

Mapping Language to Programs using Multiple Reward Components with Inverse Reinforcement Learning

Sayan Ghosh Shashank Srivastava

UNC Chapel Hill

{sayghosh, sssrivastava}@cs.unc.edu

Abstract

Mapping natural language instructions to programs that computers can process is a fundamental challenge. Existing approaches focus on likelihood-based training or using reinforcement learning to fine-tune models based on a single reward. In this paper, we pose program generation from language as Inverse Reinforcement Learning. We introduce several interpretable reward components and jointly learn (1) a reward function that linearly combines them, and (2) a policy for program generation. Fine-tuning with our approach achieves significantly better performance than competitive methods using Reinforcement Learning (RL). On the VirtualHome framework, we get improvements of up to 9.0% on the Longest Common Subsequence metric and 14.7% on recall-based metrics over previous work on this framework (Puig et al., 2018). The approach is data-efficient, showing larger gains in performance in the low-data regime. Generated programs are also preferred by human evaluators over an RL-based approach, and rated higher on relevance, completeness, and human-likeness.

1 Introduction

Mapping natural language descriptions to programs is vital for developing agents that can mimic human behavior in the real world. For example, imagine a robot that is instructed to watch television (Figure 1). The robot needs to parse the language instruction into a sequence of commands for navigating to the living room, then going to the sofa, turning on television, etc. For parsing, the robot needs to map objects mentioned in the instruction to objects in the surrounding environment, and then perform the actions required to complete the task. Parsing language into actions has been widely explored in diverse settings (Chen and Mooney, 2011; Chen et al., 2019; Anderson et al., 2018; Tellex et al., 2011). Most recent approaches leverage supervised learning with maximum likelihood estimation, followed

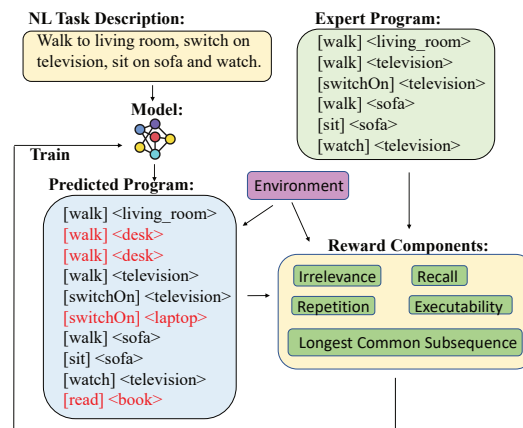


Figure 1: We frame conversion of a natural language task description into a program that can be executed in an environment as Inverse Reinforcement Learning. We design multiple interpretable reward components encoding preferred characteristics in generated programs. A reward function (optimally combining these components) and a policy for program generation are jointly learnt from expert demonstrations/programs.

by a fine-tuning phase of reinforcement learning with a single user-specified reward, encoding signals for things like task completion and executability (Misra et al., 2017; Goyal et al., 2019). To alleviate the problem of sparse rewards, reward shaping terms are often used to encode aspects like distance from goal state or deviation from labeled trajectories (Misra et al., 2017). Reward signals and shaping terms are combined using hyperparameters which need to be manually tuned, and this process becomes tedious. Thus, these approaches can't efficiently leverage multiple possible sources of supervision that might be available. This paper explores an approach to alleviate this problem.

Figure 1 illustrates the core idea of this paper. We use multiple reward components encoding various desired characteristics of a good program to drive program generation, given a natural language task description and an environment. Each compo-

ment consists of a manually defined reward (such as irrelevance, recall, and repetition) with clear semantics. The example in the figure illustrates the need for multiple reward components. We note that the predicted program mentions all the commands in the same sequence as in the expert program. However, it also generates irrelevant commands and repeats steps. Thus even though it scores optimally in the Longest Common Subsequence metric which has been used as a reward for this task (Puig et al., 2018), we need other signals to improve program generation. We investigate subsuming such reward components in an IRL framework to jointly learn: (1) a composite reward function combining the reward components with optimal weights, and (2) a policy that closely mimics the expert. The only supervision involved consists of task descriptions paired with labeled (expert) programs. Automatically learning the weights of reward components becomes increasingly vital as their number increases.

In terms of testbeds, much of previous work on instruction following has focused on block worlds (Misra et al., 2015) and navigation tasks (Chen et al., 2019; Anderson et al., 2018; Misra et al., 2018; Shridhar et al., 2020). We focus on the VirtualHome environment (Puig et al., 2018). In contrast with previous datasets, it contains stepwise instructions and programs for a large number of realistic household activities, such as making coffee or folding laundry. Steps in the programs often involving interacting with objects, or changing the state of the environment. Thus, it offers a rich set of realistic challenges including object interactions and commonsense reasoning for program generation.

Our evaluation shows that using IRL to fine-tune a model leads to significant gains over reinforcement learning. In fact, IRL leads to improved performance even on some metrics that a reinforcement learning policy directly optimizes. Additionally, our approach is data efficient. Experimenting with different dataset sizes reveals that the method generalizes better than baseline methods in low-data scenarios. More significantly, the approach can extend to other domains and provides a general framework for incorporating multiple sources of supervision or inductive biases about a task.

Our contributions are:

- We pose mapping task descriptions to programs as an IRL problem, and learn a composite reward function combining semantically interpretable

characteristics of expert programs¹.

- We achieve up to 9% increase in the Longest Common Subsequence (LCS) metric w.r.t. previous methods. Programs generated by our approach are qualitatively better and are preferred by human evaluators.
- Our approach is data-efficient in limited data scenarios compared to previous methods.

2 Related Work

Semantic Parsing and Instruction Following:

Parsing natural language to programs has been explored in diverse settings. Common semantic parsing applications include text-to-SQL (Zhong et al., 2017; Yu et al., 2018, 2019b,a) text-to-code (Yin et al., 2018; Shin et al., 2019), robot navigation and interaction tasks (Misra et al., 2016; Nyga et al., 2018; Squire et al., 2015). Other tasks involve mapping instructions to actions in simple environments (Artzi and Zettlemoyer, 2013; Chen and Mooney, 2011; Misra et al., 2015; Malmaud et al., 2014). Instruction following is also a crucial part of complex Vision and Language Navigation (VLN) tasks (Anderson et al., 2018; Misra et al., 2018; Chen et al., 2019; Nguyen and Daumé III, 2019). Significant work has explored developing models that can use additional context in this space (Fried et al., 2018; Ke et al., 2019; Wang et al., 2019; Nguyen and Daumé III, 2019; Nguyen et al., 2019; Thomason et al., 2020; Ma et al., 2019a,b). Compared to environments like Room2Room (Anderson et al., 2018), VirtualHome (Puig et al., 2018) contains realistic activities and involves dynamic state changes and interacting with the objects to complete tasks.

