Object-Aware Regularization for Addressing Causal Confusion in Imitation Learning

Anonymous Author(s)
Affiliation
Address
email

Abstract

Behavioral cloning has proven to be effective for learning sequential decision-making policies from expert demonstrations. However, behavioral cloning often suffers from the causal confusion problem where a policy relies on the noticeable *effect* of expert actions due to the strong correlation but not the *cause* we desire. This paper presents Object-aware REgularizatiOn (OREO), a simple technique that regularizes an imitation policy in an object-aware manner. Our main idea is to encourage a policy to uniformly attend to all semantic objects, in order to prevent the policy from exploiting nuisance variables strongly correlated with expert actions. To this end, we introduce a two-stage approach: (a) we extract semantic objects from images by utilizing discrete codes from a vector-quantized variational autoencoder, and (b) we randomly drop the units that share the same discrete code together, i.e., masking out semantic objects. Our experiments demonstrate that OREO significantly improves the performance of behavioral cloning on Atari environments, outperforming various other regularization and causality-based methods. We also show that our method even outperforms inverse reinforcement learning methods trained with a considerable amount of environment interaction.

1 Introduction

Imitation learning (IL) holds the promise of learning skills or behaviors directly from expert demonstrations, effectively reducing the need for costly and dangerous environment interaction [19, 40]. Its simplest and effective form is behavioral cloning (BC), which learns a policy by solving a supervised learning problem over state-action pairs from expert demonstrations. While being simple, BC has been successful in a wide range of tasks [4, 7, 26, 28] with careful designs. However, it has been recently evidenced that BC often suffers from the causal confusion problem, where the policy relies on nuisance variables strongly correlated with expert actions, instead of the true *causes* [10, 11, 48].

For example, when we train a BC policy on the Atari Pong environment (see Figure 1a), we observe that a policy relies on nuisance variables in images (i.e., scores) for predicting expert actions, instead of learning the underlying fundamental rule of the environment that experts would have used for making decisions. In particular, Table 1c shows that the policy trained using images with scores struggles to generalize to images with scores masked out (see Figure 1b). However, the policy trained with masked images could generalize to original images with scores, which shows that it successfully learned the rule of the environment. This implies that learning the policy that can identify the true cause of expert actions is important for stable performance at deployment time, where nuisance correlates usually do not hold as in expert demonstrations.

In order to address this causal confusion problem, one can consider causal discovery approaches to deduce the cause-effect relationships from observational data [22, 43]. However, it is difficult to apply these approaches to domains with high-dimensional inputs, as (i) causal discovery from observational

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

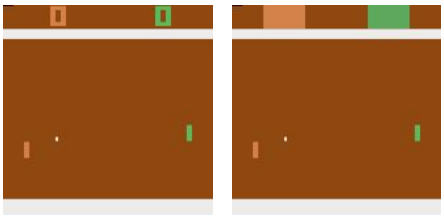

(a) Original (b) Masked

Setup		Scores
Train	Eval	
Original	Original	3.1± 1.4
	Masked	-15.6± 9.2
Masked	Original	**15.9**± 0.4
	Masked	**16.6**± 0.6

(c) Performance of behavioral cloning

Figure 1: Atari Pong environment with (a) original images and (b) images where scores are masked out. (c) Performance of behavioral cloning (BC) policy trained in Original and Masked environments, averaged over four runs. We observe that the policy trained with original images suffers in both environments, which shows that the policy exploits score information for predicting expert actions, instead of learning the underlying fundamental rule of the environment.

data is impossible in general without certain conditions[1] [33], and (ii) these domains usually do not satisfy the assumption that inputs are structured into random variables connected by a causal graph, e.g., objects in images [24, 41]. To address these limitations, de Haan et al. [11] recently proposed a method that learns a policy on top of disentangled representations from a β-VAE encoder [17] with random masking, and infers an optimal causal mask during the environment interaction by querying interactive experts [38] or environment returns. However, given that environment interaction could be dangerous and incur additional costs, we argue that it is important to develop a method for learning the policy robust to causal confusion problem without such a costly environment interaction.

In this paper, we present OREO: **O**bject-aware **RE**gularizati**O**n, a new regularization technique that addresses the causal confusion problem in imitation learning without environment interaction. The key idea of our method is to regularize a policy to attend uniformly to all semantic objects in images, in order to prevent the policy from exploiting nuisance correlates for predicting expert actions. To this end, we propose to extract semantic objects from raw images by utilizing vector-quantized variational autoencoder (VQ-VAE) [30]. In our experiments, we discover that the units of a feature map corresponding to the objects with similar semantics, e.g., backgrounds, scores, and characters, are mapped into the same or similar discrete codes (see Figure 3). Based upon this observation, we propose to regularize the policy by randomly dropping units that share the same discrete code together throughout training. Namely, our method randomly masks out semantically similar objects, which allows object-aware regularization of the policy.

We highlight the main contributions of this paper below:

- We present OREO, a simple and effective regularization method for addressing the causal confusion problem, and support the effectiveness of OREO with extensive experiments.
- We show that OREO significantly improves the performance of behavioral cloning on confounded Atari environments [5, 11], outperforming various other regularization methods [12, 13, 49, 45] and causality-based methods [11, 42].
- We show that OREO even outperforms inverse reinforcement learning methods that require a considerable amount of environment interaction [8, 18].

2 Related work

Imitation learning. Imitation learning (IL) aims to solve complex tasks where learning a policy from scratch is difficult or even impossible, by learning useful skills or behaviors from expert demonstrations [2, 16, 21, 32, 35, 34, 50]. There are two main approaches for IL: inverse reinforcement learning (IRL) methods that find a cost function under which the expert is uniquely optimal [8, 18, 29, 39, 53], and behavioral cloning methods that formulate the IL problem as a supervised learning problem that predicts expert actions from states [3, 4, 7, 26, 28, 35]. Our work employs behavioral cloning as it exhibits the benefit of avoiding costly and dangerous environment interaction, which is crucial for applying imitation learning to real-world scenarios.

