

Learning on the Job: Self-Rewarding Offline-to-Online Finetuning for Industrial Insertion of Novel Connectors from Vision

Ashvin Nair^{*1}, Brian Zhu^{*12}, Gokul Narayanan², Eugen Solowjow², Sergey Levine¹

Abstract—Learning-based methods in robotics hold the promise of generalization, but what can be done if a learned policy does not generalize to a new situation? In principle, if an agent can at least evaluate its own success (i.e., with a reward classifier that generalizes well even when the policy does not), it could actively practice the task and finetune the policy in this situation. We study this problem in the setting of industrial insertion tasks, such as inserting connectors in sockets and setting screws. Existing algorithms rely on precise localization of the connector or socket and carefully managed physical setups, such as assembly lines, to succeed at the task. But in unstructured environments such as homes or even some industrial settings, robots cannot rely on precise localization and may be tasked with previously unseen connectors. Offline reinforcement learning on a variety of connector insertion tasks is a potential solution, but what if the robot is tasked with inserting previously unseen connector? In such a scenario, we will still need methods that can robustly solve such tasks with online practice. One of the main observations we make in this work is that, with a suitable representation learning and domain generalization approach, it can be significantly easier for the reward function to generalize to a new but structurally similar task (e.g., inserting a new type of connector) than for the policy. This means that a learned reward function can be used to facilitate the finetuning of the robot’s policy in situations where the policy fails to generalize in zero shot, but the reward function generalizes successfully. We show that such an approach can be instantiated in the real world, pretrained on 50 different connectors, and successfully finetuned to new connectors via the learned reward function. Videos and visualizations can be viewed at sites.google.com/view/learningonthejob

I. INTRODUCTION

Generalizable policies require broad and diverse datasets, but for realistic applications, learning policies that can always generalize in zero shot to new objects and environments is often infeasible – indeed, even humans do not exhibit such universal generalization capabilities. Instead, when faced with a task that we don’t precisely know how to do, we can quickly learn the task by leveraging our prior knowledge and a little bit of practice. Reinforcement learning (RL) provides us with a way to implement this kind of *learning on the job*, using online finetuning in the new domain or task, and potentially even extending it into a lifelong learning system where the robot improves its generalization capacity continually with each new task it masters.

However, instantiating this concept in a practical robotics setting requires overcoming a number of obstacles. The robot must be able to combine large amounts of diverse offline data with small amounts of targeted online experience, and do so

in a way that doesn’t require revisiting previously learned tasks or domains, which means that we need an offline RL algorithm that supports online finetuning. Perhaps more importantly, the entire finetuning process must be supported by the robot’s own sensors, without privileged information or environment instrumentation, so as to retain the benefits of autonomous learning. In particular, this means that when adapting to a new task, the robot must be able to evaluate on its own whether it is making progress on the task, using a learned reward function.

We study this problem in the setting of learning a policy from vision for performing industrial insertion tasks. This family of assembly tasks, including plugging in connectors into sockets, keys into locks, screwdrivers into screw intrusions, setting screws, and so on, are found in many stages of manufacturing. When automated in factories today, these tasks are done by robots with specialized control algorithms that rely on precise localization of the socket location.

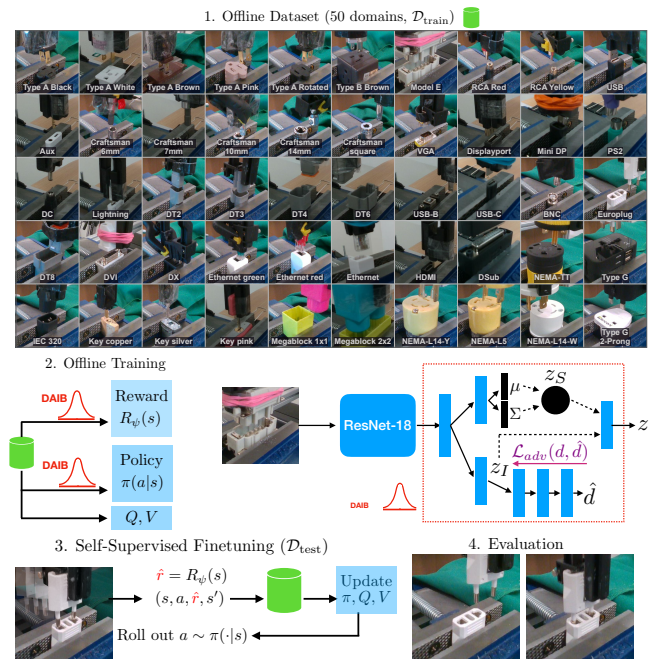


Fig. 1: We describe how a robot can learn to insert an unseen connector from prior experience under realistic conditions by evaluating its own reward. (1) We first collect a diverse dataset with 50 connectors. There is significant variation in the shape of the connectors, sockets, and background. (2) We learn a reward function, policy, and value functions using offline RL. We propose a domain adversarial information bottleneck (DAIB) in order to generalize to new domains. A domain invariance loss is applied on part of the latent representation z_I , while a domain specific latent variable z_S is constrained by an information bottleneck. (3) We finetune online with self-generated rewards to master a new domain. (4) We evaluate π in $\mathcal{D}_{\text{test}}$.

^{*} First two authors contributed equally, ¹ University of California, Berkeley. ² Siemens Corporation. Correspondence: anair17@berkeley.edu

For robots to perform these insertion tasks in industrial and warehouse settings with less human supervision, or in unstructured environments such as homes, they must rely on highly accurate state information of the external world (e.g., socket position and in-hand pose estimation). But such state estimation, using either machine learning or computer vision approaches, is brittle on unseen connectors. To solve the general problem of inserting a novel connector, one promising solution is to generalize previously collected experience of connector insertion to learn a policy to insert connectors from vision. Among these tasks, there is enough variability to require generalization and adaptation, but also enough internal structural regularity that we expect transfer between connectors. Offline RL alone enables robots to generalize to connectors very similar to those from its training dataset, but we will also expect robots to be able to perform tasks in new domains, perhaps after some practice. How can a robot insert test connectors from vision in this setting, utilizing offline RL to enable active online finetuning on a new connector?

