CAUSAL CURIOSITY: RL AGENTS DISCOVERING SELF-SUPERVISED EXPERIMENTS FOR CAUSAL REPRESENTATION LEARNING

Anonymous authors
Paper under double-blind review

ABSTRACT

Humans show an innate ability to learn the regularities of the world through interaction. By performing experiments in our environment, we are able to discern the causal factors of variation and infer how they affect the dynamics of our world. Analogously, here we attempt to equip reinforcement learning agents with the ability to perform experiments that facilitate a categorization of the rolled-out trajectories, and to subsequently infer the causal factors of the environment in a hierarchical manner. We introduce a novel intrinsic reward, called causal curiosity, and show that it allows our agents to learn optimal sequences of actions, and to discover causal factors in the dynamics. The learned behavior allows the agent to infer a binary quantized representation for the ground-truth causal factors in every environment. Additionally, we find that these experimental behaviors are semantically meaningful (e.g., to differentiate between heavy and light blocks, our agents learn to lift them), and are learnt in a self-supervised manner with approximately 2.5 times less data than conventional supervised planners. We show that these behaviors can be re-purposed and fine-tuned (e.g., from lifting to pushing or other downstream tasks). Finally, we show that the knowledge of causal factor representations aids zero-shot learning for more complex tasks.

1 INTRODUCTION

Discovering causation in environments an agent might encounter remains an open and challenging problem for causal reinforcement learning (Schölkopf (2015), Bengio et al. (2013), Schölkopf (2019)). Most approaches take the form of BAMDPs (Bayes Adaptive Markov Decision Processes) (Zintgraf et al. (2019)) or Hi-Param MDP (Hidden Parameter MDPs) (Doshi-Velez & Konidaris (2016); Yao et al. (2018); Killian et al. (2017); Perez et al. (2020)) which condition the transition $p(s_{t+1}|s_t, a_t; H)$ and/or reward function $R(r_{t+1}|s_t, a_t, s_{t+1}; H)$ of each environment on hidden parameters (also referred to as causal factors in some of the above studies). Let $s \in \mathcal{S}$, $a \in \mathcal{A}$, $r \in \mathcal{R}$, $H \in \mathcal{H}$ where $\mathcal{S}$, $\mathcal{A}$, $\mathcal{R}$, and $\mathcal{H}$ are the set of states, actions, rewards and feasible hyper-parameters. In the physical world and in the case of mechanical systems, examples of the parameter H include gravity, coefficients of friction, masses and sizes of objects. Typically, H is treated as a latent variable for which an embedding is learned during training, using variational methods (Kingma et al. (2014); Ilse et al. (2019)). Let $s_{0:T}$ be the entire state trajectory of length T. Similarly, $a_{0:T}$ is the sequence of actions applied during that trajectory by the agent that results in $s_{0:T}$. In an environment represented by z, these latent variable approaches define a probability distribution over the entire sequence of (rewards, states, actions) as $p(r_{0:T}, s_{0:T}, a_{0:T-1}; z)$ that factorizes as

$$\prod_{i=1}^{T-1} p(r_{t+1}|s_t, a_t, s_{t+1}, z)p(s_{t+1}|s_t, a_t, z)p(a_t|s_t, z) \tag{1}$$

due to the Markov assumption. At test time, the agent infers the causal factor associated with its environment by observing the trajectories produced by its initial actions that can be issued by any policy such as model-based reinforcement learning.

In practice, however, discovering causal factors in a physical environment is prone to various challenges that are caused by the disjointed nature of the influence of these factors on the produced

trajectories. More specifically, at each time step, the transition function is affected by a subset of global causal factors. This subset is implicitly defined on the basis of the current state and the action taken. For example, if a body in an environment loses contact with the ground, the coefficient of friction between the body and the ground no longer affects the outcome of any action that is taken. Likewise, the outcome of an upward force applied by the agent to a body on the ground is unaffected by the friction coefficient. We can therefore take advantage of this natural discontinuity to discern causal factors.

Without knowledge of how independent causal mechanisms affect the outcome of a particular action in a given state in an environment, it becomes impossible for the agent to conclude where the variation it encountered came from. Unsurprisingly, Hi-Param and BAMDP approaches fail to learn a disentangled embedding for the causal factors, making their behaviors uninterpretable (Perez et al. (2020)). For example, if, in an environment, a body remains stationary under a particular force, the Hi-Param or BAMDP agent may apply a higher force to achieve its goal of perhaps moving the body, but will be unable to conclude whether the "un-movability" was caused by high friction or high mass of the body. Additionally, these approaches require human-supervised reward engineering, making it difficult to apply them outside of the simulated environments they are tested in.