[1] de Haan et al. [11] showed that causal discovery methods that depend on faithfulness condition [33] are not applicable to imitation learning setup, as the condition does not hold in environments with nuisance correlates.

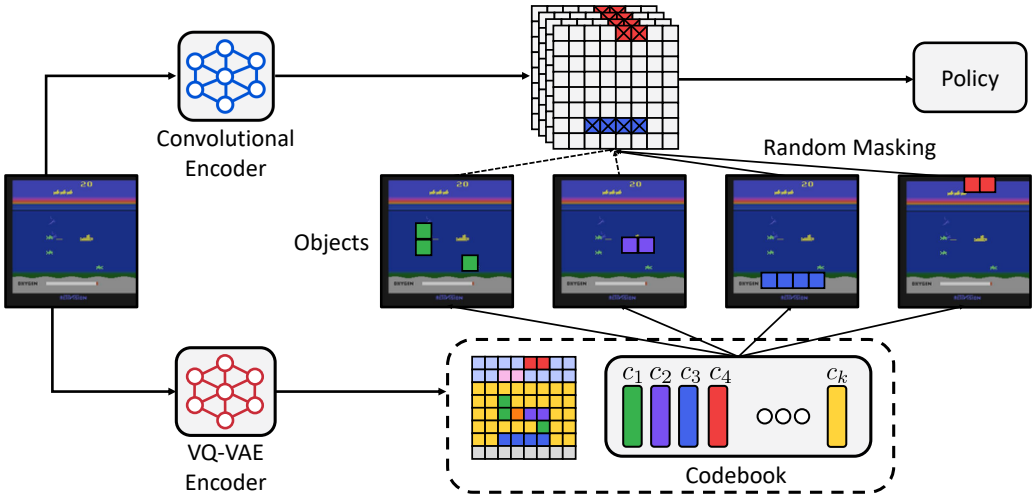

Figure 2: Overview of OREO. We first train a VQ-VAE model that encodes images into discrete codes from a codebook, where each discrete (prototype) representation represents different semantic objects in images. We then regularize a policy by randomly dropping units that share the same discrete code together, i.e., random objects, throughout training.

Distributional shift and the causal confusion problem. Despite its simplicity, BC is known to suffer from the distributional shift, where the state distribution induced by a policy gets different from the training distribution on which the policy was trained. Several approaches have been proposed for learning the policy robust to distributional shift, including interactive IL methods that query experts [37, 38, 46], and regularization techniques [4, 7]. Recently, it has been evidenced that distributional shift leads to the causal confusion problem [4, 11, 48] where a policy exploits the nuisance correlates in states for predicting expert actions. To address this problem, Bansal et al. [4] proposed to randomly drop previous samples from a sequence of samples, and Wen et al. [48] proposed an adversarial training scheme of removing information related to previous actions. The work closest to ours is de Haan et al. [11], which learns a policy with randomly masked disentangled representations and infers the best mask through during environment interaction. Our approach differs in that we regularize the policy to be robust to the causal confusion problem, without any environment interaction.

Causal discovery from observational data. Causal discovery aims to discover causal relations among variables by utilizing observational data [33]. Most prior approaches assume that inputs are structured as disentangled variables [6, 14, 22, 31, 43, 42], which often does not hold in domains with high-dimensional inputs, i.e., images. While Lopez-Paz et al. [24] demonstrated the possibility of observational causal discovery from high-dimensional images, combining causal models and representation learning in such domains still remains an open problem [41]. Hence, we instead explore the approach of regularizing a policy that operates on high-dimensional states.

3 Method

3.1 Preliminaries

We consider the standard imitation learning (IL) framework where an agent learns to solve a target task from expert demonstrations. Specifically, IL is typically defined in the context of a discrete-time Markov decision process (MDP) [47] without an explicitly-defined reward function, which is defined as a tuple $(\mathcal{S}, \mathcal{A}, p, \gamma, \rho_0)$. Here, $\mathcal{S}$ is the state space, $\mathcal{A}$ is the action space, $p(s'|s,a)$ is the transition dynamics, ρ_0 is the initial state distribution, and $\gamma \in [0, 1)$ is the discount factor. The goal of IL is to learn a policy π, mapping from states to actions, using a set of expert demonstrations $\mathcal{D} = \{(s_i, a_i)\}_{i=1}^N$. In our problem setup, an agent cannot interact with the environment, hence it should learn the policy by using only expert demonstrations.

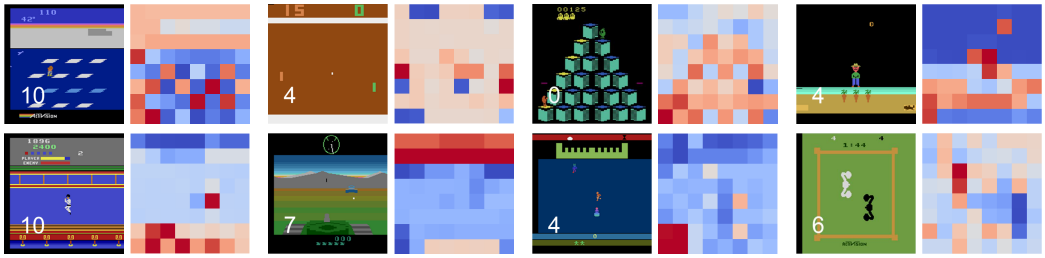

Figure 3: Visualization of the discrete codes from a VQ-VAE model trained on 8 confounded Atari environments, where previous actions are augmented to the images as nuisance variables following the setup in de Haan et al. [11]. The odd columns show images from environments, and even columns show the corresponding quantized feature maps, respectively. The discrete codes are visualized in 1D using t-SNE [25]. We observe that the units with similar semantics exhibit similar colors, i.e., mapped into the same or similar discrete codes.

