

Inverse Reinforcement Learning with Natural Language Goals

Li Zhou, Kevin Small

Amazon Alexa
{lizhouml, smakevin}@amazon.com

Abstract

Humans generally use natural language (NL) to communicate task requirements to each other. Ideally, NL should also be usable for communicating goals to autonomous machines (e.g., robots) to minimize friction in task specification. However, understanding and mapping NL goals to sequences of states and actions is challenging. Specifically, existing work along these lines has encountered difficulty in generalizing learned policies to new NL goals and environments. In this paper, we propose a novel adversarial inverse reinforcement learning algorithm to learn a language-conditioned policy and reward function. To improve generalization of the learned policy and reward function, we use a variational goal generator to relabel trajectories and sample diverse goals during training. Our algorithm outperforms multiple baselines by a large margin on a vision-based NL instruction following dataset (Room-2-Room), demonstrating a promising advance in enabling the use of NL instructions in specifying agent goals.

1 Introduction

Humans use natural language (NL) to communicate with each other. Accordingly, a desired method for conveying goals or assigning tasks to an autonomous machine (e.g., a robot) is also by NL. For example, in household tasks, when asking an agent to pick up a toolbox from the garage, we do not assign the coordinates of the toolbox to the robot. Instead, we state to the agent *retrieve my toolbox from the garage*. This requires the agent to understand the semantics of the natural language goals, associate states and actions with the natural language goals, and infer whether the natural language goals are achieved or not – all of which are very challenging tasks. Furthermore, we want to be able to learn policies that are robust to lexical variation in goals and generalize well to new goals and environments.

Inverse reinforcement learning (IRL) (Abbeel and Ng 2004; Fu, Luo, and Levine 2018), a form of imitation learning (Osa et al. 2018), is the task of learning a reward function and hence a policy based on expert demonstrations. Imitation learning has been successfully applied to a wide range of tasks including robot manipulation (Finn, Levine, and Abbeel 2016), autonomous driving (Kuefler et al. 2017), human behavior forecasting (Rhinehart and Kitani 2017),

video game AI (Tucker, Gleave, and Russell 2018). Generally, task goals are specified intrinsically by the environment and an agent is trained for each task. To generalize the learned policy to new goals, goal-conditioned imitation learning (Ding et al. 2019) and reinforcement learning (RL) algorithms (Kaelbling 1993; Schaul et al. 2015; Nair et al. 2018) have been proposed where the policy is explicitly conditioned on a goal. Normally, the goals either share the same space with the states or can be easily mapped to the state space. For example, in Ding et al. (2019), the goals and states are both coordinates in the environment and the goals are provided to the agent by specifying goal positions.

Fu et al. (2019) and Bahdanau et al. (2019) represent two recent related works from the perspective of understanding NL goal specifications, where they propose learning a language-conditioned reward function under the maximum entropy inverse reinforcement learning (Ziebart et al. 2008) and generative adversarial imitation learning frameworks (Ho and Ermon 2016). However, the NL goals are produced by a preset template-based grammar. Fu et al. (2019)’s policy is optimized exactly in a grid environment with known dynamics and when the policy is optimized by sample-based reinforcement learning algorithms such as deep Q-learning (Mnih et al. 2015), the model performance drops significantly. This reveals a sample efficiency challenge in this setting and implies that when the NL goals and the environments are complicated, generalization to new goals and environments becomes a very difficult problem.

One potential method for addressing these shortcomings are goal relabeling techniques such as hindsight experience replay (HER) (Andrychowicz et al. 2017) and latent goal relabeling (Nair et al. 2018), which have been shown to improve sample efficiency in RL settings. However, when the goals are NL and are in a different space than the state space, goal relabeling cannot be applied directly as we cannot easily relabel a state to a NL goal. Cideron et al. (2019) proposes to build such a relabeling function with a SEQ2SEQ model that takes a trajectory as input and outputs a relabeled NL goal. However, this uses the ground-truth reward function and their experiments are again based on simple template-based NL goals. Moreover, as we will show in this paper, applying HER to IRL with NL goals doesn’t significantly improve performance, as the reward function does not generalize well to relabeled goals.

In this work, we propose a sample efficient IRL algorithm with natural language (NL) goals. To the best of our knowledge, our work is the first IRL algorithm that works with real human generated NL goals (as opposed to template-based language goals (Fu, Luo, and Levine 2018; Bahdanau et al. 2019)) in a real world vision-based environment. Our contributions include: (1) proposing a NL goal-conditioned adversarial inverse reinforcement learning algorithm, (2) specifying a variational goal generator that efficiently sample diverse NL goals given a trajectory, (3) utilizing goal relabeling and sampling strategies to improve the generalization of both the policy and the reward function to new NL goals and new environments, (4) proposing a self-supervised learning mechanism to further improve the generalization in new environments. Through these innovations, we show that our algorithm significantly outperform strong baselines on the Room-2-Room dataset (Anderson et al. 2018).

2 Problem Formulation

A task is defined as a pair (E, G) , where E is an environment that the agent can interact with and G is a NL goal that the agent has to fulfill. $G = \{w_1, w_2, \dots, w_N\}$ consists of N words. For example, in our experiments, E is a realistic 3D indoor environment and G is a human-generated navigation instruction. The true reward function of each task is unknown, but there exists a set of expert demonstrations consisting of a task (E, G) and a trajectory τ . The trajectory $\tau = \{s_1, a_1, s_2, a_2, \dots, s_T, a_T\}$ consists of a sequence of T states perceived by the experts and the actions taken by the experts given the states. For example, in our experiments the states are first-person views in a 3D indoor environment and the actions are movements towards a direction in the environment. While our proposed methods can be applied to both continuous and discrete action spaces, we focus on discrete action spaces in this paper. The objective of the algorithm is to learn a policy that imitates expert demonstrations such that given a new NL goal in an existing or new environment, the agent can perform a sequence of actions to achieve that goal. In this paper, we use G to represent an arbitrary goal, \mathcal{G} to represent the set of goals in expert demonstrations, and G_i to represent the i -th goal in \mathcal{G} . The same notation style applies to E and τ too.

