. This model has virtually no knowledge of its domain or target logical formalism; it merely requires a planner and a method to reject syntactically invalid formulas. Many problems in robotics and NLP could be tackled by such an approach because of its low requirements for annotations. For example, data already exists to guide agents to reproduce the actions of customer service agents in response to queries. And in the robotic domain, future work might involve observing what humans say to one another and then acquiring a domain-specific semantic parser to guide robots on a worksite for example. Being able to adapt to variations in language use and to changes in the environment is crucial to building useful robots, because the same language may carry very different meanings in different contexts. In the long term, we hope that this line of research leads both to robots that understand us and to robotic systems that can be used to probe how children acquire language, bringing robotics and linguistics closer together.

References

[1] N. Chomsky. Approaching UG from below. In U. Sauerland and H.-M. Gärtner, editors, *Interfaces + Recursion = Language?*, pages 1–29. Mouton de Gruyter, NY, 2007.

[2] A. Brunello, A. Montanari, and M. Reynolds. Synthesis of LTL formulas from natural language texts: State of the art and research directions. In *26th International Symposium on Temporal Representation and Reasoning (TIME 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019.

[3] Y.-L. Kuo, B. Katz, and A. Barbu. Encoding formulas as deep networks: Reinforcement learning for zero-shot execution of LTL formulas. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020.

[4] J. Andreas, D. Klein, and S. Levine. Modular multitask reinforcement learning with policy sketches. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 166–175. JMLR. org, 2017.

[5] T. Winograd. Procedures as a representation for data in a computer program for understanding natural language. Technical report, Massachusetts Institute of Technology Cambrdige Project MAC, 1971.

[6] J. M. Zelle and R. J. Mooney. Learning to parse database queries using inductive logic programming. In *Proceedings of the national conference on artificial intelligence,*, page 1050–1055, 1996.

[7] L. S. Zettlemoyer and M. Collins. Learning to map sentences to logical form: Structured classification with probabilistic categorial grammars. In *UAI 2005 Proceedings of the Twenty-First Conference on Uncertainty in Artificial Intelligence*, 2005.

[8] J. Berant, A. Chou, R. Frostig, and P. Liang. Semantic parsing on freebase from question-answer pairs. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1533–1544, 2013.

[9] C. Liang, J. Berant, Q. Le, K. Forbus, and N. Lao. Neural symbolic machines: Learning semantic parsers on freebase with weak supervision. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 23–33, 2017.

[10] C. Liang, M. Norouzi, J. Berant, Q. V. Le, and N. Lao. Memory augmented policy optimization for program synthesis and semantic parsing. In *Advances in Neural Information Processing Systems*, pages 9994–10006, 2018.

[11] R. Agarwal, C. Liang, D. Schuurmans, and M. Norouzi. Learning to generalize from sparse and underspecified rewards. In *International Conference on Machine Learning*, pages 130–140, 2019.

[12] Y. Artzi and L. Zettlemoyer. Weakly supervised learning of semantic parsers for mapping instructions to actions. *Transactions of the Association for Computational Linguistics*, 1:49–62, 2013.

[13] E. C. Williams, N. Gopalan, M. Rhee, and S. Tellex. Learning to parse natural language to grounded reward functions with weak supervision. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–7. IEEE, 2018.

[14] C. Ross, A. Barbu, Y. Berzak, B. Myanganbayar, and B. Katz. Grounding language acquisition by training semantic parsers using captioned videos. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2647–2656, 2018.

[15] K. Lee, Y. Artzi, J. Dodge, and L. Zettlemoyer. Context-dependent semantic parsing for time expressions. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1437–1447. Association for Computational Linguistics, June 2014.

[16] M. B. Dwyer, G. S. Avrunin, and J. C. Corbett. Patterns in property specifications for finite-state verification. In *Proceedings of the 21st international conference on Software engineering*, pages 411–420, 1999.

[17] V. Gruhn and R. Laue. Patterns for timed property specifications. *Electronic Notes in Theoretical Computer Science*, 153(2):117–133, 2006.

[18] S. Konrad and B. H. Cheng. Real-time specification patterns. In *Proceedings of the 27th international conference on Software engineering*, pages 372–381, 2005.

[19] A. P. Nikora and G. Balcom. Automated identification of LTL patterns in natural language requirements. In *2009 20th International Symposium on Software Reliability Engineering*, pages 185–194. IEEE, 2009.

[20] O. Goldman, V. Latcinnik, E. Nave, A. Globerson, and J. Berant. 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. Association for Computational Linguistics, July 2018.

[21] L. Jehl, C. Lawrence, and S. Riezler. Learning neural sequence-to-sequence models from weak feedback with bipolar ramp loss. *Transactions of the Association for Computational Linguistics*, 7:233–248, 2019.