The key insight is that we need to (1) adapt to new tasks quickly with online finetuning if the zero-shot solution is not sufficient and (2) generalize to new domains by finding common structure between domains while preserving important domain-specific information. Ideally, a policy trained offline can generalize from vision to new tasks. But if it does not, we can still finetune in a new domain with minimal supervision as long as we have a reward function that generalizes instead. For training policies and reward functions that generalize to test domains, we propose a split representation that combines domain adversarial neural networks [1] for domain invariance and a variational information bottleneck [2] for controlling the flow of domain-specific information. This representation, which we call domain adversarial information bottleneck (DAIB) is used first for learning a robust reward function to detect successful insertions for an unseen connector. Next, we modify implicit Q-learning (IQL), an offline RL algorithm amenable to online finetuning, to use DAIB. During online finetuning, DAIB can be used in combination with online RL to enable fast learning of novel connectors.

We present two main contributions. We demonstrate a system for finetuning under realistic real-world constraints with minimal human supervision, and applied it to insert connectors robustly from vision without the need of accurate socket localization, both for observations and rewards. To accomplish this, we propose a novel representation learning method that allows better generalization of policies and reward functions to unseen domains. We outperform regression-based baselines on the same dataset that combine localizing the socket with hand-designed control policies, as well as prior RL methods. We show that new tasks can be finetuned within 200 trials (about 50 minutes of real-world interaction), given our dataset of 50 connectors across 2 robots. This system allows us to finetune IQL to a test connector, increasing performance significantly over the offline performance. Project videos and our dataset of robotic insertion will be made public at sites.google.com/view/learningonthejob

II. RELATED WORK

Reinforcement learning has been applied to a variety of robotics tasks [3]–[11]. To utilize offline datasets with diverse data in robotics, algorithms developed for offline RL [12]–[15] have been studied in the robotics setting [16]–[20]. A subset of offline RL algorithms are amenable to finetuning [14], [21]–[24]. Our work builds on the direction of offline pretraining followed by online finetuning in robotics. But beyond this line of work, we focus on finetuning from visual input in realistic settings with multiple domains and without ground truth reward functions for the new task.

In this respect, our work is closest to prior work on self-supervised RL that does not assume an external reward function and instead learns it from data. One class of self-supervised RL methods uses goal-conditioned RL with self-supervised rewards [25]–[35]. While general, this class of methods is a poor fit for industrial insertion, as high precision is required in both the policy and in evaluating rewards. Instead, we train a domain generalizing reward classifier from prior data. Prior methods have used learned rewards [36] and classifier rewards have been proposed as a scalable solution for robotics tasks previously [37], [38]. However, learned rewards have not been shown to be useful for finetuning in novel real-world robotic domains previously. Because we focus on applying offline RL and finetuning from vision in the industrial insertion setting, domain generalization of the reward function is vital for our method to work in practice.

Many aspects of robotic insertion, or peg-in-hole assembly, has been studied in prior work [39]–[44], often utilizing geometry and dynamic analysis, force control, tactile sensing, and search, but these methods can be brittle to state estimation errors. Learning-based methods, including RL, have also been applied, usually for a single connector from ground-truth state information [45]–[47]. In these cases, the RL algorithm must learn to navigate the specific dynamics of the single connector, but does not generalize across connectors. More recent work has considered using meta-learning to generalize and improve few-shot between domains [48]. Zhao et al. use offline RL and finetuning combined with meta-learning to adapt to a new connector [49]. This work assumes a known position of the socket and consistent grasping of the connector, and is robust to a small amount ($\pm 1\text{mm}$) of noise. With known socket position and small error, the learning algorithm can learn a structured noise or exploration strategy that can overcome these errors. In contrast, we initialize connectors within $\pm 20\text{mm}$ of the socket ($20\times$ the variance), which requires the robot to rely on visual feedback since blind exploration will rarely succeed.

Closest to our work is prior work that also uses pixel input for robotic insertion. Luo et al. incorporate vision alongside proprioception, using a VAE to embed pixel input [50]. InsertionNet uses a vision system to localize the object and socket, operating on a "residual policy" which is learned from state by supervised learning [51]. InsertionNet 2.0 incorporates contrastive representation learning to improve performance [52]. These prior works collect data on a single

connector and then show robust insertion of that connector. In contrast, our work focuses on what can be done to leverage prior experience for a novel connector without having access to localization of the socket for supervision. Our work also demonstrates robustness to larger variation in initial connector pose, up to 20mm error, than prior work. For visual generalization to a test connector from our offline dataset, a suitable representation learning method is vital.

Many prior methods have explored representation methods for improving the sample efficiency of RL algorithms, including reconstructive objectives, bisimulation, contrastive methods, latent space prediction, forward model prediction, and other mutual information objectives [53]–[70]. In this work, the representation learning challenge is to be able to generalize to new domains from prior domains for RL finetuning to function. Thus, most closely related to our work is domain generalization and domain adaptation. Domain adaptation methods generalize from source domains to a target domain, usually by matching the distribution of features between domains via matching statistics or using an adversarial loss [1], [71]–[74] and has been applied successfully in the sim-to-real setting [75], [76]. Successfully matching distributions makes features indistinguishable between domains. However, in many settings including in our insertion setting, domain-specific information is also important and full domain invariance is not desired. We propose a domain generalization method that can trade off between domain-invariant and domain-specific information.

III. PROBLEM STATEMENT

In RL, we consider a MDP \mathcal{M} with states $s_t \in \mathcal{S}$, actions $a_t \in \mathcal{A}$, dynamics $p(s_{t+1}|s_t, a_t)$, and reward r_t . The RL agent learns a policy $\pi(a_t|s_t)$ to maximize return $\mathcal{R}_t = \sum_{i=t}^T \gamma^{i-t} r_i$, where horizon T may be infinite, and γ is the discount factor. Many algorithms have been developed for this purpose [77]–[79]. For real-world robot learning, sample efficiency is vital, and algorithms that estimate the value function $V^\pi(s_t) = \mathbb{E}_{p_\pi(\tau)}[\mathcal{R}_t|s_t]$ or action-value $Q^\pi(s_t, a_t) = \mathbb{E}_{p_\pi(\tau)}[\mathcal{R}_t|s_t, a_t] = \mathbb{E}_{p(s_{t+1}|s_t, a_t)}[r(s_t, a_t, r_{t+1}) + V^\pi(s_{t+1})]$ better utilize off-policy data.

