

Invariant Policy Optimization: Towards Stronger Generalization in Reinforcement Learning

Anoopkumar Sonar[†]

Vincent Pacelli*

Anirudha Majumdar*

ASONAR@PRINCETON.EDU

VPACELLI@PRINCETON.EDU

ANI.MAJUMDAR@PRINCETON.EDU

[†] *Department of Computer Science, * Mechanical and Aerospace Engineering, Princeton University*

Abstract

A fundamental challenge in reinforcement learning is to learn policies that generalize beyond the operating domains experienced during training. In this paper, we approach this challenge through the following invariance principle: an agent must find a *representation* such that there exists an *action-predictor* built on top of this representation that is simultaneously optimal across all training domains. Intuitively, the resulting invariant policy enhances generalization by finding *causes of successful actions*. We propose a novel learning algorithm, Invariant Policy Optimization (IPO), that implements this principle and learns an invariant policy during training. We compare our approach with standard policy gradient methods and demonstrate significant improvements in generalization performance on unseen domains for linear quadratic regulator and grid-world problems, and an example where a robot must learn to open doors with varying physical properties.

Keywords: Reinforcement Learning, Generalization, Invariance, Causality

1. Introduction

One of the fundamental challenges with state-of-the-art reinforcement learning (RL) approaches is their limited ability to generalize beyond the specific domains they were trained on. The problem of generalization is particularly acute in complex robotics applications. Deploying an RL policy on a robot outside of the laboratory requires learning a policy that can generalize to a wide range of operating domains, especially in safety-critical applications. For example, autonomous vehicles must contend with unfamiliar obstacles, lighting, and road conditions when deployed at scale; robotic manipulators deployed in homes must deal with new objects and environments; and robots operating in close proximity to humans must be able to handle new patterns of human motion.

As a simple example, consider the problem shown in Figure 1. A robot is placed in a grid-world and must learn to navigate to a goal located in a different room. In order to do this, it must learn to first navigate to a key, use this key to open the door, and then navigate to the goal. During training, the robot is presented with environments containing red and green keys. A policy trained using standard RL techniques demonstrates strong performance when deployed in test environments with key colors seen during training. However, its performance significantly degrades in test environments with different key colors (see Section 5.2 for a thorough exploration of this problem).

Learning policies capable of such generalization remains challenging for a number of reasons. Primarily, RL algorithms have a tendency to memorize solutions to training environments, thereby achieving high training rewards with a brittle policy that will not generalize to novel environments. Moreover, learned policies often fail to ignore distractors in their sensor observations (e.g., the key colors) and are highly sensitive to changes in such irrelevant factors. The goal of this paper is to address these challenges and learn policies that achieve strong generalization across new operating domains given a limited set of training domains.

Statement of Contributions. We approach the problem of generalizing across domains (formalized in Section 2) with the following principle: a policy will generalize well if it exploits in-

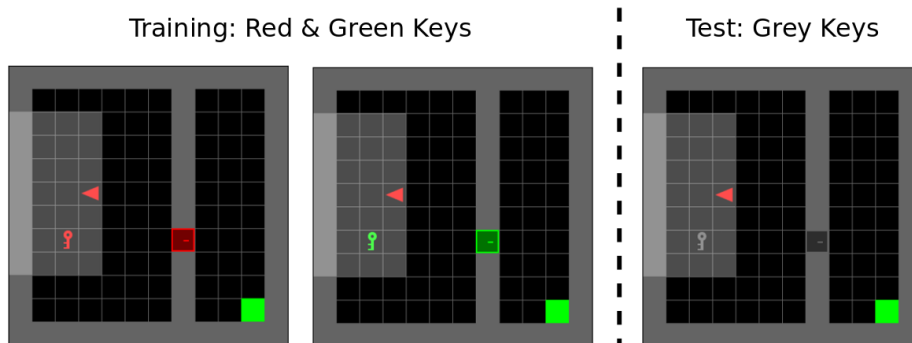


Figure 1: A depiction of the Colored-Key problem described in Section 1 on a 12×12 grid. The color of the keys in the environment corresponds to a different operating *domain*. The agent (red triangle) must learn to use the key to open the door and reach the goal (green square). The agent is trained on domains with red and green keys. At test time, the learned policy is deployed on a domain with differently-colored keys (e.g., grey keys). Our results in Section 5 demonstrate that our algorithm generalizes to this novel testing domain significantly better than one trained using standard techniques.




Figure 2: (a) Door-opening environment in DoorGym. Proximal Policy Optimization (PPO) and IPO tend to find *qualitatively* different policies when trained on domains with low door hinge friction. PPO tends to find a policy that uses the outside of the hook (b), while IPO finds a robust hooking strategy (c).

variances resulting from causal relationships present across domains (e.g. key color does not cause rewards and thus rewards are invariant to the color). To embody this principle, we leverage a close connection between causality and invariance (Section 3) in an approach we refer to as *Invariant Policy Optimization* (IPO). The key idea is to learn a representation that makes the optimal policy built on top of this representation invariant across training domains. Effectively, this approach attempts to learn and exploit the *causes of successful actions*. We demonstrate that IPO exhibits significantly stronger generalization compared to traditional on-policy methods in three different scenarios (Section 5): a linear-quadratic output feedback problem with distracting observations, an instantiation of the colored-key problem, and an example where a robot must learn to open doors with varying physical properties (Figure 2 (a)).

