Interactive Learning from Activity Description

Khanh Nguyen [1] **Dipendra Misra** [2] **Robert Schapire** [2] **Miro Dudík** [2] **Patrick Shafto** [3]

Abstract

We present a novel interactive learning protocol that enables training request-fulfilling agents by verbally describing their activities. Unlike imitation learning (IL), our protocol allows the teaching agent to provide feedback in a language that is most appropriate for them. Compared with reward in reinforcement learning (RL), the description feedback is richer and allows for improved sample complexity. We develop a probabilistic framework and an algorithm that practically implements our protocol. Empirical results in two challenging request-fulfilling problems demonstrate the strengths of our approach: compared with RL baselines, it is more sample-efficient; compared with IL baselines, it achieves competitive success rates without requiring the teaching agent to be able to demonstrate the desired behavior using the learning agent's actions. Apart from empirical evaluation, we also provide theoretical guarantees for our algorithm under certain assumptions about the teacher and the environment.

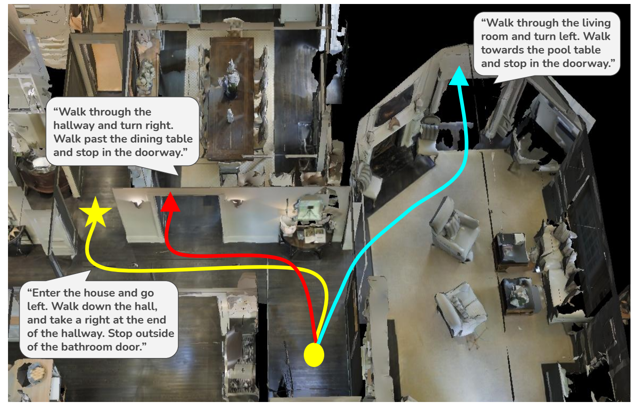

Figure 1. A real example of training an agent to fulfill a navigation request in a 3D environment (Anderson et al., 2018) using ADEL, our implementation of the ILIAD protocol. The agent receives a request *"Enter the house..."* which implies the path. Initially, it wanders far from the goal because it does not understand language. Its execution ( ) is described as *"Walk through the living room..."*. To ground language, the agent learns to generate the path conditioned on the description. After a number of interactions, its execution ( ) is closer to the optimal path. As this process iterates, the agent learns to ground diverse descriptions to executions and can execute requests more precisely.

1. Introduction

The goal of a *request-fulfilling* agent is to map a given request in a situated environment to an execution that accomplishes the intent of the request (Winograd, 1972; Chen & Mooney, 2011; Tellex et al., 2012; Artzi et al., 2013; Misra et al., 2017; Anderson et al., 2018; Chen et al., 2019; Nguyen et al., 2019; Nguyen & Daumé III, 2019; Gaddy & Klein, 2019). Request-fulfilling agents have been typically trained using *non-verbal* interactive learning protocols such as imitation learning (IL) which assumes labeled executions as feedback (Mei et al., 2016; Anderson et al., 2018; Yao et al., 2020), or reinforcement learning (RL) which uses scalar rewards as feedback (Chaplot et al., 2018; Hermann et al., 2017). These protocols are suitable for training agents with pre-collected datasets or in simulators, but they do not lend themselves easily to training by human teachers that only possess domain knowledge, but might not be able to precisely define the reward function, or provide direct demonstrations. To enable training by such teachers, we introduce a verbal interactive learning protocol called ILIAD: **I**nteractive **L**earning from **A**ctivity **D**escription, where feedback is limited to *descriptions of activities*, in a language that is appropriate for a given teacher (e.g., a natural language for humans).

Figure 1 illustrates an example of training an agent using the ILIAD protocol. Learning proceeds in episodes of interaction between a learning agent and a teacher. In each episode, the agent is presented with a request, provided in the teacher's description language, and takes a sequence of actions in the environment to execute it. After an execution is completed, the teacher provides the agent with a description of the execution, in the same description language. The agent then uses this feedback to update its policy.

[1]Department of Computer Science, University of Maryland, Maryland, USA [2]Microsoft Research, New York, USA [3]Rutgers University, New Jersey, USA. Correspondence to: Khanh Nguyen <kxnguyen@umd.edu>.

Proceedings of the 38th International Conference on Machine Learning, PMLR 139, 2021. Copyright 2021 by the author(s).

Table 1. Trade-offs between the *learning* effort of the agent and the teacher in three learning protocols. Each protocol employs a different medium for the teacher to convey feedback. If a medium is not natural to the teacher (e.g., IL-style demonstration), it must learn to express feedback using that medium (*teacher communication-learning effort*). For example, in IL, to provide demonstrations, the teacher must learn to control the agent to accomplish tasks. Similarly, if a medium is not natural to the agent (e.g., human language), it needs to learn to interpret feedback (*agent communication-learning effort*). The agent also learns tasks from information decoded from feedback (*agent task-learning effort*). The qualitative claims about the "agent learning effort" column summarize our empirical findings about the learning efficiency of algorithms that implement these protocols (Table 2).

Protocol	Feedback medium	Learning effort: Teacher (communication learning)	Learning effort: Agent (comm. & task learning)
IL	Demonstration	Highest	Lowest
RL	Scalar reward	None	Highest
ILIAD	Description	None	Medium

The agent receives *no other* feedback such as ground-truth demonstration (Mei et al., 2016), scalar reward (Hermann et al., 2017), or constraint (Miryoosefi et al., 2019). Essentially, ILIAD presents a setting where task learning is enabled by *grounded* language learning: the agent improves its request-fulfilling capability by exploring the description language and learning to ground the language to executions. This aspect distinguish ILIAD from IL or RL, where task learning is made possible by imitating actions or maximizing rewards.

The ILIAD protocol leaves two open problems: (a) *the exploration problem*: how to generate executions that elicit useful descriptions from the teacher and (b) *the grounding problem*: how to effectively ground descriptions to executions. We develop an algorithm named ADEL: **A**ctivity-**D**escription **E**xplorative **L**earner that offers practical solutions to these problems. For (a), we devise a semi-supervised execution sampling scheme that efficiently explores the description language space. For (b), we employ maximum likelihood to learn a mapping from descriptions to executions. We show that our algorithm can be viewed as density estimation, and prove its convergence in the contextual bandit setting (Langford & Zhang, 2008b), i.e., when the task horizon is 1.

Our paper does *not* argue for the primacy of one learning protocol over the others. In fact, an important point we raise is that there are multiple, possibly competing metrics for comparing learning protocols. We focus on highlighting the *complementary* advantages of ILIAD against IL and RL (Table 1). In all of these protocols, the agent and the teacher establish a communication channel that allows the teacher to encode feedback and send it to the agent. At one extreme, IL uses demonstration, an *agent-specific* medium, to encode feedback, thus placing the burden of establishing the communication channel entirely on the teacher. Concretely, in standard interactive IL (e.g., Ross et al., 2011), a demonstration can contain only actions in the agent's action space. Therefore, this protocol implicitly assumes that the teacher must be familiar with the agent's control interface. In practice, non-experts may have to spend substantial effort in order to learn to control an agent.[1] In these settings, the agent usually learns from relatively few demonstrations because it does not have to learn to interpret feedback, and the feedback directly specifies the desired behavior. At another extreme, we have RL and ILIAD, where the teacher provides feedback via *agent-agnostic* media (reward and language, respectively). RL eliminates the agent communication-learning effort by hard-coding the semantics of scalar rewards into the learning algorithm.[2] But the trade-off of using such limited feedback is that the task-learning effort of the agent increases; state-of-the-art RL algorithms are notorious for their high sample complexity (Hermann et al., 2017; Chaplot et al., 2018; Chevalier-Boisvert et al., 2019). By employing a natural and expressive medium like natural language, ILIAD offers a compromise between RL and IL: it can be more sample-efficient than RL while not requiring the teacher to master the agent's control interface as IL does. Overall, no protocol is superior in all metrics and the choice of protocol depends on users' preferences.