3 Method

3.1 Preliminary: Adversarial Inverse Reinforcement Learning (AIRL)

AIRL (Fu, Luo, and Levine 2018) is based on maximum entropy inverse reinforcement learning (MaxEntIRL) (Ziebart et al. 2008). In MaxEntIRL, the probability of a trajectory is defined as $p_\theta(\tau) = \frac{1}{Z} \exp(f_\theta(\tau))$, where f_θ is the reward function to learn and $Z = \sum_\tau \exp(f_\theta(\tau))$ is the partition function. MaxEntIRL learns the reward function from expert demonstrations by maximizing the likelihood of trajectories in the expert demonstrations: $\max_\theta \mathbb{E}_{\tau \sim \mathcal{D}}(\log p_\theta(\tau))$. Maximizing this likelihood is challenging because the partition function Z is hard to estimate. AIRL maximizes this likelihood by using a generative adversarial network

(GAN) (Goodfellow et al. 2014). The GAN generator is the policy π to learn and the GAN discriminator is defined as

$$D_{\theta, \phi}(s, a, s') = \frac{\exp\{f_{\theta, \phi}(s, a, s')\}}{\exp\{f_{\theta, \phi}(s, a, s')\} + \pi(a|s)} \quad (1)$$

where $f_{\theta, \phi}(s, a, s') = g_\theta(s, a) + \gamma h_\phi(s') - h_\phi(s)$, $g_\theta(s, a)$ is the reward function to learn, and $h_\phi(s') - h_\phi(s)$ is the reward shaping term. The policy π and the discriminator are updated alternately.

3.2 Inverse Reinforcement Learning with Natural Language Goals

Our AIRL-based solution builds on language-conditioned reward learning (LC-RL) (Fu et al. 2019) by adding template-free NL capabilities and continuous state space support. In this problem setting, the policy and the reward function are conditioned on a NL goal G , so we first extend the discriminator in Equation (1) to have $f_{\theta, \phi}(s, a, s', G) = g_\theta(s, a, G) + \gamma h_\phi(s', G) - h_\phi(s, G)$ such that

$$\begin{aligned} g_\theta(s, a, G) &= \text{MLP}([e^s; \text{Att}(s, G); e^a]) \\ h_\phi(s, G) &= \text{MLP}([e^s; \text{Att}(s, G)]) \end{aligned}$$

where $\text{MLP}(\cdot)$ is a multilayer Perceptron, e^s and e^a are the embeddings of state s and action a , respectively, and $\text{Att}(s, G)$ is an attention function. $\text{Att}(s, G) = \sum \alpha_i e_i^w$ where e_i^w is the word embedding of w_i in G , $\alpha_i = (\text{Linear}(e^s) \cdot e_i^w) / (\sum_i \text{Linear}(e^s) \cdot e_i^w)$ and $\text{Linear}(\cdot)$ is a single-layer Perceptron. We use soft actor-critic (SAC) (Haarnoja et al. 2018), one of the state-of-the-art off-policy RL algorithms, to optimize policy π given the reward function $g_\theta(s, a, G)$. SAC includes a policy network to predict an action given a state and a goal, and a Q-network that estimate the Q-value of an action given a state and a goal. We define the policy network as

$$\pi_w(s, a, G) = \text{Softmax}(e^a \cdot \text{MLP}([e^s; \text{Att}(s, G)]))$$

and we define the Q-network $q_\psi(s, a, G)$ as the same network architecture as g_θ . Compared with on-policy algorithms such as TRPO (Schulman et al. 2015), SAC utilizes a replay buffer to re-use sampled trajectories. The replay buffer is beneficial to training both the discriminator and the policy. *To update the discriminator*, we sample negative (s, a, s', G) examples from the replay buffer and sample positive (s, a, s', G) examples from the expert demonstrations. *To update the policy*, we sample a batch of (s, a, s', G) from the replay buffer and use g_θ to estimate their rewards; then we update the Q- and policy network using these reward-augmented samples. We modify SAC slightly to support discrete action spaces. Appendix B contains details about model architecture and optimization.

3.3 A Variational Goal Generator for Improved Generalization and Sample Efficiency

While SAC has better sample efficiency than many on-policy RL algorithms, as we will show in our experiments, in environments with complicated NL goals and high dimensional state spaces (e.g., vision-based instruction following),

AIRL with SAC performs only slightly better than supervised learning based behavior cloning and the learned policy and discriminator doesn't generalize well to new goals or new environments. In this section, we show that by learning a variational goal generator, we can enrich the training data for both the discriminator and the policy, which leads to large improvements in sample efficiency and generalization for both the discriminator and policy in NL settings.

A goal generator takes a trajectory τ (sequence of states and actions) as input and outputs a NL goal G . Given expert demonstrations, a straightforward way of learning a goal generator is to train an encoder-decoder model that encodes a trajectory and decodes a NL goal. However, in reality, there are many possible ways to describe a specific trajectory using NL, and the description will likely be biased due to variance in people's expression preferences. For example, consider a trajectory in which the agent goes to the kitchen and washes the dishes. Two possible goal descriptions for this trajectory can be *go to the kitchen and wash the dishes* or *clean the dishes on the dining table*. To better model variations and generate more diverse NL goals, we learn a variational encoder-decoder model as goal generator. The generative process of the variational goal generator is

$$\begin{aligned} \mu_{prior}, \sigma_{prior}^2 &= f_{prior}(\tau) \\ z &\sim \mathcal{N}(\mu_{prior}, \sigma_{prior}^2 \mathbf{I}) \\ G &= f_{dec}(z, \tau) \end{aligned}$$

where f_{prior} is a LSTM-based trajectory encoder and f_{dec} is an attention-based goal decoder. The posterior distribution of latent variable z is approximated by

$$\begin{aligned} \mu_{posterior}, \sigma_{posterior}^2 &= f_{posterior}(\tau, G) \\ q(z|\tau, G) &= \mathcal{N}(z|\mu_{posterior}, \sigma_{posterior}^2 \mathbf{I}) \end{aligned}$$

where $f_{posterior}$ is a LSTM-based trajectory and goal encoder. We then maximize the variational lower bound $-D_{KL}(q(z|\tau, G)||p(z)) + \mathbb{E}_{q(z|\tau, G)}[\log p(G|z, \tau)]$. Network architecture and optimization details are in Appendix B. Given the learned variational goal generator, we propose using three goal relabeling and sampling strategies to improve generalization as itemized below.