Training paradigms: Encoder-decoder architectures (Sutskever et al., 2014) are the dominant modeling paradigm for instruction following tasks. Models are usually pre-trained with Maximum Likelihood Estimation, and fine-tuned with Reinforcement Learning (Puig et al., 2018; Misra et al., 2017; Zhong et al., 2017) using REINFORCE (Williams, 1992). There has been limited work on using IRL (Abbeel and Ng, 2004; Ziebart et al., 2008; Finn et al., 2016; Shi et al., 2018; Li et al., 2018; Ghosh et al., 2021) in NLP. Fu et al. (2019) use IRL to learn a language-guided reward by using a deep neural network to parameterize a reward function. Our work is closest to Ghosh et al. (2021), which formulates reward components

¹Code and dataset splits for the paper are available at https://github.com/sgdgp/VirtualHome_IRL

for table-to-text generation and learns their optimal linear combination using IRL. Following similar ideas, we define multiple interpretable reward components representing desired characteristics of good programs in our context and learn their optimal linear combination while jointly learning a policy for program generation. However, we differ from Ghosh et al. (2021) by applying IRL on a different downstream task (program generation from natural language instruction) and further show that using IRL as opposed to RL can lead to higher gains in limited labelled data scenarios.

3 Method

Each example in our data consists of a natural language description of a task to be performed in an environment, along with its corresponding program (see Figure 1). In VirtualHome, a program is a sequence of commands that can be executed to complete the task being described. A command is an action-object-object triplet, where one or both of the object arguments can be empty based on the action. For example, in the command “[walk] ⟨living_room⟩” the action “walk” is followed by an object argument “living_room”. Similarly the command “[putBack] ⟨plate⟩ ⟨cupboard⟩” has two object arguments and “sleep” has none.

Let $D = \{d_i; 1 \leq i \leq m\}$ denote the natural language description of the task to be performed in VirtualHome, where d_i denotes the i^{th} token. Let $C = \{c_j; 1 \leq j \leq n\}$ denote the corresponding program, where c_j denotes j^{th} command. Our objective is to predict program C given the text description D and an environment E . Previous approaches performed likelihood-based pre-training followed by fine-tuning using reinforcement learning (Puig et al., 2018). Instead, we use IRL to learn the reward function from the expert programs and fine-tune the program generation policy simultaneously. We design a set of rewards with clear semantics that encode the desired characteristics of a good program in our context. In the rest of this section, we describe the model architecture, provide definitions of reward components, and finally describe model training using IRL.

3.1 Model architecture

We use an encoder-decoder architecture (illustrated in Figure 2). The encoder is modeled as a RNN with LSTM cells, and provides a representation of the task description. The decoder is also mod-

eled by using LSTM cells for predicting commands at each time step. We pre-compute a representation/embedding for every possible command by simply averaging the word2vec (Mikolov et al., 2013) embeddings of its action and object arguments. Let ψ denote the function that maps a command to its embedding.

Description encoder: Task descriptions are tokenized, and each token represented using its word2vec embedding. Tokens are passed through a LSTM network to get the representation of the description. The final hidden state of the encoder is used to initialize the hidden state of the decoder. **Attention over text encoder:** We attend over the sequence of encoder hidden states, h^{enc} , for every step t , using the previous hidden state of the decoder, h_{t-1} . We obtain the context vector, z_t , as:

$$\alpha^t = \text{softmax}(\psi(c_{t-1})^T W_{att}[h_{t-1}, h^{enc}]) \quad (1)$$

$$z_t = \sum_{j=1}^m \alpha_j^t h_j^{enc} \quad (2)$$

where $\psi(c_{t-1})$ is the embedding for the previous command and W_{att} is a learnable matrix.

Decoder: The decoder is an RNN with LSTM cells. The LSTM takes as input a concatenation of the previous command embedding, $\psi(c_{t-1})$, and a context vector, z_t . The operation of the decoder at step t can be described as :

$$r_t = W_{dec} \times LSTM([\psi(c_{t-1}), z_t], h_{t-1}, s_{t-1}) \quad (3)$$

where h_{t-1} and s_{t-1} are the hidden and cell states at step $t - 1$ and W_{dec} is a learnable weight matrix that defines a linear transformation from LSTM output space to command-embedding space.

Command prediction: The decoder’s output, r_t , is a vector in the command-embedding space. We calculate its cosine similarity with each possible command in the environment. The most similar command is chosen as the output at every step.

3.2 Reward components

Defining rewards for reinforcement learning requires identifying preferred characteristics of outputs. Intuitively, correct programs corresponding to a task description have multiple characteristics: they should contain objects and actions similar to those mentioned in the task description in their

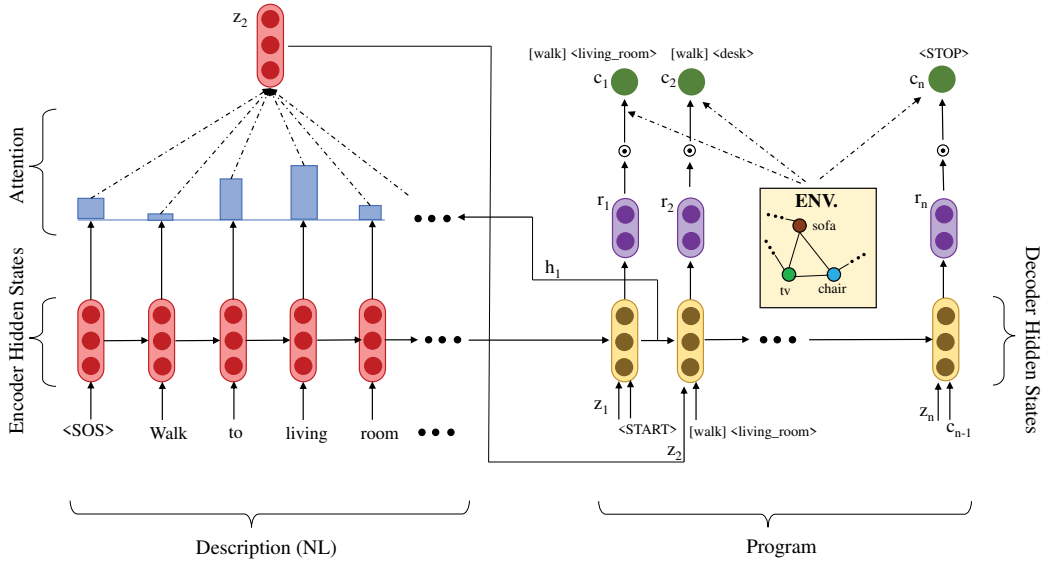


Figure 2: We follow a similar architecture to Puig et al. (2018). The final hidden state of the encoder is used to initialise the hidden state of the decoder. Each decoder cell takes in the previous step command embedding ($\psi(c_{t-1})$) and a context vector z_t (computed by attention over encoder hidden states) as input. The output of the decoder (r_t) is used to predict the next command by computing cosine similarity (denoted by \odot) to embeddings of all possible commands. During training, the most probable action is chosen. While testing, we explore two variations of inference : (1) regular beam search, and (2) beam search enforcing executability as a hard constraint.