Our goal is, instead of focusing on maximizing reward for some particular task, to allow agents to discover causal processes through exploratory interaction. During training, our agents discover self-supervised experimental behaviors which they apply to a set of training environments. These behaviors allow them to learn about the various causal mechanisms that govern the transitions in each environment. During inference in a novel environment, they perform these discovered behaviors sequentially and use the outcome of each behavior to infer the embedding for a single causal factor (Figure 1).

The main challenge while learning a disentangled representation for the causal factors of the world is that several causal factors may affect the outcome of behaviors in each environment. For example, when pushing a body on the ground, the outcome, i.e., whether the body moves, or how far the body is pushed, depends on several factors, e.g., mass, shape and size, frictional coefficients, etc. However, if, instead of pushing on the ground, the agent executes a perfect grasp-and-lift behavior, only mass will affect whether the body is lifted off the ground or not.

Thus, it is clear that not all experimental behaviors are created equal and that the outcomes of some behaviors are caused by fewer causal factors than others. Our agents learn these behaviors without supervision using *causal curiosity*, an intrinsic reward. The outcomes of these experimental behaviors are then used to infer a binary quantized embedding describing the single isolated causal factor. Even though causal factors of variation in a physical world are easily identifiable to humans, a concrete definition is required to back up our proposed method. We conjecture that the causality of a factor of variation depends on the available actions to the agent. If the set of actions that an agent can take is very limited, there is no way for it to discern a diverse set of causal factors in the environment.

Definition 1 (Causal factors)**.** *Consider the MDP ($\mathcal{S}$, $\mathcal{A}$, p, r) with state space $\mathcal{S}$, action space $\mathcal{A}$, the transition function p, and the reward function r. Let T be the length of the trajetcories during each rollout and $s_{0:T} \in \mathcal{S}^T$ denotes a trajectory. Let $d(\cdot,\cdot) : \mathcal{S}^T \times \mathcal{S}^T \to \mathbb{R}_+$ be a distance function defined on the space of trajectories of length T. The set $H = \{h_1, h_2, \ldots, h_k\}$ is called a set of $\epsilon-$causal factors if for every $h_i \in H$, there exists a sequence of actions that clusters the state trajectories into two sets S and S' such that*

$$\min\{d(s_{0:T}, s'_{0:T}) : s_{0:T} \in S, s'_{0:T} \in S'\} > \epsilon \tag{2}$$

Intuitively, a factor of variation affecting a set of environments is called causal if there exists a sequence of actions available to the agent where the resultant trajectories are clustered into two or more sets (for simplicitly here we assume binary clusters). This is analogous to the human ability to conclude whether objects are heavy or light, big or small.

Using the above, we propose an intrinsic reward, which allows our agents to discover experimental behaviors which are semantically meaningful and can be used to re-train for downstream tasks, resulting in high sample efficiency. Our work, therefore, forms an important link between structured representation learning and skill discovery, two largely disjoint fields in RL, which stand to benefit from each other.

The contributions of the work are as follows:

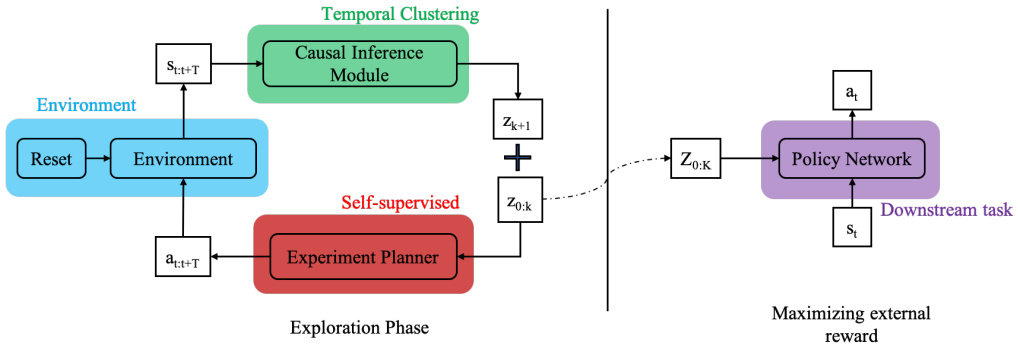

Figure 1: Overview of Inference. The exploration loop produces a series of K experiments allowing the agent to infer the representations for K causal factors. After exploration, the agent utilizes the acquired knowledge for downstream tasks.