Behavioral cloning. Behavioral cloning (BC) reduces an imitation learning problem to the supervised learning problem of training a policy that imitates expert actions. Specifically, we introduce a policy π that maps a state s_t to an action a_t, and a convolutional encoder f that maps a state s_t to a low-dimensional feature map. Then π and f are learned by minimizing the negative log-likelihood of expert actions from demonstrations as follows:

$$\mathcal{L}_{\texttt{BC}}(s_t, a_t) = -\log \pi(a_t | f(s_t)), \tag{1}$$

where π is modeled as a multinomial distribution over actions to handle discrete action spaces.

Vector quantized variational autoencoder. The VQ-VAE [30] model consists of an encoder g that compresses images into discrete latent representations, and a decoder d that reconstructs images from these discrete representations. Both encoder and decoder share a codebook C of prototype vectors which are also learned throughout training. Formally, given a state s_t, the encoder g encodes s_t into a feature map $h_t \in \mathbb{R}^{L \times D}$ that consists of a series of L latent vectors $h_{t,i} \in \mathbb{R}^D, i \in \{1, 2, ..., L\}$. Then $h_t = g(s_t)$ is quantized to discrete representations $e \in \mathbb{R}^{L \times D}$ based on the distance of latent vectors $h_{t,i}$ to the prototype vectors in the codebook $C = \{e_k\}_{k=1}^K$ as follows:

$$e_t = (e_{q(t,1)}, e_{q(t,2)}, \cdots, e_{q(t,L)}), \quad \text{where } q(t,i) = \underset{k \in [K]}{\operatorname{argmin}} \|h_{t,i} - e_k\|_2, \tag{2}$$

where $[K]$ is the set $\{1, \cdots, K\}$. Then the decoder d learns to reconstruct s_t from discrete representations e_t. The VQ-VAE is trained by minimizing the following objective:

$$\mathcal{L}_{\texttt{VQVAE}}(s_t) = \underbrace{\|s_t - d(e_t)\|_2^2}_{\mathcal{L}_{\texttt{recon}}} + \underbrace{\|sg\,[h_t] - e_t\|_2^2}_{\mathcal{L}_{\texttt{codebook}}} + \underbrace{\beta \cdot \|sg\,[e_t] - h_t\|_2^2}_{\mathcal{L}_{\texttt{commit}}}, \tag{3}$$

where the operator sg refers to a stop-gradient operator, $\mathcal{L}_{\texttt{recon}}$ is a reconstruction loss for learning representations useful for reconstructing images, $\mathcal{L}_{\texttt{codebook}}$ is a codebook loss to bring codebook representations closer to corresponding encoder outputs h, and $\mathcal{L}_{\texttt{commit}}$ is a commitment loss weighted by β to prevent encoder outputs from fluctuating frequently between different representations.

3.2 OREO: Object-aware regularization for behavioral cloning

In this section, we present OREO: **O**bject-aware **RE**gularizati**O**n that regularizes a policy in an object-aware manner to address the causal confusion problem. Our main idea is to encourage the policy to uniformly attend to all semantic objects in images, in order to prevent the policy from exploiting nuisance variables strongly correlated with expert actions. To this end, we introduce a two-stage approach: we first train a VQ-VAE model that encodes images into discrete codes, then learn the policy with our regularization scheme of randomly dropping units that share the same discrete codes (see Figure 2 and Algorithm 1 for the overview and pseudocode of OREO, respectively).

Algorithm 1 Object-aware regularization (OREO)

Initialize parameters of encoder g, decoder d, codebook C, policy π.
while not converged **do** // VQ-VAE TRAINING
  Sample $\{s_i\}_{i=1}^{B} \sim \mathcal{D}$.
  Update parameters of g, d, C by minimizing $\sum_{i=1}^{B} \mathcal{L}_{\texttt{VQVAE}}(s_i)$
end while
Initialize encoder f with parameters of g.
while not converged **do** // UPDATE POLICY VIA BEHAVIORAL CLONING
  Sample $\{s_i, a_i\}_{i=1}^{B} \sim \mathcal{D}$
  Get random masks $\{M_i\}_{i=1}^{B}$ in (4)
  Update parameters of f, π by minimizing $\sum_{i=1}^{B} \mathcal{L}_{\texttt{OREO}}(s_i, a_i, M_i)$ in (5)
end while

Extracting semantic objects. To regularize a policy in an object-aware manner, we propose to utilize discrete representations from a VQ-VAE model trained by optimizing the objective in (3) with images from expert demonstrations $\mathcal{D}$. Our motivation comes from the observation that the units of a feature map corresponding to similar objects are mapped into similar discrete codes (see Figure 3). Then, in order to extract semantic objects from images and utilize them for regularizing the policy, we propose to randomly drop the units of a feature map that share the same discrete code together throughout training. Formally, for each state s_t, we sample K binary random variables $m_k \in \{0, 1\}$, $k = 1, 2, ..., K$ from a Bernoulli distribution with probability $1 - p$, where p is the drop probability. Then, we construct a mask M_t by utilizing the discrete representations e_t in (2) as follows:

$$M_t = (m_{q(t,1)}, m_{q(t,2)}, \cdots, m_{q(t,L)}), \quad \text{where } q(t, i) = \underset{k \in [K]}{\operatorname{argmin}} \|h_{t,i} - e_k\|_2. \tag{4}$$

By considering units of a feature map with the same discrete code, we remark that our method can effectively extract semantic objects from high-dimensional images.

Behavioral cloning with OREO. Now we propose to utilize our object-aware masking scheme for the regularization of a policy. To this end, we first initialize a convolutional encoder f with the parameters of a VQ-VAE encoder g. We empirically find that employing f as our backbone encoder for π instead of a fixed encoder g is more effective, as it allows an encoder to learn useful information for predicting actions. Then, we train the policy π by minimizing the following objective:

$$\mathcal{L}_{\texttt{OREO}}\left(s_t, a_t, M_t\right) = -\log \pi\left(a_t | f(s_t) \odot M_t\right), \tag{5}$$

where $\odot$ denotes elementwise product, and the mask M_t is shared across all channels in a feature map $f(s_t)$. Here, our intuition is that our object-aware regularization scheme should be useful for enforcing the policy not to exploit specific objects strongly correlated with expert actions, as the policy should utilize all semantic objects throughout training. Additionally, following Srivastava et al. [45], we scale the masked features by a factor of $1/(1 - p)$ during the training to ensure the scale of the expected output with masked features to match the scale of outputs at test time.