In the real world, we often do not want to solve a single static environment but rather learn a policy that can generalize to a variety of new scenarios based on prior experience. To formalize this, we will consider a set of MDPs $\mathcal{M}_d \sim p(\mathcal{M})$, indexed by the domain index d , with shared observation and action spaces but potentially different dynamics and reward [80]. We will consider domains $\{0, 1, \dots, n_d\}$ to be training domains and d_{test} be a test domain. The problem is to maximize returns in a test domain $\mathcal{D}_{\text{test}}$ using the least amount of test data possible, given experience from the training domains.

From the training domains, we assume access to a prior dataset of n_d domains drawn from a distribution of domains $p(\mathcal{D})$, with each domain $\mathcal{D}_d = \{\tau_i^d\}_{i=1}^{n_d}$ containing trajectories of transitions $\tau_i^d = \{(s_t, a_t, r_t, s_{t+1})^d\}$. These trajectories can be of arbitrary quality and used within an

offline RL framework to train offline RL algorithms. The challenge is then to utilize the offline data effectively in order to maximize returns through online interaction in a test domain $\mathcal{D}_{\text{test}}$. In the real-world, obtaining a reward on a test domain can be difficult - sometimes as difficult as specifying the optimal policy - as it requires human intervention or supervision. For this reason, we do not assume we receive a reward signal in $\mathcal{D}_{\text{test}}$ during online interaction, and reward information must instead be extracted from the training data, where the data collection can be carefully controlled.

IV. ROBOT SETUP AND DATASET




Fig. 2: Full view of the robot setup, including a Rethink Sawyer robot, an Intel Realsense D435 camera, and a mount for various connectors.

We instantiate real-world learning on the job for a family of industrial insertion tasks, which exhibit real-world challenges that require learning visual policies that generalize from multiple domains. Individual connectors are pictured in Fig 1 and a third-person view of the robot setup for connector insertion is pictured in Fig 2. We use two 7 DoF Sawyer robots for collecting data and running experiments. The connector is grasped by the robot and the socket is mounted to a clamp.

The robot is commanded with an end-effector controller in Cartesian coordinates. Desired joint angles are computed via inverse kinematics, and executed with a joint-space velocity controller with force limits for safety and to prevent grasped connectors from dislodging.

We initially collect a dataset of insertions across 50 connectors, consisting of about 70k trajectories total. Each trajectory is at most 20 steps (terminating early on a successful insertion) and only sparse rewards are provided. During data collection, the robot is manually placed in a successful grasp position and a calibration procedure captures the initial pose, which is used for computing the exact position relative to the socket and for detecting ground truth success. The robot follows a predefined policy, moving towards the centerline of the socket based on the calibrated pose until within 1mm, then moving downwards. Noisy actions are executed with probability 0.2 to visit a diverse set of states and induce recovery behavior in the prior data. Sometimes the connector becomes dislodged from the robot grasp and the expert policy fails, this data is manually discarded. During this phase, the policy benefits from the exact position between the connector and socket; the key challenge is that when finetuning to a new connector, this information is not available.

We want to utilize this data to allow the robot to insert a held-out connector not seen during training. In the online practice phase, a test connector is grasped by the robot and the socket is mounted on a clamp. However, the connector position relative to the socket is not made available to the agent, as calibrating it would require human intervention. The agent can actively practice on the new task, but it must be self-supervised: it must estimate its own reward signal and operate from rich visual input.

V. LEARNING ON THE JOB WITH RL FINETUNING

Our method combines offline pretraining and online finetuning to rapidly adapt to new connectors. We will first provide an overview of the RL system, and then describe the representation learning method that we use to enable generalization for both the reward model and policy.

A. System Overview

We aim to “learn on the job” to finetune to $\mathcal{D}_{\text{test}}$ utilizing data from related training domains $\{\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_{n_d}\}$. We can interact in $\mathcal{D}_{\text{test}}$ but without rewards, so we need to extract artifacts from the training data such that we can coherently practice and master the task in the test domain.

Our solution is illustrated in Fig. 1 and can be considered in four phases. (1) We collect an offline dataset on a diverse set of connectors using the procedure described in Section IV. (2) In the offline training phase, we run offline RL to train a policy, value function, and reward model $R_\psi(s)$ on the offline dataset. These models all take images as input, and the training details are discussed in Section V-B. We discuss techniques to improve generalization to new domains in Section V-D. (3) In the online phase, with minimal human supervision, we start with the pre-trained policy and value function, and finetune the policy and value function with online data. Since we don’t have rewards in the test domain, we use $R_\psi(s)$ to obtain reward predictions. (4) In the evaluation phase, we measure the ability of the finetuned policy to succeed at the task. Success rates are measured manually, as we do not assume ground truth rewards in the test environment.

B. Offline Training Phase

In the offline training phase, we wish to recover artifacts that can be used to maximize performance on the test domain. One option is offline RL, but offline policies may not consistently generalize to a test domain. Since we can interact actively with the test domain, the alternative is to treat the offline phase as preparation for online finetuning.

A subset of offline RL algorithms are amenable to online finetuning. In this work, we use IQL [24], which learns value functions by quantile regression, optimizing:

$$\mathcal{L}_Q = \mathbb{E}_{(s,a,s') \sim \mathcal{D}} [(r(s) + \gamma \bar{V}(s') - Q(s,a))^2] \quad (1)$$

$$\mathcal{L}_V = \mathbb{E}_{(s,a) \sim \mathcal{D}} [L_2^\tau(\bar{Q}(s,a) - V(s))^2], \quad (2)$$

where $L_2^\tau(u) = |\tau - \mathbf{1}(u < 0)|u^2$, and $\bar{\cdot}$ indicates a stop gradient. The policy is extracted from the value function with advantage-weighted regression [81]:

$$\mathcal{L}_\pi = \mathbb{E}_{(s,a) \sim \mathcal{D}} [\log \pi(a|s) \exp(\bar{A}(s,a)/\beta)], \quad (3)$$

where $A(s,a) = Q(s,a) - V(s)$ and β is a temperature parameter. Each function, π , Q , and V is parametrized as a CNN that take images as state input, using independent ResNet-18 backbones [82].

The trained policy could in theory be applied to any test connector. However, when faced with a new connector, the offline-trained policy may not generalize well in zero shot.