- We equip agents with the ability to perform experiments and behave meaningfully in a set of environments in an unsupervised manner. These behaviors can expose or obfuscate specific independent causal mechanisms that occur in the world of the agent, allowing the agent to learn about each in the absence of the others, an important human behavioral trait.
- We introduce an intrinsic reward, *causal curiosity*, which allows our agents to discover these behaviors without human-engineered rewards. The outcomes of the experiments are used to learn a disentangled quantized binary representation for the causal factors of the environment, analogous to the human ability to conclude whether objects are light/heavy, big/small etc.
- Through extensive experiments, we conclude that knowledge of the causal factors aids sample efficiency in two ways - first, that the knowledge of the causal factors aids transfer learning across multiple environments, and, second, that the experimental behaviors acquired can be repurposed for downstream tasks.

2 METHOD

Consider a set of N environments $\mathcal{E}$ where the ith such environment $e^{(i)} \in \mathcal{E}$ is parameterized by a K-dimensional vector of causal factors $H^{(i)}$. The jth entry of $H^{(i)}$, $H^{(i)}_{(j)}$ represents the scalar value of the jth causal factor of the ith MDP. For each environment $e^{(i)}$, $(z^{(i)}_{(0)}, z^{(i)}_{(1)} ... z^{(i)}_{(K-1)})$ represents the disentangled embedding vector, such that $z^{(i)}_{(j)}$ encodes $H^{(i)}_{(j)}$.

2.1 LEARNING TO PERFORM EXPERIMENTS

To learn about causal processes through interaction, the agent must produce a sequence of actions $a_{0:T-1}$ that we call *experimental behavior*, which, when applied to environment $e^{(i)} \in \mathcal{E}$, produces a sequence of states (observation) $O^{(i)} = \{s^{(i)}_0, s^{(i)}_1 .. s^{(i)}_T\}$, which is then used to infer the value of the embedding for a single causal factor $z^{(i)}_{(j)}$.

According to Def. 1, a causal factor is a parameter in the environment which, when intervened on over a set of values, results in trajectories of states that are divisible into disjoint clusters under a particular sequence of actions. These clusters represent the quantized values of the causal factor. For example, mass, which is a causal factor of a body, under an action sequence of a grasping and lifting motion, results in 2 clusters, liftable and not liftable. However, such an action sequence is not known in advance. Therefore, discovering a causal factor in the environment boils down to finding a sequence of actions that makes the effect of that factor prominent by producing clustered trajectories for different values of that environmental factor. More concretely, the agent solves the following

optimization problem:

$$\arg\max_{a_{0:T}\in\mathcal{A}^T}[\min\{d(s_{0:T}, s'_{0:T}) : s_{0:T} \in S, s'_{0:T} \in S'\} - \quad \{d(s_{0:T}, s''_{0:T}) : s''_{0:T}, s_{0:T} \in S\} - \\ \{d(s'_{0:T}, s'''_{0:T}) : s'_{0:T}, s'''_{0:T} \in S'\}] \tag{3}$$

where the distance function $d(\cdot,\cdot)$ in the space of trajectories is set to be Soft Dynamic Time Warping (Cuturi & Blondel (2017)). Intuitively, this optimization finds a sequence of actions that maximizes the inter-cluster distance between two clusters of trajectories while minimizing the intra-cluster distance within each of them. The experiment planner produces a plan over a horizon of 6 control steps (approximately 2 seconds). This sequence of actions is applied to the training set of environments. The set of observations is binarily clustered and the Silhouette Score (Rousseeuw (1987)) of the clustering is returned to the planner as reward.

This procedure can also be motivated as a model selection criterion. Among the many explanations for the outcome of an interaction of the agent with an environment, the simplest model is preferred. If the description length of the environment is taken as a proxy to the complexity, Minimum Description Length can be seen as an information criterion for the underlying dynamic environment as suggested in (Mehrjou et al., 2018).

