

Object-Oriented Dynamics Learning through Multi-Level Abstraction

Guangxiang Zhu,^{1*} Jianhao Wang,^{1*} Zhizhou Ren,^{1*} Zichuan Lin,² Chongjie Zhang¹

¹Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing, China

²Department of Computer Science and Technology, Tsinghua University, Beijing, China
{zhu-gx15, wjh19, rzz16, linzc16}@mails.tsinghua.edu.cn, chongjie@tsinghua.edu.cn

Abstract

Object-based approaches for learning action-conditioned dynamics has demonstrated promise for generalization and interpretability. However, existing approaches suffer from structural limitations and optimization difficulties for common environments with multiple dynamic objects. In this paper, we present a novel self-supervised learning framework, called *Multi-level Abstraction Object-oriented Predictor* (MAOP), which employs a three-level learning architecture that enables efficient object-based dynamics learning from raw visual observations. We also design a spatial-temporal relational reasoning mechanism for MAOP to support instance-level dynamics learning and handle partial observability. Our results show that MAOP significantly outperforms previous methods in terms of sample efficiency and generalization over novel environments for learning environment models. We also demonstrate that learned dynamics models enable efficient planning in unseen environments, comparable to true environment models. In addition, MAOP learns semantically and visually interpretable disentangled representations.

Introduction

Model-based deep reinforcement learning (DRL) has recently attracted much attention for improving sample efficiency of DRL (Gu et al. 2016; Racanière et al. 2017; Finn and Levine 2017). One of the core problems for model-based DRL is to learn action-conditioned dynamics models through interacting with environments. Pixel-based approaches have been proposed for such dynamics learning from raw visual perception, achieving remarkable performance in training environments (Oh et al. 2015; Chiappa et al. 2017).

To unlock sample efficiency of model-based DRL, learning action-conditioned dynamics models that generalize over unseen environments is critical yet challenging. Finn, Goodfellow, and Levine (2016) proposed a dynamics learning method that takes a step towards generalization over object appearances. Zhu, Huang, and Zhang (2018) developed an object-oriented dynamics predictor to support generalization. However, due to structural limitations and optimization difficulties, these methods do not efficiently generalize over

environments with multiple controllable and uncontrollable dynamic objects and different static object layouts.

To address these limitations, we propose a novel three-level learning framework for self-supervised learning of object-oriented dynamics model, called *Multi-level Abstraction Object-oriented Predictor* (MAOP). This framework simultaneously learns disentangled object representations and predicts object motions conditioned on their historical states, their interactions to other objects, and an agent’s actions. To reduce the complexity of such concurrent learning and improve sample efficiency, MAOP employs a three-level learning architecture from the most abstract level of motion detection, to dynamic instance segmentation, and to dynamics learning and prediction. A more abstract learning level solves an easier problem and has lower learning complexity, and its output provides a coarse-grained guidance for a less abstract learning level, improving its speed and quality of learning.

Specifically, we perform motion detection to detect proposal regions that potentially contain dynamic instances for the follow-up dynamic instance segmentation. Then we exploit spatial-temporal information of locomotion property and appearance patterns to capture coarse region proposals of dynamic instances. Finally we use them to guide the learning of the object representations and instance localization at the level of dynamics learning. This three-level architecture is inspired by humans’ multi-level motion perception from cognitive science studies (Johansson 1975; Lu and Sperling 1995) and multi-level abstraction search in constraint optimization (Zhang and Shah 2016). In addition, we design a novel CNN-based spatial-temporal relational reasoning mechanism for MAOP, which includes a Relation Net to reason about spatial relations between objects and an Inertia Net to learn temporal effects. This mechanism offers a disentangled way to handle physical reasoning in settings with partial observability.

Our results show that MAOP significantly outperforms previous methods for learning dynamics models in terms of sample efficiency and generalization over novel settings with multiple controllable and uncontrollable dynamic objects and different object layouts. MAOP enables model learning from few interactions with environments and accurately

*Equal contribution

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

predicting the dynamics of objects as well as raw visual observations in previously unseen environments. The learned dynamics model enables an agent to directly plan in unseen environments without retraining. In addition, MAOP learns disentangled representations and gains visually and semantically interpretable knowledge, including meaningful object masks, accurate object motions, disentangled relational reasoning process, and controllable factors. Last but not least, MAOP provides a general multi-level framework for learning object-based dynamics model from raw visual observations, offering opportunities to easily leverage well-studied object detection methods (e.g., Mask R-CNN (He et al. 2017)) in the area of computer vision.

