

Learning to Generalize from Sparse and Underspecified Rewards

Rishabh Agarwal^{1*} Chen Liang¹ Dale Schuurmans^{1,2} Mohammad Norouzi¹

Abstract

We consider the problem of learning from sparse and underspecified rewards, where an agent receives a complex input, such as a natural language instruction, and needs to generate a complex response, such as an action sequence, while only receiving binary success-failure feedback. Such success-failure rewards are often underspecified: they do not distinguish between purposeful and accidental success. Generalization from underspecified rewards hinges on discounting spurious trajectories that attain accidental success, while learning from sparse feedback requires effective exploration. We address exploration by using a mode covering direction of KL divergence to collect a diverse set of successful trajectories, followed by a mode seeking KL divergence to train a robust policy. We propose Meta Reward Learning (MeRL) to construct an auxiliary reward function that provides more refined feedback for learning. The parameters of the auxiliary reward function are optimized with respect to the validation performance of a trained policy. The MeRL approach outperforms an alternative method for reward learning based on Bayesian Optimization, and achieves the state-of-the-art on weakly-supervised semantic parsing. It improves previous work by 1.2% and 2.4% on WIKITABLEQUESTIONS and WIKISQL datasets respectively.

1. Introduction

Effortlessly communicating with computers using natural language has been a longstanding goal of artificial intelligence (Winograd, 1971). Reinforcement Learning (RL) presents a flexible framework for optimizing goal oriented behavior (Sutton & Barto, 2018). As such, one can use RL to

*Work done as part of the Google AI Residency Program.

¹Google Research, Brain Team ²University of Alberta. Correspondence to: Rishabh Agarwal <rishabhagarwal@google.com>, Mohammad Norouzi <mnorouzi@google.com>.

Rank	Nation	Gold	Silver
1	USA	10	12
2	GBR	9	4
3	CHN	8	11
4	RUS	2	4
5	GER	2	2
6	JPN	2	1
7	FRA	2	1

x = “Which nation won the most silver medal?”

$R(a) = \mathbb{1}[\text{Execute}(a) = \text{“USA”}]$


$a_1 = \text{argmax}_{\text{row}}(\text{Silver}) . \text{Nation}$

$a_2 = \text{argmax}_{\text{row}}(\text{Gold}) . \text{Nation}$

$a_3 = \text{argmin}_{\text{row}}(\text{Rank}) . \text{Nation}$

$R(a_1) = R(a_2) = R(a_3) = 1$

Figure 1. Semantic parsing from question-answer pairs. An agent is presented with a natural language question x and is asked to generate a SQL-like program a . The agent receives a reward of 1 if execution of a program a on the relevant data table leads to the correct answer (e.g., USA). The reward is underspecified because spurious programs (e.g., a_2, a_3) can also achieve a reward of 1.



x = “Right Up Up Right”

$R(a) = \mathbb{1}[\text{Execute}(\bullet, a) = \star]$

$a_1 = (\rightarrow, \uparrow, \uparrow, \rightarrow)$

$a_2 = (\leftarrow, \rightarrow, \rightarrow, \uparrow, \uparrow, \rightarrow)$

$a_3 = (\uparrow, \rightarrow, \rightarrow, \uparrow)$

$R(a_1) = R(a_2) = R(a_3) = 1$

Figure 2. Instruction following in a simple maze. A blind agent is presented with a sequence of (Left, Right, Up, Down) instructions. Given the input text, the agent (\bullet) performs a sequence of actions, and only receives a reward of 1 if it reaches the goal (\star).

optimize language communication if it is expressed in terms of achieving concrete goals. In this pursuit, researchers have created a number of simulation environments where a learning agent is provided with a natural language input and asked to produce a sequence of actions for achieving a goal specified in the input text (e.g., Long et al. (2016); Hermann et al. (2017); Chaplot et al. (2018); Fu et al. (2019); Chevalier-Boisvert et al. (2018)). These tasks are typically episodic, where the agent receives sparse binary success-failure feedback indicating whether an intended goal has been accomplished. After training, the agent is placed in new contexts and evaluated based on its ability to reach novel goals, indicating the quality of its behavior policy and language interpretation skills. The emphasis on generalization in these tasks makes them suitable for benchmarking overfitting in RL (Cobbe et al., 2018; Zhang et al., 2018).

Figure 1 and 2 illustrate two examples of contextual environments with sparse and underspecified rewards. The rewards are sparse, since only a few trajectories in the combinatorial

space of all trajectories leads to a non-zero return. In addition, the rewards are *underspecified*, since the agent may receive a return of 1 for exploiting *spurious* patterns in the environment. We assert that the generalization performance of an agent trained in this setting hinges on (1) effective exploration to find successful trajectories, and (2) discounting spurious trajectories to learn a generalizable behavior.

To facilitate effective and principled exploration, we propose to disentangle combinatorial search and exploration from robust policy optimization. In particular, we use a *mode covering* direction of KL divergence to learn a high entropy exploration policy to help collect a diverse set of successful trajectories. Then, given a buffer of promising trajectories, we use a *mode seeking* direction of KL divergence to learn a robust policy with favorable generalization performance.

A key challenge in language conditional learning environments is the lack of fully specified rewards that perfectly distinguish optimal and suboptimal trajectories. Designing a rich trajectory-level reward function requires a deep understanding of the semantic relationship between the environment and the natural language input, which is not available in most real-world settings. Such a challenge arises in weakly supervised semantic parsing as depicted in Figure 1 (Pasupat & Liang, 2015). From an AI safety perspective, underspecified rewards may lead to reward hacking (Amodei et al., 2016) causing unintended and harmful behavior when deployed in real-world scenarios.

In this paper, we investigate whether one can automatically discover a rich trajectory-level reward function to help a learning agent discount spurious trajectories and improve generalization. Toward this end, we utilize both gradient-based Meta-Learning (Finn et al., 2017; Maclaurin et al., 2015) and Bayesian Optimization (Snoek et al., 2012) for reward learning. We propose to optimize the parameters of the auxiliary reward function in an outer loop to maximize generalization performance of a policy trained based on the auxiliary rewards. Our work is distinct from recent works (Bahdanau et al., 2019; Fu et al., 2019) on learning rewards for language tasks because we do not require any form of trajectory or goal demonstration.

We evaluate our overall approach (see Figure 3 for an overview) on two real world weakly-supervised semantic parsing benchmarks (Pasupat & Liang, 2015; Zhong et al., 2017) (Figure 1) and a simple instruction following environment (Figure 2). In all of the experiments, we observe a significant benefit from the proposed Meta Reward Learning (MeRL) approach, even when the exploration problem is synthetically mitigated. In addition, we achieve notable gains from the mode covering exploration strategy, which combines well with MeRL to achieve the state-of-the-art results on weakly-supervised semantic parsing.