We empirically evaluate ADEL against IL and RL baselines on two tasks: vision-language navigation (Anderson et al., 2018), and word-modification via regular expressions (Andreas et al., 2018). Our results show that ADEL significantly outperforms RL baselines in terms of both sample efficiency and quality of the learnt policies. Also, ADEL's success rate is competitive with those of the IL baselines on the navigation task and is lower by 4% on the word modification task. It takes approximately 5-9 times more training episodes than the IL baselines to reach comparable success rates, which is quite respectable considering that the algorithm has to search in an exponentially large space for the ground-truth executions whereas the IL baselines are *given* these executions. Therefore, ADEL can be a preferred algorithm whenever annotating executions with correct (agent) actions is not feasible or is substantially more expensive than describing executions in some description language. For example, in the word-modification task, ADEL teaches the agent without requiring a teacher with

[1] Third-person or observational IL (Stadie et al., 2017; Sun et al., 2019) allows the teacher to demonstrate tasks with their action space. However, this framework is *non-interactive* because the agent imitates pre-collected demonstrations and does not interact with a teacher. We consider interactive IL (Ross et al., 2011), which is shown to be more effective than non-interactive counterparts.

[2] By design, RL algorithms understand that higher reward value implies better performance.

knowledge about regular expressions. We believe the capability of non-experts to provide feedback will make ADEL and more generally the ILIAD protocol a strong contender in many scenarios. The code of our experiments is available at *https://github.com/khanhptnk/iliad*.

2. ILIAD: Interactive Learning from Activity Description

Environment. We borrow our terminology from the reinforcement learning (RL) literature (Sutton & Barto, 2018). We consider an agent acting in an environment with state space $\mathcal{S}$, action space $\mathcal{A}$, and transition function $T : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$, where $\Delta(\mathcal{S})$ denotes the space of all probability distributions over $\mathcal{S}$. Let $\mathcal{R} = \{R : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]\}$ be a set of reward functions. A *task* in the environment is defined by a tuple $(R, s_1, d^\star)$, where $R \in \mathcal{R}$ is the task's reward function, $s_1 \in \mathcal{S}$ is the start state, and $d^\star \in \mathcal{D}$ is the task's (language) request. Here, $\mathcal{D}$ is the set of all nonempty strings generated from a finite vocabulary. The agent only has access to the start state and the task request; the reward function is only used for evaluation. For example, in robot navigation, a task is given by a start location, a task request like "*go to the kitchen*", and a reward function that measures the distance from a current location to the kitchen.

Execution Episode. At the beginning of an episode, a task $q = (R, s_1, d^\star)$ is sampled from a task distribution $\mathbb{P}^\star(q)$. The agent starts in s_1 and is presented with $d^\star$ but *does not* observe R or any rewards generated by it. The agent maintains a *request-conditioned policy* $\pi_\theta : \mathcal{S} \times \mathcal{D} \rightarrow \Delta(\mathcal{A})$ with parameters θ, which takes in a state $s \in \mathcal{S}$ and a request $d \in \mathcal{D}$, and outputs a probability distribution over $\mathcal{A}$. Using this policy, it can generate an *execution* $\hat{e} = (s_1, \hat{a}_1, s_2, \cdots, s_H, \hat{a}_H)$, where H is the task horizon (the time limit), $\hat{a}_i \sim \pi_\theta(\cdot \mid s_i, d^\star)$ and $s_{i+1} \sim T(\cdot \mid s_i, \hat{a}_i)$ for every i. Throughout the paper, we will use the notation $e \sim \mathbb{P}_\pi(\cdot \mid s_1, d)$ to denote sampling an execution e by following policy π given a start state s_1 and a request d. The objective of the agent is to find a policy π with maximum value, where we define the policy value $V(\pi)$ as:

$$V(\pi) = \mathbb{E}_{q \sim \mathbb{P}^\star(\cdot), \hat{e} \sim \mathbb{P}_\pi(\cdot \mid s_1, d^\star)} \left[\sum_{i=1}^{H} R(s_i, \hat{a}_i) \right] \quad (1)$$

ILIAD protocol. Alg 1 describes the ILIAD protocol for training a request-fulfilling agent. It consists of a series of N training episodes. Each episode starts with sampling a task $q = (R, s_1, d^\star)$ from $\mathbb{P}^\star$. The agent then generates an execution $\hat{e}$ given s_1, $d^\star$, and its policy π_θ (line 4). The feedback mechanism in ILIAD is provided by a *teacher* that can describe executions in a description language. The teacher is modeled by a fixed distribution $\mathbb{P}_T : (\mathcal{S} \times \mathcal{A})^H \rightarrow \Delta(\mathcal{D})$, where $(\mathcal{S} \times \mathcal{A})^H$ is the space over H-step executions. After

Algorithm 1 ILIAD protocol. Details of line 4 and line 6 are left to specific implementations.

```
1: Initialize agent policy π_θ : S × D → Δ(A)
2: for n = 1, 2, ··· , N do
3:     World samples a task q = (R, s_1, d*) ~ P*(·)
4:     Agent generates an execution ê given s_1, d*, and π_θ
5:     Teacher generates a description d̂ ~ P_T(· | ê)
6:     Agent uses (d*, ê, d̂) to update π_θ
   return π_θ
```

generating $\hat{e}$, the agent sends it to the teacher and receives a *description* of $\hat{e}$, which is a sample $\hat{d} \sim \mathbb{P}_T(\cdot \mid \hat{e})$ (line 5). Finally, the agent uses the triplet $(d^\star, \hat{e}, \hat{d})$ to update its policy for the next round (line 6). Crucially, the agent *never* receives any other feedback, including rewards, demonstrations, constraints, or direct knowledge of the *latent* reward function. Any algorithm implementing the ILIAD protocol has to decide how to generate executions (the exploration problem, line 4) and how to update the agent policy (the grounding problem, line 6). The protocol does not provide any constraints for these decisions.

Consistency of the teacher. In order for the agent to learn to execute requests by grounding the description language, we require that the description language is similar to the request language. Formally, we define the ground-truth joint distribution over tasks and executions as follows

$$\mathbb{P}^\star(e, R, s_1, d) = \mathbb{P}_{\pi^\star}(e \mid s_1, d)\, \mathbb{P}^\star(R, s_1, d) \quad (2)$$

where $\pi^\star$ is an optimal policy that maximizes Eq 1. From this joint distribution, we derive the ground-truth execution-conditioned distribution over requests $\mathbb{P}^\star(d \mid e)$. This distribution specifies the probability that a request d can serve as a valid description of an execution e.

We expect that if the teacher's distribution $\mathbb{P}_T(d \mid e)$ is close to $\mathbb{P}^\star(d \mid e)$ then grounding the description language to executions will help with request fulfilling. In that case, the agent can treat a description of an execution as a request that is fulfilled by that execution. Therefore, the description-execution pairs $(\hat{d}, \hat{e})$ can be used as supervised-learning examples for the request-fulfilling problem.

The learning process can be sped up if the agent is able to exploit the compositionality of language. For example, if a request is "*turn right, walk to the kitchen*" and the agent's execution is described as "*turn right, walk to the bedroom*", the agent may not have successfully fulfilled the task but it can learn what "*turn right*" and "*walk to*" mean through the description. Later, it may learn to recognize "*kitchen*" through a description like "*go to the kitchen*" and compose that knowledge with its understanding of "*walk to*" to better execute "*walk to the kitchen*".