1.1. Related Work

Quantifying generalization. The problem of finding policies that generalize beyond their training domain is becoming an increasingly popular topic as reinforcement learning continues to mature and a number of recent studies have attempted to quantify and understand the generalization challenge in RL. In Song et al. (2019), the authors quantify the effects of *observational overfitting*, where learned policies are sensitive to irrelevant observational features. Benchmark suites includ-

ing Sonic Nichol et al. (2018) and Atari 2600 games Mnih et al. (2013) have also been proposed to quantify generalization. Recently, CoinRun Cobbe et al. (2019b) and the broader Procgen Benchmark Cobbe et al. (2019a) use procedural generalization of environments at controllable levels of difficulty to demonstrate that effective generalization can require an extremely large number of training environments. Another manifestation of the generalization gap is the sim-to-real problem in robotics: agents trained in simulation overfit to this domain and fail to operate in hardware Tobin et al. (2017); Peng et al. (2018); Tan et al. (2018).

Regularization and domain randomization. The most common approach for improving the out-of-domain generalization of a learning algorithm is to add different forms of regularization. Popular ones borrowed from supervised learning include L_2 regularization, dropout Srivastava et al. (2014), and batch normalization Ioffe and Szegedy (2015); each of these has been shown to improve generalization Cobbe et al. (2019b). While practical and easy to implement, these methods typically do not explicitly exploit any structure of the RL problem. Another approach is to constrain the agent’s policy to only depend on a set of learned task-relevant variables, which are found by introducing an information-theoretic regularizer Pacelli and Majumdar (2020); Goyal et al. (2019). These methods have been shown to generalize to new domains that contain task-irrelevant features not present during training. However, the task-relevant variables are not guaranteed to exploit causal relationships in the environment, which is the focus of this paper. Data augmentation and domain randomization have also been shown to be particularly useful in crossing the sim-to-real barrier Urakami et al. (2019); Peng et al. (2018); Akkaya et al. (2019). These methods are complementary to the approach presented here and could potentially be used to generate a diverse set of training domains for our method.

Distributional robustness. The PAC-Bayes Control approach Majumdar et al. (2020); Veer and Majumdar (2020) provides a way to make provable generalization guarantees under distributional shifts. This approach is particularly useful in safety-critical applications where it is important to quantify the impact of switching between training and test domains. Another approach that provides robustness guarantees is to train with adversarial perturbations to the underlying data distribution Sinha et al. (2017). However, the challenge with both of these approaches is that they require an *a priori* bound on how much the test domain differs from the training domain (e.g., in terms of an f -divergence). In contrast, the recently proposed risk-extrapolation method Krueger et al. (2020) promotes out-of-distribution generalization by encouraging robustness of hypotheses over affine combinations of training risks. This method is shown to improve performance of RL agents when their state space is augmented with noisy copies of true system states.

Causality and invariance. Recently, the task of learning causal predictors has drawn interest in the supervised learning setting. An approach formalized in Peters et al. (2016) attempts to find features that are causally linked to a target variable by exploiting the invariance of causal relationships Pearl (2009); Peters et al. (2017). This idea was expanded upon in the invariant risk minimization (IRM) approach Arjovsky et al. (2019), which formulates the problem in terms of finding a representation such that the optimal classifier built on top of this representation is invariant across domains. This results in classifiers that ignore spurious correlations that may exist in any single domain. The formulation leads to a challenging bilevel optimization problem and is tackled via a regularizer that approximates its solution. In Ahuja et al. (2020), the authors present a game-theoretic reformulation of the IRM principle and propose a new algorithm, known as IRM-Games, which offers better empirical results. The approach presented in Teney et al. (2020) uses a variance regularization scheme to find approximately-invariant classifiers. Our approach adapts ideas from causality and invariance to RL settings by learning representations that invariantly predict actions across domains. We provide more background on invariance, causality, and IRM in Section 3.

Causality in RL. Lastly, there are a number of recent methods that attempt to exploit causality in RL. For example, [Dasgupta et al. \(2019\)](#) observed that, in some instances, causal reasoning can emerge in agents trained via meta-learning. Other approaches explicitly attempt to learn causal graphs that describe the dynamics of the agent’s environment. Along these lines, [Nair et al. \(2019\)](#) proposes a two-phase training process where interactions with the environment are first used to learn the causal graph of the environment, and then a policy that exploits this graph is trained. Finally, the IRM method has recently been applied to RL problems. In [Zhang et al. \(2020\)](#), the authors attempt to learn a causal Markov decision process (MDP) that is bisimilar to the full MDP present during training. This formulation requires learning a model for both the causal and full dynamics of the system, a mapping between the two, and a causal model of the rewards. Standard RL algorithms are then used in conjunction with these causal models to produce a final policy. This approach is distinct from the one in this paper, which focuses on *identifying the causes of successful actions*, which is potentially simpler.

2. Problem Formulation

We are interested in the problem of zero-shot generalization to environments that can be significantly different from environments seen during training. We formalize this as follows. Let $s_t = (s_t^a, s_t^e)$ denote the joint state of the agent and environment. In the colored-keys example (Section 1), s_t^a corresponds to the location of the robot at time t , while s_t^e corresponds to locations of the obstacles, key, door, and goal. In our formulation, different environments correspond to different (initial) states of the environment (e.g., different configurations of obstacles, key, door, and goal). We denote the agent’s actions, observations, and rewards by $a_t \in \mathcal{A}$, $o_t \in \mathcal{O}$, and $r_t \in \mathbb{R}$ respectively.