Figure 3. Overview of the proposed approach. We employ (1) mode covering exploration to collect a diverse set of successful trajectories in a memory buffer; (2) Meta-learning or Bayesian optimization to learn an auxiliary reward function to discount spurious trajectories.

2. Formulation

2.1. Problem statement

Let \mathbf{x} denote a complex input, such as a natural language question or instruction, which places an agent in some context. Let \mathbf{a} denote a multivariate response, such as an action trajectory that the agent should produce. Let $R(\mathbf{a}|\mathbf{x}, y) \in \{0, 1\}$ denote a contextual success-failure feedback that uses some side information y to decide whether \mathbf{a} is successful in the context of \mathbf{x} and y . For instance, y may be some goal specification, e.g., the answer (denotation) in Figure 1, or the 2D coordinates of the goal in Figure 2. For simplicity of the exposition, we assume that $R(\mathbf{a}|\mathbf{x}, y)$ is deterministic, even though our results are applicable to stochastic rewards as well. To simplify the equations, we drop the conditioning of the return function on \mathbf{x} and y and express the return function as $R(\mathbf{a})$.

Our aim is to optimize the parameters of a stochastic policy $\pi(\mathbf{a}|\mathbf{x})$ according to a training set in order to maximize the empirical success rate of a policy on novel test contexts. For evaluation, the agent is required to only provide a single action trajectory $\hat{\mathbf{a}}$ for each context \mathbf{x} , which is accomplished via greedy decoding for interactive environments, and beam search for non-interactive environments to perform approximate inference:

$$\hat{\mathbf{a}} \approx \operatorname{argmax}_{\mathbf{a} \in \mathcal{A}(\mathbf{x})} \pi(\mathbf{a}|\mathbf{x}). \quad (1)$$

Let $\mathcal{A}(\mathbf{x})$ denote the combinatorial set of all plausible action trajectories for a context \mathbf{x} , and let $\mathcal{A}^+(\mathbf{x})$ denote a subset of $\mathcal{A}(\mathbf{x})$ comprising successful trajectories, i.e., $\mathcal{A}^+(\mathbf{x}) \equiv \{\mathbf{a} \in \mathcal{A}(\mathbf{x}) \mid R(\mathbf{a}|\mathbf{x}, y) = 1\}$.

2.2. Standard Objective Functions

To address the problem of policy learning from binary success-failure feedback, previous work has proposed the

following objective functions:

► **IML (Iterative Maximum Likelihood)** estimation (Liang et al., 2017; Abolafia et al., 2018) is an iterative process for optimizing a policy based on

$$O_{\text{IML}} = \sum_{\mathbf{x} \in \mathcal{D}} \frac{1}{|\mathcal{A}^+(\mathbf{x})|} \sum_{\mathbf{a}^+ \in \mathcal{A}^+(\mathbf{x})} \log \pi(\mathbf{a}^+ | \mathbf{x}). \quad (2)$$

The key idea is to replace $\mathcal{A}^+(\mathbf{x})$ in (2) with a buffer of successful trajectories collected so far, denoted $\mathcal{B}^+(\mathbf{x})$. While the policy is being optimized based on (2), one can also perform exploration by drawing *i.i.d.* samples from $\pi(\cdot | \mathbf{x})$ and adding such samples to $\mathcal{B}^+(\mathbf{x})$ if their rewards are positive.

The more general variant of this objective function for non-binary reward functions has been called Reward Augmented Maximum Likelihood (RAML) (Norouzi et al., 2016), and one can think of an iterative version of RAML as well,

$$O_{\text{RAML}} = \sum_{\mathbf{x} \in \mathcal{D}} \frac{1}{Z(\mathbf{x})} \sum_{\mathbf{a} \in \mathcal{A}(\mathbf{x})} \exp(R(\mathbf{a})/\tau) \log \pi(\mathbf{a} | \mathbf{x}), \quad (3)$$

where $Z(\mathbf{x}) \equiv \sum_{\mathbf{a} \in \mathcal{A}(\mathbf{x})} \exp(R(\mathbf{a})/\tau)$.

► **MML (Maximum Marginal Likelihood)** (Guu et al., 2017; Berant et al., 2013) is an alternative approach to parameter estimation related to the EM algorithm, which is only concerned with the *marginal* probability of successful trajectories and not with the way probability mass is distributed across $\mathcal{A}^+(\mathbf{x})$,

$$O_{\text{MML}} = \sum_{\mathbf{x} \in \mathcal{D}} \log \sum_{\mathbf{a}^+ \in \mathcal{A}^+(\mathbf{x})} \pi(\mathbf{a}^+ | \mathbf{x}). \quad (4)$$

Again, $\mathcal{A}^+(\mathbf{x})$ is approximated using $\mathcal{B}^+(\mathbf{x})$ iteratively. Dayan & Hinton (1997) also used a variant of this objective function for Reinforcement Learning.

► **RER (Regularized Expected Return)** is the common objective function used in RL

$$O_{\text{RER}} = \sum_{\mathbf{x} \in \mathcal{D}} \tau \mathcal{H}(\pi(\cdot | \mathbf{x})) + \sum_{\mathbf{a} \in \mathcal{A}(\mathbf{x})} R(\mathbf{a})\pi(\mathbf{a} | \mathbf{x}), \quad (5)$$

where $\tau \geq 0$ and \mathcal{H} denotes Shannon Entropy. Entropy regularization often helps with stability of policy optimization leading to better solutions (Williams & Peng, 1991).

Liang et al. (2018) make the important observation that the expected return objective can be expressed as a sum of two terms: a summation over the trajectories inside a context specific buffer $\mathcal{B}^+(\mathbf{x})$ and a separate expectation over the trajectories outside of the buffer:

$$O_{\text{ER}} = \sum_{\mathbf{x} \in \mathcal{D}} \underbrace{\sum_{\mathbf{a} \in \mathcal{B}^+(\mathbf{x})} R(\mathbf{a})\pi(\mathbf{a} | \mathbf{x})}_{\text{enumeration inside buffer}} + \underbrace{\sum_{\mathbf{a} \notin \mathcal{B}^+(\mathbf{x})} R(\mathbf{a})\pi(\mathbf{a} | \mathbf{x})}_{\text{expectation outside buffer}}. \quad (6)$$

Based on this observation, they propose to use enumeration to estimate the gradient of the first term on the RHS of (6) and use Monte Carlo sampling followed by rejection sampling to estimate the gradient of the second term on the RHS of (6) using the REINFORCE (Williams, 1992) estimator. This procedure is called Memory Augmented Policy Optimization (MAPO) and in its ideal form provides a low variance unbiased estimate of the gradient of (6) for deterministic $R(\cdot)$. Note that one can also incorporate entropy into MAPO (Liang et al., 2018) as the contribution of entropy can be absorbed into the reward function as $R'(\mathbf{a}) = R(\mathbf{a}) - \tau \log \pi(\mathbf{a} | \mathbf{x})$. We make heavy use of the MAPO estimator and build our code¹ on top of the open source code of MAPO generously provided by the authors.