4 Experiments

In this section, we designed our experiments to answer the following questions:

- How does OREO compare to other regularization schemes that randomly drop units from a feature map [13, 45], data augmentation schemes [12, 49], and causality-based methods [11, 42] (see Table 1)?
- How does OREO compare to inverse reinforcement learning methods that learn a policy with environment interaction [8, 18] (see Figure 5)?
- Why is regularization necessary for addressing the causal confusion problem (see Figure 6a), and why is OREO effective for addressing this problem (see Figure 8)?
- Can OREO improve BC using various sizes of expert demonstrations (see Figure 7)?

Table 1: Performance of policies trained on various confounded Atari environments without environment interaction. OREO achieves the best score on 15 out of 27 environments, and the best median and mean human-normalized score (HNS) over all environments. The results for each environment report the mean of returns averaged over eight runs. We provide standard deviations in the supplementary material. CCIL† denotes the results without environment interaction.

Environment	BC	Dropout	DropBlock	Cutout	RandomShift	CCIL†	CRLR	OREO
Alien	954.1	1003.8	926.4	973.3	806.5	820.0	82.5	**1056.2**
Amidar	95.8	89.4	110.1	**118.7**	98.0	74.9	12.0	105.7
Assault	793.8	820.4	815.0	687.6	828.9	683.3	0.0	**840.9**
Asterix	292.2	313.8	345.4	212.4	135.5	643.2	**650.0**	180.8
BankHeist	442.1	485.7	508.4	486.1	367.2	**653.5**	0.0	493.9
BattleZone	11921.2	12457.5	12025.0	11107.5	9180.0	6370.0	1468.8	**12700.0**
Boxing	18.8	20.3	32.2	20.5	**38.3**	34.8	-43.0	36.4
Breakout	**5.7**	5.4	4.8	1.0	2.0	0.5	0.0	4.2
ChopperCommand	874.2	921.4	919.4	1016.1	936.4	760.6	**1077.2**	977.4
CrazyClimber	45372.9	39501.6	38345.6	44523.2	41924.0	22616.8	112.5	**55523.4**
DemonAttack	157.2	180.5	167.8	173.1	**241.8**	171.3	0.0	224.5
Enduro	241.4	250.4	341.8	119.6	316.4	143.1	3.9	**522.8**
Freeway	32.3	32.4	32.7	32.5	33.0	**33.1**	21.4	32.7
Frostbite	116.3	124.5	128.2	**139.4**	121.6	53.3	80.0	129.9
Gopher	1713.9	1819.1	1818.2	1481.0	1995.0	1404.5	0.0	**2515.0**
Hero	11923.1	14109.7	14711.4	14896.6	12816.0	6567.8	346.2	**15219.8**
Jamesbond	419.0	451.0	473.8	381.8	428.4	387.2	0.0	**502.8**
Kangaroo	2781.5	2912.9	3217.1	2824.0	1923.9	1670.5	122.8	**3700.2**
Krull	3634.3	3892.1	3832.1	3656.4	3788.7	3090.8	0.1	**4051.6**
KungFuMaster	15074.8	14452.1	15753.0	11405.6	13389.9	13394.9	0.0	**18065.6**
MsPacman	1432.9	1733.1	1446.4	1711.0	1223.5	1084.2	105.3	**1898.4**
Pong	3.2	10.2	11.5	6.8	-0.1	-2.7	-21.0	**14.2**
PrivateEye	2681.8	2599.1	2720.6	2670.6	**3969.2**	305.3	-1000.0	3124.9
Qbert	5438.4	6469.0	6140.3	5748.6	3921.4	5138.0	125.0	**6966.4**
RoadRunner	18381.5	21470.9	22265.4	12417.1	16210.0	11834.1	1022.9	**24644.2**
Seaquest	454.4	471.3	486.8	330.1	**1016.8**	271.2	172.5	753.1
UpNDown	4221.1	4147.1	**4789.2**	4159.6	3880.2	2631.1	20.0	4577.9
Median HNS	44.1%	47.4%	49.8%	42.0%	47.6%	36.2%	-1.5%	**51.2%**
Mean HNS	73.2%	79.0%	91.7%	69.5%	88.1%	71.7%	-45.9%	**105.6%**

Environments and datasets. We evaluate OREO on 27 Atari environments [5], which are selected by following prior works [11, 44]. Following de Haan et al. [11], we consider *confounded* Atari environments, where images are augmented with previous actions (see Figure 4). We utilize a single frame as an input to a policy, to focus on the causal confusion problem from nuisance correlates in the current state[2]. In our experiments, we report two evaluation metrics: average score from environments and human-normalized score (HNS := $\frac{\text{Agent}_{\text{score}} - \text{Random}_{\text{score}}}{\text{Human}_{\text{score}} - \text{Random}_{\text{score}}}$), following Mnih et al. [27]. For expert demonstrations, we utilize DQN Replay dataset [1]. As this dataset consists of 50M transitions of each environment collected during the training of a DQN agent [27], we use the last N trajectories as expert demonstrations. We preprocess input images to grayscale images of $84 \times 84 \times 1$, by utilizing Dopamine library [9]. We provide more details in the supplementary material.

(a) Pong (b) Enduro (c) Seaquest

Figure 4: Confounded Atari environments with previous actions (white number in lower left).

Implementation. We use a single Nvidia P100 GPU and 8 CPU cores for each training run. The training time for OREO is 6 hours on the dataset of size 50000, compared to 3 hours for BC, which is because OREO additionally trains a VQ-VAE model. As for hyperparameter selection, we use the default hyperparameters from previous or similar works [30, 45], i.e., a drop probability of $p = 0.5$, a codebook size of $K = 512$, and a commitment cost of $\beta = 0.25$. We use the same hyperparameters across all environments. We report the results over 8 runs unless specified. Source code and more details on implementation are available in the supplementary material.