arguments and actions, but not contain many unmentioned objects and actions. Moreover, the order of the commands should be semantically feasible. Also, the program should not repeat cycles of commands (a common problem in conditional generation models) and be executable in the given environment. Finally, these programs should be similar to expert programs. A key insight is that these ideas can be captured quantitatively with reward components, which we describe next. In the following, C^{pred} and C^{gt} denote the predicted and ground truth (expert) program. O^{pred} denotes the set of objects in C^{pred} and O^{gt} denotes the set of objects (nouns) in the task description.

- **Recall from description:** An object in O^{pred} is similar to some object in O^{gt} if their cosine similarity is above a threshold. This reward value is the fraction of objects in O^{gt} similar to any object in O^{pred} .
- **Irrelevance:** For a command predicted at step t if the object present in the command is not relevant w.r.t. O^{gt} a penalty of -1 is given to this command. An object in the command is related to the description if its cosine similarity with at least one object from O^{gt} is greater than a threshold. The total penalty for a predicted program is normalized by its length.
- **Repetition:** We penalize the programs for rep-

etition. For every predicted command c_t (for $t \geq 2$) if the bigram (c_{t-1}, c_t) is not unique, a penalty of -1 is given to c_t . The total penalty is normalized by the length of the sequence.

- **Longest Common Subsequence (LCS):** We quantify how close the predicted program is to the expert program by finding the LCS score between C^{pred} and C^{gt} normalized by length.
- **Recall from program:** Fraction of the set of commands in C^{gt} contained in C^{pred} .
- **Executability :** Program is given reward (+1) if it is executable in the given environment.

The total reward for a program is a linear combination of the above components.

3.3 Training

We pre-train the model in a supervised manner, followed by finetuning it by using Maximum Entropy IRL. Details on implementation and hyperparameters are provided in the supplementary material.

3.3.1 Supervised training

We optimize for the maximum likelihood estimation (MLE) objective using a cross-entropy loss at each step. We use teacher-forcing (Bengio et al., 2015) during training, by using ground truth commands at a step as decoder inputs for the next step.

3.3.2 Maximum Entropy (MaxEnt) IRL

We fine-tune the model using MaxEnt IRL (Ziebart et al., 2008) to mimic salient characteristics of expert demonstrations. We formulate program generation as IRL, where at each step t , a command is generated. The reward is not observed but learned from expert demonstrations in the training data. IRL consists of two alternating steps: (1) **Reward approximation**: estimate the underlying reward function using the expert demonstration and the current policy. (2) **Reinforcement Learning**: use the estimated reward function to optimize the policy for program generation.

Reward approximation: Following the standard MaxEnt IRL framework, we assume that expert programs are drawn from a distribution $p_\phi(C|D, E)$.

$$p_\phi(C|D, E) = \frac{1}{Z} \exp(R_\phi(C|D, E))$$

$$\text{and } Z = \int_C \exp(R_\phi(C|D, E)) \quad (4)$$

where the reward function, $R_\phi(C|D, E)$ has parameters ϕ , and Z is the partition function. Also, the total reward of a program is sum of rewards at each step. Given a policy for program generation $q_\theta(C|D, E)$, our objective is to maximise the log-likelihood of the samples in the training set (Equation 5).

$$J_r(\phi) = \frac{1}{N} \sum_{n=1}^N \log(p_\phi(C_n|D_n, E_n))$$

$$= \frac{1}{N} \sum_{n=1}^N R_\phi(C_n|D_n, E_n) - \log Z \quad (5)$$

Thus, gradient w.r.t. reward parameters is given by

$$\nabla_\phi J_r(\phi) = \frac{1}{N} \sum_n \nabla_\phi R_\phi(C_n|D_n, E_n)$$

$$- \frac{1}{Z} \int_C \exp(R_\phi(C|D, E)) \nabla_\phi R_\phi(C|D, E) dC$$

$$= \mathbb{E}_{C \sim p_{data}} \nabla_\phi R_\phi(C|D, E) - \mathbb{E}_{C \sim p_\phi(C|D, E)} \nabla_\phi R_\phi(C|D, E) \quad (6)$$

We use importance sampling to approximate the gradient of the log partition function when drawing programs from the distribution of generated programs. The importance weight β_i for a generated program C_i is given by

$$\beta_i \propto \frac{\exp(R_\phi(C_i|D_i, E_i))}{q_\theta(C_i|D_i, E_i)} \quad (7)$$

Using importance sampling, the gradient of the objective function can be approximated as

$$\nabla_\phi J_r(\phi) = \frac{1}{N} \sum_{i=1}^N \nabla_\phi R_\phi(C_i|D_i, E_i) - \frac{1}{\sum_j \beta_j} \sum_{j=1}^M \beta_j \nabla_\phi R_\phi(C'_j|D'_j, E'_j) \quad (8)$$

where C_i and C'_j are drawn from training data and $q_\theta(C|D, E)$ respectively. In this work, we also assume that R_ϕ is a linear combination of the reward components defined in Section 3.2.

$$R_\phi(C|D, E) = \sum_{t=1}^{\tau} \phi^T \Psi^t \quad (9)$$

where ϕ is a weight vector, Ψ^t is the vector of reward component values at step t and τ denotes total time-steps. Owing to the linear formulation, the weight update for each reward component simply becomes a difference between the expected expert and the expected roll-out reward component. We use N expert programs and sample M programs from our policy $q_\theta(C|D, E)$. The weight update for a component ψ is:

$$\nabla_\phi J_r(\phi)_\psi = \frac{1}{N} \sum_{i=1}^N \psi_i - \frac{1}{\sum_j \beta_j} \sum_{j=1}^M \beta_j \psi'_j \quad (10)$$

where ψ_i and ψ'_j are total value of reward component over all steps for i^{th} expert program j^{th} generated program respectively. These weights are learned in a data-driven approach when the supervised model is fine-tuned using MaxEnt IRL.