Expert Goal Relabeling (EGR). As discussed previously, multiple goals can be mapped to the same trajectory. We propose to augment expert demonstrations by generating N other goals for each expert trajectory τ_i : $G_{i,n} \sim \text{GoalGenerator}(\tau_i)$, $n = \{1, 2, \dots, N\}$, where $G_{i,n}$ is the n -th generated goal for τ_i , and $\text{GoalGenerator}(\cdot)$ is the learned variational goal generator. For expository and experimental simplicity, we set $N = 2$, which works well for this dataset. Model performance under different values of N are shown in Appendix D. These newly generated tuples $\{(E_i, G_{i,n}, \tau_i)\}_{n=1}^N$ are also treated as expert demonstrations for training. We represent the set of goals in the augmented expert demonstrations by \mathbf{G}^1 where the superscript of \mathbf{G} represents the round of generation as described shortly.

Hindsight Goal Relabeling (HGR) for the Discriminator. In the AIRL framework, the quality of the discriminator is crucial to the generalization of the learned policy. A

good discriminator learns a good reward function that generalizes well to new goals and new environments such that the learned policy can also well generalize to new goals and new environments (Fu et al. 2019). To improve the discriminator, we propose to augment the positive examples of the discriminator by relabeling the goals of the sampled trajectories. More specifically, during the training, given a goal (E_j, G_j^1) , the agent interacts with the environment and samples a trajectory τ_j^1 . We then use the variational goal generator to sample a goal for τ_j^1 :

$$\tau_j^1 \sim \pi(E_j, G_j^1) \quad (2)$$

$$G_j^2 \sim \text{GoalGenerator}(\tau_j^1) \quad (3)$$

The tuples $(E, \mathbf{G}^2, \tau^1)$ are treated as positive examples for the discriminator, as a supplement to the positive examples from expert demonstrations (E, \mathbf{G}^1, τ) .

Hindsight Goal Sampling (HGS) for Policy Optimization. When optimizing the policy with SAC, we have to sample NL goals to train the policy. One natural way is to sample goals from expert demonstrations. However, expert demonstrations are relatively scarce, expensive to acquire, and it is difficult to cover goal variance encountered in the testing. Meanwhile, as we will show in our experiments, training with a diverse set of goals is beneficial to the generalization of the policy. Therefore, we propose to also sample goals from \mathbf{G}^2 so that the policy can train with goals beyond these in the expert demonstrations.

$$\tau_j^2 \sim \pi(E_j, G_j^2) \quad (4)$$

$$G_j^3 \sim \text{GoalGenerator}(\tau_j^2) \quad (5)$$

Of course, training the policy with \mathbf{G}^2 relies on the discriminator's generalization ability to provide reasonable reward estimates for states and actions sampled under \mathbf{G}^2 . This is ensured by HGR in Section 3.3, as $(E, \mathbf{G}^2, \tau^1)$ are provided as positive examples for the discriminator.

The process from Equation (2) to Equation (5) can be seen as using \mathbf{G}^1 and τ as seed, and iteratively sample τ^v from \mathbf{G}^v using the policy, and then sample \mathbf{G}^{v+1} from τ^v using the goal generator. We can go deeper in this loop to sample more diverse NL goals and trajectories for policy and discriminator training. More specifically, *to train the discriminator*, positive examples are sampled from (E, \mathbf{G}^1, τ) with probability 0.5, and are sampled from $\{(E, \mathbf{G}^{v+1}, \tau^v) | v \geq 1\}$ with probability 0.5; negative examples are sampled from $\{(E, \mathbf{G}^v, \tau^v) | v \geq 1\}$. *To train the policy*, goals are sampled from (E, \mathbf{G}^1) with probability 0.5 and are sampled from $\{(E, \mathbf{G}^{v+1}) | v \geq 1\}$ with probability 0.5. Algorithm 1 shows the overall description for LangGoalIRL.

Hindsight Experience Replay (HER) for Policy Optimization. A closely related approach, hindsight experience replay (Andrychowicz et al. 2017), has been shown to work very well in RL settings with sparse rewards. In principle, we could easily incorporate HER into our training procedure. That is, when sampling from replay buffer to optimize policy π , we sample batches not only from $\{(E, \mathbf{G}^v, \tau^v) | v \geq 1\}$, but also from $\{(E, \mathbf{G}^{v+1}, \tau^v) | v \geq 1\}$.

Algorithm 1 Inverse Reinforcement Learning with Natural Language Goals (LangGoalIRL)

Input: GoalGenerator: the variational goal generator from section 3.3; \mathcal{D} : expert demonstrations; $\mathcal{R} = \emptyset$: replay buffer; b : batch size; N : number of expert relabeling goals.