[2]We refer to Wen et al. [48] for the discussion on the causal confusion problem from stacking states.

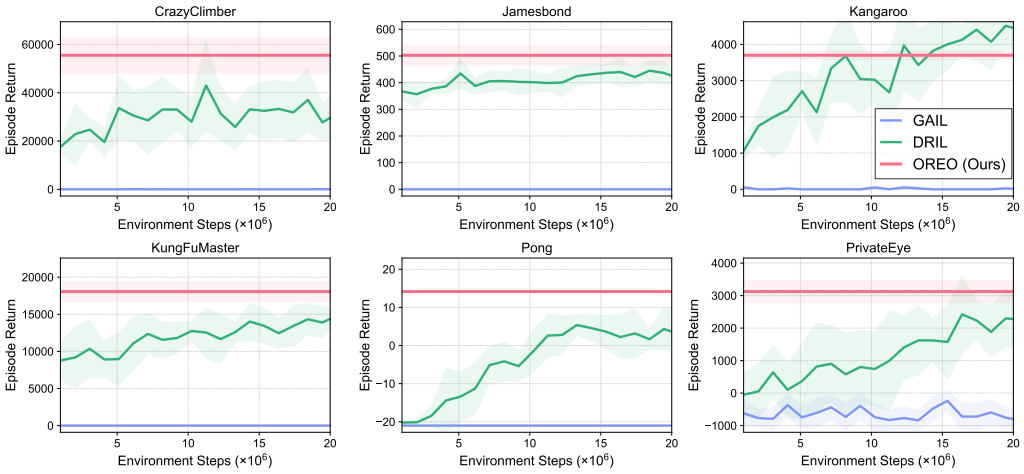

Figure 5: We compare OREO to inverse reinforcement learning methods that require environment interaction for learning a policy, on 6 confounded Atari environments. OREO outperforms baseline methods in most cases, even without using any interaction with environments. The solid line and shaded regions represent the mean and standard deviation, respectively, across eight runs.

Baselines. We consider BC as the most basic baseline method. To evaluate the effectiveness of our object-aware regularization scheme, we compare to regularization techniques that drop the randomly sampled units (i.e., Dropout [45]) or randomly sampled blocks (i.e., DropBlock [13]) from the feature map of a convolutional encoder. We also compare to data augmentation schemes, i.e., Cutout [12] that randomly masks out a square patch from images, and RandomShift [49] that randomly shifts pixels of images for regularization. We also consider the method of de Haan et al. [11] that learns a policy on top of disentangled representations from a β-VAE [17] (i.e., CCIL), and an observational causal inference method that estimates the causal contribution of each variable by confounder balancing (i.e., CRLR [42]). We provide the details for baselines in the supplementary material.

Comparative evaluation. Table 1 shows the performance of various methods that learn a policy without environment interaction. OREO significantly improves the performance of BC in most environments, outperforming other regularization techniques. In particular, OREO achieves the mean HNS of 105.6%, while the second-best method, i.e., DropBlock, achieves 91.7%. This demonstrates that our object-aware regularization scheme is indeed effective for addressing the causal confusion problem (see Figure 8 for qualitative experimental results). We found that CCIL without environment interaction does not exhibit strong performance in most environments, possibly due to the difficulty of learning disentangled representations from high-dimensional images [23]. We also provide experimental results for CCIL with environment interaction in the supplementary material, where the performance slightly improves but overall trends are similar. We observe that CRLR underperforms in most environments, which shows the difficulty of causal inference from high-dimensional images. We emphasize that OREO also outperforms other baselines in original Atari environments, which implies that our method is also effective for addressing the causal confusion that naturally occurs (see Figure 1). We provide experimental results for the original setup in the supplementary material.

Comparison with inverse reinforcement learning methods. To demonstrate that OREO can exhibit strong performance without environment interaction, we compare our method to inverse reinforcement learning (IRL) methods that first learn a reward function using expert demonstrations, and train a policy with environment interaction using learned reward function. Specifically, we consider GAIL [18], a method that learns reward function by discriminating expert states from on-policy states during environment interaction; and DRIL [8], a method that utilizes the disagreement between ensemble policies as a reward function. As shown in Figure 5, we observe that OREO exhibits superior performance to GAIL and DRIL on most confounded Atari environments, which are trained with 20M environment steps following the setup in Brantley et al. [8]. While IRL methods might outperform OREO asymptotically with more environment interaction, this result demonstrates that OREO indeed allows for achieving strong performance without interaction. We also found that GAIL exhibits almost zero performance in most environments, which is similar to the observation of previous works that GAIL suffers in environments with high-dimensional image inputs [8, 11, 36].

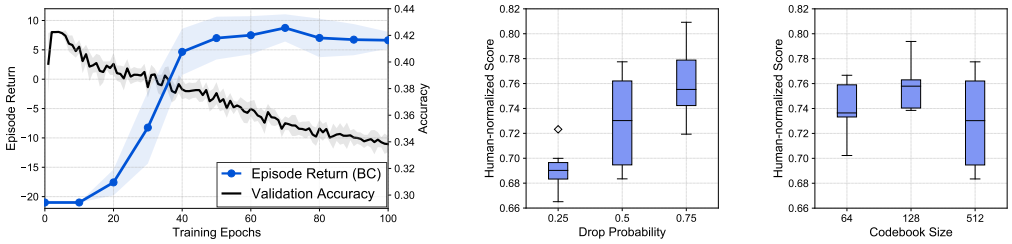

(a) Effects of validation accuracy (b) Drop probability (c) Codebook size

Figure 6: (a) Score and validation accuracy on confounded Pong environment, where validation accuracy is not aligned with the score at test time, which necessitates the use of regularization for addressing the causal confusion problem. We visualize the performance of OREO over 8 confounded Atari environments with varying (b) the drop probability of each code from a codebook and (c) codebook size. Boxplots are drawn using mean human-normalized scores obtained from eight runs.

