

Learning from Ambiguous Demonstrations with Self-Explanation Guided Reinforcement Learning

Yantian Zha*, Lin Guan*, Subbarao Kambhampati

Arizona State University
{yantian.zha,guanlin,rao}@asu.edu

Abstract

Our work aims at efficiently leveraging ambiguous demonstrations for the training of a reinforcement learning (RL) agent. An ambiguous demonstration can usually be interpreted in multiple ways, which severely hinders the RL agent from learning stably and efficiently. Since an optimal demonstration may also suffer from being ambiguous, previous works that combine RL and learning from demonstration (RLfD works) may not work well. Inspired by how humans handle such situations, we propose to use self-explanation (an agent generates explanations for itself) to recognize valuable high-level relational features as an interpretation of why a successful trajectory is successful. This way, the agent can leverage the explained important relations as guidance for its RL learning. Our main contribution is to propose the Self-Explanation for RL from Demonstrations (SERLFD) framework, which can overcome the limitations of existing RLfD works. Our experimental results show that an RLfD model can be improved by using our SERLFD framework in terms of training stability and performance. To foster further research in self-explanation-guided robot learning, we have made our demonstrations and code publicly accessible at <https://github.com/YantianZha/SERLFD>. For a deeper understanding of our work, interested readers can refer to our arXiv version at <https://arxiv.org/pdf/2110.05286.pdf>, including an accompanying appendix.

1 Introduction

Reinforcement Learning from human visual demonstrations (RLfD) has gained prominence as an approach that improves the sampling efficiency of reinforcement learning (RL) by using the demonstrations to warm-start the learning process. RLfD offers advantages over pure imitation learning methods, such as overcoming distribution drift and limited adaptivity. As a driving force in the field of robotics, RLfD has shown great promise. However, it still suffers from sampling inefficiency when dealing with highly ambiguous demonstrations.

The challenges arise from the inherent ambiguity and noisy details present in human demonstrations, which can easily distract the learning process. This issue becomes even more vexing in tasks that involve both tacit and explicit

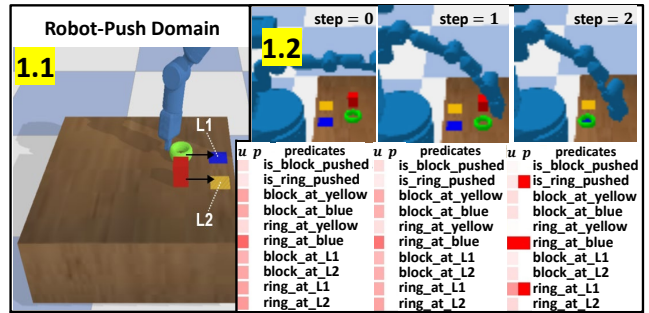


Figure 1: Fig 1.1 shows the Robot-Push domain: There are two target regions which we can index as L1 and L2. L1 and L2 are also randomly assigned with the colors yellow-and-blue or blue-and-yellow in each episode. A human user demonstrates the task of pushing the ring and block into the blue and yellow region respectively. Fig 1.2: a three-step robot execution with grounded predicates (p) and predicted self-explanations (u).

knowledge components. For example, consider a task that requires understanding specific relations between objects (e.g., putting an object in a particular spatial relation with respect to another object) as well as a tacit motor demonstration of how the object is moved. In such cases, the ambiguity in the demonstration can be considerably reduced if the robot takes into account the possible space of explicit knowledge or symbolic relations that the human teacher could be interested in. To put it simply, imagine the difficulty of playing a game of “dumb charades” when you are unsure of whether your friend is miming movie names or place names.

In this paper, we propose a framework inspired by cognitive psychology research, which suggests that humans enhance their learning by employing a self-explanation strategy. We introduce a self-explanation guided reinforcement learning approach where the robot attempts to “self-explain” the potential relations over objects that the human teacher is considering. To support this, we establish a shared vocabulary of possible predicates in the domain, and the robot takes into account this space of predicate groundings when parsing a given demonstration. By doing so, we significantly reduce the ambiguity in the demonstration, effectively addressing the sample complexity issue.

*Equal contribution; correspondence to yantian.zha@asu.edu
Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

To illustrate the challenges of learning from ambiguous demonstrations, we consider an everyday household robot scenario in Fig. 1. In this scenario, a robot needs to push two components into different target regions, denoted as L1 and L2, which are distinguished by their colors, blue and yellow, respectively. Suppose a human teleoperates the robot to push a ring object into the blue L1 region and a block object into the yellow L2 region. This demonstration can be interpreted in two ways: either the robot should push the block into the yellow region and the ring into the blue region, or it should push the block into L2 and the ring into L1.

If humans were in the robot’s position, they would employ a self-explanation strategy to learn from ambiguous demonstrations. Drawing on their background knowledge of important relations in the environment, they would form hypotheses about the relevant object relations and test these hypotheses by taking actions accordingly. If the initial hypothesis fails, they would adjust their explanations and try a different approach. In our work, the task of self-explaining involves identifying the task-relevant relations among all the domain relations at a given state, which aids in disambiguating demonstrations during the learning process.

We assume that the robot possesses human-related background knowledge, which includes relations and symbols, as well as classifiers for extracting relations in the domain. By incorporating such knowledge into the learning process, robots can benefit from the disambiguation capabilities of self-explanation. The key insight behind our approach is the explicit identification of task-relevant object-event relations, which can guide and improve the learning process.