Reinforcement Learning: For reinforcement learning stage of IRL to learn the policy for program generation we use Self Critical Sequence Training (SCST) (Rennie et al., 2017). We also perform entropy regularisation (Williams, 1992; Nachum et al., 2017) when training our model. The objective of the program generator ($J_g(\theta)$) is

$$J_g(\theta) = \mathbb{E}_{C \sim q_\theta(C|D, E)} [R_\phi(C|D, E)] + \delta H(q_\theta(C|D, E)) \quad (11)$$

where δ is a hyper-parameter (0.05 for our experiments) and $H(q_\theta(C|D, E))$ is the entropy of $q_\theta(C|D, E)$.

Summary of training process: The model training consists of an iterative process with two steps. In the first step, we fix the program generation policy, and use programs sampled from the current

policy to update the weights corresponding to each reward component. For this step, we assume that the expert programs’ distribution is exponential in their rewards. The sampled programs from the current policy are used to approximate the log partition function required in making the updates to the reward components. In the second step, we fix the composite linear reward function (i.e. the weights of each reward component), and use it to update the program generation policy for the current reward function using policy-gradient updates.

4 Experiments and Results

In this section, we describe the dataset, metrics used for quantitative evaluation followed by a comparative evaluation of the proposed method with baselines and results of human evaluation².

4.1 Data

The VirtualHome dataset consists of programs for household tasks, accompanied by task descriptions and environments with necessary preconditions to carry out the tasks. VirtualHome provides a graph-based simulator, where each node represents entities (like rooms, object etc.) and the simulator tracks changes of attributes and interactions. However, there are some programs not executable in any home scenario. We remove these, and also keep a single program when multiple programs appear against a single task description. Thus our modified dataset consists of unique executable programs which is divided into train, validation, and test splits of sizes 697, 185 and 500 samples respectively. To study the performance of our method in a low data regime, we randomly sample and form two smaller subsets of the training set of sizes 221 and 70. To calculate rewards, we extract nouns from task descriptions using spaCy’s POS tagger.

4.2 Metrics

We use the following for quantitative evaluation.

- **Normalized Longest Common Subsequence (LCS):** This is the length of the longest common subsequence of the predicted and the expert programs normalized by the maximum of their lengths. We find LCS values of actions, objects, and commands separately. The mean LCS score is the average of these.

²Code and data splits are available at https://github.com/sgdgp/VirtualHome_IRL

- **Edit Distance (ED):** This metric denotes edit distance between the predicted and the expert programs normalized by the maximum of the two programs’ lengths.
- **Recall from program :** We calculate the average reward value for recall from program reward component as described in Section 3.2.
- **Executability :** Using the available preconditions for the programs, we obtain the percentage of generated programs which are executable in the graph-based simulator of VirtualHome.

For all metrics except **ED**, a higher value is better. However, note that our metrics do not explicitly capture task completion. For example, the executability metric will be 1 even if the task is incomplete as long as the predicted program is executable in the environment.

4.3 Quantitative Evaluation

Table 1 shows a comparative evaluation of our IRL approach with baseline methods. As baselines, we use (1) Random command generation, (2) MLE approach from Puig et al. (2018), (3) RL-based fine-tuning approach from Puig et al. (2018) using LCS and execution rewards and (4) RL using all reward components (and a variant excluding executability component) weighted uniformly. We explore two variants of IRL: using all reward components, and all except the execution reward component. The table shows results with two types of inference: (1) we perform beam search with beam-size of 3 and (2) we enforce executability as a hard constraint (rows in the table marked with \diamond) while doing beam search, i.e., at any step, non-executable programs are dropped from the beam. We can make the following conclusions on the quantitative performance of IRL over baselines:

- **Significant improvement over prior work:** Table 1 shows that IRL achieves relative improvement of 9% on the mean LCS score of the test set against RL(LCS) (Puig et al., 2018) when using 697 training samples. The gains over RL (LCS) are statistically significant ($p < 0.01$). The trend remains same even with reduced training set size.
- **Better gains in low-data regime:** IRL outperforms RL significantly when not using the execution reward for all training set sizes in mean LCS and recall-based metrics. On using all the reward components the performances of RL and IRL models are comparable when the training data size is largest (697) even though IRL is slightly

MODEL	MEAN LCS	ED	RECALL PROG.	EXEC.
Random	0.075 ±0.002	0.997 ±0.000	0.009 ±0.004	0.003 ±0.003
697 training samples				
MLE	0.368 ±0.015	0.750 ±0.016	0.357 ±0.023	0.348 ±0.045
RL (LCS)	0.400 ±0.044	0.715 ±0.046	0.404 ±0.067	0.379 ±0.055
RL (LCS) [◊]	0.410 ±0.009	0.701 ±0.014	0.423 ±0.015	1.000 ±0.000
RL (LCS + exec)	0.385 ±0.055	0.731 ±0.056	0.373 ±0.074	0.373 ±0.081
RL (LCS + exec) [◊]	0.372 ±0.065	0.737 ±0.066	0.363 ±0.088	1.000 ±0.000
RL (All - Exec.)	0.431 ±0.007	0.694 ±0.013	0.460 ±0.015	0.392 ±0.025
RL (All)	0.423 ±0.016	0.698 ±0.017	0.448 ±0.022	0.394 ±0.027
Ours (IRL: All - Exec.)	0.436* ±0.012	0.687 ±0.015	0.454 ±0.009	0.402 ±0.039
Ours (IRL: All - Exec.) [◊]	0.427 ±0.019	0.683 ±0.014	0.437 ±0.019	1.000 ±0.000
Ours (IRL: All)	0.428 ±0.027	0.693 ±0.029	0.437 ±0.032	0.410 ±0.029
Ours (IRL: All) [◊]	0.424 ±0.022	0.687 ±0.023	0.432 ±0.024	1.000 ±0.000
221 training samples				
MLE	0.283 ±0.007	0.837 ±0.005	0.234 ±0.012	0.282 ±0.049
RL (LCS)	0.324 ±0.016	0.796 ±0.010	0.272 ±0.031	0.286 ±0.037
RL (LCS) [◊]	0.316 ±0.012	0.799 ±0.010	0.276 ±0.018	1.000 ±0.000
RL (LCS + exec)	0.319 ±0.018	0.798 ±0.014	0.276 ±0.025	0.309 ±0.038
RL (LCS + exec) [◊]	0.313 ±0.016	0.801 ±0.011	0.271 ±0.019	1.000 ±0.000
RL (All - Exec.)	0.322 ±0.018	0.802 ±0.014	0.290 ±0.024	0.289 ±0.054
RL (All)	0.328 ±0.009	0.796 ±0.009	0.299 ±0.016	0.273 ±0.025
Ours (IRL: All - Exec.)	0.342* ±0.013	0.780 ±0.014	0.312* ±0.018	0.299 ±0.043
Ours (IRL: All - Exec.) [◊]	0.336 ±0.010	0.780 ±0.010	0.305 ±0.012	1.000 ±0.000
Ours (IRL: All)	0.334 ±0.018	0.791 ±0.019	0.301 ±0.024	0.286 ±0.032
Ours (IRL: All) [◊]	0.327 ±0.018	0.790 ±0.015	0.294 ±0.024	1.000 ±0.000
70 training samples				
MLE	0.190 ±0.009	0.919 ±0.005	0.121 ±0.006	0.221 ±0.066
RL (LCS)	0.231 ±0.005	0.892 ±0.006	0.158 ±0.010	0.152 ±0.047
RL (LCS) [◊]	0.221 ±0.005	0.895 ±0.007	0.145 ±0.010	1.000 ±0.000
RL (LCS + exec)	0.223 ±0.009	0.900 ±0.008	0.146 ±0.011	0.191 ±0.044
RL (LCS + exec) [◊]	0.216 ±0.008	0.899 ±0.007	0.137 ±0.011	1.000 ±0.000
RL (All - Exec.)	0.226 ±0.012	0.899 ±0.007	0.156 ±0.017	0.142 ±0.035
RL (All)	0.224 ±0.014	0.896 ±0.011	0.155 ±0.018	0.182 ±0.033
Ours (IRL: All - Exec.)	0.246* ±0.008	0.881 ±0.009	0.180* ±0.009	0.152 ±0.039
Ours (IRL: All - Exec.) [◊]	0.236 ±0.008	0.883 ±0.006	0.166 ±0.004	1.000 ±0.000
Ours (IRL: All)	0.241 ±0.019	0.745* ±0.297	0.176 ±0.024	0.148 ±0.027
Ours (IRL: All) [◊]	0.234 ±0.020	0.880 ±0.014	0.165 ±0.021	1.000 ±0.000
Semi-supervised learning				
Sup. #70 + Unsup. #697	0.127 ±0.023	0.968 ±0.011	0.063 ±0.026	0.103 ±0.017