- 1: $\tilde{\mathcal{D}} = \emptyset, \mathcal{G} = \emptyset$
- 2: **for** each $(E_i, G_i, \tau_i) \in \mathcal{D}$ **do**
- 3: **for** $n = 1, 2, \dots, N$ **do**
- 4: \triangleright **Expert Goal Relabeling (EGR)**
- 5: $G_{i,n} \sim \text{GoalGenerator}(\tau_i)$
- 6: **end for**
- 7: add (E_i, G_i, τ_i) and $\{(E_i, G_{i,n}, \tau_i)\}_{n=1}^N$ to $\tilde{\mathcal{D}}$
- 8: **end for**
- 9: **while** not converged **do**
- 10: $r_1, r_2 \sim \text{Uniform}(0, 1)$
- 11: **if** $r_1 < 0.5$ **then**
- 12: Sample a goal $(E_j, G_j) \sim \tilde{\mathcal{D}}$
- 13: **else**
- 14: \triangleright **Hindsight Goal Sampling (HGS)**
- 15: Sample a goal $(E_j, G_j) \sim \mathcal{G}$
- 16: **end if**
- 17: Sample a trajectory $\tau'_j \sim \pi(E_j, G_j)$
- 18: Sample a relabeled goal $G'_j \sim \text{GoalGenerator}(\tau'_j)$
- 19: Add $(E_j, G_j, G'_j, \tau'_j)$ to replay buffer \mathcal{R}
- 20: Add (E_j, G'_j) to \mathcal{G}
- 21: **if** $r_2 < 0.5$ **then**
- 22: Sample a batch $\mathcal{P}_+ = \{(s_k^t, a_k^t, s_k^{t+1}, G_k)\}_{k=1}^b \sim \tilde{\mathcal{D}}$
- 23: **else**
- 24: \triangleright **Hindsight Goal Relabeling (HGR)**
- 25: Sample a batch $\mathcal{P}_+ = \{(s_k^t, a_k^t, s_k^{t+1}, G'_k)\}_{k=1}^b \sim \mathcal{R}$
- 26: **end if**
- 27: Sample a batch $\mathcal{P}_- = \{(s_k^t, a_k^t, s_k^{t+1}, G_k)\}_{k=1}^b \sim \mathcal{R}$
- 28: Update discriminator parameters with \mathcal{P}_+ and \mathcal{P}_- as positive and negative examples, respectively.
- 29: Sample a batch $\mathcal{Q} = \{(s_k^t, a_k^t, s_k^{t+1}, G_k)\}_{k=1}^b \sim \mathcal{R}$
- 30: Expand each entry of \mathcal{Q} with reward $r_k^t = g_\theta(s_k^t, a_k^t, G_k)$
- 31: Optimize π using Soft Actor-Critic with \mathcal{Q}
- 32: **end while**

1}. The difference between HER and HGR is that the former is goal relabeling for the policy while the latter is goal relabeling for the discriminator. HER is most efficient when rewards are sparse; however, in our setting, the rewards provided by the discriminator are not sparse and we do not observe a boost of performance after applying HER. This is discussed in greater detail in Section 4.

3.4 Self-Supervised Learning in New Environments

In this section, we specifically examine the scenario where the learned policy is deployed to a new environment. For example, after training an embodied agent to perform tasks in a set of buildings, we may deploy this agent to a new building with different floor plans. We assume that we have access to these new environments, but we do not have any expert demonstrations nor any NL goals in new environments. Note that we can not directly apply NL goals from existing environments to new environments, because goals are tied to the environments. For example, in instruction-following tasks, an example goal is *go downstairs and walk pass the living*

room. However, there may be no stairs in a new environment. As a result, we can not just sample goals from expert demonstrations to train in the new environments.

The high-level idea of our method is to first learn a policy π_p that is goal-agnostic. That is, the policy selects an action purely based on the state without conditioning on a goal. This model gives us a prior of how the agent will act given a state. We use behavior cloning to train this policy on expert demonstrations. More details are available in Appendix B. We then sample trajectories in the new environments using this policy and sample goals of these trajectories using the goal generator:

$$\begin{aligned}\tau^{new} &\sim \pi_p(E^{new}) \\ G^{new} &\sim \text{GoalGenerator}(\tau^{new})\end{aligned}$$

$(E^{new}, G^{new}, \tau^{new})$ are treated as expert demonstrations in the new environments. We then fine-tune the discriminator and the policy in new environments using Algorithm 1 (w/o expert goal relabeling). The self-supervised learning signals come from both the generated expert demonstrations for the new environments, and the hindsight goal relabeling and sampling proposed in Section 3.3. The generated expert demonstrations can be seen as seeds to bootstrap hindsight goal relabeling and sampling.

4 Experiments

While LangGoalIRL works on any natural language (NL) goal based IRL problems, our experiments focus on a challenging vision-based instruction following problem. We evaluate our model on the Room-2-Room (R2R) dataset (Anderson et al. 2018), a visually-grounded NL navigation task in realistic 3D indoor environments. The dataset contains 7,189 routes sampled from 90 real world indoor environments. A route is a sequence of viewpoints in the indoor environments with the agent’s first-person camera views. Each route is annotated by humans with 3 navigation instructions. The average length of navigation instructions is 29 words. The dataset is split into train (61 environments and 14,025 instructions), seen validation (61 environments same as train set, and 1,020 instructions), unseen validation (11 new environments and 2,349 instructions), and test (18 new environments and 4,173 instructions). We don’t use the test set for evaluation because the ground-truth routes of the test set are not released. Along with the dataset, a simulator is provided to allow the embodied agent to interact with the environments. The observation of the agent is the first-person camera view, which is a panoramic image. The action of the agent is the nearby viewpoints that the agent can move to. For details about the dataset, please see Appendix A. The R2R dataset has become a very popular testbed for language grounding in visual context (Tan, Yu, and Bansal 2019; Wang et al. 2019; Fried et al. 2018; Zhu et al. 2020).

We evaluate the model performance based on the trajectory success rate. Each navigation instruction in the R2R dataset is labeled with a goal position. Following Anderson et al. (2018), the agent is considered successfully reaching the goal if the navigation error (the distance between the stop position and the goal position) is less than 3 meters.

Note that the goal position is not available to our model, as we use navigation instructions as goals.