3. Mode Covering Exploration (MAPOX)

When it comes to using O_{IML} (2), O_{MML} (4), and O_{RER} (5) for learning from sparse feedback (*e.g.*, program synthesis) and comparing the empirical behavior of these different objective functions, there seems to be some disagreement among previous work. Abolafia et al. (2018) suggest that IML outperforms RER on their program synthesis problem, whereas Liang et al. (2017) assert that RER significantly outperforms IML on their weakly supervised semantic parsing problem. Here, we present some arguments and empirical evidence that justify the results of both of these papers, which helps us develop a novel combination of IML and RER that improves the results of (Liang et al., 2017).

Inspired by (Norouzi et al., 2016; Nachum et al., 2017), we first note that the IML objective per context \mathbf{x} can be expressed in terms of a KL divergence between an optimal policy π^* and the parametric policy π , *i.e.*, $\text{KL}(\pi^* \| \pi)$, whereas the RER objective per context \mathbf{x} can be expressed in terms of the same KL divergence, but *reversed*, *i.e.*, $\text{KL}(\pi \| \pi^*)$. It is well understood that $\text{KL}(\pi^* \| \pi)$ promotes *mode covering* behavior, whereas $\text{KL}(\pi \| \pi^*)$ promotes mode seeking behavior. In other words, $\text{KL}(\pi^* \| \pi)$ encourages all of the trajectories in \mathcal{A}^+ to have an equal probability, whereas RER, at least when $\tau = 0$, is only concerned with the *marginal* probability of successful trajectories and not with the way probability mass is distributed across $\mathcal{A}^+(\mathbf{x})$ (very much like MML). Notably, Guu et al. (2017) proposed an objective combining RER and MML to learn a robust policy that can discount spurious trajectories.

Our key intuition is that for the purpose of exploration and collecting a diverse set of successful trajectories (regardless of whether they are spurious or not) robust behavior of RER and MML should be disadvantageous. On the other hand,

¹Our open-source implementation can be found at https://github.com/google-research/google-research/tree/master/meta_reward_learning.




Figure 4. Fraction of total contexts for which at least k programs ($1 \leq k \leq 100$) are discovered during the entire course of training using the IML and MAPO (*i.e.*, RER) objectives on weakly-supervised semantic parsing datasets (a) WIKITABLEQUESTIONS and (b) WIKISQL.

the mode covering behavior of IML should encourage more exploratory behavior. We conduct some experiments to evaluate this intuition, and in Figure 3, we plot the fraction of contexts for which $|\mathcal{B}^+(\mathbf{x})| \geq k$, *i.e.*, the size of the buffer $\mathcal{B}^+(\mathbf{x})$ after convergence is larger than k as a function of k on two semantic parsing datasets.

Interestingly, we find that IML generally discovers many more successful trajectories than MAPO. For example, the fraction of context for which no plausible trajectory is found ($k = 10^0$ on the plots) is reduced by a few percent on both datasets, and for all other values of $k > 1$, the curve corresponding to IML is above the curve corresponding to MAPO, especially on WIKISQL. Examining the details of the experiments in Abolafia et al. (2018), we realize that their program synthesis tasks are primarily about discovering an individual program that is consistent with a few input-output examples. In this context, due to the presence of multiple input-output pairs, the issue of underspecified rewards poses a less serious challenge as compared to the issue of exploration. Hence, we believe that the success of IML in that context is consistent with our results in Figure 3.

Based on these findings, we develop a novel combination of IML and MAPO, which we call MAPOX (MAPO eXploratory). The key difference between MAPO and MAPOX is in the way the initial memory buffer of programs is initialized. In addition to using random search to populate an initial buffer of programs as in (Liang et al., 2018), we also use IML to find a large set of diverse trajectories, which are passed to MAPO to select from. MAPOX can be interpreted as a two-stage annealing schedule for temperature in Nachum et al. (2017), where one would use log-likelihood first (∞ temperature) and then switch to expected reward (zero temperature). In our experiments, we observe a notable gain from this form of mode covering exploration combining the benefits of IML and MAPO.

4. Learning Rewards without Demonstration

Designing a reward function that distinguishes between optimal and suboptimal behavior is critical for the use of RL in real-world applications. This problem is particularly challenging when expert demonstrations are not available. When learning from underspecified success-failure rewards, one expects a considerable benefit from a refined reward function that differentiates different successful trajectories. While a policy $\pi(\mathbf{a} | \mathbf{x})$ optimized using a robust objective function such as RER and MML learns its own internal preference between different successful trajectories, such a preference may be overly complex. This complexity arises particularly because the typical policies are autoregressive and only have limited access to trajectory level features. Learning an auxiliary reward function presents an opportunity for using trajectory level features designed by experts to influence a preference among successful trajectories.

For instance, consider the problem of weakly-supervised semantic parsing, *i.e.*, learning a mapping from natural language questions to logical programs only based on the success-failure feedback for each question-program pair. In this problem, distinguishing between purposeful and accidental success without human supervision remains an open problem. We expect that one should be able to discount a fraction of the spurious programs by paying attention to trajectory-level features such as the length of the program and the relationships between the entities in the program and the question. The key technical question is how to combine different trajectory level features to build a useful auxiliary reward function.

For the general category of problems involving learning with underspecified rewards, our intuition is that fitting a policy on spurious trajectories is disadvantageous for the policy’s generalization to unseen contexts. Accordingly, we put forward the following hypothesis: One should be able to learn an auxiliary reward function based on the performance of the policy trained with that reward function on a held-out validation set. In other words, we would like to learn reward functions that help policies generalize better. We propose two specific approaches to implement this high level idea: (1) based on gradient based Meta-Learning (MAML) (Finn et al., 2017) (Algorithm 1) (2) using BayesOpt (Snoek et al., 2012) as a gradient-free black box optimizer (Algorithm 2). Each one of these approaches has its own advantages discussed below, and it was not clear to us before running the experiments whether either of the techniques would work, and if so which would work better.

Notation. $\mathcal{D}_{\text{train}}$ and \mathcal{D}_{val} denote the training and validation datasets respectively. $\mathcal{B}_{\text{train}}^+$ represents the training memory buffer containing successful trajectories (based on underspecified rewards) for contexts in $\mathcal{D}_{\text{train}}$.