Related Work

Object-oriented reinforcement learning has received much research attention, which exploits efficient representations based on objects and their interactions. This learning paradigm is close to that of human cognition in the physical world and the learned object-level knowledge can be efficiently generalized across environments. Early work on object-oriented RL requires explicit encodings of object representations, such as *relational MDPs* (Guestrin et al. 2003), *OO-MDPs* (Diuk et al. 2008), object focused q-learning (Cobo, Isbell, and Thomaz 2013), and Schema Networks (Kansky et al. 2017). In this paper, we present an end-to-end, self-supervised neural network framework that automatically learns object representations and dynamics conditioned on actions and object relations from raw visual observations.

Action-conditioned dynamics learning aims to address one of the core problems for model-based DRL, i.e., constructing an environment dynamics model. Several pixel-based approaches have been proposed for learning how an environment changes in response to actions through unsupervised video prediction and achieve remarkable performance in training environments (Oh et al. 2015; Chiappa et al. 2017). Fragkiadaki et al. (2016) propose an object-centric prediction method to learn the dynamics model when given the object localization and tracking. Finn, Goodfellow, and Levine (2016) propose an action-conditioned video prediction method that explicitly models pixel motion and learns invariance to object appearances. Recently, Zhu, Huang, and Zhang (2018) propose an object-oriented dynamics learning paradigm. However, it focuses on environments with a single dynamic object. In this paper, we take a further step towards object-oriented dynamics modeling in more general environments with multiple dynamic objects and also demonstrate its usage for model-based planning. In addition, we design an instance-aware dynamics mechanism to support instance-level dynamics learning and handle partial observations.

Relation-based deep learning approaches make significant progress in a wide range of domains such as physical reasoning (Chang et al. 2016; Battaglia et al. 2016; van Steenkiste et al. 2018), computer vision (Watters et al. 2017; Wu et al. 2017), natural language processing (Santoro et al. 2017), and reinforcement learning (Zambaldi et al. 2018; Zhu, Huang, and Zhang 2018). Relation-based nets introduce relational inductive biases into neural networks, which

facilitate generalization over entities and relations and enables relational reasoning (Battaglia et al. 2018). This paper proposes a novel spatial-temporal relational reasoning mechanism, which includes a CNN-based Inertia Net for learning temporal effects in addition to a CNN-based Relation Net for reasoning about spatial relations.

Instance Segmentation has been one of the fundamental problems in computer vision and many approaches have been proposed (Pinheiro, Collobert, and Dollár 2015; Li et al. 2017; He et al. 2017). Instance segmentation can be regarded as the combination of semantic segmentation and object localization. Most models are supervised learning and require a large labeled training dataset. Liu, He, and Gould (2015) proposes a weakly-supervised approach to infer object instances in foreground by exploiting dynamic consistency in video. In this paper, we design a self-supervised, three-level approach for learning dynamic rigid object instances. First, foreground detection computes region proposals for potential dynamic objects. Based on these region proposals, we then learn coarse dynamic instance segmentation. This coarse instance segmentation provides a guidance for learning accurate instances at the dynamics learning level, whose instance segmentation considers not only object appearances but also motion prediction conditioned on object-to-object relations and actions.

Multi-level Abstraction Object-Oriented Predictor (MAOP)

In this section, we will present a novel self-supervised deep learning framework, aiming to learn object-oriented dynamics models that are able to efficiently generalize over unseen environments with different object layouts and multiple dynamic objects. Such a generalized object-oriented dynamics learning approach requires simultaneously learning object representations and motions conditioned on their historical states, their interactions to other objects, and an agent’s actions. This concurrent learning is challenging for an end-to-end approach in complex environments. Evidences from cognitive science studies (Johansson 1975; Lu and Sperling 1995) show that human beings are born with prior motion perception ability (in Cortical area MT) of perceiving moving and motionlessness, which enables learning more complex knowledge, such as object-level dynamics prediction. Inspired by these studies, we design a multi-level learning framework, called *Multi-level Abstraction Object-oriented Predictor* (MAOP), which incorporates motion perception levels to assist in dynamics learning. Since we focus on the seminal study towards the multi-level framework for interpretable and efficient dynamics learning, we make a basic assumption that the environment only contains rigid objects, which has also been widely adopted by many papers (Watters et al. 2017; Wu et al. 2017; Zhu, Huang, and Zhang 2018).