Normally in model selection applications, the observations are fixed and the goal is to find a model M^* that is closest to reality, as represented by:

$$M^* = \arg\min_{M}(L(M) + L(O|M)) \tag{4}$$

where $L(\cdot)$ is the description length. However, here, the situation is reversed. A simple bi-modal clustering model is fixed, motivated by Definition 1. Then, the agent is motivated to produce actions that result in observations that are best explained by this model. These discovered action sequences are the *experimental behaviors* we desire. Hence eq. (3) becomes

$$a^*_{0:T} = \arg\min_{a_{0:T}}(L(M) + L(O|M)) \tag{5}$$

where each observed trajectory $O = O(a_{0:T})$ is a function of the action sequence. As mentioned earlier, the model is fixed in this formulation; hence, the first term $L(M)$ is constant and not a function of the actions. The cost function in eq. (3) can actually be regarded as an approximation to the abstract function $-L(O|M)$ that is fed back to the RL agent as a reward function to maximize. We regard this reward function as *causal curiosity*.

Note that the since each causal factor has its own independent causal mechanism that causes O, the MDL of O will be higher if multiple causal factors cause O. On the contrary, if the agent produces actions which result in an O that is easily explained by a low-capacity bi-modal model M, then it will imply that O is caused by fewer causal factors. Consequently, the causal curiosity reward for such an action sequence, $-L(O|M)$, will be high. Therefore, causal curiosity favors experimental behaviors that result in observations caused by few causal factors - thereby allowing us to use O to infer a representation for a single causal factor.

2.2 Interventions on beliefs

By maximizing the causal curiosity reward it is possible to achieve behaviors which result in trajectories of states only caused by a single hidden parameter. However, we wish to learn these experimental behaviors for each of the hidden parameters that may vary in an environment. Optimizing the causal curiosity reward results in an implicit selection of the hidden parameter which causes the outcome of the experiment - the hidden parameter which maximizes the causal curiosity reward. Thus, to acquire a diverse set of experimental behaviors that allow the agent to learn about all of the independent mechanisms in its world, we perform an intervention on the embedding of the causal factor isolated by the initial optimization of causal curiosity. Under an intervened value of the embedding for the isolated causal factor, the agent can subsequently learn about the mechanisms produced in the environment by the remaining causal factors. For example, having learned to distinguish heavy from light objects, the agent can then use all heavy objects to learn about additional causal factors such as size, shape, etc.

To achieve this, we recursively perform binary clustering on each of the clusters obtained by optimizing (3). This process continues in the form of a tree (Figure 4), where for each cluster of

environments, a new experiment learns to split the cluster into 2 sub-clusters depending on the value of another hidden parameter. At level n, the agent produces 2^n experiments and inference models, having already intervened on the binary quantized representations of n causal factors.

3 Related Work

Doshi-Velez & Konidaris (2016) define a class Markov Decision Processes where transition probabilities $p(s_{t+1}|s_t, a_t; \theta)$ depend on a hidden parameter θ, whose value is not observed, but its effects are felt. Killian et al. (2017) and Yao et al. (2018) utilize these Hidden Parameter MDPs (Markov Decision Processes) to enable efficient policy transfer, assuming that transition probabilities across states are a function of hidden parameters. Perez et al. (2020) relax this assumption, allowing both transition probabilities and reward functions to be functions of hidden parameters. Zintgraf et al. (2019) approach the problem from a Bayes-optimal policy standpoint, defining transition probabilities and reward functions to be dependent on a hidden parameter characteristic of the MDP in consideration. We utilize this setup to define causal factors.
Substantial attempts have been made at unsupervised disentanglement, most notably, the β-VAE Higgins et al. Burgess et al. (2018), where a combination of factored priors and the information bottleneck force disentangled representations. Kim & Mnih (2018) enforce explicit factorization of the prior without compromising on the mutual information between the data and latent variables, a shortcoming of the β-VAE. Chen et al. (2018) factor the KL divergence into a more explicit form, highlighting an improved objective function and a classifier-agnostic disentanglement metric. Locatello et al. (2018) show theoretically that unsupervised disentanglement (in the absence of inductive biases) is impossible and highly unstable, susceptible to random seed values. They follow this up with Locatello et al. (2020) where they show, both theoretically and experimentally, that pair-wise images provide sufficient inductive bias to disentangle causal factors of variation. However, these works have been applied to supervised learning problems whereas we attempt to disentangle the effects of hidden variables in dynamical environments, a relatively untouched question.
Curiosity for robotics is not a new area of research. Schmidhuber (2006), Ngo et al. (2012), Pathak et al. (2017) describe curiosity as the motivation behind the behavior of an agent in an environment for which the outcome is unpredictable, i.e., an intrinsic reward that motivates the agent to explore the unseen portions of the state space (and subsequent transitions). While causal curiosity is an intrinsic reward, it differs from these traditional definitions of curiosity in that it motivates the agent to produce structure in the outcome of its behavior.