Table 1: Test set performance for different training data sizes. MLE, RL(LCS), RL(LCS + exec) models are adapted from Puig et al. (2018). [◊] denotes executability as hard constraint during inference. * denotes statistical significance ($p < 0.01$) of IRL vs RL baselines from Puig et al. (2018) using the Wilcoxon signed-rank test.

better in mean LCS score. However, we find that IRL’s improvement over RL gets increasingly larger as the training data size drops. In the low data regime (training data size of 70), we find that using IRL is more effective and leads to higher gains than RL.

- **Improvement in other rewards:** Table 1 shows that IRL helps to get better recall from program in addition to better LCS score. The gains are increased as training data gets reduced.
- **Multiple reward components leads to better programs:** IRL (without execution reward component) improves significantly ($p < 0.01$, Wilcoxon signed-rank test) on recall from pro-

gram and mean LCS scores as compared to RL baseline (using only LCS). This is an interesting result, since the RL model optimizes directly for an LCS reward. This may suggest that the composite reward function in IRL might be leading to better optimization trajectories. In many scenarios, having an exact execution reward is not feasible (often due to lack of a robust simulator). We find that the execution reward signal does not contribute much, and other forms of supervision from the expert programs lead to better performance.

- **Automatic weight learning through IRL is helpful:** To judge usefulness of the weights of

each reward component learned by IRL we implement a RL baseline, RL (All), with all the reward components weighted uniformly and compare it with the IRL models. We observe that doing IRL outperforms this baseline in mean LCS score irrespective of the training set size considered. When using all reward components except execution the relative gain achieved by IRL against RL in mean LCS score is 1.16%, 6.21% and 8.85% when trained with 697, 221 and 70 training samples respectively. Thus we see the learning the weights of each reward component proves more fruitful when we reduce the number of training samples.

- **Semi-supervised learning:** We perform an experiment to explore unsupervised reward components that don't depend on the ground truth program (last row in Table 1). We fine-tune the supervised model trained using 70 training samples by using only three unsupervised reward components - irrelevance, recall from description, and repetition. Fine-tuning is done on all 697 training samples. In this case, we note that irrelevance reward improves from -0.479 to -0.397 and recall from description reward improves from 0.104 to 0.111 as compared to values from MLE model trained on 70 samples. However, the absence of a signal for enforcing sequential structure degenerates the predicted sequence of commands, and LCS score drops significantly.

This scenario also tests generalization as many commands in test will be unseen in the ground-truth programs corresponding to the small subset of training samples (72.9% of the commands in the test set are unseen). However, to test true generalization, IRL-based models need to be tested in unseen scenarios (possibly in new environments, with new objects, unseen tasks and new compositions of individual commands). We do not explore this direction in this work.

Weights Analysis: We analyze relative weights learned by IRL and find that it assigns highest weights to the LCS (normalized weight value of 0.50) and recall from program reward components (normalized weight value of 0.35). The next largest weights are for the repetition and execution components (normalized weight values of 0.14 and 0.06, respectively). We see that the irrelevance penalty gets assigned a small negative weight, thus, allowing a small number of unrelated objects to show up in the program as long as the LCS and recall scores

are not affected. Qualitatively, in many instances, IRL fine-tuning enables identification of the correct action verb when choosing the command as opposed to RL. High reward weights specifically to LCS and recall from program help to improve ordering of commands and prevent unrelated objects from appearing in the programs.

Ablation Study: Table 2 shows performance for ablations grouping various reward components during IRL. When only irrelevance, recall from description, and repetition are used the training is often unstable but can have higher executability owing to generation of empty programs which are trivially executable. We observe that without LCS as a reward component, program generation gradually degenerates. Hence for other ablation experiments, we keep LCS and couple it with other rewards. IRL using LCS coupled with any other reward component shows improvements in the mean LCS score w.r.t RL (LCS) baseline. LCS coupled with either of the recall-based rewards performs better than other reward groups. Using all reward components we get the best scores on validation set except on executability metric.

4.4 Human Evaluation

We perform human evaluation using Amazon Mechanical Turk to explore qualitative differences between programs generated by our approach (IRL) and baseline RL approach from Puig et al. (2018). Turkers are given a task description followed by two program samples, one from each method (order is randomized). The turkers are asked to rate both the programs w.r.t a few criteria on a scale of 1 to 5 (where 5 denotes highest) and also choose a preferred program. The criteria for rating are:

- **Relatedness to description:** relatedness of objects in the program w.r.t. the task description.
- **Human-likeness:** how closely a generated program mimics human behavior.
- **Task completion:** how much of the task is accomplished by the program.