Figure 1 illustrates three levels of MAOP framework: *dynamics learning*, *dynamic instance segmentation*, and *motion detection*. Here we present them from a top-down decomposition view. The *dynamics learning* level is an end-to-end, self-supervised neural network, aiming to learn object rep-

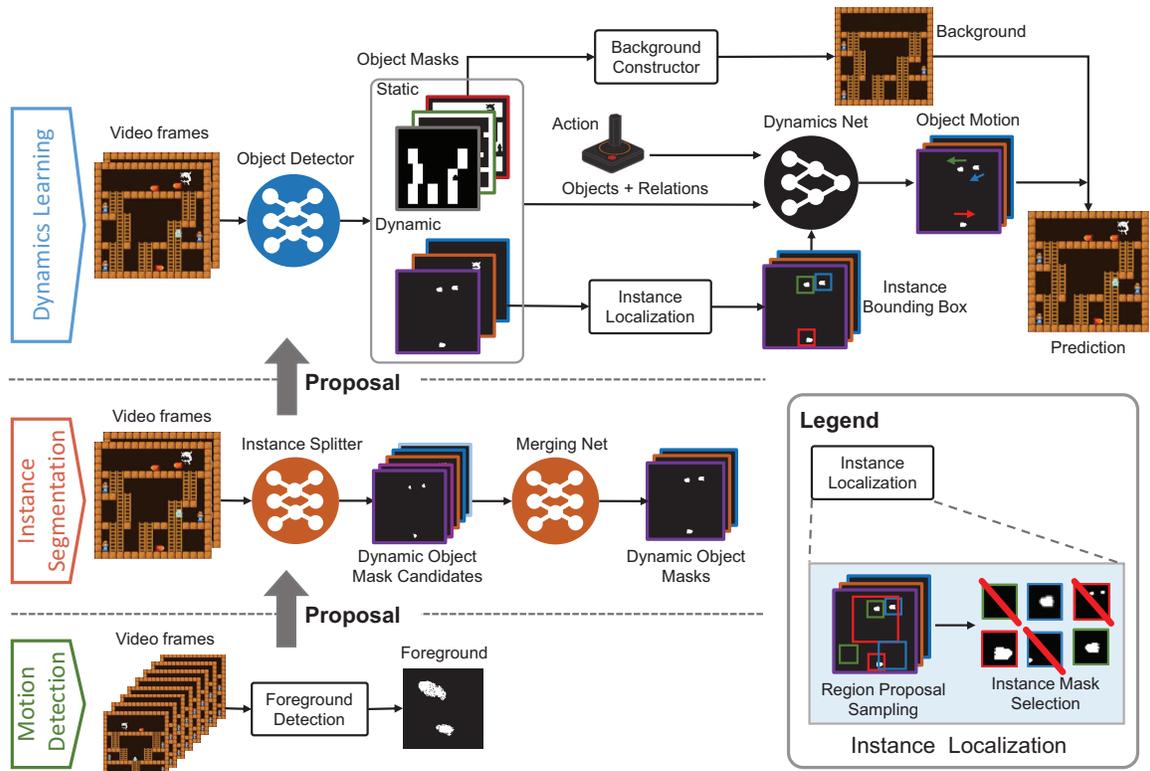


Figure 1: Multi-level dynamics learning framework. From a bottom-up view, we first perform motion detection to produce foreground masks. Then, we utilize the foreground masks as dynamic region proposals to guide the learning of dynamic instance segmentation. Finally, we use the learned dynamic instance segmentation networks (Instance Splitter and Merging Net) as a guiding network to generate region proposals of dynamic instances and guide the learning of Object Detector at the level of dynamics learning. We provide a pseudocode that sketches out this multi-level framework in Algorithm 1.