4 Experiments

Our work has 2 main thrusts - the discovered experimental behaviors and the representations obtained from the outcome of the behaviors in environments. We visualize these learnt behaviors and verify that they are indeed semantically meaningful and interpretable. We quantify the utility of the learned behaviors by using the behaviors as pre-training for a downstream task. We verify that these behaviors are indeed invariant to all other causal factors except one.
We visualize the representations obtained using these behaviors and verify that they are indeed the binary quantized representations for each of the ground truth causal factors that we manipulated in our experiments. Finally, we verify that the knowledge of the representation does indeed aid transfer learning and zero-shot generalizability in downstream tasks.

Causal World We use the Causal World Simulation (Ahmed et al. (Under submission 2020)) based on the Pybullet Physics engine to test our approach. The simulator consists of a 3-fingered robot, with 3 joints on each finger. We constrain each environment to consist of a single object that the agent can interact with. The causal factors that we manipulate for each of the objects are size, shape and mass of the blocks. The simulator allows us to capture and track the positions and velocities of each of the movable objects in an environment. While, for most experiments, the 3D position and 3D pose of the blocks is used as the state at each time step, we perform ablation studies where less information is provided to the agent.

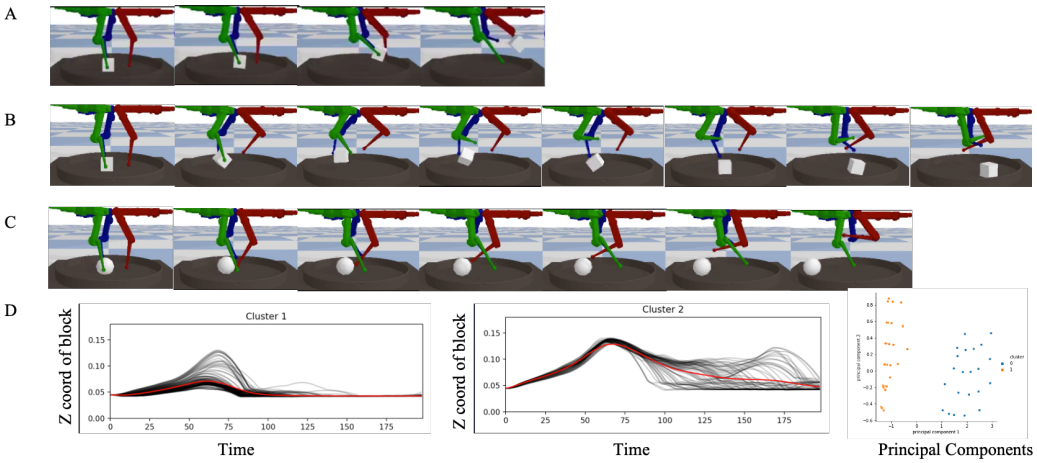

Figure 2: Examples of discovered behaviors. The agent discovers experimental behaviors that allow it to characterize each environmental object in a binary manner, e.g., heavy/light, big small, rollable/not rollable, etc. These behaviors are acquired without any external supervision by maximizing the causal curiosity reward. **A, B, C** correspond to self-discovered toss, lift-and-spin and roll behaviors respectively. **D** shows an ablation study, where the agent is only provided the z coordinate of the block in every environment. Each line corresponds to one environment and the z coordinate of the block is plotted with time when the discovered behavior is applied. It learns a lifting behavior, where cluster 1 represents the heavy blocks (z coordinate does not change much) and cluster 2 represents the light blocks (z increases as block is lifted and then falls when dropped and subsequently increases again when it bounces).

4.1 Visualizing Discovered Behaviors

We would like to analyze whether the discovered experimental behaviors are human interpretable, i.e., *are the experimental behaviors discovered in each of the setups semantically meaningful?* We find that our agents learn to perform several useful behaviors without any supervision. For instance, to differentiate between objects with varying mass, we find that they acquire a perfect grasp-and-lift behavior with an upward force. In other random seed experiments, the agents learn to lift the blocks by using the wall of the environment for support. To differentiate between cubes and spheres, the agent discovers a pushing behavior which gently rolls the spheres along a horizontal direction. Qualitatively, we find that these behaviors are stable and predictable. See videos of discovered behaviors `here` (website under construction).