We compare our model with the following baselines: (1) **LangGoalIRL_BASE**, which corresponds to our proposed model in Section 3.2 *without* the goal relabeling/sampling strategies proposed in Section 3.3. (2) **Behavior Cloning**, which is imitation learning as supervised learning. The model shares the same architecture as the policy network of LangGoalIRL and is trained to minimize cross entropy loss with actions in the expert demonstrations as labels. (3) **LC-RL (Sampling)** (Fu et al. 2019), which also uses MaxEntIRL (Ziebart et al. 2008) to learn a reward function. It optimizes the policy exactly in a grid environment, which is not scalable to our experimental setting, so we use AIRL with SAC for LC-RL to optimize the policy. In this case, LC-RL is very similar to LangGoalIRL_BASE, except that LC-RL simply concatenates state and goal embeddings as input to the policy and reward function, while LangGoalIRL_BASE uses the attention mechanism $\text{Att}(s, G)$ in Section 3.2. (4) **AGILE** (Bahdanau et al. 2019), which proposes to learn a discriminator (reward function) that predicts whether a state is the goal state for a NL goal or not. The positive examples for the discriminator are the (NL goal, goal state) pairs in expert demonstrations, and the negative examples are sampled from the replay buffer of SAC. (5) **HER** (Andrychowicz et al. 2017; Cideron et al. 2019), which uses our learned variational goal generator to relabel trajectories. This corresponds to the `final` strategy in Andrychowicz et al. (2017) applied to LangGoalIRL_BASE. (6) **Upper Bound** (Tan, Yu, and Bansal 2019). Recently, there are a few vision-language navigation models (Tan, Yu, and Bansal 2019; Wang et al. 2019; Fried et al. 2018; Zhu et al. 2020) developed and evaluated on R2R dataset. However, they assume the goal positions or the optimal actions at any states are known, and assume the environments can be exhaustively searched without sample efficiency considerations in order to do training data augmentation. As a result, their problem settings are not really comparable, but we include Tan, Yu, and Bansal (2019)’s result a potential upper bound. Note that without the additional information described above, their model performance reverts to our **Behavior Cloning** baseline. We will discuss this more in Section 5. For implementation details of our algorithms and the baselines, please refer to Appendix B.

The performance of our algorithm and baselines are shown in Figure 1, Table 1, and Table 2. From the results, we would like to highlight the following observations.

1. **LangGoalIRL outperforms baselines by a large margin.** From Table 1 and Figure 1(a) we can see that LangGoalIRL achieves a 0.530 success rate on seen validation set, a 47.63% improvement over LC-RL (sampling) and a 129% improvement over AGILE. Rewards learned by AGILE are binary and sparse, which may be one of the main reasons why AGILE performs even worse than Behavior Cloning. The difference between LC-RL (sampling) and LangGoalIRL_BASE is just the lack of attention mechanism over NL goals, which simply indicates that attention mechanism is important when NL goals are complicated. Other powerful attention mechanisms such as the Transformer architecture (Vaswani et al. 2017) could also be easily con-

	seen validation	unseen validation
Behavior Cloning	0.445 \pm 0.0122	0.300 \pm 0.0114
LC-RL (sampling)	0.359 \pm 0.0117	0.216 \pm 0.0081
AGILE	0.231 \pm 0.0173	0.253 \pm 0.0269
HER	0.464 \pm 0.0126	0.306 \pm 0.0173
LangGoalIRL_BASE	0.449 \pm 0.0130	0.308 \pm 0.0087
LangGoalIRL	0.530 \pm 0.0138	0.357 \pm 0.0089
+self-supervised	–	0.491 \pm 0.0065
Upper Bound	0.621	0.645

Table 1: Success rate after 1 million environment interactions. LangGoalIRL+self-supervised is further fine-tuned in unseen environment for 1 million environment interactions (explained in Figure 1(c)).

Expert Goal Relabeling	Hindsight Goal Relabeling	Hindsight Goal Sampling	seen validation success rate
–	–	–	0.449 \pm 0.0130
✓	–	–	0.450 \pm 0.0122
–	✓	–	0.427 \pm 0.0048
–	–	✓	0.399 \pm 0.0124
✓	✓	–	0.504 \pm 0.0185
✓	–	✓	0.444 \pm 0.0134
–	✓	✓	0.503 \pm 0.0117
✓	✓	✓	0.530 \pm 0.0138

Table 2: Ablation study of the three proposed strategies on seen validation set. A complete ablation study which includes HER is in Table 1 of Appendix D.

sidered. The success rate of LangGoalIRL is 18.04% and 15.91% higher than LangGoalIRL_BASE on seen validation and unseen validation, respectively, which shows that the three proposed strategies efficiently improve the generalization of the policy to new NL goals and new environments. From Figure 1(b) we see that, when combined with HER, LangGoalIRL converges slower and performs worse than LangGoalIRL alone. HER has been shown to efficiently deal with reward sparsity; however, the rewards learned by AIRL are not sparse. From Table 1 we see that HER alone barely outperforms LangGoalIRL_BASE. As we will discuss shortly, HER only improves the sample efficiency of the *generator (policy)*, however, the generalization of the *discriminator (reward function)* is what is crucial to the overall performance of the policy. Finally, from Figure 1(b) we can also see that LangGoalIRL_BASE consistently outperforms Behavior Cloning after about 200k interactions, which effectively indicates the benefit of using inverse reinforcement learning over behavior cloning.

2. **Hindsight goal relabeling and expert goal relabeling improve the generalization of the discriminator (reward function); such generalization is crucial to the performance of the policy.** From Table 2 and Figure 1(f), we see that if we take out HGR, the success rate drops significantly from 0.530 to 0.444. This shows that HGR plays a key role in learning a better reward function by enriching positive examples of the discriminator. Meanwhile, when applying HGS without HGR, the success rate drops from 0.503 to 0.399. This is because the discriminator cannot generalize well to sampled goals when HGR is not applied, which shows the importance of discriminator generalization. How-