In that case, we want to be able to actively finetune the policy to solve the task through additional trials with the new connector. To do so, we will also require a reward function computed from raw image observations, as the robot does not have access to the socket position (and therefore rewards) in the test domain. In the insertion dataset, the offline data contains binary rewards, so we train a reward classifier $R_\psi(s)$ with binary cross entropy: $\mathcal{L}_{\text{rew}}(R; s, r) = -r \log(R_\psi(s)) - (1-r) \log(1 - R_\psi(s))$

C. Self-Supervised Online Finetuning Phase

Next, in the online finetuning phase, we collect trajectories using the pretrained policy with additional exploration noise. In this phase, we do not assume a ground truth reward, as this requires accurate state estimation and localization. Instead, we evaluate the reward from the reward model. For a transition (s, a, s') , we compute the predicted reward $\hat{r} = R_\psi(s')$ and append (s, a, s', \hat{r}) to the replay buffer. We then run the batch gradient descent updates of the policy, Q function, and value function according to equations 1, 2, and 3. To balance the online and offline training data, we separate the data into unique replay buffers and in each mini-batch we sample 25% online data, with no data augmentations, and 75% offline data with data augmentations.

D. Domain Adversarial Information Bottleneck

This system for offline RL and online finetuning can in theory finetune to novel domains, but the degree to which it succeeds depends heavily on the generalization of the offline policy, value functions, and rewards to the test domain. As we show in our ablation experiments, this base system often fails to solve test tasks within a practical time frame on the robot. Therefore, we need a method that generalizes more effectively. In this section, we will describe how to use the domain adversarial information bottleneck to learn an intermediate representation $z = g_\phi(s)$ to improve generalization of the reward model, then explain how we incorporate this bottleneck into RL in the following section. We will overload notation in this section, referring to $R_\psi(g_\phi(s))$ as $R_\psi(z)$.

The representation should allow us to generalize to an unseen domain by capturing the common structure from previously seen domains. So what common structure does our insertion domain have? Coarsely, these tasks are peg insertion variants. At a finer level, the structural similarity of the reward function and expert policy to solve these tasks is clustered into groups of connectors — for instance, NEMA connectors or Deutsch DT series as shown in Figure 1. Within each group, the visual features for aligning the connector with the socket and the insertion depth is shared.

The most commonly proposed approach for generalizing to new domains is domain invariance [1]. In our setting, a fully domain invariant representation could not take advantage of domain-specific information about the similarity between certain connectors: at which depth the reward is obtained, connector-specific visual cues that the policy or reward classifier can take advantage of, and so on. Instead, consider a split $z = (z_I, z_S)$ consisting of a domain invariant

representation z_I and a domain specific representation z_S . Can we design an objective to factorize the two representations, such that it improves generalization to new domains?

For enforcing domain invariance of z_I , we take inspiration from domain adversarial neural networks (DANN) [1], which backpropagates the signal from a domain classifier into the representation. The domain classifier F_θ is trained to minimize negative log likelihood: $\mathcal{L}_{\text{domain}}(F; s, d) = -\log F_\theta(z_I)_d$. For training the reward model, $\mathcal{L}_{\text{adv}} = -\mathcal{L}_{\text{domain}}(F; s, d)$ is added as an auxiliary objective, adversarially optimizing z_I to worsen the domain classifier.

The auxiliary loss imposes a cost for any domain specific information in z_I , but as discussed earlier, allowing for some domain-specific information may be important for the classifier to perform well. However, if we simply concatenate a representation z_S without a domain invariance loss, z_S carries all the information and the features z_I degenerate to be completely uninformative when trained. As we show in our experiments, this leads to reduced performance, similar to ERM in theory and in practice since z_I is not used. We need to somehow limit the information carried by z_S , such that the model maximizes the domain-invariant processing captured by z_I . A natural tool to accomplish this is the variational information bottleneck (VIB) [2]. To learn a split representation, we constrain the information through z_S :

$$\mathcal{L} = \mathcal{L}_{\text{rew}}(z) + \mathcal{L}_{\text{adv}} \text{ s.t. } I(s; z_S) \leq C. \quad (4)$$

Following Alemi et al. [2], we can turn this into an unconstrained problem and compute the evidence lower bound with encoder $q_\phi(z_S|s) = \mathcal{N}(\mu_\phi(s), \sigma_\phi^2(s))$:

$$\mathcal{L}_R = \mathbb{E}_{\epsilon \sim p(\epsilon)} \mathcal{L}_{\text{rew}}(z) + \underbrace{\mathcal{L}_{\text{adv}} + \lambda [KL(q_\phi(z_S|s) || p(z_S))]}_{\mathcal{L}_{\text{DAIB}}(z)} \quad (5)$$

The prior $p(z_S)$ is taken to be an isotropic Gaussian. We use $\lambda = 0.01$ and optimize this objective to learn a reward classifier and representation with Adam [83]. The architecture and losses are pictured in Fig. 1 in the offline training section. Further details are provided on the website.

E. Reinforcement Learning with DAIB

Besides the reward generalizing to a new domain, we also need the policy to generalize at least partially to a new domain to obtain meaningful exploration data to improve from with finetuning. We incorporate a domain adversarial information bottleneck into the policy to enable domain generalization. Let $z^\pi = g_\phi^\pi(s)$ represent the output of the CNN backbone of the policy. We optimize the following policy loss: $\mathcal{L}_\pi + \mathcal{L}_{\text{DAIB}}(z^\pi)$. We find that using DAIB for the policy bottleneck enables consistent performance across connectors, as explained further in Section VI-A.

In the online finetuning phase, when we sample an action from the policy during a rollout and recompute the reward from the reward model, we use a deterministic encoding in the information bottleneck, where the mean of the output distribution is taken as the sampled value. We freeze the convolutional layers during the online finetuning phase. We refer to the overall system as “learning on the job” (LOTJ).

In our experiments, we first evaluate the effectiveness of the proposed system at inserting novel connectors, and compare it with a variety of baselines. Then, we ablate the method and DAIB in particular to isolate the contribution of the representation learning component. Visualizations of saliency maps [85] of trained networks to better understand what features the networks are paying attention to are available on the project website.