However, learning self-explanations from non-expert humans in the absence of explicit labels presents a significant challenge. Human demonstrators provide non-symbolic and potentially ambiguous demonstrations without self-explanation labels. To tackle this challenge, we propose the Self-Explanation for Reinforcement Learning from Demonstration (SERLFD) framework. We adapt Generative Adversarial Inverse Reinforcement Learning (GAN-IRL) methods, which simultaneously train a generator (RL agent) and a discriminator to distinguish between demonstrations and generated trajectories. In our framework, we enhance the discriminator in GAN-IRL with a non-linear self-explanation predictor (SE-Net). The SE-Net predicts predicate utilities as self-explanations from raw data. Unlike GAN-IRL, our discriminator distinguishes between successful and failed trajectories, similar to the reflective nature of humans. We utilize the SE-Net to assist RL agents in maximizing environment rewards by augmenting states and/or rewards with self-explanations. As a result, our SERLFD algorithm combines the benefits of RLfD and GAN-IRL.

To the best of our knowledge, the proposed SERLFD framework pioneers the direction of learning self-explanations from agents’ past experiences and human demonstrations to guide deep reinforcement learning. This opens up a broad space for future research. Our main contributions include: 1) Introducing the SERLFD framework that contrastively learns to identify task-relevant predicates while performing the task iteratively; 2) Extensive evaluation of the SERLFD framework with multiple

RL agents and various ways of using self-explanations; and 3) Demonstrating the superiority of self-explanations over traditional RLfD and GAN-IRL methods in terms of learning stability and performance, even in challenging continuous control domains.

2 Related Works

Imitation Learning from Ambiguous Demonstrations: In the realm of learning from ambiguous demonstrations, only a handful of studies have explored the field of robot learning from demonstrations (LfD). Breazeal et al. (Breazeal et al. 2006) and Bensch et al. (Bensch and Hellström 2010) approach the problem by modeling demonstration ambiguity as differences in the intentions of humans and robots. Breazeal et al. (Breazeal et al. 2006) propose a Bayesian inference framework that explicitly models human intention and belief, enabling the robot to identify conflicts, seek clarification from humans, and enhance its learning process. Bensch et al. (Bensch and Hellström 2010) view ambiguity as a hypothesis space comprising multiple categories of concepts (e.g., colors, shapes) and address disambiguation by gradually reducing the hypothesis space through concept learning with additional demonstrations. Morales et al. (Morales and De la Rosa 2013) tackle demonstration ambiguity by focusing on differences in demonstrated actions within the same or similar states. They employ clustering algorithms and propose a two-level clustering approach to categorize similar situations and ambiguous actions within each situation. Another approach by Brown et al. (Brown, Niekum, and Petrik 2020) leverages Bayesian optimization to infer reward uncertainty learned through an Inverse Reinforcement Learning (IRL) algorithm, demonstrating the potential for imitation learning from ambiguous demonstrations. However, the work primarily concentrates on the IRL aspect and evaluates its performance in simple domains. A distinct direction is taken by De et al. (de Haan, Jayaraman, and Levine 2019), who address the issue of distribution drift in imitation learning by mitigating causal misidentification resulting from ambiguous demonstrations. Although their work touches on RL, it does not fall under the RLfD paradigm since it does not utilize task rewards. In contrast, RLfD approaches, including our own, as highlighted by Rajeswaran et al. (Rajeswaran et al. 2017), offer several advantages over LfD approaches: (1) they require fewer demonstrations and (2) they naturally handle distribution drift problems by directly incorporating task rewards into the learning process.

Deep RL from Demonstrations: Numerous studies have explored the advantages of using RLfD frameworks. Salimans et al. (Salimans and Chen 2018) utilize states in demonstrations as starting points for training Deep Reinforcement Learning (DRL) models with short-term interactions. This approach focuses exploration in the local region surrounding promising states from the demonstrations. Hester et al. (Hester et al. 2017) and Vecerik et al. (Vecerik et al. 2017) propose initializing the replay buffer with demonstrations to facilitate training a Deep Q-Network (DQN) (Hester et al. 2017) or a Deep Deterministic Policy Gradient (DDPG) network (Vecerik et al. 2017). Pretraining neural

network parameters with an imitation learning objective is advocated by Cruz et al. (Cruz Jr, Du, and Taylor 2017), Rajeswaran et al. (Rajeswaran et al. 2017), and Pfeiffer et al. (Pfeiffer et al. 2018), who directly use demonstrations for initialization. In addition to pretraining, auxiliary imitation learning losses can be constructed, as demonstrated by Vecerik et al. (Vecerik et al. 2017) and Nair et al. (Nair et al. 2018). Gao et al. (Gao et al. 2018) tackle the problem of RL from imperfect demonstrations, focusing on noisy and corrupted demonstrations. Their approach incorporates reward information to normalize the Q-function over actions. However, the issue of RL from ambiguous demonstrations addressed in our work is distinct, as these demonstrations may yield the highest reward and may not necessarily be corrupted but can confuse the learner. However, reinforcement learning from ambiguous demonstrations that we address is an orthogonal issue, as these demonstrations may yield the highest reward and may not necessarily be corrupted but can confuse the learner.

3 Our Approach

3.1 Background

Reinforcement Learning and Reward Shaping: In our work, we consider a finite-horizon and discounted Markov decision process (MDP) model that can be learned by RL methods. We further assume that the reward function $r(s_t, a_t)$ is sparse, i.e., $r(s_t, a_t) = 0$ in most of the states $s \in S$. Training RL agents in sparse rewards environments could be challenging due to the delayed training signals from effective feedback. One extensively adopted way to ameliorate such training is by adding reward shaping, which provides denser training signals so that the agent could obtain valuable feedback much sooner. However, we need to be careful at providing shaped rewards. As demonstrated in (Ng, Harada, and Russell 1999), a poorly-designed reward shaping function may cause the converged optimal policy to shift as against the one under original rewards. (Ng, Harada, and Russell 1999) proves that potential-based reward shaping function, which follows the form in Eq. 1, is the only class of reward shaping function that can guarantee the invariance of optimal policies.