Algorithm 2 Simple algorithm for learning an agent's policy with access to the true marginal $\mathbb{P}^\star(e \mid s_1)$ and teacher $\mathbb{P}_T(d \mid e)$.

```
B = ∅
for i = 1, 2, ···, N do
    World samples a task q = (R, s_1, d*) ~ P*(·)
    Sample (ê, d̂) as follows: ê ~ P*(· | s_1), d̂ ~ P_T(· | ê)
    B ← B ∪ {(ê, d̂)}
Train a policy π_θ(a | s, d) via maximum log-likelihood:
    max_θ Σ_{(ê,d̂)∈B} Σ_{(s,â_s)∈ê} log π_θ(â_s | s, d̂)
    where â_s is the action taken by the agent in state s
return π_θ
```

Algorithm 3 ADEL: our implementation of the ILIAD protocol.

```
1: Input: teacher P_T(d | e), approximate marginal P_{π_ω}(e | s_1),
   mixing weight λ ∈ [0, 1], annealing rate β ∈ (0, 1)
2: Initialize π_θ : S × D → Δ(A) and B = ∅
3: for n = 1, 2, ···, N do
4:     World samples a task q = (R, s_1, d*) ~ P*(·)
5:     Agent generates ê ~ P̃(· | s_1, d*) (see Eq 4)
6:     Teacher generates a description d̂ ~ P_T(· | ê)
7:     B ← B ∪ (ê, d̂)
8:     Update agent policy:
           θ ← max_{θ'} Σ_{(ê,d̂)∈B} Σ_{(s,â_s)∈ê} log π_{θ'}(â_s | s, d̂)
       where â_s is the action taken by the agent in state s
9:     Anneal mixing weight: λ ← λ · β
10: return π_θ
```

3. ADEL: Learning from Activity Describers via Semi-Supervised Exploration

We frame the ILIAD problem as a density-estimation problem: *given that we can effectively draw samples from the distribution* $\mathbb{P}^\star(s_1, d)$ *and a teacher* $\mathbb{P}_T(d \mid e)$*, how do we learn a policy* π_θ *such that* $\mathbb{P}_{\pi_\theta}(e \mid s_1, d)$ *is close to* $\mathbb{P}^\star(e \mid s_1, d)$*?* Here, $\mathbb{P}^\star(e \mid s_1, d) = \mathbb{P}_{\pi^\star}(e \mid s_1, d)$ is the ground-truth request-fulfilling distribution obtained from the joint distribution defined in Eq 2.

If s_1 is not the start state of e, then $\mathbb{P}^\star(e \mid s_1, d) = 0$. Otherwise, by applying Bayes' rule, and noting that s_1 is included in e, we have:

$$\begin{aligned}\mathbb{P}^\star(e \mid s_1, d) &\propto \mathbb{P}^\star(e, d \mid s_1) = \mathbb{P}^\star(e \mid s_1)\mathbb{P}^\star(d \mid e, s_1), \\ &= \mathbb{P}^\star(e \mid s_1)\mathbb{P}^\star(d \mid e), \\ &\approx \mathbb{P}^\star(e \mid s_1)\mathbb{P}_T(d \mid e). \end{aligned} \tag{3}$$

As seen from the equation, the only missing piece required for estimating $\mathbb{P}^\star(e \mid s_1, d)$ is the marginal[3] $\mathbb{P}^\star(e \mid s_1)$. Alg 2 presents a simple method for learning an agent policy if we have access to this marginal. It is easy to show that the pairs $(\hat{e}, \hat{d})$ in the algorithm are approximately drawn from the joint distribution $\mathbb{P}^\star(e, d \mid s_1)$ and thus can be directly used to estimate the conditional $\mathbb{P}^\star(e \mid s_1, d)$.

Unfortunately, $\mathbb{P}^\star(e \mid s_1)$ is unknown in our setting. We present our main algorithm ADEL (Alg 3) which simultaneously estimates $\mathbb{P}^\star(e \mid s_1)$ and $\mathbb{P}^\star(e \mid s_1, d)$ through interactions with the teacher. In this algorithm, we assume access to an *approximate marginal* $\mathbb{P}_{\pi_\omega}(e \mid s_1)$ defined by an *explorative policy* $\pi_\omega(a \mid s)$. This policy can be learned from a dataset of unlabeled executions or be defined as a program that synthesizes executions. In many applications, reasonable unlabeled executions can be cheaply constructed using knowledge about the structure of the execution. For example, in robot navigation, valid executions are collision-free and non-looping; in semantic parsing, predicted parses should follow the syntax of the semantic language.

[3] We are largely concerned with the relationship between e and d, and so refer to the distribution $\mathbb{P}^\star(e \mid s_1)$ as the marginal and $\mathbb{P}^\star(e \mid s_1, d)$ as the conditional.

After constructing the approximate marginal $\mathbb{P}_{\pi_\omega}(e \mid s_1)$, we could substitute it for the true marginal in Alg 2. However, using a fixed approximation of the marginal may lead to sample inefficiency when there is a mismatch between the approximate marginal and the true marginal. For example, in the robot navigation example, if most human requests specify the kitchen as the destination, the agent should focus on generating executions that end in the kitchen to obtain descriptions that are similar to those requests. If instead, a uniform approximate marginal is used to generate executions, the agent obtains a lot of irrelevant descriptions.

ADEL minimizes potential marginal mismatch by iteratively using the estimate of the marginal $\mathbb{P}^\star(e \mid s_1)$ to improve the estimate of the conditional $\mathbb{P}^\star(e \mid s_1, d)$ and vice versa. Initially, we set $\mathbb{P}_{\pi_\omega}(e \mid s_1)$ as the marginal over executions. In each episode, we *mix* this distribution with $\mathbb{P}_{\pi_\theta}(e \mid s_1, d)$, the current estimate of the conditional, to obtain an improved estimate of the marginal (line 5). Formally, given a start state s_1 and a request $d^\star$, we sample an execution $\hat{e}$ from the following distribution:

$$\tilde{\mathbb{P}}(\cdot \mid s_1, d^\star) \triangleq \lambda \mathbb{P}_{\pi_\omega}(\cdot \mid s_1) + (1 - \lambda)\mathbb{P}_{\pi_\theta}(\cdot \mid s_1, d^\star) \tag{4}$$

where $\lambda \in [0, 1]$ is a mixing weight that is annealed to zero over the course of training. Each component of the mixture in Eq 4 is essential in different learning stages. Mixing with $\mathbb{P}_{\pi_\omega}$ accelerates convergence at the early stage of learning. Later, when π_θ improves, $\mathbb{P}_{\pi_\theta}$ skews $\tilde{\mathbb{P}}$ towards executions whose descriptions are closer to the requests, closing the gap with $\mathbb{P}^\star(e \mid s_1)$. In line 6-8, similar to Alg 2, we leverage the (improved) marginal estimate and the teacher to draw samples $(\hat{e}, \hat{d})$ and use them to re-estimate $\mathbb{P}_{\pi_\theta}$.

Theoretical Analysis. We analyze an epoch-based variant of ADEL and show that under certain assumptions, it converges to a near-optimal policy. In this variant, we run the algorithm in epochs, where the agent policy is only updated at the end of an epoch. In each epoch, we collect a fresh batch of examples $\{(\hat{e}, \hat{d})\}$ as in ADEL (line 4-7), and use them to perform a batch update (line 8). We provide a sketch of our theoretical results here and defer the full details to Appendix A.

We consider the case of $H = 1$ where an execution $e = (s_1, a)$ consists of the start state s_1 and a single action a taken by the agent. This setting while restrictive captures the non-trivial class of contextual bandit problems (Langford & Zhang, 2008b). Sequential decision-making problems where the agent makes decisions solely based on the start state can be reduced to this setting by treating a sequence of decisions as a single action (Kreutzer et al., 2017; Nguyen et al., 2017a). We focus on the convergence of the iterations of epochs, and assume that the maximum likelihood estimation problem in each epoch can be solved optimally. We also ablate the teacher learning difficulty by assuming access to a perfectly consistent teacher, i.e., $\mathbb{P}_T(d \mid e) = \mathbb{P}^\star(d \mid e)$.