During training, we assume access to multiple sets of environments $\{e_i^d\}_{i=1}^{n_d}$ from different domains $d \in D_{\text{tr}}$. We assume that the action space \mathcal{A} and observation space \mathcal{O} are shared across all domains (state spaces need not be shared). Each domain d corresponds to a partially observable Markov decision process (POMDP) [Thrun et al. \(2005\)](#) with dynamics mapping $p^d(s_{t+1}|s_t, a_t)$, observation mapping $p^d(o_t|s_t)$, and reward mapping $p^d(r_t|s_t, a_t)$. In the colored-keys example, domains differ (only) in terms of the observation mapping; in particular, each domain assigns a particular color to keys. Each domain also defines a distribution \mathcal{D}^d over environments.

Our goal is to learn a policy that generalizes to domains $D_{\text{all}} \supseteq D_{\text{tr}}$ *beyond* the training domains (e.g., generalizing to key colors not seen during training). Specifically, let $R_e(\pi)$ denote the expected cumulative reward $\sum_t r_t$ (over a finite or infinite horizon) when policy π is executed in environment e . We would then like to maximize the worst-case rewards over all domains:

$$R^{\text{all}}(\pi) := \min_{d \in D_{\text{all}}} R^d(\pi) := \min_{d \in D_{\text{all}}} \mathbb{E}_{e \sim \mathcal{D}^d} [R_e(\pi)]. \quad (1)$$

Without further assumptions on the relationship between D_{tr} and D_{all} , finding a policy that performs well on domains D_{all} may be impossible. We discuss this further in Section 4.

3. Background: Invariance and Causality

In this section, we provide a brief exposition of causality and its relationship to invariance. We refer the reader to [Peters et al. \(2017\)](#); [Arjovsky et al. \(2019\)](#); [Pearl \(2009\)](#); [Ahuja et al. \(2020\)](#) for a thorough introduction.

Definition 1 (Structural Causal Model [Peters et al. \(2017\)](#)) A structural causal model (SCM) $\mathcal{C} = (S, \eta)$ governing the random vector $x := (x_1, \dots, x_m)$ is a collection S of m assignments:

$$S_j : x_j \leftarrow f_j(\text{Pa}(x_j), \eta_j), \quad j = 1, \dots, m, \quad (2)$$

where $\text{Pa}(x_j) \subseteq \{x_1, \dots, x_m\} \setminus \{x_j\}$ are the parents of x_j , the η_1, \dots, η_m are independent noise variables, and each f_j is a mapping from these variables to x_j . The graph \mathcal{G} of an SCM is obtained by associating one vertex for each x_j and edges from each parent in $\text{Pa}(x_j)$ to x_j . We assume acyclic causal graphs. We refer to the elements of $\text{Pa}(x_j)$ as the direct causes of x_j .

Definition 2 (Intervention Peters et al. (2017)) Consider an SCM $\mathcal{C} = (S, \eta)$. An intervention d changes one or more of the structural assignments to obtain a new SCM $\mathcal{C}^d = (S^d, \eta^d)$ with assignments:

$$S_j^d : x_j \leftarrow f_j^d(\text{Pa}^d(x_j), \eta_j^d), \quad j = 1, \dots, m. \quad (3)$$

We say that the variables whose structural assignments we have changed have been intervened on.

Figure 3(a) shows the causal graph for the RL formulation in Section 2. Here, the reward r_t depends on the action a_t and a set of “reward-relevant” variables \tilde{s}_t . Thus, $\text{Pa}(r_t) = \{\tilde{s}_t\} \cup \{a_t\}$. In our running colored-keys example, \tilde{s}_t is purely a function of the agent’s state s_t^a and the goal location.

Modularity principle Bareinboim et al. (2012); Peters et al. (2017). The modularity principle establishes a close relationship between causality and *invariance*: a set of variables $\{x_1, \dots, x_k\}$ are the direct causes of y if and only if the conditional probability $p(y|x_1, \dots, x_k)$ remains invariant for all interventions where y has not been intervened on. This is also related to the notion of “autonomy” and the principle of independent mechanisms (Peters et al., 2017, Ch. 2.1). As an example, consider the reward r_t to be the variable of interest in Figure 3 (a). Then, (\tilde{s}_t, a_t) are the direct causes of r_t if and only if for all interventions where r_t has not been intervened on, $p(r_t|\tilde{s}_t, a_t)$ remains invariant. Thus, in the context of the colored-keys example, \tilde{s}_t does not contain any color-related information.




Figure 3: (a) Causal graph corresponding to our RL setting. Here, s_t is the state, o_t is the observation, and a_t is the action. The reward r_t depends only on \tilde{s}_t and a_t . (b) We seek a representation $\Phi : o_t \mapsto h_t$ such that there exists $\pi : h_t \mapsto a_t$ that is simultaneously optimal across domains.

Invariant Risk Minimization (IRM). IRM Arjovsky et al. (2019) exploits the modularity principle in the context of supervised learning. One assumes datasets $\{(x_i^d, y_i^d)\}_{i=1}^{n_d}$ from multiple training domains¹ $d \in D_{\text{tr}}$ corresponding to different interventions on the data-generating process that do not intervene on the target variable y . Here $x_i^d \in \mathcal{X} \subseteq \mathbb{R}^{n_x}$ and $y_i^d \in \mathcal{Y} \subseteq \mathbb{R}^{n_y}$. The goal is to learn a data representation $\Phi : \mathcal{X} \rightarrow \mathcal{H} \subseteq \mathbb{R}^{n_h}$ that elicits an *invariant predictor* $w \circ \Phi$ across training domains, i.e., a representation Φ such that there exists a classifier $w : \mathcal{H} \rightarrow \mathcal{Y}$ that is *simultaneously optimal* for all training domains $d \in D_{\text{tr}}$. Intuitively, the representation Φ should capture the direct causes of y and thus eliminate any features in x that spuriously correlate with y . The optimization problem associated with IRM is a challenging bi-leveled one. The authors of Arjovsky et al. (2019) propose IRM-v1, where a regularization scheme is used to find a representation Φ that leads to classifiers that are approximately locally optimal across training domains.