$$\hat{r}(s_t, a_t) = r(s_t, a_t) + \lambda\Phi(s_{t+1}) - \Phi(s_t) \quad (1)$$

where $\hat{r}(s_t, a_t)$ and $r(s_t, a_t)$ denote the shaped and original reward respectively, λ is an adjustment parameter, Φ denotes any real-valued function, and $\lambda\Phi(s_{t+1}) - \Phi(s_t)$ is the reward shaping term. Note that *we set λ to 1 in the rest of paper.*

Inverse Reinforcement Learning (IRL): IRL is about the problem that the reward function in an MDP model is unknown and needs to be discovered from expert demonstrations. Typically in an IRL framework, the unknown reward function is modeled by a parameterized function of certain features (Ziebart et al. 2008). The design of our SERLFD framework is based on GAN-IRL methods like the Guided Cost Learning (GCL (Finn, Levine, and Abbeel 2016)) optimized by Generative-Adversarial Networks (GAN-GCL (Finn et al. 2016)). GCL and GAN-GCL propose to couple IRL and RL together to do continuous control learning with

unknown rewards. Such frameworks allow learning nonlinear rewards that help under complex and unknown dynamics (essentially a model-free IRL).

GAN-GCL integrates IRL and RL by viewing the RL model as a generator and the IRL model as a discriminator that are trained in GAN formulation. The IRL model provides rewards for the RL model. The RL model is trained to gradually shift its sampling distribution to match that of demonstrations. The IRL model is trained to distinguish sampled trajectories from demonstration trajectories, by using the binary cross-entropy loss in Eq. 2. Using such a discriminator to provide rewards for policy learning can solve their Imitation Learning problems.

$$L_{irl}(D_\theta) = \mathbb{E}_{\tau \sim p}[-\log D_\theta(\tau)] + \mathbb{E}_{\tau \sim q}[-\log(1 - D_\theta(\tau))] \quad (2)$$

where D_θ denotes a discriminator model. τ denotes a trajectory. The work Adversarial IRL (Fu, Luo, and Levine 2018) further proposes to replace trajectory τ with state-action pairs, which makes the training more stable:

$$D_\theta(s, a) = \frac{\exp\{\hat{r}_\theta(s, a)\}}{\exp\{\hat{r}_\theta(s, a)\} + q(a|s)} \quad (3)$$

where $\hat{r}_\theta(s, a)$ denotes the estimation of reward feature on a pair of state and action. The $q(a|s)$ is a policy function that can be used to calculate the probability of taking the action a at state s . We refer to this version of GAN-GCL as State-Action-GAN-GCL (**SA-GAN-GCL**).

Note that our SERLFD framework is essentially solving an RLFD problem as in (Vecerik et al. 2017; Hester et al. 2017; Salimans and Chen 2018). RLFD assumes that there exists some task reward signal from the environment. By populating the replay buffer with demonstrations, RLFD methods are able to optimize the environment reward more efficiently than ordinary RL. In contrast, IRL problem assumes that there is no environment reward, so the agent has to simultaneously learn a reward function and a policy. As we will see later, although our self-explainer employs a similar computational framework to some IRL methods (Fu, Luo, and Levine 2018; Finn et al. 2016), we are not solving an IRL problem because the learned self-explanation (predicate utility weights) in SERLFD is used to *augment original state and/or reward*, as a means to mitigate the negative effects of ambiguity in demonstrations.

3.2 The SERLFD Framework

The main idea behind our SERLFD framework is to support RLFD by simultaneously learning to identify which of the domain predicates are relevant to accomplishing a task at each step, i.e. to self-explain. SERLFD interleaves the learning of two sub-tasks: 1) Train a generator (RL agent) by maximizing the accumulative environment rewards while being guided by self-explanations; and 2) Train the Self-Explainer by including it in a Discriminator to distinguish between successful and unsuccessful experiences. Fig. 2 illustrates the main components of our approach.

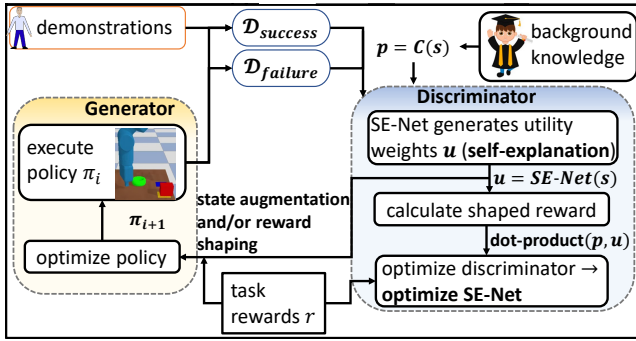


Figure 2: The SERLFD framework that couples learning Self-Explanations Networks (SE-Nets) and an RL agent. Roboticists provide predicates as human-related background domain knowledge to help robots disambiguate non-expert demonstrations for specific tasks. Buffers $\mathcal{D}_{success}$ stores successful experiences which include those from demonstrations. $\mathcal{D}_{failure}$ stores unsuccessful experiences. We train Self-Explanation Network (SE-Net) by inserting it into a Discriminator that distinguishes between successful and failed experiences in a Generative-Adversarial framework.

Predicates: In our approach, we assume that robotics experts provide robots with task-agnostic domain knowledge in the form of human-understandable predicates. This knowledge is aimed at assisting robots in disambiguating non-expert human demonstrations. A predicate is a logical expression used to describe a relationship among objects. The objects involved in the predicate are referred to as its arguments. Take for instance a predicate “pushed(obj)” in the aforementioned robot-pushing domain, “obj” is an argument that can be grounded with “block” or “ring”. When the argument is grounded, the predicate becomes a grounded predicate, e.g. “pushed(block)”. In our work, we use a binary predicate variable to represent if a grounded predicate is satisfied or not. For example, “pushed(block)” can be represented by a binary predicate variable “is_block_pushed”. When it is satisfied, we assign this binary predicate variable with value 1, otherwise -1.