representations and instance-level dynamics, and predict the next visual observation conditioned on object-to-object relations and an agent’s action. To guide the learning of the object representations and instance localization at the level of dynamics learning, the more abstracted level of *dynamic instance segmentation* learns a guiding network in a self-supervised manner, which can provide coarse mask proposals of dynamic instances. This level exploits spatial-temporal information of locomotion property and appearance patterns to capture region proposals of dynamic instances. To facilitate the learning of instance segmentation, MAOP employs the more coarse-grained level of *motion detection*, which detects changes in image sequences and provides guidance on proposing regions potentially containing dynamic instance.

Algorithm 1 shows a pseudocode that summarizes the training process of our framework. As the learning proceeds, the knowledge distilled from the more coarse-grained level are gradually refined at the more fine-grained level by considering additional information. When the training is finished, the coarse-grained levels of dynamic instance segmentation and motion detection will be removed at the testing stage. In the rest of this section, we will describe in detail the design of each level and their connections.

Motion Detection Level

At this level, we employ foreground detection to detect potential regions of dynamic objects from a sequence of image frames and provide coarse dynamic region proposals for assisting in dynamic instance segmentation. In our experiments, we use a basic unsupervised foreground detection approach (Lo and Velastin 2001). Our framework is also compatible with many advanced unsupervised foreground detection methods (Lee 2005; Maddalena and Petrosino 2008; Zhou, Yang, and Yu 2013; Guo et al. 2014) that are more efficient or more robust to moving camera. These complex foreground detection methods have the potential to improve the performance but are not the focus of this work.

Dynamic Instance Segmentation Level

This level aims to generate region proposals of dynamic instances to guide the learning of object masks and facilitate instance localization at the level of dynamics learning. The architecture is shown in the middle level of Figure 1. Instance Splitter aims to identify regions, each of which potentially contains one dynamic instance. As we focus on the motion of rigid objects, the affine transformation is approximately consistent for all pixels of each dynamic instance mask. Inspired by this, we define a discrepancy loss $\mathcal{L}_{\text{instance}}$ for a

Algorithm 1 Training process for our multi-level framework.

- 1: Initialization. Initialize the parameters of all neural networks with random weights respectively.
 - 2: Motion Detection Level. Perform foreground detection to produce dynamic region proposals, which potentially have moving objects.
 - 3: Instance Segmentation Level. Train the dynamic instance segmentation network (including Instance Splitter and Merging Net) by minimizing \mathcal{L}_{DIS} , which includes a proposal loss to focus the dynamic instance segmentation on the dynamic region proposals from Step 2.
 - 4: Dynamic learning Level. Train the dynamics learning network by minimizing \mathcal{L}_{DL} , which includes a proposal loss to utilize the dynamic instance proposals generated by the trained dynamic instance segmentation network in Step 3 to facilitate the learning of Object Detector.
-

sampled region that measures motion consistence of its pixels and use it to train Instance Splitter. To compute this loss, we first compute an average rigid transformation of a sampled region on object masks between two time steps, then apply this transformation to this region at the previous time step by Spatial Transformer Network (STN) (Jaderberg et al. 2015), and finally compared this predicted region with the region at the current time (the difference is measured by l_2 distance). Obviously, when a sampled region contains exactly one dynamic instance, this loss will be extremely small, and even zero when object masks are perfectly learned. As $\mathcal{L}_{\text{instance}}$ decreases on every sampled regions of object masks, Instance Splitter gradually learns to isolate dynamic instances from background and divide different dynamic objects onto different masks.

Considering that one object instance may be split into smaller patches on different masks, we append a Merging Net (i.e., a two-layer CNN with 1 kernel size and 1 stride) to Instance Splitter to learn to merge masks. This module uses a merging loss $\mathcal{L}_{\text{merge}}$ that aims to merge mask candidates that are adjacent and share the same motion. In addition, we add a foreground proposal loss $\mathcal{L}_{\text{foreground}}$ to encourage attentions on dynamic regions provided by the level of motion detection, which is defined similar to $\mathcal{L}_{\text{proposal}}$ at the level of dynamics learning. The total loss of this level is given by, $\mathcal{L}_{\text{DIS}} = \mathcal{L}_{\text{instance}} + \lambda_3 \mathcal{L}_{\text{merge}} + \lambda_4 \mathcal{L}_{\text{foreground}}$.