IRM Games. Inspired by IRM, the authors of Ahuja et al. (2020) demonstrate that the set of invariant predictors corresponds to the set of pure Nash equilibria of a game played among $n_d := |D_{\text{tr}}|$ players. Each player (corresponding to a training domain d) can choose its own classifier w^d and is trying to maximize its own utility function: $-L^d(w^{\text{av}} \circ \Phi)$, where $L^d(\cdot)$ is the loss on domain

1. We note that Arjovsky et al. (2019) uses the term “environment” instead of “domain”. However, we use “domain” since “environment” has a different meaning in RL contexts.

d and $w^{\text{av}}(\cdot) := \frac{1}{n_d} \sum_{d=1}^{n_d} w^d(\cdot)$. While finding Nash equilibria for continuous games is challenging in general, the game theory literature contains several heuristic schemes. In [Ahuja et al. \(2020\)](#), the authors propose a strategy based on *best response dynamics* [Barron et al. \(2010\)](#), where players take turns maximizing their utility functions. The resulting algorithm achieves similar or better empirical performance as compared to IRM-v1, with significantly reduced variance.

4. Invariant Policy Optimization

We now describe our novel reinforcement learning algorithm, which we refer to as *invariant policy optimization* (IPO). The key insight behind this algorithm is to implement the following invariance principle: learn a representation $\Phi : \mathcal{O} \mapsto \mathcal{H} \subseteq \mathbb{R}^{n_h}$ that maps observations o_t to $h_t \in \mathcal{H}$ in a manner that supports *invariant action prediction* (see Figure 3 (b)). More precisely, the goal is to learn a representation Φ such that there exists an ‘‘action-predictor’’ $\pi : \mathcal{H} \rightarrow \mathcal{A}$ built on top of this representation that is *simultaneously optimal across all training domains*.² We will refer to the resulting policy $\pi \circ \Phi$ as an *invariant policy*. This invariance principle can be formally embodied as the following optimization problem:

$$\underset{\substack{\Phi: \mathcal{O} \rightarrow \mathcal{H} \\ \pi: \mathcal{H} \rightarrow \mathcal{A}}}{\text{maximize}} \sum_{d=1}^{n_d} R^d(\pi \circ \Phi) \quad \text{subject to} \quad \pi \in \underset{\bar{\pi}: \mathcal{H} \rightarrow \mathcal{A}}{\text{argmax}} R^d(\bar{\pi} \circ \Phi), \quad \forall d = 1, \dots, n_d. \quad (4)$$

Here, R^d is the reward associated with domain d , as defined in Section 2. Intuitively, given a set of training domains, IPO attempts to learn a representation that corresponds to the ‘‘causes of successful actions’’. This interpretation elucidates the role of the different training domains; these must correspond to different interventions on the causal graph shown in Figure 3 (a) that leave optimal actions unaffected. Assuming a diverse set of training domains, one learns a representation that eliminates features that spuriously correlate with good actions (i.e., actions that achieve high rewards). For example, in the colored-keys problem, such a representation corresponds to one that eliminates color from observations. By eliminating such features, an invariant policy generalizes well to novel domains corresponding to unseen interventions on the spurious features.

Our algorithmic approach for IPO is inspired by the game-theoretic formulation of [Ahuja et al. \(2020\)](#) (see Section 3). We endow each domain d with its own policy π^d and define an overall averaged policy $\pi^{\text{av}}(h_t) := \frac{1}{n_d} \sum_{d=1}^{n_d} \pi^d(h_t)$.³ The optimization problem behind IPO then becomes:

$$\underset{\Phi, \pi^1, \dots, \pi^{n_d}}{\text{maximize}} \sum_{d=1}^{n_d} R^d \left(\underbrace{\frac{1}{n_d} \sum_{d=1}^{n_d} \pi^d \circ \Phi}_{:= \pi^{\text{av}} \circ \Phi} \right) \quad (5)$$

$$\text{subject to } \pi^d \in \underset{\bar{\pi}^d}{\text{argmax}} R^d \left(\frac{1}{n_d} \left[\bar{\pi}^d \circ \Phi + \sum_{i \neq d} \pi^i \circ \Phi \right] \right), \quad \forall d = 1, \dots, n_d.$$

Next, we relate Problem (5) to a game played between n_d players. Each player corresponds to a domain d and chooses a policy π^d to maximize its own utility function $R^d(\pi^{\text{av}} \circ \Phi)$. Since Problem (5) is identical to the one in [Ahuja et al. \(2020\)](#) for finding invariant representations (with policies playing the role of classifiers), the results from [Ahuja et al. \(2020\)](#) carry over to our setting. In particular, under mild technical assumptions on the policies, the set of pure Nash equilibria of the game correspond to the set of invariant policies. We refer the reader to [Ahuja et al. \(2020\)](#) for details

2. For the ease of exposition, we discuss the case of memoryless policies. However, it is straightforward to handle policies with memory (e.g., by augmenting observations with a memory state).

3. Further details on this averaging step for stochastic policies are provided in the appendix of the extended addition of this paper [Sonar et al. \(2020\)](#).