We make two crucial assumptions. Firstly, we make a standard realizability assumption to ensure that our policy class is expressive enough to accommodate the optimal solution of the maximum likelihood estimation. Secondly, we assume that for every start state s_1, the teacher distribution's matrix $\mathbb{P}^\star(d \mid e_{s_1})$ over descriptions and executions e_{s_1} starting with s_1, has a non-zero minimum singular value $\sigma_{min}(s_1)$. Intuitively, this assumption implies that descriptions are rich enough to help in deciphering actions. Under these assumptions, we prove the following result:

Theorem 1 (Main Result)**.** *Let $\mathbb{P}_n(e \mid s_1)$ be the marginal distribution in the n^{th} epoch. Then for any $t \in \mathbb{N}$ and any start state s_1 we have:*

$$\|\mathbb{P}^\star(e \mid s_1) - \frac{1}{t}\sum_{n=1}^{t}\mathbb{P}_n(e \mid s_1)\|_2 \leq \frac{1}{\sigma_{min}(s_1)}\sqrt{\frac{2\ln|\mathcal{A}|}{t}}.$$

Theorem 1 shows that the running average of the estimated marginal distribution converges to the true marginal distribution. The error bound depends logarithmically on the size of action space, and therefore, suitable for problems with exponentially large action space. As argued before, access to the true marginal can be used to easily learn a near-optimal policy. For brevity, we defer the proof and other details to Appendix A. Hence, our results show that under certain conditions, we can expect convergence to the optimal policy. We leave the question of sample complexity and addressing more general settings for future work.

4. Experimental Setup

In this section, we present a general method for simulating an execution-describing teacher using a pre-collected dataset (§4.1). Then we describe setups of the two problems we conduct experiments on: vision-language navigation (§4.2) and word modification (§4.3). Details about the data, the model architecture, training hyperparameters, and how the teacher is simulated in each problem are in the Appendix.

We emphasize that the ILIAD protocol or the ADEL algorithm do *not* propose learning a teacher. Similar to IL and RL, ILIAD operates with a fixed, black-box teacher that is given in the environment. Our experiments specifically simulate *human* teachers that train request-fulfilling agents by talking to them (using descriptions). We use labeled executions only to learn approximate models of human teachers.

4.1. Simulating Teachers

ILIAD assumes access to a teacher $\mathbb{P}_T(d \mid e)$ that can describe agent executions in a description language. For our experimental purposes, employing human teachers is expensive and irreproducible, thus we simulate them using pre-collected datasets. We assume availability of a dataset $\mathcal{B}_{\text{sim}} = \{(\mathcal{D}^\star_n, e^\star_n)\}_{n=1}^N$, where $\mathcal{D}^\star_n = \{d_n^{\star(j)}\}_{j=1}^M$ contains M human-generated requests that are fulfilled by execution $e^\star_n$. Each of the two experimented problems is accompanied by data that is partitioned into training/validation/test splits. We use the training split as $\mathcal{B}_{\text{sim}}$ and use the other two splits for validation and testing, respectively. Our agents do *not* have direct access to $\mathcal{B}_{\text{sim}}$. From an agent's perspective, it communicates with a black-box teacher that can return descriptions of its executions; it does not know how the teacher is implemented.

Each ILIAD episode (Alg 1) requires providing a request $d^\star$ at the beginning and a description $\hat{d}$ of an execution $\hat{e}$. The request $d^\star$ is chosen by first uniformly randomly selecting an example $(\mathcal{D}^\star_n, e^\star_n)$ from $\mathcal{B}_{\text{sim}}$, and then uniformly sampling a request $d_n^{\star(j)}$ from $\mathcal{D}^\star_n$. The description $\hat{d}$ is generated as follows. We first gather all the pairs $(d_n^{\star(j)}, e^\star_n)$ from $\mathcal{B}_{\text{sim}}$ and train an RNN-based conditional language model $\tilde{\mathbb{P}}_T(d \mid e)$ via standard maximum log-likelihood. We can then generate a description of an execution $\hat{e}$ by greedily decoding[4] this model conditioned on $\hat{e}$: $\hat{d}_{\text{greedy}} = \texttt{greedy}\big(\tilde{\mathbb{P}}_T(\cdot \mid \hat{e})\big)$. However, given limited training data, this model may not generate sufficiently high-quality descriptions. We apply two techniques to improve the quality of the descriptions. First, we provide the agent with the human-generated requests in $\mathcal{B}_{\text{sim}}$ when the executions are near optimal. Let $\texttt{perf}\,(\hat{e}, e^\star)$[5] be a performance metric that evaluates an agent's execution $\hat{e}$ against a ground-truth $e^\star$ (higher is better). An execution $\hat{e}$ is near optimal if $\texttt{perf}\,(\hat{e}, e^\star) \geq \tau$, where τ is a constant threshold. Second, we apply prag-

[4]Greedily decoding an RNN-based model refers to stepwise choosing the highest-probability class of the output softmax. In this case, the classes are words in the description vocabulary.

[5]The metric perf is only used in simulating the teachers and is not necessarily the same as the reward function R.

matic inference (Andreas & Klein, 2016; Fried et al., 2018a), leveraging the fact that the teacher has access to the environment's simulator and can simulate executions of descriptions. The final description given to the agent is

$$\hat{d} \sim \begin{cases} \texttt{Unif}\left(\mathcal{D}^{\star}_{n}\right) & \text{if } \texttt{perf}\left(\hat{e}, e^{\star}_{n}\right) \geq \tau, \\ \texttt{Unif}\left(\mathcal{D}_{\text{prag}} \cup \{\emptyset\}\right) & \text{otherwise} \end{cases} \quad (5)$$

where $\texttt{Unif}(\mathcal{D})$ is a uniform distribution over elements of $\mathcal{D}$, $e^{\star}_{n}$ is the ground-truth execution associated with $\mathcal{D}^{\star}_{n}$, $\mathcal{D}_{\text{prag}}$ contains descriptions generated using pragmatic inference (which we will describe next), and $\emptyset$ is the empty string.

Improved Descriptions with Pragmatic Inference. Pragmatic inference emulates the teacher's ability to mentally simulate task execution. Suppose the teacher has its own execution policy $\pi_T(a \mid s, d)$, which is learned using the pairs $\left(e^{\star}_{n}, d^{\star(j)}_{n}\right)$ of $\mathcal{B}_{\text{sim}}$, and access to a simulator of the environment. A *pragmatic* execution-describing teacher is defined as $\mathbb{P}^{\text{prag}}_{T}(d \mid e) \propto \mathbb{P}_{\pi_T}(e \mid s_1, d)$. For this teacher, the more likely that a request d causes it to generate an execution e, the more likely that it describes e as d.

In our problems, constructing the pragmatic teacher's distribution explicitly is not feasible because we would have to compute a normalizing constant that sums over all possible descriptions. Instead, we follow Andreas et al. (2018), generating a set of candidate descriptions and using $\mathbb{P}_{\pi_T}(e \mid s_1, d)$ to re-rank those candidates. Concretely, for every execution $\hat{e}$ where $\texttt{perf}\left(\hat{e}, e^{\star}_{n}\right) < \tau$, we use the learned language model $\tilde{\mathbb{P}}_T$ to generate a set of candidate descriptions $\mathcal{D}_{\text{cand}} = \{\hat{d}_{\text{greedy}}\} \cup \{\hat{d}^{(k)}_{\text{sample}}\}^{K}_{k=1}$. This set consists of the greedily decoded description $\hat{d}_{\text{greedy}} = \texttt{greedy}\left(\tilde{\mathbb{P}}_T(\cdot \mid \hat{e})\right)$ and K descriptions $\hat{d}^{(k)}_{\text{sample}} \sim \tilde{\mathbb{P}}_T(\cdot \mid \hat{e})$. To construct $\mathcal{D}_{\text{prag}}$, we select descriptions in $\mathcal{D}_{\text{cand}}$ from which π_T generates executions that are similar enough to $\hat{e}$:

$$\mathcal{D}_{\text{prag}} = \left\{ d \mid d \in \mathcal{D}_{\text{cand}} \wedge \texttt{perf}\left(e^{d}, \hat{e}\right) \geq \tau \right\} \quad (6)$$

where $e^{d} = \texttt{greedy}\left(\mathbb{P}_{\pi_T}(\cdot \mid s_1, d)\right)$ and s_1 is the start state of $\hat{e}$.