A. Finetuning Comparisons

For insertion of a novel connector, we first run offline training to obtain a policy, Q function, and value function. Then, we run online finetuning of RL on the novel connector. We evaluate four connector insertion tasks: Deutsch DT 12-way, Megablock, NEMA 15-5, and Europlug. In practice, we would simply train offline once and apply it to any new connector, but for statistically rigorous results on several test connectors, we train offline on a set of connectors excluding that connector or close variants (e.g., same connector on a different robot), similar to hold-one-out cross-validation. We evaluate in two settings: an easier scenario where the initial location of the connector is centered around the socket with $\pm 10\text{mm}$ noise, and a harder scenario where the initial location of the connector is offset by 10 – 20mm.

In the easier $\pm 10\text{mm}$ setting, we compare against four hand-engineered baselines and two RL baselines. **Localize**: represents a method with learned perception and hand-designed control, where we train a model (on the same offline data) to predict the socket position. The controller regresses onto the state positions in the same offline data at each step and executes a hand-engineered policy to reach that goal location. **Straight down**: moves straight down from the starting position. **Random search**: executes a stochastic search policy as described in [44], moving straight down from the initial starting position and then move to 5 randomly sampled positions while pressing down. **Spiral search**: moves in a spiral while pressing downward, as described in [44]. **SAC**: uses soft actor critic from vision [79] as the underlying RL algorithm, similar to [46], and does not use DAIB. **State-IQL**: trains IQL from state instead of images during offline training, then uses the localization model for state estimation during online training. Each method including ours is executed for 20 time steps per trajectory for evaluation. For the RL methods, we report offline performance after offline training and then online performance after finetuning. The policies were finetuned for 200 trials, or about 50 minutes.

The results are reported in Table I. We see that the performance of most methods is inconsistent, but LOTJ online successfully solves the task to $> 75\%$ success for all four new connectors. In the harder 10-20mm setting, we compare against similar baselines. However, we found that the simpler straight down, random search, and spiral search methods always fail from the initial position, as it is too far from the socket. Instead, we first execute the localization model and move to that goal position, then run the corresponding strategy. The results for the harder insertion scenario are

Connector	Localize	Straight	Random	Spiral	SAC		State-IQL		LOTJ, full	
		Down	Search	Search	(offline)	(online)	(offline)	(online)	(offline)	(online)
DT12-1	3/20	1/20	7/20	4/20	0/20	0/20	5/20	5/20	7/20	19/20
Mega 1x1	3/20	1/20	1/20	15/20	0/20	0/20	0/20	2/20	5/20	16/20
NEMA 15	0/20	1/20	0/20	3/20	0/20	0/20	4/20	1/20	11/20	15/20
Europlug	6/20	4/20	6/20	4/20	0/20	0/20	6/20	4/20	10/20	15/20

TABLE I: Insertion scenario I: the initial position of the gripper is randomly offset by ± 10 mm from above the socket. The best performing method in each row is bolded. (Statistical significance is computed according to Barnard’s exact test [84] for independent binomials at 95% confidence level.) Only our method online after fine-tuning solves all four test tasks the majority of the time.

Connector	Localize	Straight	Random	Spiral	SAC		State-IQL		LOTJ, full	
		Down	Search	Search	(offline)	(online)	(offline)	(online)	(offline)	(online)
DT12-1	0/20	0/20	0/20	0/20	0/20	0/20	0/20	3/20	5/20	18/20
Mega 1x1	0/20	1/20	0/20	17/20	0/20	0/20	0/20	11/20	0/20	20/20
NEMA 15	0/20	0/20	0/20	0/20	0/20	0/20	3/20	1/20	6/20	17/20
Europlug	0/20	0/20	0/20	0/20	0/20	0/20	0/20	5/20	2/20	15/20

TABLE II: Insertion scenario II: the initial position of the gripper is sampled from a box 10-20mm from the socket. This scenario is more difficult as moving straight down almost never solves the task. The baselines in this series of experiments use the localization model initially, then follow the baseline strategy. Our method significantly improves from offline performance on all connectors.

	LOTJ, no DAIB	LOTJ, DAIB reward	LOTJ, full
DT12-1	0/20 \rightarrow 5/20	4/20 \rightarrow 0/20	0/20 \rightarrow 18/20
Mega 1x1	3/20 \rightarrow 0/20	3/20 \rightarrow 20/20	0/20 \rightarrow 20/20
NEMA 15	0/20 \rightarrow 0/20	0/20 \rightarrow 2/20	6/20 \rightarrow 17/20
Europlug	0/20 \rightarrow 15/20	1/20 \rightarrow 12/20	2/20 \rightarrow 15/20

TABLE III: Ablation of using the domain adversarial information bottleneck (DAIB) for the reward and for the policy across all four test connectors. The only consistent setting where finetuning occurs across all four connectors is using the DAIB for both reward and policy. Removing the bottleneck for the policy significantly reduces finetuning performance on two tasks, but still finetunes on the other two. Removing the bottleneck for the reward prevents finetuning except for the Europlug connector.

reported in Table II. In this case, the performance of most methods, including LOTJ offline, on most connectors, is poor. However, LOTJ is able to solve these tasks after 200 trials of online finetuning. Importantly, even when the initial performance is 0/20, as in the DT 12-way connector where the initial policy is poor and does not get close to the socket, the finetuned process quickly improves the policy and makes contact with the socket within a few trials. Further trials are required for finetuning to actually observe successes through stochastic exploration to perfect the policy.

B. Domain Adversarial Information Bottleneck Ablations

Reward classification. To learn a generalizable reward function and to evaluate the proposed domain generalization method, we first need to learn a reward classifier from the offline data. The accuracy of this reward classifier has a significant impact on whether online finetuning with the learned reward successfully solves a task; if the reward is inaccurate, the policy can adversarially learn to visit regions of incorrect high reward. We train the reward classifier on a subset of 25 connectors and evaluate on a set of five held-out connectors. During training and evaluation, the data is rebalanced to be 50% positive and 50% negative samples per connector. In Table IV, we report the mean accuracy of each method averaged across the five held-out connectors.

We compare the following methods. Our method, **DAIB**,

uses a domain adversarial objective in combination with a variational information bottleneck as detailed in section V-D. **DANN** enforces domain invariance using a domain adversarial neural network [1]. **DAIB**, $\lambda = 0$ concatenates z_S to the representation as in DAIB, but without enforcing an information bottleneck. **VIB** enforces a variational information bottleneck only without domain invariance [2]. **ERM** has no auxiliary representation learning objectives.