Algorithm 1 Invariant Policy Optimization (IPO)

```

1: for iter = 1, 2, ... itermax do
2:   if Fixed- $\Phi$  then
3:      $\Phi = I$  (identity map)
4:   else
5:      $\Phi \leftarrow \text{PPO}\left(\sum_{d=1}^{n_d} R^d(\pi_{\text{cur}}^{\text{av}} \circ \Phi)\right)$  // Update  $\Phi$  via an iteration of proximal policy optimization
6:   end if
7:   for  $i = 1, \dots, K, d = 1, \dots, n_d$  do
8:     for  $d = 1, \dots, n_d$  do
9:        $\pi^d \leftarrow \text{PPO}\left(R^d(\pi^{\text{av}} \circ \Phi)\right)$  // Update  $\pi^d$  while keeping  $\pi^i$  with  $i \neq d$  fixed.
10:      Define  $\pi^{\text{av}}(\cdot) := \frac{1}{n_d} \sum_{d=1}^{n_d} \pi^d(\cdot)$  // Averaging
11:    end for
12:  end for
13: end for

```

on the technical assumptions, but note that these are satisfied by a wide range of function classes (e.g., ReLU networks with arbitrary depth, linear functions, and functions in L^p spaces).

While finding Nash equilibria for continuous games such as the one above is difficult in general, the game theory literature has developed several approximate approaches that demonstrate good performance in practice. Here, we adapt the strategy based on *best response dynamics* Barron et al. (2010) proposed in Ahuja et al. (2020) to our setting. The resulting IPO training procedure is presented in Algorithm 1. The for-loop in lines 8–11 implement the best-response dynamics; the players (corresponding to the different domains) take turns choosing π^d in order to optimize their own objective $R^d(\pi^{\text{av}} \circ \Phi)$. We choose to implement the updates using proximal policy optimization (PPO) Schulman et al. (2017). However, this choice is not fundamental and one may implement the updates using other policy gradient methods. Algorithm 1 can also accommodate actor-critic methods. In this version, each domain d has both an actor π^d and a critic v^d . In the policy-update steps, one updates both the actor and the critic using PPO.

Line 5 of the algorithm periodically updates the representation Φ . However, as demonstrated in Ahuja et al. (2020), simply fixing $\Phi = I$ (identity) is an effective approach and can (under certain conditions) recover invariant predictors that extract causal features (and ignore non-causal factors such as color in the colored-key example). The intuition is that the burden of extracting causal features from observations is simply shifted to π (the portion of the overall policy that maps the output of Φ to actions). In our numerical experiments (Section 5), we did not find significant benefits to periodically updating Φ .⁴

5. Examples

Next, we demonstrate the generalization benefits afforded by IPO on three examples. Code for all our examples is available on GitHub: <https://github.com/irom-lab/Invariant-Policy-Optimization.git>.

5.1. Linear Quadratic Regulator with Distractors

We first apply our approach to the linear-quadratic regulator (LQR) problem Anderson and Moore (2007) modified to include high-dimensional “distractor” observations. There has been a growing interest in LQR as a simplified surrogate for deep RL problems Dean et al. (2019); Tu and Recht (2018); Fazel et al. (2018); Agarwal et al. (2019). Here we consider the output-feedback control

4. Further discussion on the choice of $\Phi = I$ is provided in the appendix of the extended edition of this paper Sonar et al. (2020).

Number of training domains	2	3	4	5	10
Gradient descent	97.7±5.4	90.0±9.2	82.6±4.3	78.6±5.4	68.8±3.9
Overparameterization	86.2±3.0	75.3±4.4	69.4±1.8	64.5±1.9	51.4±1.0
IPO (ours) [Fixed- Φ]	78.8±3.5	64.8±2.3	57.7±1.0	52.3±1.8	43.2±1.1
IPO (ours) [Variable- Φ]	71.3±2.1	62.7±2.4	57.8±1.0	53.9±1.1	45.3±0.45
LQR oracle	32.1	32.1	32.1	32.1	32.1

Table 1: LQR with distractors: comparison of IPO with two baselines (gradient descent and overparameterization) with distractor dimension $n_y = 1000$ and varying number of training domains. IPO demonstrates stronger generalization (lower costs) compared to the baselines. The mean and std. dev. are across 10 seeds.

problem proposed in Song et al. (2019) as a benchmark for assessing generalization with respect to changes in the observation model. The dynamics of the system are described by $s_{t+1} = As_t + Ba_t$, where $s_t \in \mathbb{R}^{n_s}$, $a_t \in \mathbb{R}^{n_a}$, and A, B are fixed matrices. The agent receives a high-dimensional observation vector $o_t = (W_c s_t, W_d s_t)$, where W_c and W_d are semi-orthogonal matrices. This ensures that the portion of the observation corresponding to $W_c s_t \in \mathbb{R}^{n_s}$ contains full information about the state, while $W_d s_t \in \mathbb{R}^{n_y}$ is a high-dimensional “distractor.” The goal is to choose policies of the form $a_t = K o_t$ in order to minimize the infinite-horizon LQR cost $\sum_{t=0}^{\infty} s_t^T Q s_t + a_t^T R a_t$.

Here, a domain corresponds to a particular choice of W_d ; all other parameters (A, B, Q, R, W_c) are shared across domains and unknown to the agent. During training time, one learns a policy $K o_t$ using n_d domains. At test time, the learned policy is assessed on a new domain. In the case where there is a single domain (used for both training and test) and $o_t = s_t$, one can find the globally optimal policy via gradient descent (despite the non-convexity of the problem) Fazel et al. (2018). However, as demonstrated in Song et al. (2019), simple policy gradient using the combined costs of multiple training domains finds a policy that overfits to the training domains in the more general setting considered here. Intuitively, this is because the learned policy fails to ignore the distractors.

