Structure Mapping for Transferability of Causal Models

Purva Pruthi¹ Javier González^{*2} Xiaoyu Lu³ Madalina Fiterau¹

Abstract

Human beings learn causal models and constantly use them to transfer knowledge between similar environments. We use this intuition to design a transfer-learning framework using objectoriented representations to learn the causal relationships between objects. A learned causal dynamics model can be used to transfer between variants of an environment with exchangeable perceptual features among objects but with the same underlying causal dynamics. We adapt continuous optimization for structure learning techniques (Zheng et al., 2018) to explicitly learn the cause and effects of the actions in an interactive environment and transfer to the target domain by categorization of the objects based on causal knowledge. We demonstrate the advantages of our approach in a gridworld setting by combining causal modelbased approach with model-free approach in reinforcement learning.

1. Introduction

In reinforcement learning, two kind of approaches are common - *model-free* learning and *model-based* learning (Sutton & Barto, 1999). In model-free learning, the agent learns to directly estimate the future reward for different states without building an explicit model of the system while in model-based learning the agent explicitly learns the state transition and reward distributions of the system. In modelbased learning, the model can be prediction-based, where the goal is to accurately predict the future state or it can be causal, where the goal is to explicitly model the variables (causes) responsible for the state transition and rewards. Experimental work in cognitive science has shown that humans and animals use a combination of model-based and model-free algorithms.(Dolan & Dayan, 2013)

Building explicit causal models for the systems provide the advantage of explicit understanding of the dynamics of the system and re-using the causal knowledge to transfer to the variations of the system. For example, while learning a new video game, we first learn about different objects within the game, their perceptual features (e.g. color, shape, material etc.), causal characteristics (like what will happen if you hit a wall) and causal relationships between different objects (e.g. switching on the switch causes bulb to light.) Now, if we are presented with a new variant of the game with different looking objects but similar causal characteristics then we tend to transfer knowledge by mapping the objects seen in the previous game with the objects in the new game based on similar causal behavior. On the other hand, if we are presented with a new variant of the game with similar looking objects but with different causal characteristics, then we need to augment the causal model for those objects with new causal characteristics.

In either case, multiple characterization of the objects in terms of their perceptual features as well as causal properties provide several advantages. First, this characterization allows the flexibility to decouple the perceptual representation of an object from their causal properties, allowing combinatorial generalization over systems with different combination of these features. Second, it allows to build abstract causal knowledge in the form of high-level causal concepts. (e.g. Hitting a wall causes no movement even if the wall were red, blue or black in color). [(Kemp et al., 2010)].

Object-oriented MDP (Diuk et al., 2008) provides an efficient way of modeling an environment as it factorizes the state space over different objects and thus, reducing the large state-space. Also, it provides an abstraction by attributing the effect of the objects to the object's characteristics rather than it's specific location in the state space, which is the case in pixel-based representation (assuming the model is trained on images).

In this work, we combine ideas from object-oriented MDPs (Diuk et al., 2008), structure learning (Zheng et al., 2018), and categorization as causal reasoning (Rehder, 2003), (Kemp et al., 2010). Our main contribution can be summarized as below:

^{*}Work done while working at Amazon, Cambridge, UK. ¹College of Information and Computer Sciences, University of Massachusetts Amherst, USA ²Microsoft Research, Cambridge, United Kingdom ³Amazon, Cambridge, United Kingdom . Correspondence to: Purva Pruthi cpruthi@cs.umass.edu>.

Proceedings of the 37^{th} International Conference on Machine Learning, Vienna, Austria, PMLR 108, 2020. Copyright 2020 by the author(s).

- We formulate and evaluate the learning of the state transition model in OO-MDP as a parametric structure learning problem using continuous optimization of score-based structure learning techniques.
- We formulate and evaluate the causal dynamics model combined with model-free approach on a toy gridworld problem as a proof-of-concept.
- We formulate transfer learning as a structure mapping problem where we learn mapping from objects in the source domain to the target domain based on the causal structure.



Figure 1. High-level schematic diagram to illustrate the overall approach for transfer learning. (a) Object based segmentation of source environment consisting of an agent (in yellow), two keys (in blue-green) and two locks (in red). Each object is described by it's attributes {x-position (O.x), y-position (O.y), color (O.c)} attributes (b) Causal Dynamics model of source environment for taking action 'UP'. (c) Target environment with same rules as source environment but different (inverted) colors of keys and locks. (d) Causal Dynamics model of target environment for taking action 'UP'.