[22] K. Guu, P. Pasupat, E. Liu, and P. Liang. From language to programs: Bridging reinforcement learning and maximum marginal likelihood. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1051–1062, 2017.

[23] H. Sahni, S. Kumar, F. Tejani, and C. Isbell. Learning to compose skills. *arXiv preprint arXiv:1711.11289*, 2017.

[24] A. Camacho, R. Toro Icarte, T. Q. Klassen, R. Valenzano, and S. A. McIlraith. LTL and beyond: Formal languages for reward function specification in reinforcement learning. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 6065–6073. International Joint Conferences on Artificial Intelligence Organization, 7 2019.

[25] M. Alshiekh, R. Bloem, R. Ehlers, B. Könighofer, S. Niekum, and U. Topcu. Safe reinforcement learning via shielding. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.

[26] R. Patel, R. Pavlick, and S. Tellex. Learning to ground language to temporal logical form. In *NAACL 2019, SpLU RoboNLP Workshop*, 2019.

[27] D. Dong, H. Wu, W. He, D. Yu, and H. Wang. Multi-task learning for multiple language translation. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1723–1732, 2015.

[28] L. Dong and M. Lapata. Language to logical form with neural attention. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 33–43, 2016.

[29] J. Pennington, R. Socher, and C. D. Manning. Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, 2014.

[30] D. Bahdanau, K. Cho, and Y. Bengio. Neural machine translation by jointly learning to align and translate. In Y. Bengio and Y. LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.

[31] J. Bastings. The annotated encoder-decoder with attention, 2018.

[32] G. De Giacomo and M. Y. Vardi. Linear temporal logic and linear dynamic logic on finite traces. In *Twenty-Third International Joint Conference on Artificial Intelligence*, 2013.

[33] R. S. Sutton, D. A. McAllester, S. P. Singh, and Y. Mansour. Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems*, pages 1057–1063, 2000.

[34] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pages 1928–1937, 2016.

[35] R. J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.

[36] S. Horie and O. Watanabe. Hard instance generation for sat. In *International Symposium on Algorithms and Computation*, pages 22–31. Springer, 1997.

[37] Z. Manna and A. Pnueli. A hierarchy of temporal properties (invited paper, 1989). In *Proceedings of the Ninth Annual ACM Symposium on Principles of Distributed Computing*, PODC '90, page 377–410. Association for Computing Machinery, 1990. ISBN 089791404X.

[38] R. Jia and P. Liang. Data recombination for neural semantic parsing. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12–22. Association for Computational Linguistics, Aug. 2016.

[39] A. Duret-Lutz, A. Lewkowicz, A. Fauchille, T. Michaud, E. Renault, and L. Xu. Spot 2.0—a framework for LTL and ω-automata manipulation. In *International Symposium on Automated Technology for Verification and Analysis*, pages 122–129. Springer, 2016.

[40] S. Dutta and M. Y. Vardi. Assertion-based flow monitoring of systemc models. In *2014 Twelfth ACM/IEEE Conference on Formal Methods and Models for Codesign (MEMOCODE)*, pages 145–154. IEEE, 2014.

[41] P. Anderson, Q. Wu, D. Teney, J. Bruce, M. Johnson, N. Sünderhauf, I. Reid, S. Gould, and A. van den Hengel. Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3674–3683, 2018.

A Grammar

The grammar used to produce our machine generated commands is shown below:

$$
\begin{aligned}
\text{BINOP} &\rightarrow \text{`and'} \mid \text{`or'} \\
\text{UOP} &\rightarrow \text{`do not'} \mid \text{`you should not'} \\
\text{ITEM} &\rightarrow \text{`apple'} \mid \text{`orange'} \mid \text{`pear'} \\
\text{LANDMARK} &\rightarrow \text{`flag'} \mid \text{`house'} \mid \text{`tree'} \\
\text{PREDICATE} &\rightarrow \text{`be around the' LANDMARK} \mid \text{`be near the' LANDMARK} \\
&\quad \mid \text{`go to the' LANDMARK} \mid \text{`hold the' Item} \\
&\quad \mid \text{`take the' ITEM} \mid \text{`possess the' ITEM} \\
\text{P} &\rightarrow \text{PREDICATE} \mid \text{UOP PREDICATE} \mid \text{PREDICATE BINOP PREDICATE} \mid \text{UOP P} \\
\text{S} &\rightarrow \text{SAFETY} \mid \text{GUARANTEE} \mid \text{OBLIGATION} \mid \text{RECURRENCE} \mid \\
&\quad \mid \text{PERSISTENCE} \mid \text{REACTIVITY} \\
\text{SPREFIX} &\rightarrow \text{`always'} \mid \text{`at all times,'} \\
\text{SSUFFIX} &\rightarrow \text{`forever'} \mid \text{`at all times'} \mid \text{`all the time'} \\
\text{SAFETY} &\rightarrow \text{SPREFIX P} \mid \text{P SSUFFIX} \mid \text{SAFETY BINOP SAFETY} \\
\text{GPREFIX} &\rightarrow \text{`eventually'} \mid \text{`at some point'} \\
\text{NOTPREDICATE} &\rightarrow \text{UOP PREDICATE} \\
\text{GUARANTEE} &\rightarrow \text{GPREFIX P} \mid \text{`guarantee that you will' PREDICATE} \\
&\quad \mid \text{`guarantee that you'NOTPREDICATE} \mid \text{GUARANTEE BINOP GUARANTEE} \\
\text{OBLIGATION} &\rightarrow \text{SAFETY BINOP GUARANTEE} \mid \text{OBLIGATION BINOP SAFETY} \\
&\quad \mid \text{OBLIGATION BINOP GUARANTEE} \\
\text{RECURRENCE} &\rightarrow \text{`eventually,' P `and do this repeatedly'} \mid \text{RECURRENCE BINOP RECURRENCE} \\
\text{PERSISTENCE} &\rightarrow \text{`at some point, start to' P `and keep doing it'} \\
&\quad \mid \text{PERSISTENCE BINOP PERSISTENCE} \\
\text{REACTIVITY} &\rightarrow \text{RECURRENCE BINOP PERSISTENCE} \mid \text{REACTIVITY BINOP RECURRENCE} \\
&\quad \mid \text{REACTIVITY BINOP PERSISTENCE}
\end{aligned}
$$

B Sample output predictions

B.1 Grammar-generated sentences

Table 5: Predicted output on the best model parameters for the test set of the grammar-generated commands

Input	*eventually, be around the tree or go to the flag and do this repeatedly*
Target	$(\square\lozenge((\text{ TREE }) \vee (\text{ FLAG })))$
RL+gen	$(\square\lozenge((\text{ FLAG }) \vee (\text{ TREE })))$
IML+gen	$(\square\lozenge((\lozenge(\text{ TREE })) \vee (\square\lozenge((\square(\text{ TREE })) \vee (\square\lozenge(\text{ FLAG }))))))$
Input	*at some point, start to be around the house and go to the flag and keep doing it*
Target	$(\lozenge\square((\text{ HOUSE }) \wedge (\text{ FLAG })))$
RL+gen	$(\square\lozenge((\text{ HOUSE }) \vee (\text{ FLAG })))$
IML+gen	$(\lozenge\square((\lozenge(\text{ HOUSE })) \vee (\square(\text{ FLAG }))))$
Input	*at all times, possess the pear or be around the house and guarantee that you will go to the tree*
Target	$((\square((\text{ PEAR }) \vee (\text{ HOUSE }))) \wedge (\lozenge(\text{ TREE })))$
RL+gen	$(\square\lozenge((\text{ HOUSE }) \vee (\text{ TREE })))$
IML+gen	$(\square\lozenge((\square\lozenge(\text{ PEAR })) \vee (\square\lozenge((\square\lozenge(\text{ PEAR })) \vee (\square(\text{ PEAR }))))))$
Input	*at some point, start to be around the house and possess the orange and keep doing it*
Target	$(\lozenge\square((\text{ HOUSE }) \wedge (\text{ ORANGE })))$
RL+gen	$(\square\lozenge(\text{ HOUSE }))$
IML+gen	$(\square\lozenge((\square\lozenge(\text{ HOUSE })) \vee (\square\lozenge((\lozenge(\text{ HOUSE })) \vee (\square\lozenge(\text{ HOUSE }))))))$

B.2 Human-generated sentences

Table 6: Predicted output on the best model parameters for the test set of the human-generated sentences

Input : *get your hands on the orange or the apple some time.*
RL+gen: $(\square\lozenge(\text{ ORANGE }))$
IML+gen: $(\square\lozenge((\lozenge((\lozenge\square(\text{ APPLE })) \vee (\lozenge(\text{ ORANGE })))) \vee (\square\lozenge(\text{ TREE }))))$
Input: *guarantee that you snatch the pear*
RL+gen: $(\square\lozenge(\text{ PEAR }))$
IML+gen: $(\square\lozenge((\square\lozenge((\lozenge(\text{ APPLE })) \vee (\square(\text{ TREE })))) \vee (\square\lozenge(\text{ PEAR }))))$
Input: *guarantee that you approach the trash can and visit the tree*
RL+gen: $(\square\lozenge(\text{ TREE }))$
IML+gen: $(\square\lozenge((\square\lozenge((\lozenge(\text{ TREE })) \vee (\lozenge(\text{ HOUSE })))) \vee (\square\lozenge(\text{ TREE }))))$
Input: *ensure that you pick up the peach and then take it to the trash can*
RL+gen: $(\square\lozenge(\text{ ORANGE }))$
IML+gen: $(\square\lozenge((\lozenge((\square\lozenge(\text{ ORANGE })) \wedge (\square(\text{ ORANGE })))) \vee (\square(\text{ ORANGE }))))$