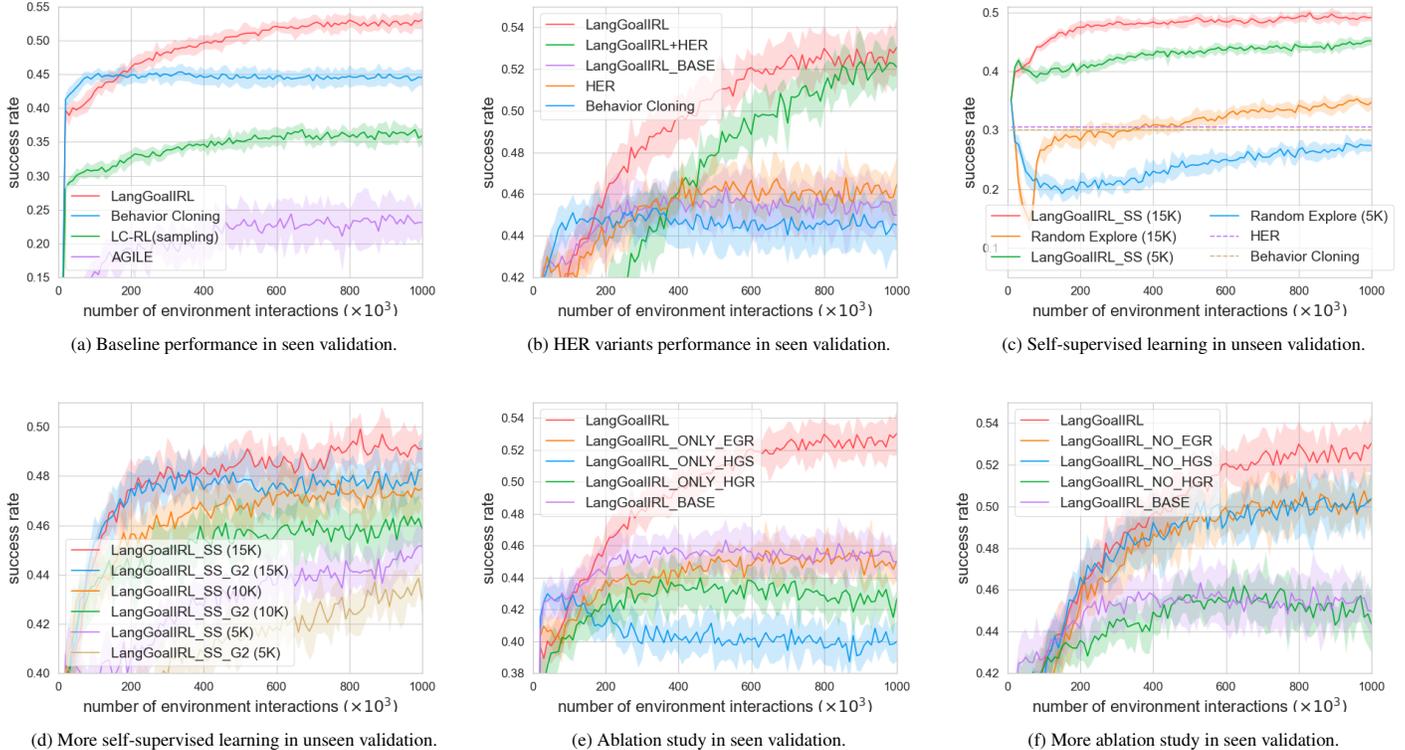


Figure 1: Model performance on both seen and unseen validation. LangGoalIRL_SS in Figure (c) and (d) represents self-supervised learning in unseen environments. LangGoalIRL_SS (15K) means that 15,000 demonstrations are sampled by the goal-agnostic policy π_p . Random Explore means that a random policy instead of π_p is used to sample demonstrations. LangGoalIRL_SS_G2 means that the policy does not iteratively sample goals from $\{(\mathbf{E}, \mathbf{G}^v) | v > 2\}$ as described in Section 3.3.3.

ever, from Table 2 and Figure 1(e) we can also see that if we only keep HGR, the success rate drops to 0.427 which is even lower than LangGoalIRL_BASE. We observe that in this case the goals generated by the variational goal generator only appear in the positive examples of the discriminator, so the discriminator easily identify these goals, and simply assigns high rewards to these goals regardless of what action is taken. When EGR and HGR are applied together, re-labeled goals appear in both positive and negative examples, and the success rate improves from 0.427 to 0.504.

3. Self-supervised learning largely improves performance in unseen environments. Table 1 shows that with self-supervised learning, LangGoalIRL achieves a success rate of 0.491, a 59.42% improvement over LangGoalIRL_BASE and a 37.54% improvement over LangGoalIRL w/o self-supervised learning. This shows that the proposed self-supervised learning algorithm can significantly improve policy performance even though neither expert demonstrations nor NL goals in new environments are given. The number of trajectories sampled by the goal-agnostic policy π_p in new environments by default is 15,000. In Figure 1(c), we show the performance of a baseline model where the embodied agent randomly explores the new environments (randomly selects an action at any states, but does not visit any visited viewpoints) rather than using π_p to sam-

ple trajectories. We set the number of sampled trajectory to 15,000 and 5,000. We see that the agent performance drops during early stages, as randomly explored trajectories are not good seeds to bootstrap Algorithm 1 in new environments. After early stages, the model performance gradually increases as HGR and HGS provide some self-supervised signals but the overall the performance remains lower than LangGoalIRL. More results are in Table 2 of Appendix D. Finally, self-supervised learning can also be applied to existing environments by augmenting the expert demonstrations in seen validation set. While this is not our target setting, we include these results in Table 3 of Appendix D.

4. Hindsight goal sampling improves the generalization of the policy by enabling the policy to explore a more diverse set of goals. From Table 2 and Figure 1(f) we can see that HGS further improves the policy performance from 0.504 to 0.530 on seen validation set. Figure 1(d) shows the impact of HGS on self-supervised learning in new environments. The baseline, LangGoalIRL_SS_G2 samples goals only from \mathbf{G}^1 and \mathbf{G}^2 defined in Section 3.3 and does not iteratively sample goals from $\{(\mathbf{E}, \mathbf{G}^v) | v > 2\}$. As goals sampled are less diverse, we see that LangGoalIRL_SS_G2 performs worse than LangGoalIRL_SS given 5,000, 10,000 and 15,000 generated expert demonstrations in new environments. More experiments are in Table 2 of Appendix D.