For our numerical experiments, we choose $n_x = n_a = 20$ and $Q = R = I_{20 \times 20}$. The matrices A and W_c are random orthogonal matrices, B is $I_{20 \times 20}$, and the W_d are random semi-orthogonal matrices (different for each domain). For IPO, we employ a policy $n_d^{-1} \sum_{i=1}^{n_d} K_i^d$ that averages policies K^d corresponding to the training domains. Instead of PPO, we simply use gradient descent to perform policy updates.⁵

We compare our approach with two baselines: (i) gradient descent on K using the combined cost of n_d training domains, and (ii) gradient descent using an *overparameterized* class of policies with two layers (i.e., $K = K_1 K_2$) and hidden dimension of $10n_a$. Interestingly, Song et al. (2019) found that this form of overparameterization induces an implicit regularization towards “simpler” policies (i.e., ones that are less “dependent” on the distractors). We present results using both the Fixed- Φ version of IPO (ref. Algorithm 1) and the version where Φ is optimized. In the latter version, we choose Φ to be linear (with output dimension $10n_a$). Table 1 compares the generalization performance of the learned policies to new domains as we vary the number of training domains. Here, the distractors have dimension 1000. Consistent with Song et al. (2019), we find that overparameterization forms a strong baseline for this problem. However, IPO significantly outperforms both baselines. As expected, performance improves with increasing number of training domains and tends towards the performance achieved by an “oracle” policy that has access to the full state s_t on the test domain. Note that the results using Fixed- Φ are similar to the version with optimized Φ , which is consistent with the arguments provided in Section 4. Table 2 assesses the impact of the distractor dimension. Here, we fix the number of training domains to five. Again, IPO demonstrates significantly improved performance.

5. Hyperparameters and implementation details for all examples are provided in the appendix of the extended edition of this paper Sonar et al. (2020).

Distractor dimension	100	500	1000	1500	2000
Gradient descent	46.3±2.4	70.3±4.5	78.6±5.4	88.7±5.6	94.7±8.2
Overparameterization	36.3±0.5	54.3±2.0	64.5±1.9	71.6±2.0	79.1±3.4
IPO (ours) [Fixed- Φ]	42.5±2.0	50.8±1.1	52.3±1.8	55.0±1.7	59.4±2.2
IPO (ours) [Variable- Φ]	38.5±0.6	48.0±1.1	53.9±1.1	58.1±1.5	62.0±1.6
LQR oracle	32.1	32.1	32.1	32.1	32.1

Table 2: LQR with distractors: comparison of IPO with two baselines (gradient descent and overparameterization) with varying dimensionality of distractors and $n_d = 5$. IPO demonstrates stronger generalization (i.e., lower costs) compared to the two baselines. The reported mean and std. dev. are across 10 different seeds.

5.2. Colored-Key Domains

We now consider the colored-keys problem introduced in Section 1. In this example, a robot is placed in a grid-world that contains a goal (located in a room), a door, and a key (see Figure 1). The robot is presented with a reward if it reaches the goal. Using this sparse reward signal, it must learn to first navigate to the key, use this to open the door, and then navigate to the goal. In this setting, an environment corresponds to a particular configuration of the key, door, goal, and obstacles. Different domains correspond to different key colors.

We implement our approach on 5×5 grid-worlds using MiniGrid [Chevalier-Boisvert et al. \(2018\)](#). Observations in MiniGrid correspond to $5 \times 5 \times 3$ values; the three channels encode the object type (e.g. door), object color, and object state (e.g., open/closed) for a 5×5 neighborhood around the robot. The robot receives a sparse reward of $1 - 0.9t/T$, where t is the time taken to reach the goal and $T = 250$ is the time-limit for completing the task. During training, the robot has access to environments from two domains corresponding to red and green keys. We use 48 training environments split evenly between these domains. At test-time, the robot is placed in environments with grey keys. This color choice is motivated by the fact that in MiniGrid, colors are encoded using integers (e.g., red: 0, green: 1), and grey corresponds to the color that is “furthest away” in terms of this encoding (grey: 5). We implement IPO with an actor-critic architecture and the Fixed- Φ option (given the results in Section 5.1 and argument in Section 4). Table 3 reports the average rewards on 50 test environments from the training and test domains. We compare our approach to PPO [Schulman et al. \(2017\)](#) trained to maximize rewards combined across training environments. IPO achieves better generalization to the new domain and is also more consistent across training seeds.

Key color	Red (training)	Green (training)	Grey (testing)
PPO	0.94±0.004	0.94±0.005	0.80±0.12
IPO (ours)	0.94±0.003	0.94±0.003	0.85±0.03

Table 3: Colored-key domains: comparison of the average reward on 50 test environments drawn from different domains. The reported mean and std. dev. are across 10 different seeds.

Friction	1.1	1.2	1.3	1.4	1.5
PPO	96.6±2.1	94.0±3.6	89.8±8.1	83.4±12.0	78.2±13.7
IPO (ours)	99.0±1.3	96.0±2.7	93.8±5.9	87.4±9.8	81.0±12.6

Table 4: Door-opening environments. Training was performed using two domains with 0.0 and 0.1 friction respectively. The table compares the door-opening success rates for PPO and IPO on 100 test environments with higher friction. The reported mean and std. dev. are across 5 seeds.