Algorithm 1 Meta Reward-Learning (MeRL)

Input: $\mathcal{D}_{\text{train}}, \mathcal{D}_{\text{val}}, \mathcal{B}_{\text{train}}^+, \mathcal{B}_{\text{val}}^+$
for step $t = 1, \dots, T$ **do**
 Sample a mini-batch of contexts $\mathcal{X}_{\text{train}}$ from $\mathcal{D}_{\text{train}}$
 and \mathcal{X}_{val} from \mathcal{D}_{val}
 Generate n_{explore} trajectories using π_{θ} for each context in $\mathcal{X}_{\text{train}}, \mathcal{X}_{\text{val}}$ and save successful trajectories to $\mathcal{B}_{\text{train}}^+, \mathcal{B}_{\text{val}}^+$ respectively
 Compute $\theta' = \theta - \alpha \nabla_{\theta} O_{\text{train}}(\pi_{\theta}, R_{\phi})$ using samples from $(\mathcal{B}_{\text{train}}^+, \mathcal{X}_{\text{train}})$
 Compute $\phi' = \phi - \beta \nabla_{\phi} O_{\text{val}}(\pi_{\theta'})$ using samples from $(\mathcal{B}_{\text{val}}^+, \mathcal{X}_{\text{val}})$
 Update $\phi \leftarrow \phi', \theta \leftarrow \theta'$
end for

In this work, we employ a feature-based terminal reward function R_{ϕ} parameterized by the weight vector ϕ . For a given context \mathbf{x} , the auxiliary reward is only non-zero for successful trajectories. Specifically, for a feature vector $\mathbf{f}(\mathbf{a}, \mathbf{x})$ for the context \mathbf{x} and trajectory \mathbf{a} and the underspecified rewards $R(\mathbf{a} | \mathbf{x}, y)$:

$$R_{\phi}(\mathbf{a} | \mathbf{x}, y) = \phi^T \mathbf{f}(\mathbf{a}, \mathbf{x}) R(\mathbf{a} | \mathbf{x}, y). \quad (7)$$

Learning the auxiliary reward parameters determines the relative importance of features, which is hard to tune manually. Refer to the supplementary material for more details about the auxiliary reward features used in this work.

4.1. Meta Reward-Learning (MeRL)

An overview of MeRL is presented in Algorithm 1. At each iteration of MeRL, we simultaneously update the policy parameters θ and the auxiliary reward parameters ϕ . The policy π_{θ} is trained to maximize the training objective O_{train} (8) computed using the training dataset and the auxiliary rewards R_{ϕ} while the auxiliary rewards are optimized to maximize the meta-training objective O_{val} (9) on the validation dataset:

$$\begin{aligned}
 O_{\text{train}}(\pi_{\theta}, R_{\phi}) &= \sum_{\mathbf{x} \in \mathcal{D}_{\text{train}}} \sum_{\mathbf{a} \in \mathcal{B}_{\text{train}}^+(\mathbf{x})} R_{\phi}(\mathbf{a}) \pi_{\theta}(\mathbf{a} | \mathbf{x}) \\
 &+ \sum_{\mathbf{x} \in \mathcal{D}_{\text{train}}} \tau \mathcal{H}(\pi_{\theta}(\cdot | \mathbf{x})), \quad (8)
 \end{aligned}$$

$$O_{\text{val}}(\pi) = \sum_{\mathbf{x} \in \mathcal{D}_{\text{val}}} \sum_{\mathbf{a} \in \mathcal{B}_{\text{val}}^+(\mathbf{x})} R(\mathbf{a}) \pi(\mathbf{a} | \mathbf{x}). \quad (9)$$

The auxiliary rewards R_{ϕ} are not optimized directly to maximize the rewards on the validation set but optimized such that a policy learned by maximizing R_{ϕ} on the training set attains high underspecified rewards $R(\mathbf{a} | \mathbf{x}, y)$ on the validation set. This indirect optimization is robust and less susceptible to spurious sequences on the validation set.

Algorithm 2 Bayesian Optimization Reward-Learning (BoRL)

Input: $\mathcal{D}_{\text{train}}, \mathcal{D}_{\text{val}}, \mathcal{B}_{\text{train}}^+$
for trial $k = 1, \dots, K$ **do**
 Sample a parameter vector ϕ_k for R_{ϕ_k} by optimizing the acquisition function a_M over Bayesian model M i.e. $\phi_k \leftarrow \underset{\phi}{\text{argmax}} a_M(\phi | \mathcal{V}_{1:k-1})$
 Create a memory buffer \mathcal{B}_k^+ containing only the highest ranked trajectories in $\mathcal{B}_{\text{train}}^+$ based on R_{ϕ_k}
for step $t = 1, \dots, T$ **do**
 Sample batch of contexts $\mathcal{X}_{\text{train}}$ from $\mathcal{D}_{\text{train}}$
for context \mathbf{c} in $\mathcal{X}_{\text{train}}$ **do**
 Generate n_{explore} trajectories $\mathcal{S}_{\mathbf{c}}$ using π_{θ}
 Save successful trajectories in $\mathcal{S}_{\mathbf{c}}$ ranked higher than any trajectory in $\mathcal{B}_k^+(\mathbf{c})$ based on R_{ϕ_k}
end for
 Update $\theta \leftarrow \theta - \alpha \nabla_{\theta} O_{\text{train}}(\pi_{\theta})$ using samples from $(\mathcal{B}_k^+, \mathcal{X}_{\text{train}})$
end for
 Evaluate v_k , the accuracy of policy π on \mathcal{D}_{val}
 Augment $\mathcal{V}_{1:k} = \{\mathcal{V}_{1:k-1}, (\phi_k, v_k)\}$ and update the model M
end for

MeRL requires O_{val} to be a differentiable function of ϕ . To tackle this issue, we compute O_{val} using only samples from the buffer $\mathcal{B}_{\text{val}}^+$ containing successful trajectories for contexts in \mathcal{D}_{val} . Since we don't have access to ground-truth programs, we use beam search in non-interactive environments and greedy decoding in interactive environments to generate successful trajectories using policies trained with the underspecified rewards. Note that $\mathcal{B}_{\text{val}}^+$ is also updated during training by collecting new successful samples from the current policy at each step.