Effects of expert demonstration size. To investigate the effectiveness of OREO with various sizes of expert demonstrations, we evaluate the performance of OREO with a varying number of expert demonstrations $N \in \{5, 10, 20, 35, 50\}$. Specifically, we report the mean HNS over 8 confounded Atari environments, which are randomly selected due to the high computation cost of running experiments for all environments. As shown in Figure 7, OREO consistently improves the performance of BC across a wide range of dataset sizes. As for the comparison with other baselines, OREO achieves superior performance to Dropout and DropBlock except for the extreme case of $N = 5$, which is because learning a VQ-VAE model with a limited number of data could be unstable. We also observe that DropBlock and OREO consistently outperform Dropout, which supports our intuition that dropping individual units from a feature map is not sufficient for effective regularization to address the causal confusion problem.

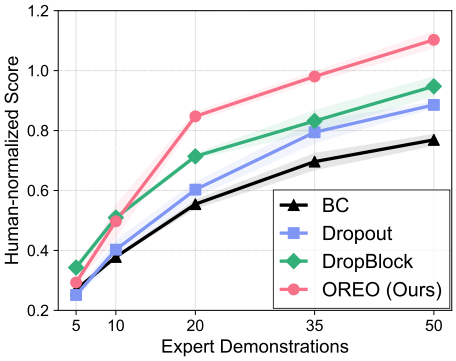

Figure 7: Mean human-normalized score over 8 confounded Atari environments with a varying number of expert demonstrations. The solid line and shaded regions represent the mean and standard deviation, respectively, across four runs.

Why is regularization necessary in confounded environments? A simple and widely used approach to address the overfitting problem in supervised learning is a model selection with a validation dataset. To see how this works in our setup, we first introduce a validation dataset consisting of 5 expert demonstrations on confounded Pong environment, and visualize the scores and validation accuracies measured by a policy learned with BC in Figure 6a. We observe that this simple scheme is not helpful for confounded Atari environments, i.e., validation accuracy is not aligned with the score at test time, because the distribution of the validation dataset could be significantly different from the distribution induced by a learned policy. As the evaluation of the policy in environments during training could be dangerous or even impossible, this result implies that regularizing the policy is necessary for successful imitation learning in confounded environments.

How does OREO improve the performance of BC? To understand how OREO improves the performance of BC, we visualize spatial attention maps from the last convolutional layer of a policy encoder in Figure 8. Specifically, following prior works [20, 51], we compute a spatial attention map by averaging the absolute values of a feature map along the channel dimension. We then apply 2-dimensional spatial softmax and multiply the upscaled attention map with images for visualization. We observe that the activations from the encoder trained with OREO are capturing all important objects of the environment (e.g., car that an agent controls and approaching cars), while the activations from BC are missing information of approaching cars by only focusing on the small region around a car and a scoreboard. This shows that our regularization scheme that encourages a policy to uniformly attend to all semantic objects allows for learning the policy that attends to important objects.

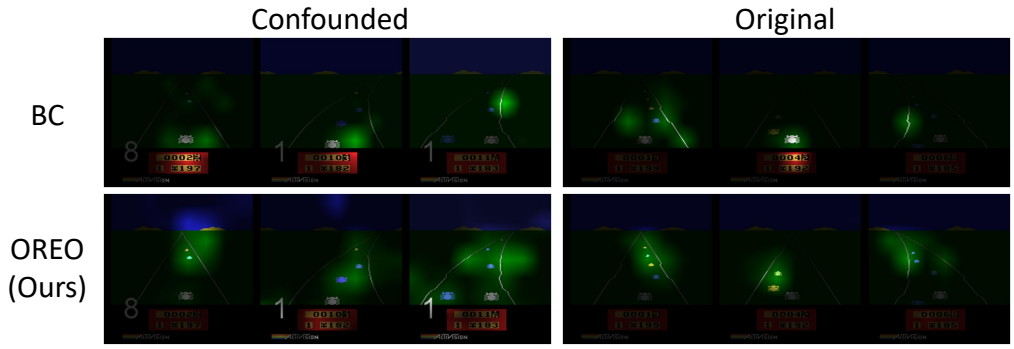

Figure 8: We visualize the spatial attention map from a convolutional encoder trained with BC and OREO in confounded (left) and original (right) Enduro environment. We observe that the encoder trained with OREO attends to important objects from images (e.g., approaching cars), while the encoder trained with BC only attends the small region around a car.

Effects of hyperparameters. We investigate the effect of two major hyperparameters, i.e., $p \in \{0.25, 0.5, 0.75\}$ for the drop probability in (4), and $K \in \{64, 128, 512\}$ for the codebook size in (2). Figure 6b shows that the performance improves as p increases, which implies that more strong regularization is effective for addressing the causal confusion problem on confounded Atari environments. Figure 6c shows that too small or large codebook size could be harmful to the performance. We remark that our experiments used default hyperparameters $p = 0.5$ and $K = 512$ for reporting the results, so the performance of OREO could be further improved with more tuning.

5 Discussion

In this paper, we present OREO, a simple regularization method to address the causal confusion problem in imitation learning. OREO regularizes a policy in an object-aware manner, by randomly dropping the units of a feature map that share the same discrete codes from a VQ-VAE model. Our experimental results demonstrate that OREO improves the performance of behavioral cloning without costly environment interaction, which is crucial for safe and successful imitation learning.

Limitations. While our method effectively regularizes a policy on confounded Atari environments, our work might be limited in that we only focused on the causal confusion problem from nuisance correlates in the current state. An interesting future direction would be to apply our regularization scheme to address the problem from a sequence of observations [4, 48]. We expect that OREO can also be effective for the setup, as our method could regularize the policy consistently over multiple timesteps, e.g., OREO could extract semantic objects with the same discrete code from consecutive images. Another limitation is that we do not deduce the cause-effect relations for addressing the causal confusion problem, but instead regularizes the policy to prevent it from exploiting nuisance correlates. However, given that it is an open problem to infer the structured disentangled variables and discover the causal relations among the variables [41], we believe encouraging the policy to attend to *all* semantic objects is a reasonable and promising direction for addressing this problem.