4.2. Vision-Language Navigation (NAV)

Problem and Environment. An agent executes natural language requests (given in English) by navigating to locations in environments that photo-realistically emulate residential buildings (Anderson et al., 2018). The agent successfully fulfills a request if its final location is within three meters of the intended goal location. Navigation in an environment is framed as traversing in a graph where each node represents a location and each edge connects two nearby unobstructed locations. A state s of an agent represents its location and the direction it is facing. In the beginning, the agent starts in state s_1 and receives a navigation request $d^{\star}$. At every time step, the agent is not given the true state s but only receives an observation o, which is a real-world RGB image capturing the panoramic view at its current location.

Agent Policy. The agent maintains a policy $\pi_\theta(a \mid o, d)$ that takes in a current observation o and a request d, and outputs an action $a \in V_{\text{adj}}$, where V_{adj} denotes the set of locations that are adjacent to the agent's current location according to the environment graph. A special <stop> action is taken when the agent wants to terminate an episode or when it has taken H actions.

Simulated Teacher. We simulate a teacher that does not know how to control the navigation agent and thus cannot provide demonstrations. However, the teacher can verbally describe navigation paths taken by the agent. We follow §4.1, constructing a teacher $\mathbb{P}_T(d \mid e)$ that outputs language descriptions given executions $e = (o_1, a_1, \cdots, o_H)$.

4.3. Word Modification (REGEX)

Problem. A human gives an agent a natural language request (in English) $d^{\star}$ asking it to modify the characters of a word w^{inp}. The agent must execute the request and outputs a word $\hat{w}^{\text{out}}$. It successfully fulfills the request if $\hat{w}^{\text{out}}$ exactly matches the expected output w^{out}. For example, given an input word *embolden* and a request "*replace all n with c*", the expected output word is *emboldec*. We train an agent that solves this problem via a semantic parsing approach. Given w^{inp} and $d^{\star}$, the agent generates a regular expression $\hat{a}_{1:H} = (\hat{a}_1, \cdots, \hat{a}_H)$, which is a sequence of characters. It then uses a regular expression compiler to apply the regular expression onto the input word to produce an output word $\hat{w}^{\text{out}} = \texttt{compile}\left(w^{\text{inp}}, \hat{a}_{1:H}\right)$.

Agent Policy and Environment. The agent maintains a policy $\pi_\theta(a \mid s, d)$ that takes in a state s and a request d, and outputs a distribution over characters $a \in V_{\text{regex}}$, where V_{regex} is the regular expression (character) vocabulary. A special <stop> action is taken when the agent wants to stop generating the regular expression or when the regular expression exceeds the length limit H. We set the initial state $s_1 = \left(w^{\text{inp}}, \emptyset\right)$, where $\emptyset$ is the empty string. A next state is determined as follows

$$s_{t+1} = \begin{cases} \left(\hat{w}^{\text{out}}, \hat{a}_{1:t}\right) & \text{if } \hat{a}_t = \texttt{<stop>}, \\ \left(w^{\text{inp}}, \hat{a}_{1:t}\right) & \text{otherwise} \end{cases} \quad (7)$$

where $\hat{w}^{\text{out}} = \texttt{compile}\left(w^{\text{inp}}, \hat{a}_{1:t}\right)$.

Simulated Teacher. We simulate a teacher that does not have knowledge about regular expressions. Hence, instead of receiving full executions, which include regular expressions $\hat{a}_{1:H}$ predicted by the agent, the teacher generates

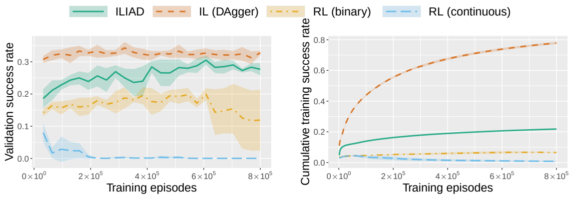

(a) NAV

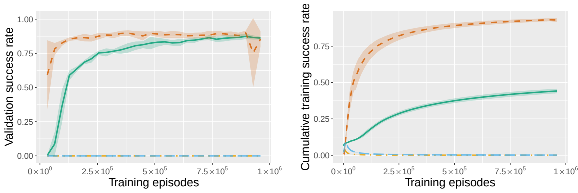

(b) REGEX

Figure 2. Validation success rate (average held-out return) and cumulative training success rate (average training return) over the course of training. For each algorithm, we report means and standard deviations over five runs with different random seeds.

descriptions given only pairs $(w_j^{\text{inp}}, \hat{w}_j^{\text{out}})$ of an input word w_j^{inp} and the corresponding output generated by the agent $\hat{w}_j^{\text{out}}$. In addition, to reduce ambiguity, the teacher requires multiple word pairs generate a description. This issue is better illustrated in the following example. Suppose the agent generates the pair *embolden* $\rightarrow$ *emboldec* by predicting a regular expression that corresponds to the description "*replace all n with c*". However, because the teacher does not observe the agent's regular expression (and cannot understand the expression even if it does), it can also describe the pair as "*replace the last letter with c*". Giving such a description to the agent would be problematic because the description does not correspond to the predicted regular expression. Observing multiple word pairs increases the chance that the teacher's description matches the agent's regular expression (e.g. adding *now* $\rightarrow$ *cow* help clarify that "*replace all n with c*" should be generated). In the end, the teacher is a model $P_T\left(d \mid \{w_j^{\text{inp}}, \hat{w}_j^{\text{out}}\}_{j=1}^{J}\right)$ that takes as input J word pairs. To generate J word pairs, in every episode, in addition to the episode's input word, we sample $J-1$ more words from the dictionary and execute the episode's request on the J words. We do *not* use any regular expression data in constructing the teacher. To train the teacher's policy π_T for pragmatic inference (§4.1), we use a dataset that consists of tuples $\left(\mathcal{D}_n^\star, (w_n^{\text{inp}}, w_n^{\text{out}})\right)$ which are not annotated with ground-truth regular expressions. π_T directly generates an output word instead of predicting a regular expression like the agent policy π_θ.

4.4. Baselines and Evaluation Metrics

We compare interactive learning settings that employ different teaching media:

- *Learning from activity description* (ILIAD): the teacher returns a language description $\hat{d}$.
- *Imitation learning* (IL): the teacher demonstrates the correct actions in the states that the agent visited, returning $e^\star = (s_1, a_1^\star, \cdots, a_H^\star, s_H)$, where s_i are the states in the agent's execution and $a_i^\star$ are the optimal actions in those states.
- *Reinforcement learning* (RL): the teacher provides a scalar reward that evaluates the agent's execution. We consider a special case when rewards are provided only at the end of an episode. Because such feedback is cheap to collect (e.g., star ratings) (Nguyen et al., 2017b; Kreutzer et al., 2018; 2020), this setting is suitable for large-scale applications. We experiment with both binary reward that indicates task success, and continuous reward that measures normalized distance to the goal (see Appendix D).

We use ADEL in the ILIAD setting, DAgger (Ross et al., 2011) in IL, and REINFORCE[6] (Williams, 1992) in RL. We report the *success rates* of these algorithms, which are the fractions of held-out (validation or test) examples on which the agent successfully fulfills its requests. All agents are initialized with random parameters.

5. Results

We compare the learning algorithms on not only success rate, but also the effort expended by the teacher. While task success rate is straightforward to compute, teacher effort is hard to quantify because it depends on many factors: the type of knowledge required to teach a task, the cognitive and physical ability of a teacher, etc. For example, in REGEX, providing demonstrations in forms of regular expressions may be easy for a computer science student, but could be challenging for someone who is unfamiliar with programming. In NAV, controlling a robot may not be viable for an individual with motor impairment, whereas generating language descriptions may infeasible for someone with a verbal-communication disorder. Because it is not possible cover all teacher demographics, our goal is to *quantitatively* compare the learning algorithms on learning effectiveness and efficiency, and *qualitatively* compare them on the teacher effort to learn to express feedback using the protocol's communication medium. Our overall findings (Table 1) highlight the strengths and weaknesses of each learning algorithm and can potentially aid practitioners in selecting algorithms that best suit their applications.

[6]We use a moving-average baseline to reduce variance. We also experimented with A2C (Mnih et al., 2016) but it was less stable in this sparse-reward setting. At the time this paper was written, we were not aware of any work that successfully trained agents using RL without supervised-learning bootstrapping in the two problems we experimented on.