The results are presented in Table IV. We see that our method, DAIB, achieves 88% accuracy, significantly higher than other methods. DANN, which enforces domain invariance, is the worst performing. This poor performance shows that the domain invariance assumption is too strong and connector-specific information is likely necessary to identify successful insertions. The DAIB, $\lambda = 0$ method achieves

Method	Acc.
DAIB (Ours)	88%
DANN	71%
DAIB, $\lambda = 0$	77%
VIB	79%
ERM	76%

TABLE IV: Comparison of test accuracy on reward classification.

similar accuracy as ERM, as the network can make the domain invariant features degenerate and bypass the domain adversarial objective. The regularizing effect of VIB alone slightly improves performance over ERM. But all methods achieve significantly worse accuracy than DAIB. Next, armed with an accurate reward model, we investigate finetuning connectors using the learned reward model.

RL ablation. Next, we ablate using LOTJ for the reward and policy. The results are presented in Table III on the harder insertion scenario (10-20mm). The results vary somewhat connector to connector, but using DAIB for both the reward and policy consistently finetunes successfully while removing the bottleneck sometimes fails to finetune.

VII. ACKNOWLEDGEMENTS

We thank Ilya Kostrikov, Kuan Fang, and Patrick Yin for useful discussions. This research was partially supported by Siemens and the Office of Naval Research.

REFERENCES

- [1] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky, "Domain-adversarial training of neural networks," *JMLR*, 5 2016.
- [2] A. A. Alemi, I. S. Fischer, J. V. Dillon, and K. P. Murphy, "Deep variational information bottleneck," *ICLR*, 2017.
- [3] J. Peters and S. Schaal, "Reinforcement learning of motor skills with policy gradients," *Neural Networks*, vol. 21, no. 4, pp. 682–697, 2008.
- [4] J. Kober and J. Peter, "Policy search for motor primitives in robotics," in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 97, 2008, pp. 83–117.
- [5] M. P. Deisenroth and C. E. Rasmussen, "PILCO: A model-based and data-efficient approach to policy search," in *International Conference on Machine Learning (ICML)*, 2011, pp. 465–472.
- [6] S. Levine, C. Finn, T. Darrell, and P. Abbeel, "End-to-End Training of Deep Visuomotor Policies," *Journal of Machine Learning Research (JMLR)*, vol. 17, no. 1, pp. 1334–1373, 2016.
- [7] S. Levine, P. Pastor, A. Krizhevsky, and D. Quillen, "Learning Hand-Eye Coordination for Robotic Grasping with Deep Learning and Large-Scale Data Collection," *International Journal of Robotics Research*, 2017.
- [8] H. Zhu, A. Gupta, A. Rajeswaran, S. Levine, and V. Kumar, "Dexterous Manipulation with Deep Reinforcement Learning: Efficient, General, and Low-Cost," in *IEEE International Conference on Robotics and Automation (ICRA)*, vol. 2019-May, Institute of Electrical and Electronics Engineers Inc., oct 2019, pp. 3651–3657.
- [9] A. Giusti, J. J. Guzzi, D. C. Cirean *et al.*, "A Machine Learning Approach to Visual Perception of Forest Trails for Mobile Robots," in *IEEE Robotics and Automation Letters (RAL)*, vol. 1, no. 2, 2015, pp. 2377–3766.
- [10] J. Nakanishi, J. Morimoto, G. Endo, G. Cheng, S. Schaal, and M. Kawato, "Learning from demonstration and adaptation of biped locomotion," in *Robotics and Autonomous Systems*, vol. 47, no. 2-3, 2004, pp. 79–91.
- [11] M. Kalakrishnan, J. Buchli, P. Pastor, and S. Schaal, "Learning Locomotion over Rough Terrain using Terrain Templates," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2009.
- [12] S. Fujimoto, D. Meger, and D. Precup, "Off-Policy Deep Reinforcement Learning without Exploration," in *International Conference on Machine Learning (ICML)*, dec 2019.
- [13] A. Kumar, A. Zhou, G. Tucker, and S. Levine, "Conservative Q-Learning for Offline Reinforcement Learning," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [14] A. Nair, M. Dalal, A. Gupta, and S. Levine, "Accelerating Online Reinforcement Learning with Offline Datasets," jun 2020.
- [15] Y. Wu, G. Tucker, and O. Nachum, "Behavior Regularized Offline Reinforcement Learning," nov 2020.
- [16] A. Singh, A. Yu, J. Yang, J. Zhang, A. Kumar, and S. Levine, "Cog: Connecting new skills to past experience with offline reinforcement learning," 2020.
- [17] A. Singh, H. Liu, G. Zhou, A. Yu, N. Rhinehart, and S. Levine, "Parrot: Data-driven behavioral priors for reinforcement learning," 2020.
- [18] Y. Chebotar, K. Hausman, Y. Lu *et al.*, "Actionable models: Unsupervised offline reinforcement learning of robotic skills," in *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, ser. Proceedings of Machine Learning Research, M. Meila and T. Zhang, Eds., vol. 139. PMLR, 2021, pp. 1518–1528.
- [19] D. Kalashnikov, J. Varley, Y. Chebotar, B. Swanson, R. Jonschkowski, C. Finn, S. Levine, and K. Hausman, "Mt-opt: Continuous multi-task robotic reinforcement learning at scale," 4 2021.