	LangGoalIRL	Language Goal-Conditioned IRL (Fu et al. 2019) (Bahdanau et al. 2019)	Language Goal-Conditioned RL (Cideron et al. 2019) (Jiang et al. 2019)	State Goal-Conditioned IRL (Ding et al. 2019)	State Goal-Conditioned RL (Andrychowicz et al. 2017) (Plappert et al. 2018)	SoTA models on R2R dataset (Tan, Yu, and Bansal 2019) (Fried et al. 2018) (Zhu et al. 2020)
Require Ground-Truth Reward function	No	No	Yes	No	Yes	No (Yes if use RL fine-tuning)
Goal Type	Natural Language	Natural Language	Natural Language	States	States	Natural Language
Use Student-forcing	No	No	No	No	No	Yes
Goals/Paths Augmentation	EGR, HGS, HGR	No	HER	HER	HER	All shortest-paths between any two locations in environments

Table 3: We compare our algorithm with closely related works that fall into 5 different problem settings. In the setting we are interested in, the goals of the tasks are NL goals, and neither the ground-truth reward function (required by RL) nor the optimal action at any state (required by student-forcing) is available.

5 Related Works

Our paper focuses on proposing a sample efficient algorithm for IRL with natural language (NL) goals. Table 3 summarizes the differences between our LangGoalIRL algorithm and recent closely related works on goal-conditioned RL and imitation learning (IL). Goal-conditioned RL and IL have been explored by many prior works (Schaul et al. 2015; Florensa et al. 2017; Nair et al. 2018; Plappert et al. 2018; Ding et al. 2019). Usually, the task is to learn a policy that is parameterized by a goal and can generalize to new goals. In most prior works, the goals and states are in a same space, such as positions in a Cartesian coordinate system (Plappert et al. 2018; Ding et al. 2019). This corresponds to column 5 and 6 in Table 3. However, in this paper we investigate a different setting where the goals are NL goals, which corresponds to column 3 and 4 in Table 3. The closest prior works to our algorithm are language-conditioned IRL and reward learning (Fu et al. 2019; Bahdanau et al. 2019; MacGlashan et al. 2015; Williams et al. 2018; Goyal, Niekum, and Mooney 2019). Goyal, Niekum, and Mooney (2019) propose to train a model to predict whether a language instruction describes an action in a trajectory, and use the predictions to perform reward shaping. Bahdanau et al. (2019) propose to use GAN (Goodfellow et al. 2014; Ho and Ermon 2016) to learn a discriminator that predicts whether a state is the goal state of a NL instruction. Fu et al. (2019) propose to use the MaxEntIRL (Ziebart et al. 2008) framework to learn a language-conditioned reward function for instruction following tasks. The language instructions in Bahdanau et al. (2019) and Fu et al. (2019)’s experiments are generated by templates and much easier to understand compared with our dataset. Meanwhile, LangGoalIRL demonstrates significantly better generalization than these two prior works. There are also many prior works on goal relabeling to improve sample efficiency, but none of them can be directly applied to the IRL with NL goals setting. Andrychowicz et al. (2017) propose hindsight experience replay that samples additional goals for each transition in the replay buffer and can efficiently deal with the reward sparsity problem. Nair et al. (2018) propose to sample additional diverse goals from a learned latent space. Ding et al. (2019) propose to relabel trajectories using the states within the trajectories. These algorithms do not work with NL goals. Cideron et al. (2019) and Jiang et al. (2019) generalize HER to language setting and relabel trajectories by hindsight language instructions. However, in this paper we show that, when doing IRL, simply applying HER to policy optimization does not improve

policy performance, and it is critical to improve the generalization of the reward function. Finally, there are multiple recent works (Tan, Yu, and Bansal 2019; Wang et al. 2019; Fried et al. 2018; Zhu et al. 2020) on R2R dataset that focus on solving vision-language navigation (Anderson et al. 2018). This corresponds to the last column in Table 3. These works do not focus on sample efficient IRL algorithms and use student-forcing (Anderson et al. 2018) to train their supervised learning models or pre-train their RL models. With student-forcing, the policy is assumed to always have access to the optimal action at any state it encounters during training, which is even a stronger assumption than knowing the ground-truth reward function. Fried et al. (2018) and Tan, Yu, and Bansal (2019) propose to learn a speaker model that back-translate a given trajectory to a NL goal, which is similar to our goal generator without latent variables. However, the purpose of the speaker model is to augment training data before the training, rather than relabeling goals during training. To augment training data, they generate the shortest path between any two viewpoints. This requires knowing the topological graph of viewpoints before the training. Moreover, they augment training set with every shortest path that is not included in the original training set. This assumes the environments can be exhaustively searched without sampling efficiency considerations. Other proposed techniques in these papers such as environment dropout (Tan, Yu, and Bansal 2019) and progress monitoring (Zhu et al. 2020) are orthogonal to our work. For a survey of reinforcement learning with NL, please see Luketina et al. (2019).

6 Conclusion

In this work, we propose a sample efficient algorithm for IRL with NL goals. Observing limitations of existing works regarding language-conditioned IRL, LangGoalIRL emphasizes template-free NL goal specification and sample efficiency when generalizing to new goals and environments. Specifically, we propose learning a variational goal generator that can relabel trajectories and sample diverse goals. Based on this variational goal generator, we describe three strategies to improve the generalization and sample efficiency of the language-conditioned policy and reward function. Empirical results demonstrate that LangGoalIRL outperforms existing baselines by a large margin and generalizes well to new natural language goals and new environments, thus increasing flexibility of expression and domain transfer in providing instructions to autonomous agents.