The validation objective is computed using the policy obtained after one gradient step update on the training objective and therefore, the auxiliary rewards affect the validation objective via the updated policy parameters θ' as shown in equations (10) and (11):

$$\theta'(\phi) = \theta - \alpha \nabla_{\theta} O_{\text{train}}(\pi_{\theta}, R_{\phi}), \quad (10)$$

$$\nabla_{\phi} O_{\text{val}}(\pi_{\theta'}) = \nabla_{\theta'} O_{\text{val}}(\pi_{\theta'}) \nabla_{\phi} \theta'(\phi). \quad (11)$$

4.2. Bayesian Optimization Reward-Learning (BoRL)

An overview of BoRL is presented in Algorithm 2. At each trial in BoRL, we sample auxiliary reward parameters by maximizing the acquisition function computed using the posterior distribution over the validation objective. After sampling the reward parameters, we optimize the O_{train} objective on the training dataset for a fixed number of iterations. Once the training is finished, we evaluate the policy

on the validation dataset, which is used to update the posterior distribution. BoRL is closely related to the previous work on learning metric-optimized example weights (Zhao et al., 2018) for supervised learning.

BoRL does not require the validation objective O_{val} to be differentiable with respect to the auxiliary reward parameters, therefore we can directly optimize the evaluation metric we care about. For example, in non-interactive environments, the reward parameters are optimized using the beam search accuracy on the validation set \mathcal{D}_{val} . In this work, we use Batched Gaussian Process Bandits (Desautels et al., 2014) employing a Matérn kernel with automatic relevance determination (Rasmussen, 2004) and the expected improvement acquisition function (Moćkus, 1975).

4.3. MeRL vs. BoRL

BoRL offers more flexibility than MeRL since we can optimize any non-differentiable objective on the validation set using BoRL but MeRL can only be used for differentiable objectives. Another advantage of BoRL over MeRL is that it performs global optimization over the reward parameters as compared to the local gradient based optimization in MeRL. Notably, the modular nature of Bayesian optimization and the widespread availability of open source libraries for black box optimization makes BoRL easier to implement than MeRL. However, MeRL is much more computationally efficient than BoRL due to having access to the gradients of the objective to optimize. Additionally, MeRL has the ability to adapt the auxiliary rewards throughout the course of policy optimization while BoRL can only express reward functions that remain fixed during policy optimization.

5. Related Work

The problem we study in this work as well as the proposed approach intersect with many subfields of machine learning and natural language processing discussed separately below.

Reward learning. Reinforcement learning (RL) problems are specified in terms of a reward function over state-action pairs, or a trajectory return function for problems with sparse feedback. A key challenge in applying RL algorithms to real world problems is the limited availability of a rich and reliable reward function. Prior work has proposed to learn the reward function (1) from expert demonstrations using inverse reinforcement learning (Abbeel & Ng, 2004; Ziebart et al., 2008) or adversarial imitation learning (Ho & Ermon, 2016) and (2) from human feedback (Christiano et al., 2017; Leike et al., 2018; Ibarz et al., 2018). Recently, these ideas have been applied to the automatic discovery of goal specifications (Xie et al., 2018; Bahdanau et al., 2019), text generation tasks (Wang et al., 2018; Wu et al., 2017; Bosselut et al., 2018) and the optimization of reward functions (*e.g.*, Gleave

& Habryka (2018); Fu et al. (2019); Shi et al. (2018)) via inverse RL. By contrast, we aim to learn a reward function through meta-learning to enhance underspecified rewards without using any form of trajectory or goal demonstrations. Another relevant work is LIRPG (Zheng et al., 2018), which learns a parametric intrinsic reward function that can be added to the extrinsic reward to improve the performance of policy gradient methods. While the intrinsic reward function in LIRPG is trained to optimize the extrinsic reward, our reward function is trained to optimize the validation set performance through meta-learning, because our main concern is generalization.

Meta-learning. Meta-learning aims to design learning algorithms that can quickly adapt to new tasks or acquire new skills, which has shown recent success in RL (Finn et al., 2017; Duan et al., 2016; Wang et al., 2016; Nichol & Schulman, 2018). There has been a recent surge of interest in the field of meta-reinforcement learning with previous work tackling problems such as automatically acquiring intrinsic motivation (Zheng et al., 2018), discovering exploration strategies (Gupta et al., 2018; Xu et al., 2018b), and adapting the nature of returns in RL (Xu et al., 2018c). It has also been applied to few-shot inverse reinforcement learning (Xu et al., 2018a), online learning for continual adaptation (Nagabandi et al., 2018), and semantic parsing by treating each query as a separate task (Huang et al., 2018). Concurrent work (Zou et al., 2019) also dealt with the problem of learning shaped rewards via meta-learning. Recent work has also applied meta-learning to reweight learning examples (Ren et al., 2018) to enable robust supervised learning with noisy labels, learning dynamic loss functions (Wu et al., 2018) and predicting auxiliary labels (Liu et al., 2019) for improving generalization performance in supervised learning. In a similar spirit, we use meta optimization to learn a reward function by maximizing the generalization accuracy of the agent’s policy. Our hypothesis is that the learned reward function will weight correct trajectories more than the spurious ones leading to improved generalization.

Semantic parsing. Semantic parsing has been a long-standing goal for language understanding (Winograd, 1972; Zelle & Mooney, 1996; Chen & Mooney, 2011). Recently, weakly supervised semantic parsing (Berant et al., 2013; Artzi & Zettlemoyer, 2013) has been proposed to alleviate the burden of providing gold programs or logical forms as annotations. However, learning from weak supervision raises two main challenges (Berant et al., 2013; Pasupat & Liang, 2016a; Guu et al., 2017): (1) how to explore an exponentially large search space to find gold programs; (2) how to learn robustly given spurious programs that accidentally obtain the right answer for the wrong reason. Previous work (Pasupat & Liang, 2016b; Mudrakarta et al., 2018; Krishnamurthy et al., 2017) has shown that efficient exploration of the search space and pruning the spurious programs

by collecting more human annotations has a significant impact on final performance. Some recent work (Berant et al., 2019; Cho et al., 2018) augments weak supervision with other forms supervisions, such as user feedback or intermediate results. Recent RL approaches (Liang et al., 2017; 2018) rely on maximizing expected reward with a memory buffer and performing systematic search space exploration to address the two challenges. This paper takes such an approach a step further, by learning a reward function that can differentiate between spurious and correct programs, in addition to improving the exploration behavior.

Language grounding. Language grounding is another important testbed for language understanding. Recent efforts includes visual question answering (Antol et al., 2015) and instruction following in simulated environments (Hermann et al., 2017; Chevalier-Boisvert et al., 2018). These tasks usually focus on the integration of visual and language components, but the language inputs are usually automatically generated or simplified. In our experiments, we go beyond simplified environments, and also demonstrate significant improvements in real world semantic parsing benchmarks that involve complex language inputs.