2. Preliminaries and Problem Statement

2.1. Notation

The traditional formalism for the reinforcement learning problem is the Markov Decision Process (MDP). An MDP is a five-tuple $(S, A, T, \mathcal{R}, \gamma)$, where S is a set of states, A is the set of actions, $T(s^{(t+1)}|s^{(t)}, a^{(t)})$ is the probability of transforming from state $s^{(t)} \in S$ to $s^{(t+1)} \in S$ after action $a^{(t)} \in A$, $\mathcal{R}(r^{(t+1)}|s^{(t)}, a^{(t)})$ is the probability of receiving the reward $r^{(t+1)} \in \mathbb{R}$ after executing action $a^{(t)}$ while in state $s^{(t)}$, and $\gamma \in [0, 1]$ is the rate at which future rewards are exponentially discounted. In object-oriented MDP, we assume that the state consists of N objects (or entities) and each object has M attributes. O_i refers to the i^{th} object and let $\alpha_{i,j}^{(t)}$ refer to the j^{th} attribute value of the i^{th} object at time t. $O_i^{(t)} = (\alpha_{i,1}^{(t)}, \ldots, \alpha_{i,M}^{(t)})$ refers to the state of the i^{th} object at time t. The complete state of the MDP at time t is $s^{(t)} = (O_1^{(t)} \ldots O_N^{(t)})$. Example of an attribute α_{ij} is $\{color, x-position, y-position$ etc}. Attributes which remain constant over time in a given environment characterize the object. (e.g. color characterizes locks, keys, agent and walls in a particular environment.)

2.2. Problem Statement

Let's assume that there are two *deterministic* environments, source S and target T. In this work, we assume that the two environments have same causal dynamics model but different perceptual features. In S, agent can experiment with the environment *more freely* by taking random actions while it is *more expensive* to experiment in T. Our goal is to transfer knowledge from S to T with minimum possible interventions in T and learn an optimal policy for T.

2.3. Toy Environment: Triggers

We provide a proof-of-concept for our approach using a toy-environment (known as Triggers), first introduced in (Ferret et al., 2019). In this gridworld based environment, the agent can take one of the four possible actions {north, south, east or west}. Hitting black walls doesn't change the agent's position and moving in the free space increases or decreases the agent's x or y position by 1, depending on the action. In source environment (Figure 1(a)), green colored boxes represent the keys and red-colored boxes represent the locks. The agent should collect all the keys before it can attempt opening one of the locks. If the agent attempts to open the lock when at least one of the keys is still present, the agent receives a negative reward of -1. It receives +1 reward for opening each lock successfully. The key boxes disappear when collected and the lock boxes disappear when opened successfully. In the target environment, (Figure 1(d)), we change the colors of the keys and locks. The goal is to transfer the knowledge from source to target by mapping the causal behavior of the objects (keys and doors) with corresponding colors (perceptual features).

3. Our approach

Our approach mainly consists of three steps which are also summarized in Figure 1. We assume that the object-oriented representation of the environment is learned from input images using unsupervised visual-object reasoning modules (Eslami et al., 2016) and is already provided to us. We learn the causal dynamics model M^S of the source environment from the agent's interactions in the source environment using a random policy and use the dynamics model for planning in S. The causal dynamics model consists of the causal structure between the object attributes at the current time-step and the next time-step, parametrized by Φ . In the target environment, we start with the assumption that it follows the same dynamics model M^S . We simulate the initial trajectories for T_0 timesteps and also observe the state transition and reward by actually experimenting in the target environment. This will lead to the inaccurate prediction using M^S when interacting with some set of objects $O_1^T \dots O_s^T$. We learn a mapping from attributes α'_{ij} of such objects in the target environment O^T to attributes α_{ii} of objects in source environment O^S such that M^S is invariant in T.

3.1. Structure Learning in Object Oriented-MDPs

The problem of structure learning in OO-MDPs can be formulated as learning a directed acyclic graph $G^a = (V, E)$ with object attributes as vertices $V \in \mathbb{R}^{N \times M}$ and $E \in$ $\mathbb{R}^{NM \times \tilde{N}M}$. We assume that the dependencies encoded by DAG represent direct causal dependencies. This assumption is reasonable in our context because we discover the causal graph by performing interventions (random policy) in the environment rather than by using observational data. Note, that in OO-MDPs each G^a is specific to each value $a \in$ \mathcal{A} of an action because different actions might result in different causes for the state transitions, resulting in different graphs. For example, if an agent takes an 'UP' action, then the only attributes corresponding to the object in the 'UP' direction might be valid causes for the state transition of the agent. Due to the determinstic and stationary nature of the environment, we obtained one fixed graph G^a for each action a.