In our framework, we establish a collection of N binary predicate variables, denoted as $P = \langle p_1, p_2, \dots, p_N \rangle$, which corresponds to a domain D capable of supporting various tasks, where p_i is the i -th binary predicate variable in the set P . We also assume to have a set of N classifiers¹ $C = \langle c_{p_1}, c_{p_2}, \dots, c_{p_N} \rangle$, where c_{p_i} is the classifier that can determine whether the i -th predicate is satisfied given a state: $p_i = c_{p_i}(s) = -1$ or 1 .

Self-Explanation: The self-explanation mechanism operates by training a SE-Net to predict a set of predicate utility weights, denoted as u . These weights, estimated without relying on ground-truth self-explanations, serve to determine

¹We note that some robot LfD works rely on marker tracking to simplify the complexity of the vision part of their theories, e.g. (Konidaris et al. 2012; Niekum et al. 2012). We take a cue from this and provide a procedural way to build each classifier. For details, please consult Appendix.B in the arXiv version (see Abstract).

the relevance of predicates to the task at hand in a given state. Mathematically, we express this as: $u = \langle u_1, u_2, \dots, u_N \rangle = SE-Net(s)$. Each predicate utility weight u_i (where $1 \leq i \leq N$) represents a quantitative estimation of the utility of the corresponding predicate variable in elucidating the viability of a decision. The values encompassed within u may accentuate certain predicate variables while diminishing the significance of others.

Self-Explanation Guided Learning of RL agent (Generator): Since the predicates may provide meaningful abstraction of the environment states and may potentially speed up RL training, we append the predicate values P to the environment state (RL state), resulting in the transformation $\langle s \rangle \rightarrow \langle s, p \rangle$. The SE-Net and RL agent mentioned earlier are trained iteratively. The yellow area in Fig. 2 illustrates the training process of the RL agent, leveraging self-explanations u predicted by the SE-Net. Notably, the SE-Net learns a sophisticated non-linear function to identify task-relevant predicates within a continuous state, providing detailed guidance (u) beyond mere numerical reward predictions. The self-explanations u can be used to assist RL agents in two aspects, namely by forming shaping rewards (reward augmentation) or by concatenating them to RL states (state augmentation).

In a similar fashion to solving an RLfD problem, the objective of SERLFD is to leverage demonstrations for training the RL agent to maximize task (environment) rewards. To accomplish this, we utilize self-explanations to construct a reward shaping value that augments the rewards. The reward shaping value, denoted as $h_\theta(s, u)$, is computed as follows:

$$h_\theta(s, u) = \sum_{i=0}^{k-1} u_i \cdot c_{p_i}(s) \quad (4)$$

where $h_\theta(s, u)$ is a scalar value used as the reward shaping value, $c_{p_i}(s)$ represents the application of the i -th predicate detector c_{p_i} to obtain the grounded predicate value from the state s , and k denotes the number of predicates. Then the shaped rewards \hat{r} can be calculated as follows:

$$\hat{r}_\theta(s_t, u_t, a_t, s_{t+1}, u_{t+1}) = r(s_t, a_t) + h_\theta(s_{t+1}, u_{t+1}) - h_\theta(s_t, u_t) \quad (5)$$

where the function \hat{r}_θ models the prediction of shaped rewards, denoted as $\hat{r}(s_t, a_t)$. The value of $h_\theta(s, u)$, computed in Eq. 4, is included in this computation. The term $r(s_t, a_t)$ represents the task rewards in the RL environment. Eq. 5 is formulated based on Eq. 1 and shares similarities with the prediction of shaped reward in work (Fu, Luo, and Levine 2018). However, unlike (Fu, Luo, and Levine 2018), our approach incorporates the task reward r to train the agent to accomplish the task instead of solely imitating.

The reward augmentation way of using self-explanations provides more direct learning signals for RL agents. However, we could also use self-explanations u to augment RL states: $\langle s, p \rangle \rightarrow \langle s, p, u \rangle$. The benefit is that it provides more detailed guidance for RL agents. Depending on whether we augment task rewards with self-explanations, the learning objective of RL agents could be either maximize the accumulative rewards of the environment rewards r , or the augmented rewards \hat{r}_θ computed by using Eq. 5.

Training the Self-Explanation Network (SE-Net) in a Discriminator: To describe the training procedure for generating self-explanations u and training the SE-Net, we follow these steps. First, we sample a batch of experiences and grounded predicate values from the buffers $\mathcal{D}_{success}$ and $\mathcal{D}_{failure}$, which store experiences from successful and unsuccessful trajectories. These samples serve as the input to the Discriminator, represented by the blue area in Fig. 2. Since the SE-Net is integrated into the Discriminator, we can optimize a binary logistic regression loss (King 2008) to train the SE-Net. The specific neural network architecture of the SE-Net can be customized to match the characteristics of the state space. For instance, in experiments where the state space involves features such as object poses, a multilayer perceptron can be employed to implement the SE-Net.