Potential negative impacts. Real-world applications of behavioral cloning, e.g., autonomous driving [4], require a large amount of data that often contain sensitive information, therefore raising privacy concerns. As our method is built upon a variational autoencoder, it could be exposed to privacy violation attacks that infer training data information, such as model inversion [52], and membership inference [15]. For example, the facial information of pedestrians may be reconstructed via membership inference attack. To address this vulnerability to privacy violation attacks, a differentially private variational autoencoder would be required for real applications. Also, a behavioral cloning policy will imitate whatever demonstrations one specifies. If some bad actions are included in expert demonstrations, the policy would perform dangerous actions to users. For these reasons, in addition to developing algorithms for better performance, it is also important to consider safe adaptation.

References

[1] Agarwal, Rishabh, Schuurmans, Dale, and Norouzi, Mohammad. An optimistic perspective on offline reinforcement learning. In *International Conference on Machine Learning*, 2020.

[2] Aytar, Yusuf, Pfaff, Tobias, Budden, David, Paine, Tom Le, Wang, Ziyu, and de Freitas, Nando. Playing hard exploration games by watching youtube. In *Advances in Neural Information Processing Systems*, 2018.

[3] Bain, Michael and Sammut, Claude. A framework for behavioural cloning. In *Machine Intelligence 15*, 1996.

[4] Bansal, Mayank, Krizhevsky, Alex, and Ogale, Abhijit. Chauffeurnet: Learning to drive by imitating the best and synthesizing the worst. In *Proceedings of Robotics: Science and Systems*, 2019.

[5] Bellemare, Marc G, Naddaf, Yavar, Veness, Joel, and Bowling, Michael. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, 2013.

[6] Bengio, Yoshua, Deleu, Tristan, Rahaman, Nasim, Ke, Rosemary, Lachapelle, Sébastien, Bilaniuk, Olexa, Goyal, Anirudh, and Pal, Christopher. A meta-transfer objective for learning to disentangle causal mechanisms. In *International Conference on Learning Representations*, 2020.

[7] Bojarski, Mariusz, Del Testa, Davide, Dworakowski, Daniel, Firner, Bernhard, Flepp, Beat, Goyal, Prasoon, Jackel, Lawrence D, Monfort, Mathew, Muller, Urs, Zhang, Jiakai, et al. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016.

[8] Brantley, Kiante, Sun, Wen, and Henaff, Mikael. Disagreement-regularized imitation learning. In *International Conference on Learning Representations*, 2020.

[9] Castro, Pablo Samuel, Moitra, Subhodeep, Gelada, Carles, Kumar, Saurabh, and Bellemare, Marc G. Dopamine: A research framework for deep reinforcement learning. *arXiv preprint arXiv:1812.06110*, 2018.

[10] Codevilla, Felipe, Santana, Eder, López, Antonio M, and Gaidon, Adrien. Exploring the limitations of behavior cloning for autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019.

[11] de Haan, Pim, Jayaraman, Dinesh, and Levine, Sergey. Causal confusion in imitation learning. In *Advances in Neural Information Processing Systems*, 2019.

[12] DeVries, Terrance and Taylor, Graham W. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017.

[13] Ghiasi, Golnaz, Lin, Tsung-Yi, and Le, Quoc V. Dropblock: A regularization method for convolutional networks. In *Advances in Neural Information Processing Systems*, 2018.

[14] Goyal, Anirudh, Lamb, Alex, Hoffmann, Jordan, Sodhani, Shagun, Levine, Sergey, Bengio, Yoshua, and Schölkopf, Bernhard. Recurrent independent mechanisms. *arXiv preprint arXiv:1909.10893*, 2019.

[15] Hayes, Jamie, Melis, Luca, Danezis, George, and De Cristofaro, Emiliano. Logan: Membership inference attacks against generative models. In *Proceedings on Privacy Enhancing Technologies*, 2017.

[16] Hester, Todd, Vecerik, Matej, Pietquin, Olivier, Lanctot, Marc, Schaul, Tom, Piot, Bilal, Horgan, Dan, Quan, John, Sendonaris, Andrew, Osband, Ian, et al. Deep q-learning from demonstrations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018.

[17] Higgins, Irina, Matthey, Loic, Pal, Arka, Burgess, Christopher, Glorot, Xavier, Botvinick, Matthew, Mohamed, Shakir, and Lerchner, Alexander. beta-vae: Learning basic visual concepts with a constrained variational framework. In *International Conference on Learning Representations*, 2017.

[18] Ho, Jonathan and Ermon, Stefano. Generative adversarial imitation learning. In *Advances in Neural Information Processing Systems*, 2016.

[19] Hussein, Ahmed, Gaber, Mohamed Medhat, Elyan, Eyad, and Jayne, Chrisina. Imitation learning: A survey of learning methods. *ACM Computing Surveys (CSUR)*, 50(2):1–35, 2017.

[20] Laskin, Michael, Lee, Kimin, Stooke, Adam, Pinto, Lerrel, Abbeel, Pieter, and Srinivas, Aravind. Reinforcement learning with augmented data. In *Advances in Neural Information Processing Systems*, 2020.

[21] Le, Hoang M, Yue, Yisong, Carr, Peter, and Lucey, Patrick. Coordinated multi-agent imitation learning. In *International Conference on Machine Learning*, 2017.

[22] Le, Thuc Duy, Hoang, Tao, Li, Jiuyong, Liu, Lin, Liu, Huawen, and Hu, Shu. A fast pc algorithm for high dimensional causal discovery with multi-core pcs. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 16(5):1483–1495, 2016.

[23] Locatello, Francesco, Bauer, Stefan, Lucic, Mario, Raetsch, Gunnar, Gelly, Sylvain, Schölkopf, Bernhard, and Bachem, Olivier. Challenging common assumptions in the unsupervised learning of disentangled representations. In *International Conference on Machine Learning*, 2019.