5.3. DoorGym

Finally, we consider a more challenging task where a robot must learn to open doors (Figure 2 (a)). We implement this example using DoorGym [Urakami et al. \(2019\)](#), which uses MuJoCo [Todorov](#)

et al. (2012) as its simulation engine. Each environment in Doorgym corresponds to a particular setting of physical parameters including door position, height, width, mass, handle location, door hinge spring constant, and hinge friction. Observations correspond to the pose (and velocity) of the gripper (hook), along with the position of the door handle. Actions correspond to torques applied on the gripper. During training, we use environments from two domains. The first domain corresponds to environments with *zero* friction at the door hinge, while the second domain corresponds to the friction set to 0.1; all other environment parameters (e.g., door position, height, etc.) are randomized. We use 32 training environments, with a 75%-25% split between the two domains.

We implement IPO using policies with an actor-critic architecture and the Fixed- Φ option. We compare IPO with PPO implemented using the same architecture and trained using environments pooled from both domains. We use the reward function specified in Urakami et al. (2019) for training. Table 4 presents the door-opening success rates of both methods on test environments with the friction set to higher values than seen during training. IPO demonstrates improved generalization to these test domains. Perhaps more interesting than the quantitative difference between the two methods, we observe a *qualitative* difference between the policies learned using IPO and PPO. Across five seeds, PPO learns a policy that uses the *outside of the hook* on three seeds (Figure 2 (b)). On one seed, PPO learns to use the hook as one would expect (Figure 2 (c)) and demonstrates a mixture of the two behaviors on the other seed. While using the outside of the hook is a viable strategy for lower friction values, this strategy is not as robust for higher values. In contrast, IPO learns to perform the hooking maneuver shown in Figure 2 (c) on all five seeds.⁶

6. Discussion and Conclusions

We have considered the problem of learning policies with strong generalization beyond training domains. The key idea behind Invariant Policy Optimization (IPO) is to learn representations that support *invariant action prediction* across different domains. We implemented the proposed techniques on: (i) linear quadratic regulator (LQR) problems with distractor observations, (ii) an example where an agent must learn to navigate to a goal by opening a door using different colored keys in its environment, and (iii) an example where a robot must learn to open doors with varying physical properties. We compared our approach with standard policy gradient methods (e.g., PPO) and demonstrated significant improvements in generalization performance on unseen domains.

Future work. As IRM is a recently developed technique for training robust models, the theory and methods surrounding it are still evolving (see e.g. Gulrajani and Lopez-Paz (2020); Rosenfeld et al. (2020); Nagarajan et al. (2020) for recent developments). Due to the structural similarities between our method and IRM, improvements to IRM can be readily transferred to IPO. This includes both theoretical developments surrounding IRM and improved bi-level optimization schemes (here best-response was chosen for simplicity and empirical performance).

There are also a number of specific directions for improvements to IPO. On the theoretical front, an important direction for future work is to provide rigorous guarantees on generalization to novel domains. One potential avenue is to combine the algorithmic techniques presented here with recent results on PAC-Bayes generalization theory applied to control and RL settings Majumdar et al. (2020); Veer and Majumdar (2020). On the algorithmic front, an interesting direction is to use domain randomization techniques to *automatically generate* new training domains that can be used to improve invariant policy learning (e.g., automatically generating domains with different colored keys in the colored-keys example). Finally, a particularly promising future direction is to explore the application of IPO to robotics problems involving sim-to-real transfer, where one considers simulation and reality as different domains to learn a policy that is invariant across them.

6. A video of both policies is available at <https://youtu.be/J3O474yZ2Tc>.

Acknowledgements

The authors are grateful to Kartik Ahuja for helpful clarifications on the training procedure for IRM-Games. The authors would like to thank Richard Song and Behnam Neyshabur for providing the code from Song et al. (2019) for the LQR example in Section 5.1. The authors also thank Damien Teney for helpful comments on an early version of this paper.

This work is partially supported by the Office of Naval Research [Award Number: N00014-18-1-2873], the Google Faculty Research Award, the Amazon Research Award, the National Science Foundation [IIS1755038], and the Alfred Rheinsein Faculty Award from Princeton University.

References

- Naman Agarwal, Brian Bullins, Elad Hazan, Sham M Kakade, and Karan Singh. Online control with adversarial disturbances. *arXiv preprint arXiv:1902.08721*, 2019.
- Kartik Ahuja, Karthikeyan Shanmugam, Kush Varshney, and Amit Dhurandhar. Invariant risk minimization games. *arXiv preprint arXiv:2002.04692*, 2020.
- Ilge Akkaya, Marcin Andrychowicz, Maciek Chociej, Mateusz Litwin, Bob McGrew, Arthur Petron, Alex Paino, Matthias Plappert, Glenn Powell, Raphael Ribas, et al. Solving Rubik’s cube with a robot hand. *arXiv preprint arXiv:1910.07113*, 2019.
- Brian Anderson and John Moore. *Optimal Control: Linear Quadratic Methods*. Courier Corporation, 2007.
- Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019.
- Elias Bareinboim, Carlos Brito, and Judea Pearl. Local characterizations of causal bayesian networks. In *Graph Structures for Knowledge Representation and Reasoning*, pages 1–17. Springer, 2012.
- E. Barron, Rafal Goebel, and R. Jensen. Best response dynamics for continuous games. *Proceedings of the American Mathematical Society*, 138(3):1069–1083, 2010.
- Maxime Chevalier-Boisvert, Lucas Willems, and Suman Pal. Minimalistic gridworld environment for openai gym. <https://github.com/maximecb/gym-minigrid>, 2018.
- Karl Cobbe, Christopher Hesse, Jacob Hilton, and John Schulman. Leveraging procedural generation to benchmark reinforcement learning. *arXiv preprint arXiv:1912.01588*, 2019a.
- Karl Cobbe, Oleg Klimov, Chris Hesse, Taehoon Kim, and John Schulman. Quantifying generalization in reinforcement learning. In *Proceedings of the International Conference on Machine Learning*, pages 1282–1289, 2019b.
- Ishita Dasgupta, Jane Wang, Silvia Chiappa, Jovana Mitrovic, Pedro Ortega, David Raposo, Edward Hughes, Peter Battaglia, Matthew Botvinick, and Zeb Kurth-Nelson. Causal reasoning from meta-reinforcement learning. *arXiv preprint arXiv:1901.08162*, 2019.
- Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. On the sample complexity of the linear quadratic regulator. *Foundations of Computational Mathematics*, pages 1–47, 2019.