Our Discriminator design aligns with the Generator in terms of utilizing self-explanations. The purpose of our self-explanation u is to enhance the task rewards and facilitate faster learning of tasks by RL agents, rather than solely performing imitation learning from ambiguous demonstrations. Consequently, our Discriminator is trained to differentiate between successful and unsuccessful trajectories to identify task-relevant predicates as self-explanations. As depicted in Fig. 2, we store a sampled trajectory into either $\mathcal{D}_{success}$ or $\mathcal{D}_{failure}$ depending on whether it accomplishes the task or not. The discriminator loss function L_{SE} , responsible for training the SE-Net, is defined in Eq. 6.

$$L_{SE}(s_t, a_t) = \mathbb{E}_{(s,a) \sim \mathcal{D}_{success}} [-\log D(\hat{r}_\theta(s_t, u_t, a_t, s_{t+1}, u_{t+1}))] + \mathbb{E}_{(s,a) \sim \mathcal{D}_{failure}} [-\log(1 - D(\hat{r}_\theta(s_t, u_t, a_t, s_{t+1}, u_{t+1})))] \quad (6)$$

where the discriminator function $D(\cdot)$ is formulated in Eq. 3. The shaped reward $\hat{r}_\theta(s_t, u_t, a_t, s_{t+1}, u_{t+1})$ is computed by using Eqs. 4 and 5. The predicate utility weights (self-explanation) u in the function $\hat{r}(\cdot)$ is predicted by the Self-Explainer Network (SE-Net).

The SERLFD learning is summarized in **Algorithm. 1**.

4 Evaluation

Since self-explanations should play a general role in improving an RL agent that is trained with ambiguous demonstrations, we evaluate our SERLFD framework in multiple domains and use different candidate deep RL models as the RL agent. We design our evaluation to answer the following questions: 1) Can SERLFD outperform RLFD? 2) What are promising ways of using self-explanations to guide an RL agents (i.e., state augmentation, reward shaping, or both)? 3) Do self-explanations play a general role in supporting an RL agent to learn from ambiguous demonstrations or self-explanations are only effective for certain of the RL models? 4) Since our SERLFD combines the benefits of RLFD and GAN-IRL, does our SERLFD outperform a state-of-the-art GAN-IRL for Imitation Learning method? 5) Does SERLFD help in both continuous and discrete domains?

To answer 1, we compare the performance of an RLFD model and the same one supported by our SE-Net. To answer 2, we investigate different ways of using self-explanations

– RLFD+SE: using both reward and state augmentation, RLFD+SE+nu: without using the predicate utility weights u to augment RL states $\langle s, p \rangle$, and RLFD+SE+nrs: without adding reward shaping. To answer 3, we investigate a diverse set of RL agents. To answer 4, we compare SERLFD with the SA-GAN-GCL (Fu, Luo, and Levine 2018) (for Imitation Learning). To answer 5, we evaluate our models in both continuous robotic control domains and a discrete Pacman domain. Finally, in our **supplemental video** and Appendix.D, we show the self-explanations produced by our SE-Net. While the self-explanations are for guiding an RL agent itself, they do reveal some interesting patterns showing robot’s understanding of tasks.

4.1 Experiments in Continuous Domain

Robots and Domains: We use a Fetch Mobile Manipulator (Wise et al. 2016) that has a 7-DoF arm in Py-Bullet simulator (Coumans and Bai 2016). We also fixed its mobile base in experiments. We consider continuous robot control domains that are increasingly complex: Robot-Push-Simple, Robot-Push, and Robot-Remove-and-Push. The Robot-Push domain (Fig 1.1) has two objects, a block and ring, and two target regions that are colored in yellow and blue. By saying a “region”, we mean a square that each side is approximately 0.1m. The colors yellow and blue are exchangeable in different episodes. In the Robot-Push-Simple domain, however, the colors yellow and blue are always fixed. We assume that there is a roboticist who gives the predicates {is_block_pushed, is_ring_pushed, block_at_yellow, block_at_blue, ring_at_yellow, ring_at_blue, block_at_L1, block_at_L2, ring_at_L1, ring_at_L2} to the Robot-Push domain (L1 and L2 are the indices of target regions and are always fixed), and gives the predicates {is_block_pushed, is_ring_pushed, block_at_yellow, block_at_blue, ring_at_yellow, ring_at_blue} to the Robot-Push-Simple domain. Robot-Remove-and-Push (in Fig. 3.1) is a robot domain that goes beyond “Pushing”. A roboticist provides 20 predicates (as listed in Fig. 5) for this domain. Similar to the Robot-Push domain, there are two target regions indexed with L1 and L2 which are fixed and could be assigned with either blue-yellow or yellow-blue colors. In each episode, either a block or a cylinder would show up and both have a removable black cover at the top.

Tasks: While each of the above domains and the corresponding background domain knowledge (predicates) support the learning of multiple tasks like pushing an object to a specific region in a specific color or with a specific index, we uniformly assume that human users provide demonstrations for the task of pushing objects (e.g. a block and a ring) into regions in specific colors (yellow or blue), no matter which region is indexed by L1 or L2. In the more complex Robot-Remove-and-Push domain, if the block or cylinder’s initial poses are on the left side of the table, the robot needs to push them to their target regions *with the black cover removed*.

Settings: We collected demonstrations by using a keyboard to input control commands. Specifically, we gathered 8 trajectories (average length: 17 steps) for Robot-Push-Simple, 15 trajectories (average length: 19 steps) for Robot-Push, and 16 trajectories (average length: 5 steps)

Algorithm 1: The SERLFD Learning Algorithm

INPUT: A dataset of human demonstrations, an environment with reward function r , state space S , action space A , and all hyper-parameters