References

- Abbeel, P.; and Ng, A. Y. 2004. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*, 1.
- Anderson, P.; Wu, Q.; Teney, D.; Bruce, J.; Johnson, M.; Sünderhauf, N.; Reid, I.; Gould, S.; and van den Hengel, A. 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*, 3674–3683.
- Andrychowicz, M.; Wolski, F.; Ray, A.; Schneider, J.; Fong, R.; Welinder, P.; McGrew, B.; Tobin, J.; Abbeel, O. P.; and Zaremba, W. 2017. Hindsight experience replay. In *Advances in neural information processing systems*, 5048–5058.
- Bahdanau, D.; Hill, F.; Leike, J.; Hughes, E.; Hosseini, A.; Kohli, P.; and Grefenstette, E. 2019. Learning to understand goal specifications by modelling reward. *International Conference on Learning Representations*.
- Cideron, G.; Seurin, M.; Strub, F.; and Pietquin, O. 2019. Self-Educated Language Agent With Hindsight Experience Replay For Instruction Following. *arXiv preprint arXiv:1910.09451*.
- Ding, Y.; Florensa, C.; Abbeel, P.; and Phielipp, M. 2019. Goal-conditioned imitation learning. In *Advances in Neural Information Processing Systems*, 15298–15309.
- Finn, C.; Levine, S.; and Abbeel, P. 2016. Guided cost learning: Deep inverse optimal control via policy optimization. In *International conference on machine learning*, 49–58.
- Florensa, C.; Held, D.; Geng, X.; and Abbeel, P. 2017. Automatic goal generation for reinforcement learning agents. *arXiv preprint arXiv:1705.06366*.
- Fried, D.; Hu, R.; Cirik, V.; Rohrbach, A.; Andreas, J.; Morency, L.-P.; Berg-Kirkpatrick, T.; Saenko, K.; Klein, D.; and Darrell, T. 2018. Speaker-follower models for vision-and-language navigation. In *Advances in Neural Information Processing Systems*, 3314–3325.
- Fu, J.; Korattikara, A.; Levine, S.; and Guadarrama, S. 2019. From language to goals: Inverse reinforcement learning for vision-based instruction following. *arXiv preprint arXiv:1902.07742*.
- Fu, J.; Luo, K.; and Levine, S. 2018. Learning Robust Rewards with Adversarial Inverse Reinforcement Learning. In *International Conference on Learning Representations*.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *Advances in neural information processing systems*, 2672–2680.
- Goyal, P.; Niekum, S.; and Mooney, R. J. 2019. Using Natural Language for Reward Shaping in Reinforcement Learning. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*, 2385–2391.
- Haarnoja, T.; Zhou, A.; Hartikainen, K.; Tucker, G.; Ha, S.; Tan, J.; Kumar, V.; Zhu, H.; Gupta, A.; Abbeel, P.; et al. 2018. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*.
- Ho, J.; and Ermon, S. 2016. Generative adversarial imitation learning. In *Advances in neural information processing systems*, 4565–4573.
- Jiang, Y.; Gu, S. S.; Murphy, K. P.; and Finn, C. 2019. Language as an abstraction for hierarchical deep reinforcement learning. In *Advances in Neural Information Processing Systems*, 9414–9426.
- Kaelbling, L. P. 1993. Learning to achieve goals. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, 1094–1098.
- Kuefler, A.; Morton, J.; Wheeler, T.; and Kochenderfer, M. 2017. Imitating driver behavior with generative adversarial networks. In *2017 IEEE Intelligent Vehicles Symposium (IV)*, 204–211. IEEE.
- Luketina, J.; Nardelli, N.; Farquhar, G.; Foerster, J.; Andreas, J.; Grefenstette, E.; Whiteson, S.; and Rocktäschel, T. 2019. A survey of reinforcement learning informed by natural language. *arXiv preprint arXiv:1906.03926*.
- MacGlashan, J.; Babes-Vroman, M.; desJardins, M.; Littman, M.; Muresan, S.; Squire, S.; Tellex, S.; Arumugam, D.; and Yang, L. 2015. Grounding English Commands to Reward Functions. In *Proceedings of Robotics: Science and Systems*. Rome, Italy.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Hiedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; et al. 2015. Human-level control through deep reinforcement learning. *Nature* 518(7540): 529–533.
- Nair, A. V.; Pong, V.; Dalal, M.; Bahl, S.; Lin, S.; and Levine, S. 2018. Visual reinforcement learning with imagined goals. In *Advances in Neural Information Processing Systems*, 9191–9200.
- Osa, T.; Pajarinen, J.; Neumann, G.; Bagnell, J. A.; Abbeel, P.; Peters, J.; et al. 2018. An algorithmic perspective on imitation learning. *Foundations and Trends® in Robotics* 7(1-2): 1–179.
- Plappert, M.; Andrychowicz, M.; Ray, A.; McGrew, B.; Baker, B.; Powell, G.; Schneider, J.; Tobin, J.; Chociej, M.; Welinder, P.; et al. 2018. Multi-goal reinforcement learning: Challenging robotics environments and request for research. *arXiv preprint arXiv:1802.09464*.
- Rhinehart, N.; and Kitani, K. M. 2017. First-person activity forecasting with online inverse reinforcement learning. In *Proceedings of the IEEE International Conference on Computer Vision*, 3696–3705.
- Schaul, T.; Horgan, D.; Gregor, K.; and Silver, D. 2015. Universal value function approximators. In *International conference on machine learning*, 1312–1320.
- Schulman, J.; Levine, S.; Abbeel, P.; Jordan, M.; and Moritz, P. 2015. Trust region policy optimization. In *International conference on machine learning*, 1889–1897.

Tan, H.; Yu, L.; and Bansal, M. 2019. Learning to Navigate Unseen Environments: Back Translation with Environmental Dropout. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics*, 2610–2621.

Tucker, A.; Gleave, A.; and Russell, S. 2018. Inverse reinforcement learning for video games. *arXiv preprint arXiv:1810.10593*.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in neural information processing systems*, 5998–6008.

Wang, X.; Huang, Q.; Celikyilmaz, A.; Gao, J.; Shen, D.; Wang, Y.-F.; Wang, W. Y.; and Zhang, L. 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*, 6629–6638.

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

Zhu, F.; Zhu, Y.; Chang, X.; and Liang, X. 2020. Vision-language navigation with self-supervised auxiliary reasoning tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10012–10022.

Ziebart, B. D.; Maas, A. L.; Bagnell, J. A.; and Dey, A. K. 2008. Maximum entropy inverse reinforcement learning. In *Aaai*, volume 8, 1433–1438. Chicago, IL, USA.