(Zheng et al., 2018) (Zheng et al., 2020) introduced an approach (NOTEARS) in which structure learning can be formulated as a continuous optimization problem, making the score-based structure learning approaches efficient for high-dimensional input variables.

Let us represent the current state and next state of attributes as $X = \{\alpha_{11}^{(t)}, \alpha_{12}^{(t)} \dots \alpha_{NM}^{(t)}, \alpha_{11}^{(t+1)}, \alpha_{12}^{(t+1)}, \dots \alpha_{NM}^{(t+1)}\}$ with dimension d = 2NM, Causal generative process of each attribute X_j can be written as: $X_j = f_j(Pa(X_j))$. We assume that the attributes at time t + 1 cannot be parents of attributes at time t. Though, we allow the edges between the attributes corresponding to same timestamp (e.g. $\alpha_{ip}^{(t)}$ and $\alpha_{jq}^{(t)}$) which are interpreted as relational dependencies rather than a causal dependency (e.g. spatial relationship between x-position of the up neighbor and x-position of the agent). Assuming agent as reference, dependence between x-y position of neighboring objects and agent can be encoded as a directed edge.

Our goal is to learn a DAG $G^a(\Phi)$ by learning $f = (f_1, f_2, f_3 \dots f_d)$. We assume that f_j is a function approximated by neural networks using the parameters Φ_j . Thus, the optimization problem becomes

$$\min_{\Phi} L(\Phi) = \frac{1}{n} \sum_{j=1}^{d} l(x_j, f(X, \Phi_j))$$

subject to

$$h(W(\Phi)) = 0$$

where $W(f) = W(f_1, f_2 \dots f_d) \in \mathbb{R}^{d \times d}$ encodes the graph edges $G^a(\Phi)$, i.e. $[w(f)]_{kj} = \|\frac{\partial f_j}{\partial X_k}\|_{L_2} = 0$ if $X_k \notin Pa(X_j)$ and $h(W(\Phi)) = tr(e^{W \cdot W}) - d$ encodes the acyclicity constraint. Refer (Zheng et al., 2018) and (Zheng et al., 2020) for more details.

Previously, (Kansky et al., 2017) have used a similar approach called Schema Networks for zero-shot learning by building an explicit causal generative model. Our approach to learn causal structure using NOTEARS provides several advantages compared to the integer-programming based structure learning used in Schema Networks. First, the binary representation for attributes used by the schema networks limits the scope of their approach as binary representation makes the approach computationally challenging for high-dimensional attributes. NOTEARS with parametric DAGs, on the other hand allows using continuous and discrete variables for the structure learning. Second, NOTEARS with object-oriented representation can learn state-space abstraction by capturing relational dependencies. For example, if two objects interact when they are adjacent to each other, then the interaction between attributes of different objects can be learned in the functional form of $f_k(j)$ regardless of the absolute positions of the objects.

3.2. Structure Mapping from Source to Target

We describe our structure mapping approach in Algorithm 1. Intuitively, there are two ways to do structure mapping, either by updating $G(\Phi)$ to $G(\Theta)$ or by mapping the source object attributes to target object attributes and categorize objects in the same group if causal behaviors are the same. The first approach requires learning the functional relationships Θ from scratch, assuming the graph structure remains the same. If we assume underlying causal dynamics remain the same, then we can adopt the second approach where the goal is to learn a mapping from source objects to target objects. (See Algorithm 1)

Algorithm 1 Structure Mapping from $G(\Phi)$ to $G(\Theta)$

Input: Causal Dynamics Model $G^{a}(\Phi)$, Optimal Qvalues $Q^*(S, A)$ from source environment repeat $S \leftarrow \text{Current state in object-oriented representation.}$ $A \leftarrow \epsilon - greedy(S,Q)$ Execute action A, observe next state S' and reward R $R^G, S'^G \leftarrow G^a(S, \Phi)$ if $R \neq R^G$ OR $S' \neq S'^G$ then $O^{A}.attr \leftarrow O^{A}.color$ (Pick attribute which characterize object) Identify O^S such that $\Theta^{reward}(O^A.attr)$ $\Phi^{reward}(O^S.attr) = R \text{ and } \Theta^{state}(O^A.attr) =$ $\Phi^{state}(O^S.attr) = S'$ Categorize O^S and O^A in same group and map attribute $O^A.attr$ to $O^S.attr$ end if until T_0 steps