[24] Lopez-Paz, David, Nishihara, Robert, Chintala, Soumith, Scholkopf, Bernhard, and Bottou, Léon. Discovering causal signals in images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017.

[25] Maaten, Laurens van der and Hinton, Geoffrey. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(Nov):2579–2605, 2008.

[26] Mahler, Jeffrey and Goldberg, Ken. Learning deep policies for robot bin picking by simulating robust grasping sequences. In *Conference on Robot Learning*, 2017.

[27] Mnih, Volodymyr, Kavukcuoglu, Koray, Silver, David, Rusu, Andrei A, Veness, Joel, Bellemare, Marc G, Graves, Alex, Riedmiller, Martin, Fidjeland, Andreas K, Ostrovski, Georg, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.

[28] Muller, Urs, Ben, Jan, Cosatto, Eric, Flepp, Beat, and Cun, Yann L. Off-road obstacle avoidance through end-to-end learning. In *Advances in Neural Information Processing Systems*, 2006.

[29] Ng, Andrew Y, Russell, Stuart J, et al. Algorithms for inverse reinforcement learning. In *International Conference on Machine Learning*, 2000.

[30] Oord, Aaron van den, Vinyals, Oriol, and Kavukcuoglu, Koray. Neural discrete representation learning. In *Advances in Neural Information Processing Systems*, 2017.

[31] Parascandolo, Giambattista, Kilbertus, Niki, Rojas-Carulla, Mateo, and Schölkopf, Bernhard. Learning independent causal mechanisms. In *International Conference on Machine Learning*, 2018.

[32] Pathak, Deepak, Mahmoudieh, Parsa, Luo, Guanghao, Agrawal, Pulkit, Chen, Dian, Shentu, Yide, Shelhamer, Evan, Malik, Jitendra, Efros, Alexei A, and Darrell, Trevor. Zero-shot visual imitation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2018.

[33] Pearl, Judea. *Causality*. Cambridge University Press, 2009.

[34] Pohlen, Tobias, Piot, Bilal, Hester, Todd, Azar, Mohammad Gheshlaghi, Horgan, Dan, Budden, David, Barth-Maron, Gabriel, Van Hasselt, Hado, Quan, John, Večerík, Mel, et al. Observe and look further: Achieving consistent performance on atari. *arXiv preprint arXiv:1805.11593*, 2018.

[35] Pomerleau, Dean. Alvinn: An autonomous land vehicle in a neural network. In *Advances in Neural Information Processing Systems*, 1989.

[36] Reddy, Siddharth, Dragan, Anca D, and Levine, Sergey. Sqil: Imitation learning via reinforcement learning with sparse rewards. In *International Conference on Learning Representations*, 2020.

[37] Ross, Stephane and Bagnell, J Andrew. Reinforcement and imitation learning via interactive no-regret learning. In *Advances in Neural Information Processing Systems*, 2014.

[38] Ross, Stéphane, Gordon, Geoffrey, and Bagnell, Drew. A reduction of imitation learning and structured prediction to no-regret online learning. In *International Conference on Artificial Intelligence and Statistics*, 2011.

[39] Russell, Stuart. Learning agents for uncertain environments. In *Annual Conference on Computational Learning Theory*, 1998.

[40] Schaal, Stefan et al. Learning from demonstration. In *Advances in Neural Information Processing Systems*, 1997.

[41] Schölkopf, Bernhard. Causality for machine learning. *arXiv preprint arXiv:1911.10500*, 2019.

[42] Shen, Zheyan, Cui, Peng, Kuang, Kun, Li, Bo, and Chen, Peixuan. Causally regularized learning with agnostic data selection bias. In *ACM international conference on Multimedia*, 2018.

[43] Spirtes, Peter, Glymour, Clark N, Scheines, Richard, and Heckerman, David. *Causation, prediction, and search*. MIT Press, 2000.

[44] Srinivas, Aravind, Laskin, Michael, and Abbeel, Pieter. Curl: Contrastive unsupervised representations for reinforcement learning. In *International Conference on Machine Learning*, 2020.

[45] Srivastava, Nitish, Hinton, Geoffrey, Krizhevsky, Alex, Sutskever, Ilya, and Salakhutdinov, Ruslan. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.

[46] Sun, Wen, Venkatraman, Arun, Gordon, Geoffrey J, Boots, Byron, and Bagnell, J Andrew. Deeply aggrevated: Differentiable imitation learning for sequential prediction. In *International Conference on Machine Learning*, 2017.

[47] Sutton, Richard S and Barto, Andrew G. *Reinforcement learning: An introduction*. MIT Press, 2018.

[48] Wen, Chuan, Lin, Jierui, Darrell, Trevor, Jayaraman, Dinesh, and Gao, Yang. Fighting copycat agents in behavioral cloning from observation histories. In *Advances in Neural Information Processing Systems*, 2020.

[49] Yarats, Denis, Kostrikov, Ilya, and Fergus, Rob. Image augmentation is all you need: Regularizing deep reinforcement learning from pixels. In *International Conference on Learning Representations*, 2021.

[50] Ye, Gu and Alterovitz, Ron. Guided motion planning. In *Robotics Research*, pp. 291–307. Springer, 2017.

[51] Zagoruyko, Sergey and Komodakis, Nikos. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. In *International Conference on Learning Representations*, 2017.

[52] Zhang, Yuheng, Jia, Ruoxi, Pei, Hengzhi, Wang, Wenxiao, Li, Bo, and Song, Dawn. The secret revealer: Generative model-inversion attacks against deep neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.

[53] Ziebart, Brian D, Maas, Andrew L, Bagnell, J Andrew, and Dey, Anind K. Maximum entropy inverse reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2008.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See Section 5.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] See Section 5.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] See Section 4, and the supplementary material.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Section 4 and the supplementary material.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See Section 4 and the supplementary material.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Section 4 and the supplementary material.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]