- [20] A. Kumar, A. Singh, S. Tian, C. Finn, and S. Levine, "A workflow for offline model-free robotic reinforcement learning," 9 2021.
- [21] A. Villafior, J. Dolan, and J. Schneider, "Fine-tuning offline reinforcement learning with model-based policy optimization," 2020.
- [22] L. Meng, M. Wen, Y. Yang, C. Le, X. Li, W. Zhang, Y. Wen, H. Zhang, J. Wang, and B. Xu, "Offline pre-trained multi-agent decision transformer: One big sequence model tackles all smac tasks," 2021.
- [23] S. Lee, Y. Seo, K. Lee, P. Abbeel, and J. Shin, "Offline-to-online reinforcement learning via balanced replay and pessimistic q-ensemble," 2021.
- [24] I. Kostrikov, A. Nair, and S. Levine, "Offline reinforcement learning with implicit q-learning," *CoRR*, vol. abs/2110.06169, 2021.
- [25] L. P. Kaelbling, "Learning to achieve goals," in *International Joint Conference on Artificial Intelligence (IJCAI)*, vol. vol.2, 1993, pp. 1094 – 8.
- [26] T. Schaul, D. Horgan, K. Gregor, and D. Silver, "Universal Value Function Approximators," in *International Conference on Machine Learning (ICML)*, 2015, pp. 1312–1320.
- [27] A. Baranes and P.-Y. Oudeyer, "Active Learning of Inverse Models with Intrinsically Motivated Goal Exploration in Robots," *Robotics and Autonomous Systems*, vol. 61, no. 1, pp. 49–73, 2012.
- [28] M. Andrychowicz, F. Wolski, A. Ray, J. Schneider, R. Fong, P. Welinder, B. McGrew, J. Tobin, P. Abbeel, and W. Zaremba, "Hindsight Experience Replay," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [29] A. Nair, V. Pong, M. Dalal, S. Bahl, S. Lin, and S. Levine, "Visual Reinforcement Learning with Imagined Goals," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [30] O. Nachum, S. S. Gu, H. Lee, and S. Levine, "Data-Efficient Hierarchical Reinforcement Learning," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [31] D. Held, X. Geng, C. Florensa, and P. Abbeel, "Automatic Goal Generation for Reinforcement Learning Agents," in *International Conference on Machine Learning (ICML)*, 2018.
- [32] A. Péré, S. Forestier, O. Sigaud, and P.-Y. Oudeyer, "Unsupervised Learning of Goal Spaces for Intrinsically Motivated Goal Exploration," in *International Conference on Learning Representations (ICLR)*, 2018.
- [33] D. Warde-Farley, T. Van De Wiele, T. Kulkarni, C. Ionescu, S. Hansen, and M. Volodymyr, "Unsupervised Control Through Non-Parametric Discriminative Rewards," in *International Conference on Learning Representations (ICLR)*, 2019.
- [34] V. H. V. Pong, M. Dalal, S. Lin, A. Nair, S. Bahl, and S. Levine, "Skew-Fit: State-Covering Self-Supervised Reinforcement Learning," in *International Conference on Machine Learning (ICML)*, 2020.
- [35] A. Khazatsky, A. Nair, D. Jing, and S. Levine, "What Can I Do Here? Learning New Skills by Imagining Visual Affordances," in *International Conference on Robotics and Automation (ICRA)*, jun 2021.
- [36] V. H. Pong, A. Nair, L. Smith, C. Huang, and S. Levine, "Offline meta-reinforcement learning with online self-supervision," *International Conference on Machine Learning (ICML)*, 2022.
- [37] J. Fu, A. Singh, D. Ghosh, L. Yang, and S. Levine, "Variational Inverse Control with Events: A General Framework for Data-Driven Reward Definition," in *Advances in Neural Information Processing Systems (NeurIPS)*, may 2018.
- [38] A. Singh, L. Yang, K. Hartikainen, C. Finn, and S. Levine, "End-to-End Robotic Reinforcement Learning without Reward Engineering," in *Robotics: Science and Systems (RSS)*, apr 2019.
- [39] D. E. Whitney, "Quasi-static assembly of compliantly supported rigid parts," *Journal of Dynamic Systems, Measurement, and Control*, vol. 104, pp. 65–77, 3 1982.
- [40] Y. Xia, Y. Yin, and Z. Chen, "Dynamic analysis for peg-in-hole assembly with contact deformation," *International Journal of Advanced Manufacturing Technology*, vol. 30, pp. 118–128, 8 2006.
- [41] G. Zhang, A. Bell, H. Zhang, J. He, J. Wang, and C. Martinez, "On-robot assembly parameter optimization," *Proceedings of the World Congress on Intelligent Control and Automation (WCICA)*, pp. 547–552, 2008.
- [42] B. Zhang, D. Gravel, G. Zhang, J. Wang, and A. Bell, "Robotic force control assembly parameter optimization for adaptive production," *Proceedings - IEEE International Conference on Robotics and Automation*, pp. 464–469, 2011.
- [43] R. Li, R. Platt, W. Yuan, A. Ten Pas, N. Roscup, M. A. Srinivasan, and E. Adelson, "Localization and Manipulation of Small Parts Using GelSight Tactile Sensing," in *International Conference on Intelligent Robots and Systems (IROS)*, 2014.
- [44] J. A. Marvel, R. Bostelman, and J. Falco, "Multi-Robot Assembly Strategies and Metrics," *ACM computing surveys*, vol. 51, no. 1, p. 14, 2018.
- [45] J. Luo, E. Solowjow, C. Wen, J. A. Ojea, A. M. Agogino, A. Tamar, and P. Abbeel, "Reinforcement learning on variable impedance controller for high-precision robotic assembly," 2019.
- [46] G. Schoettler, A. Nair, J. Luo, S. Bahl, J. Aparicio Ojea, E. Solowjow, and S. Levine, "Deep Reinforcement Learning for Industrial Insertion