Table 2. Main results. We report means and standard deviations of success rates (%) over five runs with different random seeds. RL-Binary and RL-Cont refer to the RL settings with binary and continuous rewards, respectively. Sample complexity is the number of training episodes (or number of teacher responses) required to reach a validation success rate of at least c. Note that the teaching efforts are not comparable across the learning settings: providing a demonstration can be more or less tedious than providing a language description depending on various characteristics of the teacher. Hence, even though ADEL requires more episodes to reach the same performance as DAgger, we do not draw any conclusions about the primacy of one algorithm over the other in terms of teaching effort.

				Sample complexity ↓		
Learning setting	Algorithm	Val success rate (%) ↑	Test success rate (%) ↑	# Demonstrations	# Rewards	# Descriptions
Vision-language navigation					(c = 30.0%)	
IL	DAgger	35.6 ± 1.35	32.0 ± 1.63	45K $\pm$ 26K	-	-
RL-Binary	REINFORCE	22.4 ± 1.15	20.5 ± 0.58	-	$+\infty$	-
RL-Cont	REINFORCE	11.1 ± 2.19	11.3 ± 1.25	-	$+\infty$	-
ILIAD	ADEL	32.2 ± 0.97	31.9 ± 0.76	-	-	406K $\pm$ 31K
Word modification					(c = 85.0%)	
IL	DAgger	92.5 ± 0.53	93.0 ± 0.37	118K $\pm$ 16K	-	-
RL-Binary	REINFORCE	0.0 ± 0.00	0.0 ± 0.00	-	$+\infty$	-
RL-Cont	REINFORCE	0.0 ± 0.00	0.0 ± 0.00	-	$+\infty$	-
ILIAD	ADEL	88.1 ± 1.60	89.0 ± 1.30	-	-	573K $\pm$ 116K

Main results. Our main results are in Table 2. Overall, results in both problems match our expectations. The IL baseline achieves the highest success rates (on average, 35.6% on NAV and 92.5% on REGEX). This framework is most effective because the feedback directly specifies ground-truth actions. The RL baseline is unable to reach competitive success rates. Especially, in REGEX, the RL agent cannot learn the syntax of the regular expressions and completely fails at test time. This shows that the reward feedback is not sufficiently informative to guide the agent to explore efficiently in this problem. ADEL's success rates are slightly lower than those of IL (3-4% lower than) but are substantially higher than those of RL (+9.8% on NAV and +88.1% on REGEX compared to the best RL results).

To measure learning efficiency, we report the number of training episodes required to reach a substantially high success rate (30% for NAV and 85% for REGEX). We observe that all algorithms require hundreds of thousands of episodes to attain those success rates. The RL agents cannot learn effectively even after collecting more than 1M responses from the teachers. ADEL attains reasonable success rates using 5-9 times more responses than IL. This is a decent efficiency considering that ADEL needs to find the ground-truth executions in exponentially large search spaces, while IL directly communicates these executions to the agents. As ADEL lacks access to ground-truth executions, its average training returns are 2-4 times lower than those of IL (Figure 2).

Ablation. We study the effects of mixing with the approximate marginal ($\mathbb{P}_{\pi_\omega}$) in ADEL (Table 3). First of all, we observe that learning cannot take off without using the approximate marginal ($\lambda = 0$). On the other hand, using only the approximate marginal to generate executions ($\lambda = 1$) degrades performance, in terms of both success rate and sample efficiency. This effect is more visible on REGEX where the success rate drops by 33% (compared to a 3% drop in NAV), indicating that the gap between the approximate marginal and the true marginal is larger in REGEX than in NAV. This matches our expectation as the set of unlabeled executions that we generate to learn π_ω in REGEX covers a smaller portion of the problem's execution space than that in NAV. Finally, mixing the approximate marginal and the agent-estimated conditional ($\lambda = 0.5$) gives the best results.

Table 3. Effects of mixing execution policies in ADEL.

Mixing weight	Val success rate (%) ↑	Sample complexity ↓
Vision-language navigation		
$\lambda = 0$ (no marginal)	0.0	$+\infty$
$\lambda = 1$	29.4	$+\infty$
$\lambda = 0.5$ (final)	32.0	384K
Word modification		
$\lambda = 0$ (no marginal)	0.2	$+\infty$
$\lambda = 1$	55.7	$+\infty$
$\lambda = 0.5$ (final)	88.0	608K

6. Related Work

Learning from Language Feedback. Frameworks for learning from language-based communication have been previously proposed. Common approaches include: reduction to reinforcement learning (Goldwasser & Roth, 2014; MacGlashan et al., 2015; Ling & Fidler, 2017; Goyal et al., 2019; Fu et al., 2019; Sumers et al., 2020), learning to ground language to actions (Chen & Mooney, 2011; Misra et al., 2014; Bisk et al., 2016; Liu et al., 2016; Wang et al., 2016; Li et al., 2017; 2020a;b), or devising EM-based algorithms to parse language into logical forms (Matuszek et al., 2012; Labutov et al., 2018). The first approach may discard useful learning signals from language feedback and inherits the limitations of RL algorithms. The second requires extra effort from the teacher to provide demonstrations. The third approach has to bootstrap the language parser with labeled executions. ADEL enables learning from a specific type of language feedback (language description) without reducing it to reward, requiring demonstrations, or assuming access to labeled executions.

Description Feedback in Reinforcement Learning. Recently, several papers have proposed using language description feedback in the context of reinforcement learning (Jiang et al., 2019; Chan et al., 2019; Colas et al., 2020; Cideron et al., 2020). These frameworks can be viewed as extensions of hindsight experience replay (HER; Andrychowicz et al., 2017) to language goal generation. While the teacher in ILIAD can be considered as a language goal generator, an important distinction between ILIAD and these frameworks is that ILIAD models a completely *reward-free* setting. Unlike in HER, the agent in ILIAD does not have access to a reward function that it can use to compute the reward of any tuple of state, action, and goal. With the feedback coming solely from language descriptions, ILIAD is designed so that task learning relies only on extracting information from language. Moreover, unlike reward, the description language in ILIAD does not contain information that *explicitly* encourages or discourages actions of the agent. The formalism and theoretical studies of ILIAD presented in this work are based on a probabilistic formalism and do not involve reward maximization.

Description Feedback in Vision-Language Navigation. Several papers (Fried et al., 2018b; Tan et al., 2019) apply back-translation to vision-language navigation (Anderson et al., 2018). While also operating with an output-to-input translator, back-translation is a single-round, offline process, whereas ILIAD is an iterative, online process. Zhou & Small (2021) study a test-time scenario that is similar to ILIAD but requires labeled demonstrations to learn the execution describer and to initialize the agent. The teacher in ILIAD is more general: it can be automated (i.e., learned from labeled data), but it can also be a human. Our experiments emulate applications where non-expert humans teach agents new tasks by only giving them verbal feedback. We use labeled demonstrations to simulate human teachers, but it is part of the experimental setup, not part of our proposed protocol and algorithm. Our agent does not have access to labeled demonstrations; it is initialized with *random parameters* and is trained with only language-description feedback. Last but not the least, we provide theoretical guarantees for ADEL, while these works only present empirical studies.

Connection to Pragmatic Reasoning. Another related line of research is work on the rational speech act (RSA) or pragmatic reasoning (Grice, 1975; Golland et al., 2010; Monroe & Potts, 2015; Goodman & Frank, 2016; Andreas & Klein, 2016; Fried et al., 2018a), which is also concerned with transferring information via language. It is important to point out that RSA is a mental *reasoning* model whereas ILIAD is an *interactive* protocol. In RSA, a speaker (or a listener) constructs a pragmatic message-encoding (or decoding) scheme by building an internal model of a listener (or a speaker). Importantly, during that process, one agent *never* interacts with the other. In contrast, the ILIAD agent learns through interaction with a teacher. In addition, RSA focuses on encoding (or decoding) a single message while ILIAD defines a process consisting of multiple rounds of message exchanging. We employ pragmatic inference to improve the quality of the simulated teachers but in our context, the technique is used to set up the experiments and is not concerned about communication between the teacher and the agent.