6. Experiments

We evaluate our approach on two weakly-supervised semantic parsing benchmarks, WIKITABLEQUESTIONS (Pasupat & Liang, 2015) and WIKISQL (Zhong et al., 2017). Note that we only make use of weak-supervision in WIKISQL and therefore, our methods are not directly comparable to methods trained using strong supervision in the form of (question, program) pairs on WIKISQL. Additionally, we demonstrate the negative effect of under-specified rewards on the generalization ability of an agent in the instruction following task (refer to section 6.1). For all our experiments, we report the mean accuracy and standard deviation based on 5 runs with identical hyperparameters.

6.1. Instruction Following Task

We experiment with a simple instruction following environment in the form of a simple maze of size $N \times N$ with K deadly traps distributed randomly over the maze. A goal located in one of the four corners of the maze (see Figure 2). An agent is provided with a language instruction, which outlines an optimal path that the agent can take to reach the goal without being trapped. The agent receives a reward of 1 if it succeeds in reaching the goal within a certain number of steps, otherwise 0. To increase the difficulty of this task, we reverse the instruction sequence that the agent receives, *i.e.*, the command “Left Up Right” corresponds to the optimal trajectory of actions ($\rightarrow, \uparrow, \leftarrow$).

We use a set of 300 randomly generated environments with

Table 1. Performance of the trained MAPO agent with access to different type of rewards in the instruction following task.

Reward structure	Dev	Test
Underspecified	73.0 (± 3.4)	69.8 (± 2.5)
Underspecified + Auxiliary (BoRL)	75.3 (± 1.6)	72.3 (± 2.2)
Underspecified + Auxiliary (MeRL)	83.0 (± 3.6)	74.5 (± 2.5)
Oracle Reward	95.7 (± 1.3)	92.6 (± 1.0)

(N, K) = (7, 14) with training and validation splits of 80% and 20% respectively. The agent is evaluated on 300 unseen test environments from the same distribution. To mitigate the issues due to exploration, we train the agent using a fixed replay buffer containing the gold trajectory for each environment. For more details, refer to the supplementary material. We compare the following setups for a MAPO agent trained with the same neural architecture in Table 1:

- **Oracle Reward:** This agent is trained using the replay buffer containing only the gold trajectories.
- **Underspecified Reward:** For each environment, we added a fixed number of additional spurious trajectories (trajectories which reach the goal without following the language instruction) to the oracle memory buffer.
- **Underspecified + Auxiliary Reward:** In this case, we use the memory buffer with spurious trajectories similar to the underspecified reward setup, however, we additionally learn an auxiliary reward function using MeRL and BoRL (see Algorithm 1 and 2 respectively).

All the agents trained with different types of reward signal achieve an accuracy of approximately 100% on the training set. However, the generalization performance of Oracle rewards $>$ Underspecified + Auxiliary rewards $>$ Underspecified rewards. Using our Meta Reward-Learning (MeRL) approach, we are able to bridge the gap between Underspecified and Oracle rewards, which confirms our hypothesis that the generalization performance of an agent can serve as a reasonable proxy to reward learning.

6.2. Weakly-Supervised Semantic Parsing

On WIKISQL and WIKITABLEQUESTIONS benchmarks, the task is to generate an SQL-like program given a natural language question such that when the program is executed on a relevant data table, it produces the correct answer. We only have access to weak supervision in the form of question-answer pairs (see Figure 1). The performance of an agent trained to solve this task is measured by the number of correctly answered questions on a held-out test set.

6.2.1. COMPARISON TO STATE-OF-THE-ART RESULTS

We compare the following variants of our technique with the current state-of-the-art in weakly supervised se-

Table 2. Results on WIKITABLEQUESTIONS.

Method	Dev	Test	Improvement on MAPO
MAPO	42.2 (± 0.6)	42.9 (± 0.5)	-
MAPOX	42.6 (± 0.5)	43.3 (± 0.4)	+0.4
BoRL	42.9 (± 0.6)	43.8 (± 0.2)	+0.9
MeRL	43.2 (± 0.5)	44.1 (± 0.2)	+1.2

Table 3. Results on WIKISQL using only weak supervision.

Method	Dev	Test	Improvement on MAPO
MAPO	71.8 (± 0.4)	72.4 (± 0.3)	-
MAPOX	74.5 (± 0.4)	74.2 (± 0.4)	+1.8
BoRL	74.6 (± 0.4)	74.2 (± 0.2)	+1.8
MeRL	74.9 (± 0.1)	74.8 (± 0.2)	+2.4
MAPO (Ens. of 5)	-	74.2	-
MeRL (Ens. of 5)	-	76.9	+2.7

semantic parsing, Memory Augmented Policy Optimization (MAPO) (Liang et al., 2018):

► **MAPOX**: Combining the exploration ability of IML with generalization ability of MAPO, MAPOX runs MAPO starting from a memory buffer $\mathcal{B}_{\text{train}}^+$ containing all the high reward trajectories generated during the training of IML and MAPO using underspecified rewards only.

► **BoRL** (MAPOX + Bayesian Optimization Reward-Learning): As opposed to MAPOX, BoRL optimizes the MAPO objective only on the highest ranking trajectories present in the memory buffer $\mathcal{B}_{\text{train}}^+$ based on a parametric reward function learned using BayesOpt (see Algorithm 2).

► **MeRL** (MAPOX + Meta Reward-Learning): Similar to BoRL, MeRL optimizes the MAPO objective with an auxiliary reward function simultaneously learned with the agent’s policy using meta-learning (see Algorithm 1).

Results. We present the results on weakly-supervised semantic parsing in Table 2 and Table 3. We observe that MAPOX noticeably improves upon MAPO on both datasets by performing better exploration. In addition, MeRL and BoRL both improve upon MAPOX in WIKITABLEQUESTIONS demonstrating that even when a diverse set of candidates from IML are available, one still benefits from our automatic reward learning framework. On WIKISQL, we do not see any gain from BoRL on top of MAPOX, however, MeRL improves upon MAPOX by 0.6% accuracy. Table 3 also shows that even with ensembling 5 models, MeRL significantly outperforms MAPO. Finally, Table 4 compares our approach with previous works on WIKITABLEQUESTIONS. Note that the learned auxiliary reward function matches our intuition, e.g. it prefers programs with more entity matches and shorter length.

Table 4. Comparison to previous approaches for WIKITABLEQUESTIONS

Method	Ensemble Size	Test
Pasupat & Liang (2015)	-	37.1
Neelakantan et al. (2016)	15	37.7
Haug et al. (2018)	15	38.7
Zhang et al. (2017)	-	43.7
MAPO (Liang et al., 2018)	10	46.3
MeRL	10	46.9

6.2.2. UTILITY OF META-OPTIMIZATION

We compare MeRL’s meta-optimization approach to post-hoc “fixing” the policy obtained after training using underspecified rewards. Specifically, we learn a linear re-ranking function which is trained to maximize rewards on the validation set by rescaling the beam search samples on the set. The re-ranker is used to rescore sequences sampled from the learned policy at test time. We implemented two variants of this baseline: 1) Baseline 1 uses the same trajectory-level features as our auxiliary reward function, 2) Baseline 2 uses the policy’s probability in addition to the auxiliary reward features in the ranking function. We use the policies learned using MAPOX for these baselines and evaluate them on WIKITABLEQUESTIONS.