Although the network structure of this level is similar to Object Detector in the level of dynamics learning, we do not integrated them together as a whole network because concurrent learning of both object representations and dynamics model is not stable. Instead, we first learn the coarse object representations only based on the spatial-temporal consistency of locomotion and appearance pattern, and then use them as proposal regions to guide object-oriented dynamics learning at the more fine-grained level, which in turn fine-tunes the object representations. In addition, MAOP is also readily to incorporate Mask R-CNN (He et al. 2017) or other off-the-shelf supervised or unsupervised object detection methods (Liu et al. 2018; Xu et al. 2019) as a plug-and-play module into our framework to generate region proposals of dynamic instances.

Object-Oriented Dynamics Learning Level

The semantics of this level is formulated as learning an object-based dynamics model with region proposals generated from the more abstracted level of dynamic instance

segmentation. Its architecture is shown at the top level of Figure 1, which is an end-to-end neural network and can be trained in a self-supervised manner. It takes a sequence of video frames and an agent’s actions as input, learns disentangled representations (including objects, relations and effects) and dynamics of controllable and uncontrollable dynamic object instances conditioned on actions and object relations, and produces predictions of future frames. The whole architecture includes four major components: A) an Object Detector that learns to decompose the input image into objects; B) an Instance Localization module responsible for localizing dynamic instances; C) a Dynamics Net for learning the motion of each dynamic instance conditioned on the effects from actions and object-level spatial-temporal relations; and D) a Background Constructor that computes a background image from learned static object masks. In addition to Figure 1, we further provide Supplementary Algorithm S1 (<https://arxiv.org/abs/1904.07482>) to describe interactions of these components and the learning paradigm of object-based dynamics, which is a general framework and agnostic to the concrete form of each component. In the following paragraphs, we describe detailed design of each component.

Object Detector and Instance Localization Module.

Object Detector is a CNN module aiming to learn object masks from a sequence of input images. An object mask describes a spatial distribution of a class of objects, which forms the fundamental building block of our object-oriented framework. Considering that instances of the same class are likely to have different motions, we append an Instance Localization Module to Object Detector to localize each dynamic instance to support instance-level dynamics learning. Class-specific object masks in conjunction with instance localization bridge visual perception (Object Detector) with dynamics learning (Dynamics Net), which allows learning objects based on both appearances and dynamics.

Specifically, Object Detector receives image $I^t \in \mathbb{R}^{H \times W \times 3}$ at timestep t and outputs object masks $O^t \in [0, 1]^{H \times W \times n_O}$, including dynamic object masks $D^t \in [0, 1]^{H \times W \times n_D}$ and static object masks $S^t \in [0, 1]^{H \times W \times n_S}$, where H , W denote the height and width of the input image, n_D and n_S denotes the maximum possible number of dynamic and static object classes respectively, and $n_O = n_D + n_S$. Note that our framework does not require the actual number of object classes, but needs to set a maximum number (usually 10 is enough). When they do not match, some learned object masks may be redundant, which