Concurrent with the objective they are trained on, we find that the acquired behaviors impose structure on the outcome when applied to each of the training environments. The outcome of each experimental behavior on the set of training environments results in dividing it into 2 subsets. These subsets correspond to the binary quantized values of a single factor, e.g., large or small, while being invariant to the values of other causal factors of the environments. We also perform ablation studies where instead of providing the full state vector, we provide only one coordinate (e.g., only x, y or z coordinate of the block). We find that causal curiosity results in behaviors that differentiate the environments based on outcomes along the direction provided. For example, when only the x coordinate was provided, the agent learned to evaluate mass by applying a pushing behavior along the x direction. Similarly, a lifting behavior was obtained when only the z coordinate was supplied to the curiosity module (Figure 2).

4.2 Utility of Learned Behaviors for Downstream Tasks

While the behaviors acquired are semantically meaningful, we would like to quantify their utility as pre-training for downstream tasks. We analyze the performance on `Lifting` where the agent must grasp and lift a block to a predetermined height and `Travel`, where the agent must impart a velocity to the block along a predetermined direction. We re-train the learnt planner using an external reward for these tasks. We implement a baseline vanilla Cross Entropy Method optimized Model Predictive Control Planner (De Boer et al. (2005)) trained using the identical reward function

and compare the rewards per trajectory during training. We find high zero-shot generalizability and quicker convergence as compared to the vanilla CEM planner (Figure 3).

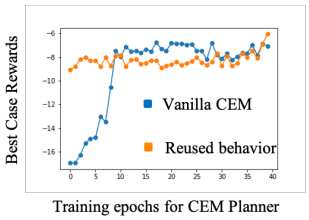

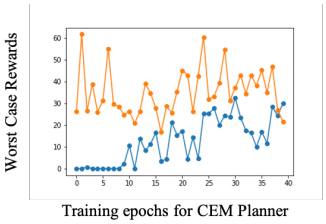

Discovered Experiments as Pre-training	
Method	**Time Steps of training**
Vanilla CEM Planner	$15x10^5$
Reused Behavior	$6x10^5$

Figure 3: Utility of discovered behaviors. We find that the behaviors discovered by the agents while optimizing causal curiosity show high zero-shot generalizability and converge to the same performance as conventional planners for downstream tasks. We also analyze the worst case performance and find that the pre-training ensures better performance than random initialization. We compare the time-steps required by both methods to achieve the zero-shot reward of the reused behavior.

4.3 Visualization of hierarchical binary latent space

Our agents discover a disentangled latent space such that they are able to isolate the sources of causation of the variability they encounters in their environments. For every environment, they learn a disentangled embedding vector which describes each of the causal factors.

To show this, we use 3 separate experimental setups - `Mass`, `SizeMass` and `ShapeSizeMass` where each of the causal factors are allowed to vary over a range of discrete values. During `Mass`, the agent is allowed access to 6 environments with objects having the same shape (cuboids) and size but differing only in mass. During `SizeMass`, the agent has access to 24 environments with cuboids having masses and sizes ranging over 6 and 4 values respectively. Finally, during `ShapeSizeMass`, the agent has access to 48 environments with objects having masses, sizes and shapes ranging over 6, 4, and 2 values respectively.

During training, the agent discovers a hierarchical binary latent space (Figure 4), where each level of hierarchy corresponds to a single causal factor. The binary values at each level of hierarchy correspond to the high/low values of the causal factor in question. To our knowledge, we obtain the first interpretable latent space describing the various causal processes in the environment of an agent. This implies that it learns to quantify each physical attribute of the blocks it encounters in a completely unsupervised manner.

4.4 Knowledge of causal factors aids transfer

Next, we test whether knowledge of the causal factors does indeed aid transfer and zero-shot generalizability. To this end, we supply the representations obtained by the agent during the experimental behavior phase as input to a policy network in addition to the state of the simulator, and train it for a place-and-orient downstream task (Figure 1). We define 2 experimental setups - `Mass` and `SizeMass`. In `Mass`, the agent is given access to 10 environments, with 10 varying values of mass. In `SizeMass`, the agent is allowed access to 10 environments, with 5 and 2 values of mass and size respectively. In both setups, the agent learns about the varying causal mechanisms by optimizing causal curiosity. Subsequently, using the causal representation along with the state for each environment, it is trained to maximize external reward.

After training, the agents are exposed to a set of unseen test environments, where we analyze their zero-shot generalizability. These test environments consist of unseen masses and sizes and their unseen combinations. This corresponds to "Strong Generalization" as defined by Perez et al. (2020). We report results averaged over 10 random seeds.