Connection to Emergent Language. Finally, our work also fundamentally differs from work on (RL-based) emergent language (Foerster et al., 2016; Lazaridou et al., 2017; Havrylov & Titov, 2017; Das et al., 2017; Evtimova et al., 2018; Kottur et al., 2017) in that we assume the teacher speaks a fixed, well-formed language, whereas in these works the teacher begins with no language capability and learns a language over the course of training.

7. Conclusion

The communication protocol of a learning framework places natural boundaries on the learning efficiency of any algorithm that instantiates the framework. In this work, we illustrate the benefits of designing learning algorithms based on a natural, descriptive communication medium like human language. Employing such expressive protocols leads to ample room for improving learning algorithms. Exploiting compositionality of language to improve sample efficiency, and learning with diverse types of feedback are interesting areas of future work. Extending the theoretical analyses of ADEL to more general settings is also an exciting open problem.

Acknowledgement

We would like to thank Hal Daumé III, Kianté Brantley, Anna Sotnikova, Yang Cao, Surbhi Goel, Akshay Krishnamurthy, and Cyril Zhang for their insightful comments on the paper. We thank Jacob Andreas for the discussion on pragmatic inference and thank Huyen Nguyen for useful conversations about human behavior. We also thank Microsoft GCR team for providing computational resources.

References

Agarwal, A., Kakade, S., Krishnamurthy, A., and Sun, W. Flambe: Structural complexity and representation learning of low rank mdps. In *Proceedings of Advances in Neural Information Processing Systems*, 2020.

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

Andreas, J. and Klein, D. Reasoning about pragmatics with neural listeners and speakers. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 1173–1182, Austin, Texas, November 2016. Association for Computational Linguistics. doi: 10.18653/v1/D16-1125. URL *https://www.aclweb.org/anthology/D16-1125*.

Andreas, J., Klein, D., and Levine, S. Learning with latent language. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 2166–2179, New Orleans, Louisiana, June 2018. Association for Computational Linguistics. doi: 10.18653/v1/N18-1197. URL *https://www.aclweb.org/anthology/N18-1197*.

Andrychowicz, M., Wolski, F., Ray, A., Schneider, J., Fong, R., Welinder, P., McGrew, B., Tobin, J., Abbeel, P., and Zaremba, W. Hindsight experience replay. *arXiv preprint arXiv:1707.01495*, 2017.

Artzi, Y., FitzGerald, N., and Zettlemoyer, L. Semantic parsing with combinatory categorial grammars. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Tutorials)*, pp. 2, Sofia, Bulgaria, August 2013. Association for Computational Linguistics. URL *https://www.aclweb.org/anthology/P13-5002*.

Bisk, Y., Yuret, D., and Marcu, D. Natural language communication with robots. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 751–761, San Diego, California, June 2016. Association for Computational Linguistics. doi: 10.18653/v1/N16-1089. URL *https://www.aclweb.org/anthology/N16-1089*.

Chan, H., Wu, Y., Kiros, J., Fidler, S., and Ba, J. Actrce: Augmenting experience via teacher's advice for multi-goal reinforcement learning. *arXiv preprint arXiv:1902.04546*, 2019.

Chaplot, D. S., Sathyendra, K. M., Pasumarthi, R. K., Rajagopal, D., and Salakhutdinov, R. Gated-attention architectures for task-oriented language grounding. In *Association for the Advancement of Artificial Intelligence*, 2018.

Chen, D. L. and Mooney, R. J. Learning to interpret natural language navigation instructions from observations. In *Association for the Advancement of Artificial Intelligence*, 2011.

Chen, H., Suhr, A., Misra, D., and Artzi, Y. Touchdown: Natural language navigation and spatial reasoning in visual street environments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019.

Chevalier-Boisvert, M., Bahdanau, D., Lahlou, S., Willems, L., Saharia, C., Nguyen, T. H., and Bengio, Y. Babyai: A platform to study the sample efficiency of grounded language learning. In *Proceedings of the International Conference on Learning Representations*, 2019.

Cideron, G., Seurin, M., Strub, F., and Pietquin, O. Higher: Improving instruction following with hindsight generation for experience replay. In *2020 IEEE Symposium Series on Computational Intelligence (SSCI)*, pp. 225–232. IEEE, 2020.

Colas, C., Karch, T., Lair, N., Dussoux, J.-M., Moulin-Frier, C., Dominey, P. F., and Oudeyer, P.-Y. Language as a cognitive tool to imagine goals in curiosity-driven exploration. In *Proceedings of Advances in Neural Information Processing Systems*, 2020.

Das, A., Kottur, S., Moura, J. M., Lee, S., and Batra, D. Learning cooperative visual dialog agents with deep reinforcement learning. In *International Conference on Computer Vision*, 2017.

Evtimova, K., Drozdov, A., Kiela, D., and Cho, K. Emergent communication in a multi-modal, multi-step referential game. In *Proceedings of the International Conference on Learning Representations*, 2018.

Foerster, J. N., Assael, Y. M., Freitas, N. D., and Whiteson, S. Learning to communicate with deep multi-agent reinforcement learning. In *NIPS*, 2016.

Fried, D., Andreas, J., and Klein, D. Unified pragmatic models for generating and following instructions. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 1951–1963, New Orleans, Louisiana, June 2018a. Association for Computational Linguistics. doi: 10.18653/v1/N18-1177. URL *https://www.aclweb.org/anthology/N18-1177*.

Fried, D., Hu, R., Cirik, V., Rohrbach, A., Andreas, J., Morency, L.-P., Berg-Kirkpatrick, T., Saenko, K., Klein, D., and Darrell, T. Speaker-follower models for vision-and-language navigation. In *Proceedings of Advances in Neural Information Processing Systems*, 2018b.

Fu, J., Korattikara, A., Levine, S., and Guadarrama, S. From language to goals: Inverse reinforcement learning for vision-based instruction following. In *Proceedings of the International Conference on Learning Representations*, 2019.

Gaddy, D. and Klein, D. Pre-learning environment representations for data-efficient neural instruction following. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 1946–1956, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1188. URL *https://www.aclweb.org/anthology/P19-1188*.

Goldwasser, D. and Roth, D. Learning from natural instructions. *Machine learning*, 94(2):205–232, 2014.

Golland, D., Liang, P., and Klein, D. A game-theoretic approach to generating spatial descriptions. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pp. 410–419, Cambridge, MA, October 2010. Association for Computational Linguistics. URL *https://www.aclweb.org/anthology/D10-1040*.

Goodman, N. D. and Frank, M. C. Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*, 20(11):818–829, 2016.

Goyal, P., Niekum, S., and Mooney, R. J. Using natural language for reward shaping in reinforcement learning. In *International Joint Conference on Artificial Intelligence*, 2019.

Grice, H. P. Logic and conversation. In *Speech acts*, pp. 41–58. Brill, 1975.

Havrylov, S. and Titov, I. Emergence of language with multi-agent games: Learning to communicate with sequences of symbols. In *Proceedings of Advances in Neural Information Processing Systems*, 2017.

Hermann, K. M., Hill, F., Green, S., Wang, F., Faulkner, R., Soyer, H., Szepesvari, D., Czarnecki, W., Jaderberg, M., Teplyashin, D., Wainwright, M., Apps, C., Hassabis, D., and Blunsom, P. Grounded language learning in a simulated 3D world. *CoRR*, abs/1706.06551, 2017.

Jiang, Y., Gu, S., Murphy, K., and Finn, C. Language as an abstraction for hierarchical deep reinforcement learning. In *Proceedings of Advances in Neural Information Processing Systems*, 2019.

Kottur, S., Moura, J., Lee, S., and Batra, D. Natural language does not emerge 'naturally' in multi-agent dialog. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pp. 2962–2967, Copenhagen, Denmark, September 2017. Association for Computational Linguistics. doi: 10.18653/v1/D17-1321. URL *https://www.aclweb.org/anthology/D17-1321*.