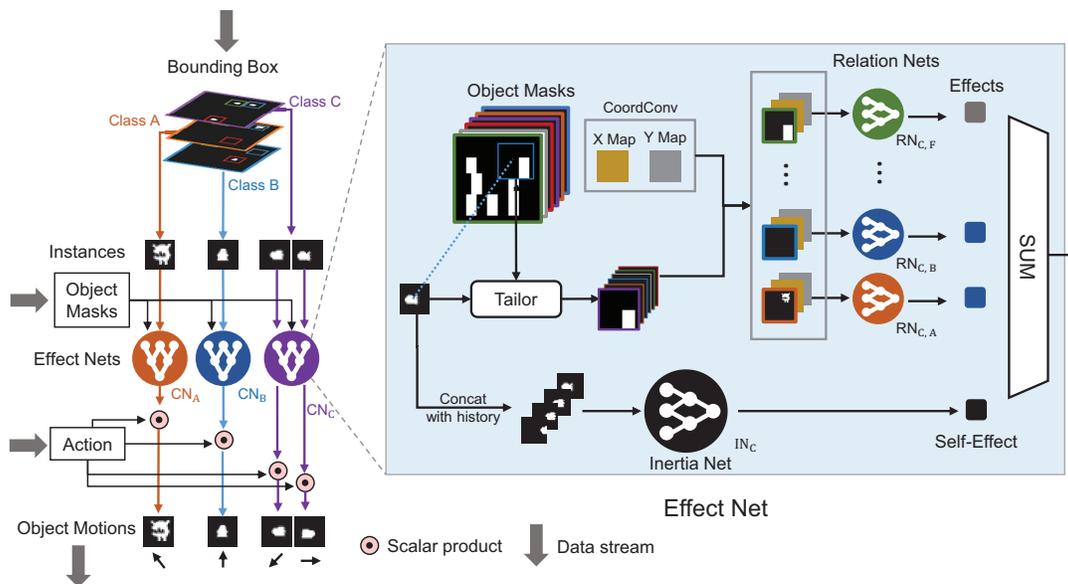


Figure 2: Architecture of Dynamics Net (left) and its component Effect Net (right). Different classes of objects are distinguished by different letters (e.g., A, B, ..., F). Dynamics Net has one Effect Net for each class of objects. An Effect Net consists of one Inertia Net and several Relation Nets.

does not affect the accuracy of predictions. We have conducted experiments to confirm this and will add the results. Entry $O_{u,v,i}$ indicates the probability that pixel $I_{u,v,:}$ belongs to the i -th object class. The Instance Localization module uses learned dynamic object masks to identify each object instance mask $X_{:::,i}^t \in [0, 1]^{H_M \times W_M}$ ($1 \leq i \leq n_M$), where H_M, W_M denote the height and width of the bounding box of this instance and n_M denotes the maximum possible number of localized instances. As shown in Figure 1, Instance Localization first samples a number of bounding boxes on dynamic object masks and then select the regions, each of which contains only one dynamic instance. We use the discrepancy loss $\mathcal{L}_{\text{instance}}$ described in Section of Dynamic Instance Segmentation Level as a selection score for selecting instance masks. More details of region proposal sampling and instance mask selection can be found in Supplementary Section 2 (<https://arxiv.org/abs/1904.07482>).

Dynamics Net. Dynamics Net is designed to learn instance-level motion effects of actions, object-to-object spatial relations (Relation Net) and temporal relations of spatial states (Inertia Net), and to reason about the motion of each dynamic instance based on these effects. Its architecture is illustrated as Figure 2, where the motion of each dynamic instance is individually computed. We take as an example the computation of the motion of the i -th instance $X_{:::,i}^t$ to illustrate the detailed structure of the Effect Net.

As shown in the right subfigure of Figure 2, Effect Net first uses a sub-differentiable tailor module introduced by Zhu, Huang, and Zhang (2018) to enable the inference of dynamics focusing on the relations with neighbour objects. This module crops a w -size ‘‘horizon’’ window from the concatenated masks of all objects O^t centered on the expected location of $X_{:::,i}^t$, where w denotes the maximum effective range of rela-

tions. Then, the cropped object masks are concatenated with constant x-coordinate and y-coordinate meshgrid map (to make networks more sensitive to the spatial information) and fed into corresponding Relation Nets (RN) according to their classes. We use $C_{:::,i,j}^t$ to denote the cropped mask that crops the j -th object class $O_{:::,j}^t$ centered on the expected location of the i -th dynamic instance (the class it belongs to is denoted as $c_i, 1 \leq c_i \leq n_D$). The effect of object class j on class c_i , $E^t(c_i, j) = \text{RN}_{c_i,j}(\text{concat}(C_{:::,i,j}^t, \text{Xmap}, \text{Ymap}))$. Note that there are total $n_D \times n_O$ RNs for $n_D \times n_O$ pairs of object classes that share the same architecture but not their weights. To handle the partial observation problem, we add an Inertia Nets (IN) to learn spatio-temporal self-effects of an object class through historical states, $E_{\text{self}}^t(c_i) = \text{IN}_{c_i}(\text{concat}(X_{:::,i}^t, X_{:::,i}^{t+1}, \dots, X_{:::,i}^{t+h}))$, where h is the history length. There are total n_D INs for n_D dynamic object classes, which share the same architecture but not their weights. To predict the motion vector m_i^t for the i -th dynamic instance, all these effects are summed up and then multiplied by the coded action vector a^t , that is, $m_i^t = ((\sum_j E^t(c_i, j)) + E_{\text{self}}^t(c_i)) \cdot a^t$.