- 1: Initialize the weights of an RL agent and a Self-Explainer (SE-Net)
- 2: Initialize buffers $\mathcal{D}_{success}$ and $\mathcal{D}_{failure}$ for training the Self-Explainer and D_{RL} for RL from Demonstrations as in the RLfD works (Hester et al. 2017; Vecerik et al. 2017)
- 3: Store expert experiences into $\mathcal{D}_{success}$ and \mathcal{D}_{RL} . Pretrain the RL agent with experiences sampled from \mathcal{D}_{RL}
- 4: Sample K trajectories with a random policy and add them to D_{RL} . Also add successful and unsuccessful trajectories to $\mathcal{D}_{success}$ and $\mathcal{D}_{failure}$ respectively.
- 5: **for** $episode = 1; episode \leq N; episode ++$ **do**
- 6: Sample experiences (including grounded predicate values) from $\mathcal{D}_{success}$ and $\mathcal{D}_{failure}$ **▷ Train Self-Explainer**
- 7: Compute utility weights u and shaped reward prediction $\hat{r}_\theta(s_t, u_t, a_t, s_{t+1}, u_{t+1})$ by using SE-Net, Eqs. 4 and 5
- 8: Update the SE-Net via binary cross entropy loss L_{SE} (Eq. 6) to distinguish successful experiences from unsuccessful experiences
- 9: Sample experiences with grounded predicate values from \mathcal{D}_{RL} **▷ Train RL agent**
- 10: Run SE-Net on sampled states to obtain utility weights u
- 11: Augment input states with grounded predicate values and utility weight values
- 12: Augment rewards with predicted shaped reward by using Eq. 5
- 13: Update RL agent with the augmented experiences
- 14: Use RL agent to sample a new trajectory and add it to D_{RL} . If the trajectory is successful, it would also be added to $\mathcal{D}_{success}$. Otherwise, it would also be added to $\mathcal{D}_{failure}$ **▷ Sample a new trajectory**
- 15: **end for**
- 16: **return** a trained RL agent and SE-Net

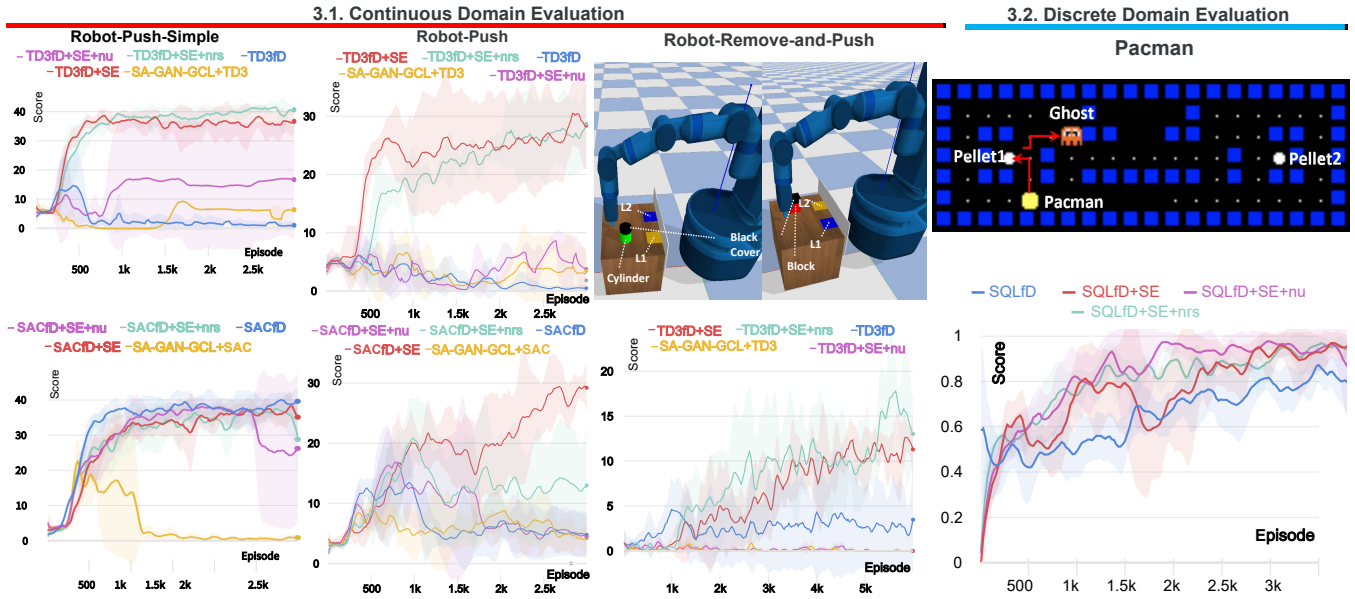


Figure 3: Learning curves of training the baseline RLfD agents (TD3fD/SACfD), TD3fD/SACfD with SE-Nets that uses self-explanation to augment states (TD3fD/SACfD+SE+nrs), augment rewards (TD3fD/SACfD+SE+nu), or both (TD3fD/SACfD+SE), and an Imitation Learning agent built by using RL in the original SA-GAN-GCL framework (Fu, Luo, and Levine 2018); The blue, red, aqua, magenta, and gold curves are the results of baseline TD3fD/SACfD, TD3fD/SACfD+SE, TD3fD/SACfD+SE+nrs, TD3fD/SACfD+SE+nu, and SA-GAN-GCL respectively. For each curve, we run three times of each algorithm and report the mean and standard-variance, which are plotted in bold and lighter color regions respectively. y-axis values are scores that each is measured as an average of over 100 episodes. x-axis values are episodes.

for Robot-Remove-and-Push domains. We use two state-of-the-art RLfD baselines for continuous control tasks: Twin-Delayed DDPG (Fujimoto, Van Hoof, and Meger 2018) from Demonstrations (TD3fD), and Soft-Actor Critic (Haarnoja et al. 2018) from Demonstrations (SACfD). Each

training episode had a maximum of 50 steps. For comprehensive details regarding the state space, action space, termination conditions, and reward function, kindly consult Appendix.C in the arXiv version (refer to Abstract).