Results. Baseline 1 and 2 resulted in -3.0% drop and +0.2% improvement in test accuracy respectively, as opposed to +0.8% improvement by MeRL over MAPOX. MeRL’s improvement is significant as the results are averaged across 5 trials. These results demonstrate the efficacy of the end-to-end approach of MeRL as compared to the two stage approach of learning a policy followed by reranking to fix it. Additionally, the learned auxiliary rewards for MeRL only have to distinguish between spurious and non-spurious programs while the post-hoc reranker has to differentiate between correct and incorrect programs too.

7. Conclusion & Future Work

In this paper, we identify the problem of learning from *sparse* and *underspecified* rewards. We tackle this problem by employing a mode covering exploration strategy and meta learning an auxiliary terminal reward function without using any expert demonstrations.

As future work, we’d like to extend our approach to learn non-terminal auxiliary rewards as well as replace the linear reward model with more powerful models such as neural networks. Another interesting direction is to improve upon the local optimization behavior in MeRL via random restarts, annealing etc.

Acknowledgments We thank Chelsea Finn, Kelvin Guu and anonymous reviewers for their review of the early draft of the paper and Christian Howard for helpful discussions.

References

- Abbeel, P. and Ng, A. Y. Apprenticeship learning via inverse reinforcement learning. *Proceedings of the twenty-first international conference on Machine learning*, 2004.
- Abolafia, D. A., Norouzi, M., Shen, J., Zhao, R., and Le, Q. V. Neural program synthesis with priority queue training. *arXiv:1801.03526*, 2018.
- Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., and Mané, D. Concrete problems in ai safety. *arXiv:1606.06565*, 2016.
- Antol, S., Agrawal, A., Lu, J., Mitchell, M., Batra, D., Lawrence Zitnick, C., and Parikh, D. Vqa: Visual question answering. *Proceedings of the IEEE international conference on computer vision*, 2015.
- Artzi, Y. and Zettlemoyer, L. Weakly supervised learning of semantic parsers for mapping instructions to actions. *Transactions of the Association of Computational Linguistics*, 2013.
- Bahdanau, D., Hill, F., Leike, J., Hughes, E., Hosseini, A., Kohli, P., and Grefenstette, E. Learning to understand goal specifications by modelling reward. *ICLR*, 2019.
- Berant, J., Chou, A., Frostig, R., and Liang, P. Semantic parsing on freebase from question-answer pairs. *EMNLP*, 2013.
- Berant, J., Deutch, D., Globerson, A., Milo, T., and Wolfson, T. Explaining relational queries to non-experts. *International Conference on Data Engineering (ICDE)*, 2019.
- Bosselut, A., Celikyilmaz, A., He, X., Gao, J., Huang, P.-S., and Choi, Y. Discourse-aware neural rewards for coherent text generation. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, volume 1, pp. 173–184, 2018.
- Chaplot, D. S., Sathyendra, K. M., Pasumarthi, R. K., Rajagopal, D., and Salakhutdinov, R. Gated-attention architectures for task-oriented language grounding. *AAAI*, 2018.
- Chen, D. L. and Mooney, R. J. Learning to interpret natural language navigation instructions from observations. *AAAI*, 2011.
- Chevalier-Boisvert, M., Bahdanau, D., Lahlou, S., Willems, L., Saharia, C., Nguyen, T. H., and Bengio, Y. Babyai: First steps towards grounded language learning with a human in the loop. *arXiv preprint arXiv:1810.08272*, 2018.
- Cho, M., Amplayo, R. K., Hwang, S.-w., and Park, J. Adversarial tableqa: Attention supervision for question answering on tables. *arXiv:1810.08113*, 2018.
- Christiano, P. F., Leike, J., Brown, T., Martic, M., Legg, S., and Amodei, D. Deep reinforcement learning from human preferences. *Advances in Neural Information Processing Systems*, 2017.
- Cobbe, K., Klimov, O., Hesse, C., Kim, T., and Schulman, J. Quantifying generalization in reinforcement learning. *arXiv:1812.02341*, 2018.
- Dayan, P. and Hinton, G. E. Using expectation-maximization for reinforcement learning. *Neural Comput.*, 1997.
- Desautels, T., Krause, A., and Burdick, J. W. Parallelizing exploration-exploitation tradeoffs in gaussian process bandit optimization. *The Journal of Machine Learning Research*, 2014.
- Duan, Y., Schulman, J., Chen, X., Bartlett, P. L., Sutskever, I., and Abbeel, P. RL²: Fast reinforcement learning via slow reinforcement learning. *arXiv:1611.02779*, 2016.
- Finn, C., Abbeel, P., and Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. *International Conference on Machine Learning*, 2017.
- Fu, J., Korattikara, A., Levine, S., and Guadarrama, S. From language to goals: Inverse reinforcement learning for vision-based instruction following. *ICLR*, 2019.
- Gleave, A. and Habryka, O. Multi-task maximum entropy inverse reinforcement learning. *arXiv:1805.08882*, 2018.
- Gupta, A., Mendonca, R., Liu, Y., Abbeel, P., and Levine, S. Meta-reinforcement learning of structured exploration strategies. *arXiv:1802.07245*, 2018.
- Guu, K., Pasupat, P., Liu, E., and Liang, P. From language to programs: Bridging reinforcement learning and maximum marginal likelihood. *ACL*, 2017.
- Haug, T., Ganea, O.-E., and Grnarova, P. Neural multi-step reasoning for question answering on semi-structured tables. *ECIR*, 2018.
- Hermann, K. M., Hill, F., Green, S., Wang, F., Faulkner, R., Soyer, H., Szepesvari, D., Czarnecki, W. M., Jaderberg, M., Teplyashin, D., et al. Grounded language learning in a simulated 3d world. *arXiv:1706.06551*, 2017.
- Ho, J. and Ermon, S. Generative adversarial imitation learning. *Advances in Neural Information Processing Systems*, 2016.