4. Experimental Evaluation

4.1. Structure Learning

We trained the structure learning algorithm on the source Trigger environment with different size of the state space (by varying grid's width and height from 5 to 75) and different numbers and locations of the keys and locks. Input X is 19-dimensional vector containing (x-position, y-position, color) attributes of agent, four neighbors and reward, and agent's x,y positions at t+1. We also include num_keys as the variable which represents the number of keys present in the environment at each time step. In Figure 2, we observe



Figure 2. L_2 norm $\|\Phi_j(X)\|_{L_2}$ of Weight Values for learned bayesian graph G^{up} for action 'UP'. Attributes on y-axis are potential parents (causes) of the corresponding j_{th} attribute on x-axis (X_j) , with parameters Φ_j . Positive values of weights indicate edge between parents and children.

that the attribute future reward r^t has total number of keys (num_keys) and color of the 'up' neighbor $(up.c^t)$ as the parents, as expected. For agent's next positions x^{t+1} and y^{t+1} , it's current x and y positions and color of the 'up' neighbor are the causes. We also see edges between x and y positions of the neighbors and the agent due to spatial dependency between them.

4.2. Agent's Performance in Source Environment

To evaluate the efficiency of learned causal dynamics model, we use the model to do planning alongwith exploration in the source environment. We use random shooting planning algorithm (Nagabandi et al., 2017) to select the optimal action sequence. Figure 3 compares the performance of trained DQN agent and DQN agent alongwith online causal planning agent (Causal Model +DQN). Exploration parameter ϵ is reduced to 0.05 in DQN after planning steps due to which less than 50% less random actions were taken by combined model-free and modelbased approach. We observe that combined agent converges faster than the purely model-free agent, with less variance. (See Appendix for more details.) Code required to reproduce our experiments is available online at https://github.com/Information-Fusion -Lab-Umass/causal_transfer_learning.



Figure 3. Planning using Causal-Dynamics Model

5. Discussion and Future Work

In this work, we built an explicit causal dynamics model for deterministic and discrete environments with the goal to transfer the knowledge in domains with similar dynamics. We observe that this is a flexible and interepretable way to model perceptually changing environments. We are in the process of evaluating our structure mapping on Triggers and extending it to real-world environments. The main challenge is in learning object-oriented representation for objects in an unsupervised manner for complex domains with continuous action and state spaces. Limited intervention setting and use of largely available observational data for structure learning are some interesting future directions for our work.

References

- Diuk, C., Cohen, A., and Littman, M. L. An object-oriented representation for efficient reinforcement learning. In *Proceedings of the 25th international conference on Machine learning*, pp. 240–247, 2008.
- Dolan, R. J. and Dayan, P. Goals and habits in the brain. *Neuron*, 80(2):312–325, 2013.
- Eslami, S. A., Heess, N., Weber, T., Tassa, Y., Szepesvari, D., Hinton, G. E., et al. Attend, infer, repeat: Fast scene understanding with generative models. In *Advances in Neural Information Processing Systems*, pp. 3225–3233, 2016.
- Ferret, J., Marinier, R., Geist, M., and Pietquin, O. Selfattentional credit assignment for transfer in reinforcement learning, 2019.
- Kansky, K., Silver, T., Mély, D. A., Eldawy, M., Lázaro-Gredilla, M., Lou, X., Dorfman, N., Sidor, S., Phoenix, S., and George, D. Schema networks: Zero-shot transfer with a generative causal model of intuitive physics, 2017.
- Kemp, C., Goodman, N. D., and Tenenbaum, J. B. Learning to learn causal models. *Cognitive science*, 34 7:1185–243, 2010.
- Nagabandi, A., Kahn, G., Fearing, R. S., and Levine, S. Neural network dynamics for model-based deep reinforcement learning with model-free fine-tuning, 2017.
- Rehder, B. Categorization as causal reasoning. *Cognitive Science*, 27(5):709–748, 2003.
- Sutton, R. S. and Barto, A. G. Reinforcement learning. *Journal of Cognitive Neuroscience*, 11(1):126–134, 1999.
- Zheng, X., Aragam, B., Ravikumar, P., and Xing, E. P. DAGs with NO TEARS: Continuous Optimization for Structure Learning. In Advances in Neural Information Processing Systems, 2018.
- Zheng, X., Dan, C., Aragam, B., Ravikumar, P., and Xing, E. P. Learning sparse nonparametric DAGs. In *International Conference on Artificial Intelligence and Statistics*, 2020.