Results and Analysis: The results in Fig. 3.1 highlight

which exhibit increased stability and achieve higher scores in a shorter timeframe. Notably, in the Robot-Push-Simple domain with SACfD as the baseline RL agent, the use of self-explanations becomes less crucial due to the lower complexity of the task and the efficacy of entropy-driven exploration. Our investigation focuses on determining the optimal approach for RL agents to leverage self-explanations. The findings emphasize the significance of augmenting states with self-explanations, as this approach provides comprehensive information that enables RL agents to effectively utilize self-explanation details. It is important to note that TD3fD/SACfD+SE, TD3fD/SACfD+SE+**nu**, and TD3fD/SACfD+SE+**nrs** employ the same SE-Nets, trained in a similar manner, with the key distinction lying in how RL agents utilize self-explanations.

Visualizing and Analyzing Self-Explanations: We visually inspect the self-explanations generated by our SE-Nets on RL states. Figures 4 and 5 depict the self-explanation predictions in the Robot-Push and Robot-Remove-and-Push domains, respectively. For Robot-Push-Simple and Pacman domains, please refer to Figures 8 and 11 in Appendix.D in the arXiv version (see Abstract).

To visualize the self-explanations over time, at each step, we show the original input frame with predicates, their groundings (p column), and predicted utility weights (u column). Each cell in the p columns is either red or white – meaning whether a predicate is satisfied or not; Each cell in the u columns has a color ranging from white to red – meaning the increasing relevance between the predicate and a successful decision-making that the robot hypothesizes. Each utility weight is normalized by dividing over the sum of all utility weights across a trajectory.

The self-explanations generated by TD3fD+SE and SACfD+SE, alongside input frames and predicate groundings, reveal consistent patterns – They assign higher weights to predicates associated with colors rather than predicates related to specific locations. In Fig. 4, when the block is pushed to the yellow region (note the color exchange of the target regions), the robot demonstrates certainty that “block_at_yellow” is more relevant than “block_at_L1/L2” at step 2. Interestingly, after step 2, since there is only one free target region remaining, the importance of either “ring_at_blue” or “ring_at_L1/L2” becomes inconsequential. Thus, at step 5, SACfD+SE assumes that “ring_at_L1” is more important than “ring_at_blue”, while TD3fD+SE continues to emphasize the significance of “ring_at_blue”.

We now show the produced self-explanations in the domain Robot-Remove-and-Push in Fig. 5. In the two examples, we can observe that the robot understands which predicates are more task-relevant – hypothesizing higher utility weights on “is_cube_pushedTo_blue” in the TD3fD+SE case, and “is_cylinder_pushedTo_yellow” in the TD3fD+SE+**nrs** case – which also guides the following behavior. For more visualizations and analysis of self-explanations, including those predicted by trained TD3fD/SACfD+SE+**nu** and TD3fD/SACfD+SE+**nrs** models, please consult our supplemental video².

²<https://youtu.be/w5nGYOdVMiA?si=TBvUq3pOpRbca4w4>

4.2 Experiments in Discrete Domain

To evaluate SE-Nets in discrete domains, we conduct experiments in a Pacman domain, and the results are depicted in Fig. 3.2. In this domain, the Pacman’s objective is to eat ghosts within a fixed time frame after consuming a power pellet. The task is considered completed when the Pacman successfully eats all randomly-wandering ghosts, resulting in a reward of 1 and the end of the episode. Otherwise, if the task is not completed, the reward is 0. Two predicate variables, {ghost_nearby, eat_capsule}, are used to indicate the proximity of a ghost to the Pacman and whether the Pacman has consumed a power pellet. The challenge lies in determining the optimal time to eat the pellet. During the demonstrations, the agent tended to hastily consume the pellet when the ghost happened to be nearby. However, in normal circumstances, it is more advantageous for the agent to wait until the ghost approaches before eating the pellet. For our experiments, we employed Soft Q-Learning (Haarnoja et al. 2017) from Demonstration (SQLfD) as our RL agent. We collected 5 demonstrations, each consisting of an average of 18 steps. In our training, each episode was capped at 2000 steps. The results show that RL agents with SE-Nets outperform others in this domain. For more details, see Appendix.D and Fig.11 in our arXiv paper linked in Abstract.

5 Discussion and Future Work

Our contribution introduces the SERLfD framework, utilizing self-explanations to elevate the efficacy of RLfD learning by simulating human-like introspection. The initial version, accessible on arXiv, pioneers the integration of self-explanation into robot learning using deep neural networks. The SERLfD framework demonstrates substantial advantages through effective handling of ambiguous demonstrations, where only specific underlying factors within each observation pertain to the task. Notably, SERLfD holds potential for practical application given its relevance to real-world scenarios where robots collaborate with non-expert humans. A potential constraint lies in the reliance on optimal object trackers for accurate relation detection, as real-world object tracking and predicate detection can be subject to noise. Future endeavors could investigate the performance of advanced visual models in attaining comparable results and implement our framework on real robotic systems.

Future research can explore a broader range of self-explanation techniques for diverse robot learning challenges. Furthermore, refining self-explanation mechanisms in demanding settings shows promise. Recent strides in using foundation models like LLMs for improved robot learning (see (Liu, Bahety, and Song 2023)) highlight opportunities for their use in self-explanation-guided robot learning, aiding adaptability to new domains. However, persistent challenges include the inherent difficulty for LLMs to self-explain and discover causally important factors (Jin et al. 2023). Additionally, aligning LLMs’ high-dimensional outputs with robot-friendly learning signals remains an ongoing hurdle. Tackling these obstacles is pivotal for the effective incorporation of foundation models into self-explanation-guided robot learning, with potential insights from SERLfD.

Acknowledgements

This research was supported by ONR grants N00014-18-1-2442, N00014-18-1-2840, N00014-19-1-2119 and N00014-23-1-2409, and a JP Morgan AI Faculty Research Grant to Kambhampati.