- Tasks with Visual Inputs and Natural Rewards,” *arXiv preprint arXiv:1906.05841*, 2019.
- [47] W. Lian, T. Kelch, D. Holz, A. Norton, and S. Schaal, “Benchmarking Off-The-Shelf Solutions to Robotic Assembly Tasks,” in *IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2021.
- [48] G. Schoettler, A. Nair, J. Ojea, S. Levine, and E. Solowjow, “Meta-reinforcement learning for robotic industrial insertion tasks,” 2020.
- [49] T. Z. Zhao, J. Luo, O. Sushkov, R. Pevceviciute, N. Heess, J. Scholz, S. Schaal, and S. Levine, “Offline meta-reinforcement learning for industrial insertion,” 10 2022.
- [50] J. Luo, O. Sushkov, R. Pevceviciute, W. Lian, C. Su, M. Vecerik, N. Ye, S. Schaal, and J. Scholz, “Robust Multi-Modal Policies for Industrial Assembly via Reinforcement Learning and Demonstrations: A Large-Scale Study,” in *Robotics: Science and Systems (RSS)*, mar 2021.
- [51] O. Spector and D. D. Castro, “InsertionNet - A Scalable Solution for Insertion,” *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5509–5516, jul 2021.
- [52] O. Spector, V. Tchuiev, and D. D. Castro, “Insertionnet 2.0: Minimal contact multi-step insertion using multimodal multiview sensory input,” *Proceedings - IEEE International Conference on Robotics and Automation*, pp. 6330–6336, 3 2022.
- [53] S. Lange and M. A. Riedmiller, “Deep learning of visual control policies,” in *European Symposium on Artificial Neural Networks (ESANN)*. Citeseer, 2010.
- [54] S. Lange, M. Riedmiller, A. Voigtlander, and A. Voigtländer, “Autonomous reinforcement learning on raw visual input data in a real world application,” in *International Joint Conference on Neural Networks (IJCNN)*, no. June. IEEE, 2012, pp. 1–8.
- [55] C. Finn, X. Y. Tan, Y. Duan, T. Darrell, S. Levine, and P. Abbeel, “Deep spatial autoencoders for visuomotor learning,” in *IEEE International Conference on Robotics and Automation (ICRA)*, vol. 2016-June. IEEE, 2016, pp. 512–519.
- [56] N. Ferns, P. Panangaden, and D. Precup, “Metrics for finite Markov decision processes,” in *Uncertainty in Artificial Intelligence (UAI)*, 2004, pp. 162–169.
- [57] P. S. Castro, “Scalable methods for computing state similarity in deterministic Markov decision processes,” in *Association for the Advancement of Artificial Intelligence (AAAI)*, 2020.
- [58] A. Zhang, R. T. McAllister, R. Calandra, Y. Gal, and S. Levine, “Learning invariant representations for reinforcement learning without reconstruction,” in *International Conference on Learning Representations*, 2021.
- [59] A. van den Oord DeepMind, Y. Li DeepMind, O. Vinyals DeepMind, A. van den Oord, Y. Li, and O. Vinyals, “Representation Learning with Contrastive Predictive Coding,” 2018.
- [60] M. Laskin, A. Srinivas, and P. Abbeel, “CURL: Contrastive unsupervised representations for reinforcement learning,” in *International Conference on Machine Learning (ICML)*. PMLR, 2020, pp. 5639–5650.
- [61] T. Nguyen, R. Shu, T. Pham, H. Bui, and S. Ermon, “Temporal predictive coding for model-based planning in latent space,” 6 2021.
- [62] M. Schwarzer, A. Anand, R. Goel, R. D. Hjelm, A. Courville, and P. Bachman, “Data-efficient reinforcement learning with self-predictive representations,” 7 2020.
- [63] D. Hafner, T. Lillicrap, I. Fischer, R. Villegas, D. Ha, H. Lee, and J. Davidson, “Learning latent dynamics for planning from pixels,” *36th International Conference on Machine Learning, ICML 2019*, vol. 2019-June, pp. 4528–4547, 11 2018.
- [64] B. Mazouze, R. T. des Combes, T. Doan, P. Bachman, and R. D. Hjelm, “Deep reinforcement and infomax learning,” *Advances in Neural Information Processing Systems*, vol. 2020-December, 6 2020.
- [65] K. Gregor, D. Rezende, and D. W. Deepmind, “Variational Intrinsic Control.”
- [66] D. Hafner, P. A. Ortega, J. Ba, T. Parr, K. Friston, and N. Heess, “Action and Perception as Divergence Minimization,” 2020.
- [67] K. Rakelly, A. Gupta, C. Florensa, and S. Levine, “Which mutual-information representation learning objectives are sufficient for control?” 2021.
- [68] R. Jonschkowski, R. Hafner, J. Scholz, and M. Riedmiller, “Pves: Position-velocity encoders for unsupervised learning of structured state representations,” 5 2017.
- [69] D. Ghosh, A. Gupta, and S. Levine, “Learning actionable representations with goal conditioned policies,” in *International Conference on Learning Representations*, 2019.
- [70] A. Sax, B. Emi, A. Zamir, L. Guibas, S. Savarese, and J. Malik, “Mid-Level Visual Representations Improve Generalization and Sample Efficiency for Learning Visuomotor Policies,” in *Conference on Robot Learning (CoRL)*, 2019.
- [71] E. Tzeng, J. Hoffman, N. Zhang, K. Saenko, and T. Darrell, “Deep domain confusion: Maximizing for domain invariance,” 12 2014.
- [72] B. Sun and K. Saenko, “Deep coral: Correlation alignment for deep domain adaptation,” 7 2016.
- [73] M. Long, Y. Cao, J. Wang, and M. I. Jordan, “Learning transferable features with deep adaptation networks,” 2 2015.
- [74] K. Bousmalis, G. Trigeorgis, N. Silberman, D. Krishnan, and D. Erhan, “Domain Separation Networks,” *Advances in Neural Information Processing Systems (NeurIPS)*, no. Nips, 2016.
- [75] K. Bousmalis, A. Irpan, P. Wohlhart *et al.*, “Using simulation and domain adaptation to improve efficiency of deep robotic grasping,” 2017.
- [76] S. James, P. Wohlhart, M. Kalakrishnan, D. Kalashnikov, A. Irpan, J. Ibarz, S. Levine, R. Hadsell, and K. Bousmalis, “Sim-to-real via sim-to-sim: Data-efficient robotic grasping via randomized-to-canonical adaptation networks,” 12 2018.
- [77] J. Schulman, S. Levine, P. Moritz, M. I. Jordan, and P. Abbeel, “Trust Region Policy Optimization,” in *International Conference on Machine Learning (ICML)*, 2015.
- [78] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal Policy Optimization Algorithms,” 2017.
- [79] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, “Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor,” in *International Conference on Machine Learning (ICML)*, 2018.
- [80] R. Kirk, A. Zhang, E. Grefenstette, and T. Rocktäschel, “A survey of generalisation in deep reinforcement learning,” 11 2021.
- [81] X. B. Peng, A. Kumar, G. Zhang, and S. Levine, “Advantage-Weighted Regression: Simple and Scalable Off-Policy Reinforcement Learning,” sep 2019.
- [82] K. He, X. Zhang, S. Ren, and J. Sun, “Deep Residual Learning for Image Recognition,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [83] D. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *International Conference on Learning Representations (ICLR)*, 2015.
- [84] D. V. Mehrotra, I. S. F. Chan, and R. L. Berger, “A cautionary note on exact unconditional inference for a difference between two independent binomial proportions,” 2003.
- [85] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, “Grad-cam: Visual explanations from deep networks via gradient-based localization,” *International Journal of Computer Vision*, vol. 128, pp. 336–359, 10 2016.