For each setup, we train a PPO-optimized Actor-Critic Policy (referred to as **Causally-curious agent**) with access to the causal representations and a 56 dimensional state vector from the environment (thus, a total of 57 dimensional input for `Mass`, and a 58 dimensional for `SizeMass`). Similar to Perez et al. (2020), we implement 2 baselines - the **Generalist** and the **Specialist**. The **Specialist**

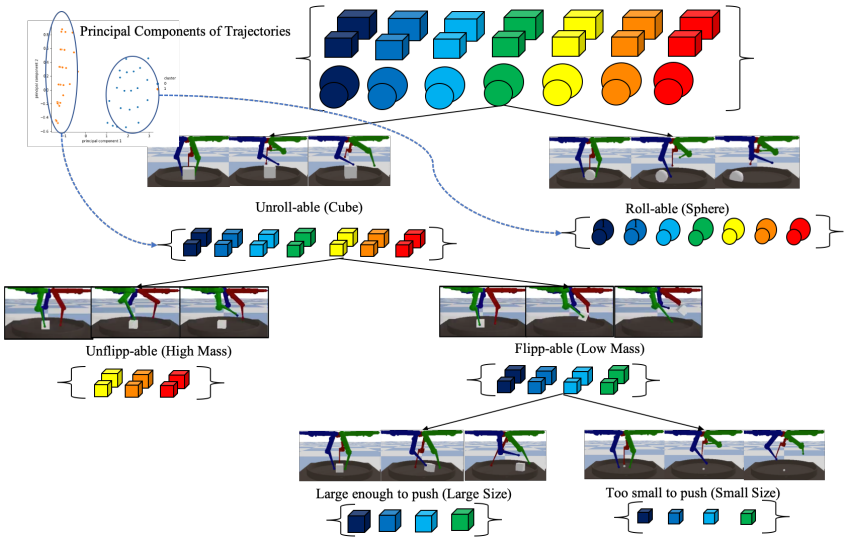

Figure 4: Discovered hierarchical latent space. The agent learns experiments that differentiate the full set of blocks in `ShapeSizeMass` into hierarchical binary clusters. At each level, the environments are divided into 2 clusters on the basis of the value of a single causal factor. We also show the principal components of the trajectories in the top left.

consists of an agent with identical architecture as **Causally-curious agent**, but without access to causal representations (i.e., receives a 56 dimensional state vector). It is initialized randomly and is trained only on the test environments, serving as a benchmark for complexity of the test tasks. It performs poorly, indicating that the test tasks are complex. The architecture of the **Generalist** is identical to the **Specialist**. Like the **Specialist**, the **Generalist** also does not have access to the causal representations, but is trained on the same set of training environments that the Causally-curious agent is trained on. The poor performance of the generalist indicates that the tasks distribution of training and test tasks differs significantly and that memorization of behaviors does not yield good transfer. We find that causally-curious agents significantly outperform the both baselines indicating that indeed, knowledge of the causal representation does aid zero-shot generalizability.

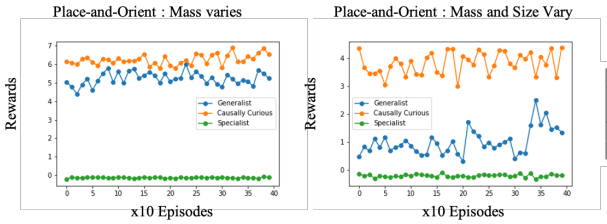

Extra time steps needed by baseline to match CC performance	
Method	**Extra time steps of training**
Generalist - Mass	4.8×10^5
Generalist - Mass and Size	8.5×10^5

Figure 5: Knowledge of causal factors aids transfer. We find that knowledge of the causal representation allows agents to generalize to unseen environments with high zero-shot performance. We also study the extra timesteps required to match the zero-shot performance of causally-curious agent.

5 Conclusion

We introduce *causal curiosity*, an intrinsic reward that allows agents to discover binary quantized representations for the causal factors that affect environments an RL agent may encounter. We show that optimizing causal curiosity rewards results in the agent performing self-supervised experiments. We find that these experiments happen to be semantically meaningful and can be used as pre-training for downstream tasks. While our work learns binary quantized causal representations, a dense encoding may improve the amount of encoded information about the causal mechanisms of the environments. We leave this to future work.

REFERENCES

Ossama Ahmed, Frederik Träuble, Anirudh Goyal, Alexander Neitz, Manuel Wüthrich, Yoshua Bengio, Bernhard Schölkopf, and Stefan Bauer. Causalworld: A robotic manipulation benchmark for causal structure and transfer learning. Under submission 2020.

Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.