We rate batches of 50 programs by 5 turkers. The programs are generated by RL(LCS) and IRL (All - Exec.) models fine-tuned on 697 training examples. Executability is not enforced during inference for both the models. Table 3 shows the human evaluation results. IRL programs are overall rated better for all three criteria. In general, IRL generates more relevant commands and prevents repetitions due to irrelevance and repetition reward compo-

MODEL	MEAN LCS	ED	RECALL PROG.	EXEC.
RL (LCS)	0.265 ±0.006	0.869 ±0.005	0.198 ±0.008	0.112 ±0.024
Irrel. + Recall desc.	0.081 ±0.015	0.984 ±0.005	0.033 ±0.008	0.210 ±0.032
Irrel. + Recall desc. + Rep.	0.147 ±0.034	0.947 ±0.014	0.074 ±0.030	0.301 ±0.301
LCS + Irrel.	0.267 ± 0.005	0.868 ± 0.010	0.192 ± 0.007	0.131 ±0.031
LCS + Recall desc.	0.271 ±0.010	0.862 ±0.011	0.201 ±0.016	0.139 ±0.012
LCS + Rep.	0.271 ±0.010	0.868 ±0.008	0.202 ±0.013	0.130 ±0.041
LCS + Recall from prog	0.274 ±0.009	0.864 ±0.008	0.209 ±0.013	0.095 ±0.020
All reward comp. - Exec	0.273 ±0.007	0.865 ±0.003	0.213 ±0.010	0.105 ±0.033
All reward comp.	0.282 ±0.003	0.857 ±0.006	0.217 ±0.005	0.114 ±0.018

Table 2: Ablation results on validation set when trained on training set of size 70

CRITERION	RL	Ours (IRL)
Relatedness	3.09	3.18
Human-likeness	3.41	3.50
Task completion	2.84	2.95
Pref. count (out of 250)	106	144

Table 3: Human evaluation results

nents. Also, programs generated by the IRL model are preferred in 57.6% of the cases (statistically significant at $p < 0.01$, Binomial test).

5 Conclusion

We explored an approach for incorporating diverse reward components in instruction following tasks. Such components can often be defined by a domain expert, and encode inductive biases about a problem. Since reward weights are learned, these models can be robust to spurious reward components. However, the issue of possibly adversarial reward components remains to be explored. While the approach requires access to expert examples, since we focus on scenarios involving RL-based fine-tuning, these are presumed to be already available. The approach can potentially generalize to other domains and applications, and can be fertile ground for directions of future research.

References

Pieter Abbeel and Andrew Y Ng. 2004. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*, page 1.

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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3674–3683.

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

Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. 2015. Scheduled sampling for sequence prediction with recurrent neural networks. In *Advances in Neural Information Processing Systems*, pages 1171–1179.

David L Chen and Raymond J Mooney. 2011. Learning to interpret natural language navigation instructions from observations. In *Twenty-Fifth AAAI Conference on Artificial Intelligence*.

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 Conference on Computer Vision and Pattern Recognition*, pages 12538–12547.

Chelsea Finn, Sergey Levine, and Pieter Abbeel. 2016. Guided cost learning: Deep inverse optimal control via policy optimization. In *International conference on machine learning*, pages 49–58.

Daniel Fried, Ronghang Hu, Volkan Cirik, Anna Rohrbach, Jacob Andreas, Louis-Philippe Morency, Taylor Berg-Kirkpatrick, Kate Saenko, Dan Klein, and Trevor Darrell. 2018. Speaker-follower models for vision-and-language navigation. In *Advances in Neural Information Processing Systems*, pages 3314–3325.

Justin Fu, Anoop Korattikara, Sergey Levine, and Sergio Guadarrama. 2019. From language to goals: Inverse reinforcement learning for vision-based instruction following. In *International Conference on Learning Representations*.

Sayan Ghosh, Zheng Qi, Snigdha Chaturvedi, and Shashank Srivastava. 2021. How helpful is inverse reinforcement learning for table-to-text generation? 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 71–79, Online. Association for Computational Linguistics.

- Prasoon Goyal, Scott Niekum, and Raymond J Mooney. 2019. Using natural language for reward shaping in reinforcement learning. *arXiv preprint arXiv:1903.02020*.
- Liyiming Ke, Xiujun Li, Yonatan Bisk, Ari Holtzman, Zhe Gan, Jingjing Liu, Jianfeng Gao, Yejin Choi, and Siddhartha Srinivasa. 2019. Tactical rewind: Self-correction via backtracking in vision-and-language navigation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6741–6749.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *ICLR (Poster)*.
- Zichao Li, Xin Jiang, Lifeng Shang, and Hang Li. 2018. [Paraphrase generation with deep reinforcement learning](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3865–3878, Brussels, Belgium. Association for Computational Linguistics.
- Chih-Yao Ma, Jiasen Lu, Zuxuan Wu, Ghassan Al-Regib, Zsolt Kira, Richard Socher, and Caiming Xiong. 2019a. Self-monitoring navigation agent via auxiliary progress estimation. *arXiv preprint arXiv:1901.03035*.
- Chih-Yao Ma, Zuxuan Wu, Ghassan AlRegib, Caiming Xiong, and Zsolt Kira. 2019b. The regretful agent: Heuristic-aided navigation through progress estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6732–6740.
- Jonathan Malmaud, Earl Wagner, Nancy Chang, and Kevin Murphy. 2014. Cooking with semantics. In *Proceedings of the ACL 2014 Workshop on Semantic Parsing*, pages 33–38.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Dipendra Misra, Andrew Bennett, Valts Blukis, Eyvind Niklasson, Max Shatkhin, and Yoav Artzi. 2018. Mapping instructions to actions in 3d environments with visual goal prediction. *arXiv preprint arXiv:1809.00786*.
- Dipendra Misra, John Langford, and Yoav Artzi. 2017. Mapping instructions and visual observations to actions with reinforcement learning. *arXiv preprint arXiv:1704.08795*.
- Dipendra Misra, Kejia Tao, Percy Liang, and Ashutosh Saxena. 2015. Environment-driven lexicon induction for high-level instructions. 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 992–1002.
- Dipendra K Misra, Jaeyong Sung, Kevin Lee, and Ashutosh Saxena. 2016. Tell me dave: Context-sensitive grounding of natural language to manipulation instructions. *The International Journal of Robotics Research*, 35(1-3):281–300.
- Ofir Nachum, Mohammad Norouzi, Kelvin Xu, and Dale Schuurmans. 2017. Bridging the gap between value and policy based reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 2775–2785.
- Khanh Nguyen and Hal Daumé III. 2019. Help, anna! visual navigation with natural multimodal assistance via retrospective curiosity-encouraging imitation learning. *arXiv preprint arXiv:1909.01871*.
- Khanh Nguyen, Debadeepta Dey, Chris Brockett, and Bill Dolan. 2019. Vision-based navigation with language-based assistance via imitation learning with indirect intervention. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 12527–12537.
- Daniel Nyga, Subhro Roy, Rohan Paul, Daehyung Park, Mihai Pomarlan, Michael Beetz, and Nicholas Roy. 2018. Grounding robot plans from natural language instructions with incomplete world knowledge. In *Conference on Robot Learning*, pages 714–723.
- Travis E Oliphant. 2006. *A guide to NumPy*, volume 1. Trelgol Publishing USA.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc.
- Xavier Puig, Kevin Ra, Marko Boben, Jiaman Li, Tingwu Wang, Sanja Fidler, and Antonio Torralba. 2018. Virtualhome: Simulating household activities via programs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8494–8502.
- Steven J Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, and Vaibhava Goel. 2017. Self-critical sequence training for image captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7008–7024.
- Zhan Shi, Xinchu Chen, Xipeng Qiu, and Xuanjing Huang. 2018. Toward diverse text generation with inverse reinforcement learning. *arXiv preprint arXiv:1804.11258*.
- Eui Chul Shin, Miltiadis Allamanis, Marc Brockschmidt, and Alex Polozov. 2019. Program synthesis and semantic parsing with learned