References

- Bensch, S.; and Hellström, T. 2010. On ambiguity in robot learning from demonstration. In *Intelligent autonomous systems*, 47–56. Citeseer.
- Breazeal, C.; Berlin, M.; Brooks, A.; Gray, J.; and Thomaz, A. L. 2006. Using perspective taking to learn from ambiguous demonstrations. *Robotics and autonomous systems*, 54(5): 385–393.
- Brown, D. S.; Niekum, S.; and Petrik, M. 2020. Bayesian Robust Optimization for Imitation Learning. *arXiv preprint arXiv:2007.12315*.
- Coumans, E.; and Bai, Y. 2016. Pybullet, a python module for physics simulation for games, robotics and machine learning.
- Cruz Jr, G. V.; Du, Y.; and Taylor, M. E. 2017. Pre-training neural networks with human demonstrations for deep reinforcement learning. *arXiv preprint arXiv:1709.04083*.
- de Haan, P.; Jayaraman, D.; and Levine, S. 2019. Causal confusion in imitation learning. In *Advances in Neural Information Processing Systems*, 11698–11709.
- Finn, C.; Christiano, P.; Abbeel, P.; and Levine, S. 2016. A connection between generative adversarial networks, inverse reinforcement learning, and energy-based models. *arXiv preprint arXiv:1611.03852*.
- Finn, C.; Levine, S.; and Abbeel, P. 2016. Guided cost learning: Deep inverse optimal control via policy optimization. In *International conference on machine learning*, 49–58.
- Fu, J.; Luo, K.; and Levine, S. 2018. Learning Robust Rewards with Adversarial Inverse Reinforcement Learning. In *International Conference on Learning Representations*.
- Fujimoto, S.; Van Hoof, H.; and Meger, D. 2018. Addressing function approximation error in actor-critic methods. *arXiv preprint arXiv:1802.09477*.
- Gao, Y.; Xu, H.; Lin, J.; Yu, F.; Levine, S.; and Darrell, T. 2018. Reinforcement learning from imperfect demonstrations. *arXiv preprint arXiv:1802.05313*.
- Haarnoja, T.; Tang, H.; Abbeel, P.; and Levine, S. 2017. Reinforcement learning with deep energy-based policies. In *ICML'17 Proceedings of the 34th International Conference on Machine Learning - Volume 70*, 1352–1361.
- Haarnoja, T.; Zhou, A.; Abbeel, P.; and Levine, S. 2018. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290*.
- Hester, T.; Vecerik, M.; Pietquin, O.; Lanctot, M.; Schaul, T.; Piot, B.; Horgan, D.; Quan, J.; Sendonaris, A.; Dulac-Arnold, G.; et al. 2017. Deep q-learning from demonstrations. *arXiv preprint arXiv:1704.03732*.
- Jin, Z.; Liu, J.; Lyu, Z.; Poff, S.; Sachan, M.; Mihalcea, R.; Diab, M.; and Schölkopf, B. 2023. Can Large Language Models Infer Causation from Correlation? *arXiv preprint arXiv:2306.05836*.
- King, J. E. 2008. Binary logistic regression. *Best practices in quantitative methods*, 358–384.
- Konidaris, G.; Kuindersma, S.; Grupen, R.; and Barto, A. 2012. Robot learning from demonstration by constructing skill trees. *The International Journal of Robotics Research*, 31(3): 360–375.
- Liu, Z.; Bahety, A.; and Song, S. 2023. REFLECT: Summarizing Robot Experiences for Failure Explanation and Correction. In *Conference on Robot Learning*, 3468–3484. PMLR.
- Morales, O. C. F.; and De la Rosa, R. F. 2013. Ambiguity analysis in learning from demonstration applications for mobile robots. In *2013 16th International Conference on Advanced Robotics (ICAR)*, 1–6. IEEE.
- Nair, A.; McGrew, B.; Andrychowicz, M.; Zaremba, W.; and Abbeel, P. 2018. Overcoming exploration in reinforcement learning with demonstrations. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 6292–6299. IEEE.
- Ng, A. Y.; Harada, D.; and Russell, S. 1999. Policy invariance under reward transformations: Theory and application to reward shaping. In *ICML*, volume 99, 278–287.
- Niekum, S.; Osentoski, S.; Konidaris, G.; and Barto, A. G. 2012. Learning and generalization of complex tasks from unstructured demonstrations. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 5239–5246. IEEE.
- Pfeiffer, M.; Shukla, S.; Turchetta, M.; Cadena, C.; Krause, A.; Siegwart, R.; and Nieto, J. 2018. Reinforced imitation: Sample efficient deep reinforcement learning for mapless navigation by leveraging prior demonstrations. *IEEE Robotics and Automation Letters*, 3(4): 4423–4430.
- Rajeswaran, A.; Kumar, V.; Gupta, A.; Vezzani, G.; Schulman, J.; Todorov, E.; and Levine, S. 2017. Learning complex dexterous manipulation with deep reinforcement learning and demonstrations. *arXiv preprint arXiv:1709.10087*.
- Salimans, T.; and Chen, R. 2018. Learning Montezuma’s Revenge from a Single Demonstration. *arXiv preprint arXiv:1812.03381*.
- Vecerik, M.; Hester, T.; Scholz, J.; Wang, F.; Pietquin, O.; Piot, B.; Heess, N.; Rothörl, T.; Lampe, T.; and Riedmiller, M. 2017. Leveraging demonstrations for deep reinforcement learning on robotics problems with sparse rewards. *arXiv preprint arXiv:1707.08817*.
- Wise, M.; Ferguson, M.; King, D.; Diehr, E.; and Dymesich, D. 2016. Fetch and freight: Standard platforms for service robot applications. In *Workshop on autonomous mobile service robots*.
- Ziebart, B. D.; Maas, A. L.; Bagnell, J. A.; and Dey, A. K. 2008. Maximum entropy inverse reinforcement learning. In *Aaai*, volume 8, 1433–1438. Chicago, IL, USA.