- Maryam Fazel, Rong Ge, Sham M Kakade, and Mehran Mesbahi. Global convergence of policy gradient methods for the linear quadratic regulator. *arXiv preprint arXiv:1801.05039*, 2018.
- Anirudh Goyal, Riashat Islam, Daniel Strouse, Zafarali Ahmed, Matthew Botvinick, Hugo Larochelle, Yoshua Bengio, and Sergey Levine. Infobot: Transfer and exploration via the information bottleneck. *arXiv preprint arXiv:1901.10902*, 2019.
- Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. *arXiv preprint arXiv:2007.01434*, 2020.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- David Krueger, Ethan Caballero, Joern-Henrik Jacobsen, Amy Zhang, Jonathan Binas, Remi Le Priol, and Aaron Courville. Out-of-distribution generalization via risk extrapolation (REx). *arXiv preprint arXiv:2003.00688*, 2020.
- Anirudha Majumdar, Alec Farid, and Anoopkumar Sonar. PAC-Bayes Control: Learning policies that provably generalize to novel environments. *International Journal of Robotics Research (IJRR)*, October 2020.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing Atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- Vaishnavh Nagarajan, Anders Andreassen, and Behnam Neyshabur. Understanding the failure modes of out-of-distribution generalization. *arXiv preprint arXiv:2010.15775*, 2020.
- Suraj Nair, Yuke Zhu, Silvio Savarese, and Li Fei-Fei. Causal induction from visual observations for goal directed tasks. *arXiv preprint arXiv:1910.01751*, 2019.
- Alex Nichol, Vicki Pfau, Christopher Hesse, Oleg Klimov, and John Schulman. Gotta learn fast: A new benchmark for generalization in reinforcement learning. *arXiv preprint arXiv:1804.03720*, 2018.
- Vincent Pacelli and Anirudha Majumdar. Learning task-driven control policies via information bottlenecks. In *Proceedings of Robotics: Science and Systems (RSS)*, 2020.
- Judea Pearl. *Causality*. Cambridge University Press, 2009.
- Xue Bin Peng, Marcin Andrychowicz, Wojciech Zaremba, and Pieter Abbeel. Sim-to-real transfer of robotic control with dynamics randomization. In *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018.
- Jonas Peters, Peter Bühlmann, and Nicolai Meinshausen. Causal inference by using invariant prediction: identification and confidence intervals. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 78(5):947–1012, 2016.
- Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of Causal Inference: Foundations and Learning Algorithms*. MIT Press, 2017.
- Elan Rosenfeld, Pradeep Ravikumar, and Andrej Risteski. The risks of invariant risk minimization. *arXiv preprint arXiv:2010.05761*, 2020.

- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Aman Sinha, Hongseok Namkoong, and John Duchi. Certifying some distributional robustness with principled adversarial training. *arXiv preprint arXiv:1710.10571*, 2017.
- Anoopkumar Sonar, Vincent Pacelli, and Anirudha Majumdar. Invariant policy optimization: Towards stronger generalization in reinforcement learning. *arXiv preprint arXiv:2006.01096*, 2020.
- Xingyou Song, Yiding Jiang, Yilun Du, and Behnam Neyshabur. Observational overfitting in reinforcement learning. *arXiv preprint arXiv:1912.02975*, 2019.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- Jie Tan, Tingnan Zhang, Erwin Coumans, Atil Iscen, Yunfei Bai, Danijar Hafner, Steven Bohez, and Vincent Vanhoucke. Sim-to-real: Learning agile locomotion for quadruped robots. *arXiv preprint arXiv:1804.10332*, 2018.
- Damien Teney, Ehsan Abbasnejad, and Anton van den Hengel. Unshuffling data for improved generalization. *arXiv preprint arXiv:2002.11894*, 2020.
- Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic Robotics*. MIT Press, 2005.
- Joshua Tobin, Wojciech Zaremba, and Pieter Abbeel. Domain randomization and generative models for robotic grasping. *arXiv preprint arXiv:1710.06425*, 2017.
- Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5026–5033. IEEE, 2012.
- Stephen Tu and Benjamin Recht. The gap between model-based and model-free methods on the linear quadratic regulator: An asymptotic viewpoint. *arXiv preprint arXiv:1812.03565*, 2018.
- Yusuke Urakami, Alec Hodgkinson, Casey Carlin, Randall Leu, Luca Rigazio, and Pieter Abbeel. Doorgym: A scalable door opening environment and baseline agent. *arXiv preprint arXiv:1908.01887*, 2019.
- Sushant Veer and Anirudha Majumdar. Probably approximately correct vision-based planning using motion primitives. In *Proceedings of the Conference on Robot Learning (CoRL)*, 2020.
- Amy Zhang, Clare Lyle, Shagun Sodhani, Angelos Filos, Marta Kwiatkowska, Joelle Pineau, Yarín Gal, and Doina Precup. Invariant causal prediction for block MDPs. *arXiv preprint arXiv:2003.06016*, 2020.