- code idioms. In *Advances in Neural Information Processing Systems*, pages 10825–10835.
- 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 *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Shawn Squire, Stefanie Tellex, Dilip Arumugam, and Lei Yang. 2015. Grounding english commands to reward functions. In *Robotics: Science and Systems*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.
- Stefanie Tellex, Thomas Kollar, Steven Dickerson, Matthew R. Walter, Ashis Gopal Banerjee, Seth Teller, and Nicholas Roy. 2011. Understanding natural language commands for robotic navigation and mobile manipulation. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*, AAAI'11, page 1507–1514. AAAI Press.
- Jesse Thomason, Michael Murray, Maya Cakmak, and Luke Zettlemoyer. 2020. Vision-and-dialog navigation. In *Conference on Robot Learning*, pages 394–406.
- Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, CJ Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. 2020. [SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python](#). *Nature Methods*, 17:261–272.
- Xin Wang, Qiuyuan Huang, Asli Celikyilmaz, Jianfeng Gao, Dinghan Shen, Yuan-Fang Wang, William Yang Wang, and Lei Zhang. 2019. Reinforced cross-modal matching and self-supervised imitation learning for vision-language navigation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6629–6638.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256.
- Pengcheng Yin, Bowen Deng, Edgar Chen, Bogdan Vasilescu, and Graham Neubig. 2018. [Learning to mine aligned code and natural language pairs from stack overflow](#). In *International Conference on Mining Software Repositories*, MSR, pages 476–486. ACM.
- Tao Yu, Rui Zhang, He Yang Er, Suyi Li, Eric Xue, Bo Pang, Xi Victoria Lin, Yi Chern Tan, Tianze Shi, Zihan Li, et al. 2019a. Cosql: A conversational text-to-sql challenge towards cross-domain natural language interfaces to databases. *arXiv preprint arXiv:1909.05378*.
- Tao Yu, Rui Zhang, Kai Yang, Michihiro Yasunaga, Dongxu Wang, Zifan Li, James Ma, Irene Li, Qingning Yao, Shanelle Roman, Zilin Zhang, and Dragomir Radev. 2018. Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing and text-to-sql task. In *"Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing"*, "Brussels, Belgium". "Association for Computational Linguistics".
- Tao Yu, Rui Zhang, Michihiro Yasunaga, Yi Chern Tan, Xi Victoria Lin, Suyi Li, Irene Li Heyang Er, Bo Pang, Tao Chen, Emily Ji, Shreya Dixit, David Proctor, Sungrok Shim, Vincent Zhang Jonathan Kraft, Caiming Xiong, Richard Socher, and Dragomir Radev. 2019b. Sparc: Cross-domain semantic parsing in context. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Florence, Italy. Association for Computational Linguistics.
- Victor Zhong, Caiming Xiong, and Richard Socher. 2017. Seq2sql: Generating structured queries from natural language using reinforcement learning. *arXiv preprint arXiv:1709.00103*.
- Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, and Anind K Dey. 2008. Maximum entropy inverse reinforcement learning. In *Aaai*, volume 8, pages 1433–1438. Chicago, IL, USA.

A Appendix

A.1 Training Details

Model parameters The task description is tokenized into words. We do not remove stop words or lemmatize words. Embeddings for words are obtained by using pretrained word2vec (Mikolov et al., 2013) vectors (300 dimensional). Next, these embeddings are passed into an encoder RNN made of LSTM cells. The LSTM network is unidirectional with hidden dimension of 100. The decoder RNN is also unidirectional with hidden dimension as 100. The last hidden state of the encoder RNN is used to initialise the hidden state of the decoder. We train each model for 400 epochs using the Adam optimizer (Kingma and Ba, 2015). We choose the hyperparameters and best epoch for each model by obtaining results on the validation set using beam

size of 3 and not enforcing executability. Since we adapt the model from Puig et al. (2018) the size of network is still same with around 3M parameters.

Hyper-parameter tuning We use five different random seeds for five trials of each experiment: 42, 101, 123, 2020 and 2021. Batch size and learning rate are manually tuned in the range {64, 128, 256} and {0.001, 0.0005, 0.0001} respectively. Based on the results on validation splits we chose batch size as 256 and learning rate as 0.001 to report the results. We use a weight of 0.05 for entropy regularization during policy gradient after trying out weights in the range {0.001, 0.005, 0.01, 0.05, 0.1, 0.5, 10}. We chose 0.05 as weight for entropy regularization on the basis of model performance on validation set which peaks for 0.05 when the model is trained on 70 samples. For Adam optimizer the value of coefficients used for computing running averages of gradient and its square are 0.9, 0.999 (default values as per Pytorch) respectively. We do not use weight decay during optimization. All the models are trained for 400 epochs. For Max-Ent IRL (Ziebart et al., 2008), we sample a subset of the training set for reward approximation stage (learning reward component weights). When we have a training set size of 70, we sample 5 expert and generated programs for this stage. Similarly we sample 50 for training set size of 221 and 150 for training set of 697 examples.

Software and hardware specifications All the models are coded using Pytorch 1.4.0³ (Paszke et al., 2019) and related libraries like numpy (Oliphant, 2006), scipy (Virtanen et al., 2020) etc. The graphical simulator for VirtualHome (Puig et al., 2018) used in the paper is publicly available⁴. We run all experiments on GeForce RTX 2080 GPU of size 12 GB. The system has 256 GB RAM and 40 CPU cores. The inference process is run in parallel on all the cores. For IRL fine-tuning on training set of size 70 it takes around 80 minutes for 400 epochs, which increases to 10 hours when fine-tuning on the training set of 697 samples. Doing RL fine-tuning also takes similar amount of time since the time required for just reward weight approximation is quite less.