Christopher P Burgess, Irina Higgins, Arka Pal, Loic Matthey, Nick Watters, Guillaume Desjardins, and Alexander Lerchner. Understanding disentangling in β-vae. arxiv 2018. *arXiv preprint arXiv:1804.03599*, 2018.

Tian Qi Chen, Xuechen Li, Roger B Grosse, and David K Duvenaud. Isolating sources of disentanglement in variational autoencoders. In *Advances in Neural Information Processing Systems*, pp. 2610–2620, 2018.

Marco Cuturi and Mathieu Blondel. Soft-dtw: a differentiable loss function for time-series. *arXiv preprint arXiv:1703.01541*, 2017.

Pieter-Tjerk De Boer, Dirk P Kroese, Shie Mannor, and Reuven Y Rubinstein. A tutorial on the cross-entropy method. *Annals of operations research*, 134(1):19–67, 2005.

Finale Doshi-Velez and George Konidaris. Hidden parameter markov decision processes: A semiparametric regression approach for discovering latent task parametrizations. In *IJCAI: proceedings of the conference*, volume 2016, pp. 1432. NIH Public Access, 2016.

Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework.

Maximilian Ilse, Jakub M Tomczak, Christos Louizos, and Max Welling. Diva: Domain invariant variational autoencoders. *arXiv preprint arXiv:1905.10427*, 2019.

Taylor W Killian, Samuel Daulton, George Konidaris, and Finale Doshi-Velez. Robust and efficient transfer learning with hidden parameter markov decision processes. In *Advances in neural information processing systems*, pp. 6250–6261, 2017.

Hyunjik Kim and Andriy Mnih. Disentangling by factorising. *arXiv preprint arXiv:1802.05983*, 2018.

Durk P Kingma, Shakir Mohamed, Danilo Jimenez Rezende, and Max Welling. Semi-supervised learning with deep generative models. In *Advances in neural information processing systems*, pp. 3581–3589, 2014.

Francesco Locatello, Stefan Bauer, Mario Lucic, Gunnar Rätsch, Sylvain Gelly, Bernhard Schölkopf, and Olivier Bachem. Challenging common assumptions in the unsupervised learning of disentangled representations. *arXiv preprint arXiv:1811.12359*, 2018.

Francesco Locatello, Ben Poole, Gunnar Rätsch, Bernhard Schölkopf, Olivier Bachem, and Michael Tschannen. Weakly-supervised disentanglement without compromises. *arXiv preprint arXiv:2002.02886*, 2020.

Arash Mehrjou, Friedrich Solowjow, Sebastian Trimpe, and Bernhard Schölkopf. A local information criterion for dynamical systems. *arXiv preprint arXiv:1805.10615*, 2018.

Hung Ngo, Matthew Luciw, Alexander Forster, and Juergen Schmidhuber. Learning skills from play: artificial curiosity on a katana robot arm. In *The 2012 international joint conference on neural networks (IJCNN)*, pp. 1–8. IEEE, 2012.

Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. Curiosity-driven exploration by self-supervised prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 16–17, 2017.

Christian F Perez, Felipe Petroski Such, and Theofanis Karaletsos. Generalized hidden parameter mdps: Transferable model-based rl in a handful of trials. *AAAI Conference On Artifical Intelligence*, 2020.

Peter J Rousseeuw. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 20:53–65, 1987.

Jürgen Schmidhuber. Developmental robotics, optimal artificial curiosity, creativity, music, and the fine arts. *Connection Science*, 18(2):173–187, 2006.

B. Schölkopf. Artificial intelligence: Learning to see and act (News & Views). *Nature*, 518(7540): 486–487, 2015.

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

Jiayu Yao, Taylor Killian, George Konidaris, and Finale Doshi-Velez. Direct policy transfer via hidden parameter markov decision processes. In *LLARLA Workshop, FAIM*, volume 2018, 2018.

Luisa Zintgraf, Kyriacos Shiarlis, Maximilian Igl, Sebastian Schulze, Yarin Gal, Katja Hofmann, and Shimon Whiteson. Varibad: A very good method for bayes-adaptive deep rl via meta-learning. *arXiv preprint arXiv:1910.08348*, 2019.

A Appendix

Acknowledgments

This work was supported by C-BRIC (one of six centers in JUMP, a Semiconductor Research Corporation (SRC) program sponsored by DARPA), and the Intel and CISCO Corporations. The authors affirm that the views expressed herein are solely their own, and do not represent the views of the United States government or any agency thereof.