Kreutzer, J., Sokolov, A., and Riezler, S. Bandit structured prediction for neural sequence-to-sequence learning. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1503–1513, Vancouver, Canada, July 2017. Association for Computational Linguistics. doi: 10.18653/v1/P17-1138. URL *https://www.aclweb.org/anthology/P17-1138*.

Kreutzer, J., Uyheng, J., and Riezler, S. Reliability and learnability of human bandit feedback for sequence-to-sequence reinforcement learning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1777–1788, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi: 10.18653/v1/P18-1165. URL *https://www.aclweb.org/anthology/P18-1165*.

Kreutzer, J., Riezler, S., and Lawrence, C. Learning from human feedback: Challenges for real-world reinforcement learning in nlp. In *Proceedings of Advances in Neural Information Processing Systems*, 2020.

Labutov, I., Yang, B., and Mitchell, T. Learning to learn semantic parsers from natural language supervision. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 1676–1690, Brussels, Belgium, October-November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1195. URL *https://www.aclweb.org/anthology/D18-1195*.

Langford, J. and Zhang, T. The epoch-greedy algorithm for multi-armed bandits with side information. In *Advances in neural information processing systems*, pp. 817–824, 2008a.

Langford, J. and Zhang, T. The epoch-greedy algorithm for multi-armed bandits with side information. In Platt, J., Koller, D., Singer, Y., and Roweis, S. (eds.), *Advances in Neural Information Processing Systems*, volume 20, pp. 817–824. Curran Associates, Inc., 2008b. URL *https://proceedings.neurips.cc/paper/2007/file/4b04a686b0ad13dce35fa99fa4161c65-Paper.pdf*.

Lazaridou, A., Peysakhovich, A., and Baroni, M. Multi-agent cooperation and the emergence of (natural) language. In *Proceedings of the International Conference on Learning Representations*, 2017.

Li, T. J.-J., Li, Y., Chen, F., and Myers, B. A. Programming iot devices by demonstration using mobile apps. In *International Symposium on End User Development*, pp. 3–17. Springer, 2017.

Li, T. J.-J., Chen, J., Mitchell, T. M., and Myers, B. A. Towards effective human-ai collaboration in gui-based interactive task learning agents. *Workshop on Artificial Intelligence for HCI: A Modern Approach (AI4HCI)*, 2020a.

Li, T. J.-J., Radensky, M., Jia, J., Singarajah, K., Mitchell, T. M., and Myers, B. A. Interactive task and concept learning from natural language instructions and gui demonstrations. In *The AAAI-20 Workshop on Intelligent Process Automation (IPA-20)*, 2020b.

Ling, H. and Fidler, S. Teaching machines to describe images via natural language feedback. In *Proceedings of Advances in Neural Information Processing Systems*, 2017.

Liu, C., Yang, S., Saba-Sadiya, S., Shukla, N., He, Y., Zhu, S.-C., and Chai, J. Jointly learning grounded task structures from language instruction and visual demonstration. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 1482–1492, 2016.

MacGlashan, J., Babes-Vroman, M., desJardins, M., Littman, M. L., Muresan, S., Squire, S., Tellex, S., Arumugam, D., and Yang, L. Grounding english commands to reward functions. In *Robotics: Science and Systems*, 2015.

Magalhaes, G. I., Jain, V., Ku, A., Ie, E., and Baldridge, J. General evaluation for instruction conditioned navigation using dynamic time warping. In *Proceedings of Advances in Neural Information Processing Systems*, 2019.

Matuszek, C., FitzGerald, N., Zettlemoyer, L., Bo, L., and Fox, D. A joint model of language and perception for grounded attribute learning. In *Proceedings of the International Conference of Machine Learning*, 2012.

Mei, H., Bansal, M., and Walter, M. R. Listen, attend, and walk: Neural mapping of navigational instructions to action sequences. In *Association for the Advancement of Artificial Intelligence (AAAI)*, 2016.

Miryoosefi, S., Brantley, K., Daume III, H., Dudik, M., and Schapire, R. E. Reinforcement learning with convex constraints. In *Proceedings of Advances in Neural Information Processing Systems*, 2019.

Misra, D., Langford, J., and Artzi, Y. Mapping instructions and visual observations to actions with reinforcement learning. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 2017.

Misra, D. K., Sung, J., Lee, K., and Saxena, A. Tell Me Dave: Context-sensitive grounding of natural language to mobile manipulation instructions. In *Robotics: Science and Systems (RSS)*, 2014.

Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., and Kavukcuoglu, K. Asynchronous methods for deep reinforcement learning. In *Proceedings of the International Conference of Machine Learning*, 2016.

Monroe, W. and Potts, C. Learning in the Rational Speech Acts model. In *Proceedings of 20th Amsterdam Colloquium*, 2015.

Nguyen, K. and Daumé III, H. Help, anna! visual navigation with natural multimodal assistance via retrospective curiosity-encouraging imitation learning. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, November 2019. URL *https://arxiv.org/abs/1909.01871*.

Nguyen, K., Daumé III, H., and Boyd-Graber, J. Reinforcement learning for bandit neural machine translation with simulated human feedback. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pp. 1464–1474, Copenhagen, Denmark, September 2017a. Association for Computational Linguistics. doi: 10.18653/v1/D17-1153. URL *https://www.aclweb.org/anthology/D17-1153*.

Nguyen, K., Daumé III, H., and Boyd-Graber, J. Reinforcement learning for bandit neural machine translation with simulated human feedback. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pp. 1464–1474, Copenhagen, Denmark, September 2017b. Association for Computational Linguistics. doi: 10.18653/v1/D17-1153. URL *https://www.aclweb.org/anthology/D17-1153*.

Nguyen, K., Dey, D., Brockett, C., and Dolan, B. Vision-based navigation with language-based assistance via imitation learning with indirect intervention. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. URL *https://arxiv.org/abs/1812.04155*.

Ross, S., Gordon, G., and Bagnell, D. A reduction of imitation learning and structured prediction to no-regret online learning. In *Artificial Intelligence and Statistics (AISTATS)*, 2011.

Stadie, B. C., Abbeel, P., and Sutskever, I. Third-person imitation learning. In *Proceedings of the International Conference on Learning Representations*, 2017.

Sumers, T. R., Ho, M. K., Hawkins, R. D., Narasimhan, K., and Griffiths, T. L. Learning rewards from linguistic feedback. In *Association for the Advancement of Artificial Intelligence*, 2020.

Sun, W., Vemula, A., Boots, B., and Bagnell, J. A. Provably efficient imitation learning from observation alone. In *Proceedings of the International Conference of Machine Learning*, June 2019.

Sutton, R. S. and Barto, A. G. *Reinforcement learning: An introduction*. MIT press, 2018.

Tan, H., Yu, L., and Bansal, M. 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: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 2610–2621, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1268. URL *https://www.aclweb.org/anthology/N19-1268*.

Tellex, S., Thaker, P., Joseph, J., and Roy, N. Toward learning perceptually grounded word meanings from unaligned parallel data. In *Proceedings of the Second Workshop on Semantic Interpretation in an Actionable Context*, pp. 7–14, Montréal, Canada, June 2012. Association for Computational Linguistics. URL *https://www.aclweb.org/anthology/W12-2802*.

Wang, S. I., Liang, P., and Manning, C. D. Learning language games through interaction. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 2368–2378, Berlin, Germany, August 2016. Association for Computational Linguistics. doi: 10.18653/v1/P16-1224. URL *https://www.aclweb.org/anthology/P16-1224*.

Williams, R. J. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8, 1992.

Winograd, T. Understanding natural language. *Cognitive Psychology*, 3(1):1–191, 1972.

Yao, Z., Tang, Y., Yih, W.-t., Sun, H., and Su, Y. An imitation game for learning semantic parsers from user interaction. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 6883–6902, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.559. URL *https://www.aclweb.org/anthology/2020.emnlp-main.559*.

Zhou, L. and Small, K. Inverse reinforcement learning with natural language goals. In *AAAI*, 2021.