- Huang, P.-S., Wang, C., Singh, R., tau Yih, W., and He, X. Natural language to structured query generation via meta-learning. *CoRR*, abs/1803.02400, 2018.
- Ibarz, B., Leike, J., Pohlen, T., Irving, G., Legg, S., and Amodei, D. Reward learning from human preferences and demonstrations in atari. In *Advances in Neural Information Processing Systems*, pp. 8011–8023, 2018.
- Krishnamurthy, J., Dasigi, P., and Gardner, M. Neural semantic parsing with type constraints for semi-structured tables. *EMNLP*, 2017.
- Leike, J., Krueger, D., Everitt, T., Martic, M., Maini, V., and Legg, S. Scalable agent alignment via reward modeling: a research direction. *arXiv:1811.07871*, 2018.
- Liang, C., Berant, J., Le, Q., Forbus, K. D., and Lao, N. Neural symbolic machines: Learning semantic parsers on freebase with weak supervision. *ACL*, 2017.
- Liang, C., Norouzi, M., Berant, J., Le, Q. V., and Lao, N. Memory augmented policy optimization for program synthesis and semantic parsing. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 31*, pp. 9994–10006. 2018.
- Liu, S., Davison, A. J., and Johns, E. Self-supervised generalisation with meta auxiliary learning. *arXiv preprint arXiv:1901.08933*, 2019.
- Long, R., Pasupat, P., and Liang, P. Simpler context-dependent logical forms via model projections. *ACL*, 2016.
- Maclaurin, D., Duvenaud, D., and Adams, R. Gradient-based hyperparameter optimization through reversible learning. *ICML*, 2015.
- Moćkus, J. On bayesian methods for seeking the extremum. *Optimization Techniques IFIP Technical Conference*, 1975.
- Mudrakarta, P. K., Taly, A., Sundararajan, M., and Dhamdhere, K. It was the training data pruning too! *arXiv:1803.04579*, 2018.
- Nachum, O., Norouzi, M., and Schuurmans, D. Improving policy gradient by exploring under-appreciated rewards. *ICLR*, 2017.
- Nagabandi, A., Finn, C., and Levine, S. Deep online learning via meta-learning: Continual adaptation for model-based rl. *arXiv:1812.07671*, 2018.
- Neelakantan, A., Le, Q. V., Abadi, M., McCallum, A. D., and Amodei, D. Learning a natural language interface with neural programmer. *arXiv:1611.08945*, 2016.
- Nichol, A. and Schulman, J. Reptile: a scalable metalearning algorithm. *arXiv:1803.02999*, 2018.
- Norouzi, M., Bengio, S., Jaitly, N., Schuster, M., Wu, Y., Schuurmans, D., et al. Reward augmented maximum likelihood for neural structured prediction. *NIPS*, 2016.
- Pasupat, P. and Liang, P. Compositional semantic parsing on semi-structured tables. *ACL*, 2015.
- Pasupat, P. and Liang, P. Inferring logical forms from denotations. *ACL*, 2016a.
- Pasupat, P. and Liang, P. Inferring logical forms from denotations. *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2016b.
- Rasmussen, C. E. Gaussian processes in machine learning. *Advanced lectures on machine learning*, 2004.
- Ren, M., Zeng, W., Yang, B., and Urtasun, R. Learning to reweight examples for robust deep learning. *arXiv:1803.09050*, 2018.
- Shi, Z., Chen, X., Qiu, X., and Huang, X. Towards diverse text generation with inverse reinforcement learning. *arXiv preprint arXiv:1804.11258*, 2018.
- Snoek, J., Larochelle, H., and Adams, R. P. Practical bayesian optimization of machine learning algorithms. *NIPS*, 2012.
- Sutton, R. S. and Barto, A. G. *Reinforcement learning: An introduction*. MIT press, 2018.
- Wang, J. X., Kurth-Nelson, Z., Tirumala, D., Soyer, H., Leibo, J. Z., Munos, R., Blundell, C., Kumaran, D., and Botvinick, M. Learning to reinforcement learn. *arXiv:1611.05763*, 2016.
- Wang, X., Chen, W., Wang, Y.-F., and Wang, W. Y. No metrics are perfect: Adversarial reward learning for visual storytelling. *arXiv preprint arXiv:1804.09160*, 2018.
- Williams, R. J. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 1992.
- Williams, R. J. and Peng, J. Function optimization using connectionist reinforcement learning algorithms. *Connection Science*, 1991.
- Winograd, T. Procedures as a representation for data in a computer program for understanding natural language. Technical report, Thesis, 1971.
- Winograd, T. Understanding natural language. *Cognitive psychology*, 1972.

- Wu, L., Zhao, L., Qin, T., Lai, J., and Liu, T.-Y. Sequence prediction with unlabeled data by reward function learning. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pp. 3098–3104. AAAI Press, 2017.
- Wu, L., Tian, F., Xia, Y., Fan, Y., Qin, T., Jian-Huang, L., and Liu, T.-Y. Learning to teach with dynamic loss functions. In *Advances in Neural Information Processing Systems*, pp. 6467–6478, 2018.
- Xie, A., Singh, A., Levine, S., and Finn, C. Few-shot goal inference for visuomotor learning and planning. *CoRL*, 2018.
- Xu, K., Ratner, E., Dragan, A., Levine, S., and Finn, C. Few-shot intent inference via meta-inverse reinforcement learning. *arXiv:1805.12573*, 2018a.
- Xu, T., Liu, Q., Zhao, L., and Peng, J. Learning to explore via meta-policy gradient. *ICML*, 2018b.
- Xu, Z., van Hasselt, H., and Silver, D. Meta-gradient reinforcement learning. *arXiv:1805.09801*, 2018c.
- Zelle, M. and Mooney, R. J. Learning to parse database queries using inductive logic programming. *Association for the Advancement of Artificial Intelligence (AAAI)*, 1996.
- Zhang, C., Vinyals, O., Munos, R., and Bengio, S. A study on overfitting in deep reinforcement learning. *arXiv:1804.06893*, 2018.
- Zhang, Y., Pasupat, P., and Liang, P. Macro grammars and holistic triggering for efficient semantic parsing. *ACL*, 2017.
- Zhao, S., Fard, M. M., and Gupta, M. Metric-optimized example weights. *arXiv preprint arXiv:1805.10582*, 2018.
- Zheng, Z., Oh, J., and Singh, S. On learning intrinsic rewards for policy gradient methods. *Advances in Neural Information Processing Systems*, 2018.
- Zhong, V., Xiong, C., and Socher, R. Seq2sql: Generating structured queries from natural language using reinforcement learning. *arXiv:1709.00103*, 2017.
- Ziebart, B. D., Maas, A. L., Bagnell, J. A., and Dey, A. K. Maximum entropy inverse reinforcement learning. *AAAI*, 2008.
- Zou, H., Ren, T., Yan, D., Su, H., and Zhu, J. Reward shaping via meta-learning. *arXiv preprint arXiv:1901.09330*, 2019.