³<https://pytorch.org/>

⁴<https://github.com/xavierpuigf/virtualhome>

A.2 Validation Set Results

We choose the best performing model given a training paradigm and also the set of parameter based on the performance in the validation set. Table 4 shows the results on the validation set. We use the same set of metrics being used for test set. We train and test five runs for each model. We perform significance testing using the Wilcoxon signed rank test. The values in the tables marked with * are statistically significant ($p < 0.05$).

A.3 Example of generated programs by IRL

Table 5 shows two sample generated programs from the IRL model trained with 697 samples using all rewards. The “good” program has higher values of LCS and recall from program values.

MODEL	MEAN LCS	ED	RECALL PROG.	EXEC.
Random	0.069 ±0.001	0.997 ±0.002	0.007 ±0.002	0.004 ±0.005
697 training samples				
MLE	0.412 ±0.009	0.719 ±0.010	0.407 ±0.013	0.373 ±0.039
RL (LCS)	0.453 ±0.012	0.678 ±0.018	0.472 ±0.010	0.419 ±0.060
RL (LCS) [◇]	0.454 ±0.011	0.668 ±0.017	0.467 ±0.016	1.000 ±0.000
RL (LCS + exec)	0.449 ±0.022	0.680 ±0.027	0.456 ±0.029	0.420 ±0.039
RL (LCS + exec) [◇]	0.438 ±0.016	0.679 ±0.023	0.444 ±0.018	1.000 ±0.000
RL (All - Exec.)	0.453 ±0.020	0.681 ±0.020	0.491 ±0.029	0.443 ±0.013
RL (All)	0.453 ±0.017	0.680 ±0.020	0.492 ±0.021	0.404 ±0.046
Ours (IRL: All - Exec.)	0.460 ±0.008	0.670 ±0.009	0.481 ±0.007	0.423 ±0.04
Ours (IRL: All - Exec.) [◇]	0.457 ±0.008	0.662 ±0.009	0.458 ±0.017	1.000 ±0.000
Ours (IRL: All)	0.458 ±0.022	0.672 ±0.027	0.480 ±0.031	0.431 ±0.038
Ours (IRL: All) [◇]	0.459 ±0.018	0.662 ±0.024	0.475 ±0.023	1.000 ±0.000
221 training samples				
MLE	0.312 ±0.010	0.818 ±0.010	0.275 ±0.013	0.252 ±0.029
RL (LCS)	0.362 ±0.006	0.771 ±0.008	0.339 ±0.007	0.308 ±0.043
RL (LCS) [◇]	0.356 ±0.010	0.779 ±0.014	0.337 ±0.014	1.000 ±0.000
RL (LCS + exec)	0.358 ±0.008	0.772 ±0.012	0.329 ±0.005	0.304 ±0.042
RL (LCS + exec) [◇]	0.359 ±0.008	0.762 ±0.011	0.331 ±0.007	1.000 ±0.000
RL (All - Exec.)	0.353 ±0.013	0.779 ±0.010	0.340 ±0.021	0.312 ±0.035
RL (All)	0.359 ±0.007	0.776 ±0.008	0.345 ±0.013	0.271 ±0.039
Ours (IRL: All - Exec.)	0.376* ±0.006	0.757 ±0.007	0.362* ±0.014	0.362 ±0.034
Ours (IRL: All - Exec.) [◇]	0.368 ±0.007	0.754 ±0.006	0.354 ±0.015	1.000 ±0.000
Ours (IRL: All)	0.370 ±0.006	0.767 ±0.010	0.354 ±0.011	0.279 ±0.039
Ours (IRL: All) [◇]	0.370 ±0.010	0.756 ±0.010	0.357 ±0.016	1.000 ±0.000
70 training samples				
MLE	0.233 ±0.015	0.866 ±0.009	0.160 ±0.011	0.232 ±0.038
RL (LCS)	0.265 ±0.006	0.869 ±0.005	0.198 ±0.008	0.112 ±0.024
RL (LCS) [◇]	0.260 ±0.010	0.865 ±0.014	0.181 ±0.017	1.000 ±0.000
RL (LCS + exec)	0.261 ±0.006	0.875 ±0.009	0.185 ±0.007	0.146 ±0.023
RL (LCS + exec) [◇]	0.258 ±0.006	0.868 ±0.007	0.173 ±0.005	1.000 ±0.000
RL (All - Exec.)	0.267 ±0.010	0.870 ±0.011	0.198 ±0.019	0.100 ±0.050
RL (All)	0.265 ±0.006	0.871 ±0.005	0.201 ±0.008	0.113 ±0.046
Ours (IRL: All - Exec.)	0.273 ±0.007	0.865 ±0.003	0.213 ±0.010	0.105 ±0.033
Ours (IRL: All - Exec.) [◇]	0.276 ±0.011	0.854 ±0.017	0.213 ±0.015	1.000 ±0.000
Ours (IRL: All)	0.282 ±0.003	0.857 ±0.006	0.217* ±0.005	0.114 ±0.018
Ours (IRL: All) [◇]	0.283* ±0.008	0.848* ±0.013	0.207 ±0.014	1.000 ±0.000
Semi-supervised learning				
Sup. #70 + Unsup. #697	0.145 ±0.025	0.963 ±0.015	0.075 ±0.027	0.130 ±0.035

Table 4: Validation set performance for different training data sizes. MLE, RL(LCS), RL(LCS + exec) models are adapted from Puig et al. (2018). [◇] denotes executability as hard constraint during inference. * denotes statistical significance ($p < 0.05$) of IRL vs RL baselines from Puig et al. (2018) using the Wilcoxon signed-rank test.

Good Example**Task description:**

I pick up groceries. I open fridge. I put groceries in the fridge.

Program:

[walk] ⟨kitchen⟩
[walk] ⟨fridge⟩
[find] ⟨fridge⟩
[open] ⟨fridge⟩
[find] ⟨groceries⟩
[grab] ⟨groceries⟩
[putin] ⟨groceries⟩ ⟨fridge⟩
[close] ⟨fridge⟩

Rewards: LCS(0.50), Rec_{desc}(1.00), Rec_{prog}(0.71), Irrel.(−0.12), Exec.(1)

Bad example:**Task description:**

Turn on laptop. Connect to internet and check for new email messages.

Program:

[walk] ⟨home_office⟩
[walk] ⟨computer⟩
[find] ⟨computer⟩
[switchon] ⟨computer⟩

Rewards: LCS(0.15), Rec_{desc}(0.25), Rec_{prog}(0.08), Irrel.(0.0), Exec.(1)

Table 5: Two sample generated programs with high and low values of LCS and Recall from program.