C Qualitative failure analysis

From observing our predicted output, we see the limitations of our approach (Appendix B). While it is encouraging to see that the model generally learns to produce the correct predicates, there are some situational ambiguities that are hard to overcome by observing execution demonstrations. Namely, it is difficult to distinguish *eventually* $\lozenge$ and *always* $\square$ by observing whether a formula accepts the execution trajectories. For example, consider the input command *always go to the tree* and a corresponding execution trace $\mathbf{y}$. Then, a hypothesis such as $\lozenge$TREE i.e., *eventually go to the tree*,

will receive a non-zero reward, because it is able to accept the demonstration **y**. This is because a trace that shows the robot *always* going to the tree must necessarily show the robot *eventually* going to the tree. In general, the model can always receive a non-zero reward for proposing $\lozenge$ in a situation which calls for $\square$. In the same way, any commands involving disjunction are difficult to learn. That is, the model can always receive a non-zero reward by proposing $p \vee q$ in a situation which requires $p \wedge q$.

These ambiguities would not hinder training if the planner always gave a higher likelihood to the correct formula. But we find that this is not the case. Table 7 shows that on average, the ground truth formula receives a lower reward than the predicted formula, which indicates that further efforts to maximize the score will not result in more accurate parses. This problem is likely to become less severe if future advances produce more accurate planners that place higher likelihood on paths which more efficiently execute a formula.

	Avg. Reward
RL + gen	0.392
IML + gen	0.373
Ground truth	0.369

Table 7: The average rewards received by the predicted formulas for the best performing parameters on the machine-generated data.
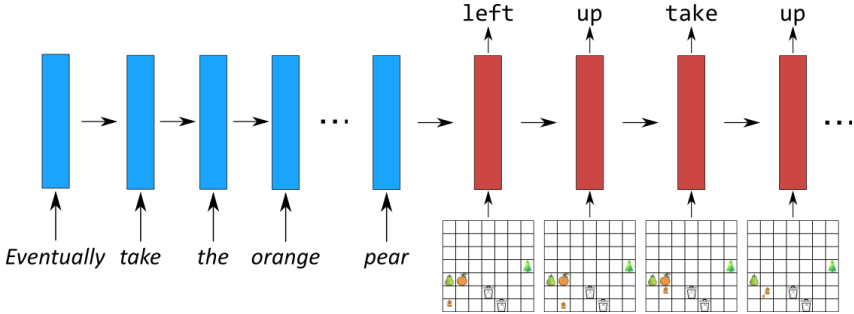
D Comparison to baselines

Figure 3: The baseline sequence-to-sequence architecture which encodes natural language into a high dimensional embedding and then decodes the embedding into a distribution over actions. The recurrent units that serve as the encoder and decoder are both stacked ($n = 2$) LSTMs with hidden-size of 1000. A feature representation of the environment is extracted using a linear layer.

To determine the effectiveness of LTL as an intermediate representation, we compare to a simple end-to-end language-to-action neural network model. We adapt the architecture from [41] as a baseline for comparison (Figure 3). During training, the model takes a natural language command as input and generates a sequence of actions. The loss is simply the likelihood of the ground truth action sequence. We train on the same set of human- and machine-generated sentences as described in Section 4. Teacher-forcing is used when decoding the sequences.

This baseline produces acceptable actions 22% of the time on our human-generated test set and 18.7% of the time on machine-generated test set. These measures fall below our model's performance of 41.3% and 43.3% respectively, as seen in the **Plan** columns of Table 3.

A comparison to a more sophisticated language-to-action planner, a variant of which serves as our executor, can be obtained from [3], which finds the correct execution 50% of the time on the test set in a related, but not identical, environment.

Outside of our model's end-to-end performance, we note that the presence of any intermediate representation is useful since it makes the behavior of the system as a whole more interpretable, which is crucial in ensuring correctness. Moreover, learning a language-to-formalism mapping allows us to generalize to different execution environments, whereas a language-to-action model is only useful in the specific execution context for which it was trained.