Background Constructor. This module constructs the static background of an input image based on the static object masks learned by Object Detector. Since Object Detector can decompose its observation into objects in an unseen environment with different layouts, Background Constructor is able to generate a corresponding static background and support the visual observation prediction in novel environments. Specifically, Background Constructor maintains a background memory $B \in \mathbb{R}^{H \times W \times 3}$ which is continuously

updated with the static object mask learned by Object Detector. Denoting α as the decay rate, the updating formula is given by, $B^t = \alpha B^{t-1} + (1 - \alpha) I^t \sum_i S_{:, :, i}^t$, and $B^0 = 0$.

Prediction and Training Loss. Based on the learned masks and motions of the object instances, we propose an object-oriented prediction loss, $\mathcal{L}_{\text{pred-object}} = \sum_i \left\| \text{STN} \left((\bar{u}_i, \bar{v}_i)^t, m_i^t \right) - (\bar{u}_i, \bar{v}_i)^{t+1} \right\|_2^2$, where $(\bar{u}_i, \bar{v}_i)^t$ is the expected location of i -th instance mask $X_{:, :, i}^t$. To utilize the information of ground-true future frames, we also use a conventional image prediction loss. Our prediction of the next frame is produced by merging the learned object motions and the background B^t . The pixels of a dynamic instance can be calculated by masking the raw image with the corresponding instance mask and we can use STN to apply the learned instance motion vector m_i^t on these pixels. First, we transform all the dynamic instances according to the learned instance-level motions. Then, we merge all the transformed dynamic instances with the background image calculated from Background Constructor to generate the prediction of the next frame. We use the pixel-wise l_2 loss to restrain image prediction error, denoted as $\mathcal{L}_{\text{pred-image}}$. In addition, we add a proposal loss to utilize the dynamic instance proposals for guiding the learning, $\mathcal{L}_{\text{proposal}} = \left\| \sum_i (D_{:, :, i}^t - P_{:, :, i}^t) \right\|_2^2$, where P denotes the dynamic instance region proposals provided by the level of dynamic instance segmentation. Therefore, the total loss of the dynamics learning level is given by, $\mathcal{L}_{\text{DL}} = \mathcal{L}_{\text{pred-object}} + \lambda_1 \mathcal{L}_{\text{pred-image}} + \lambda_2 \mathcal{L}_{\text{proposal}}$.

Experiments

We compare MAOP with state-of-the-art action-conditioned dynamics learning baselines, AC Model (Oh et al. 2015), CDNA (Finn, Goodfellow, and Levine 2016), and OODP (Zhu, Huang, and Zhang 2018). AC Model adopts an encoder-LSTM-decoder structure, which performs transformations in hidden space and constructs pixel predictions. CDNA explicitly models pixel motions to achieve invariance to appearance. OODP tries to simultaneously learn object-based representations, relations and motion effects. MAOP adopts a multi-level abstraction framework to decompose raw images into objects and predict instance-level dynamics based on actions and object relations. OODP and MAOP both aim at learning object-level dynamics through an object-oriented learning paradigm, which decomposes raw images into objects and perform prediction based on object-level relations. OODP is only designed for class-level dynamics, while MAOP is able to learn instance-level dynamics. See Supplementary Material (<https://arxiv.org/abs/1904.07482>) for more implementation details.

Generalization Ability and Sample Efficiency

We first evaluate zero-shot generalization and sample efficiency on *Monster Kong* from Pygame Learning Environment (Tasfi 2016), which allows us to test generalization ability over various scenes with different layouts. It is the advanced version of that used by Zhu, Huang, and Zhang (2018), which has a more general and complex setting. The monster wanders around and breathes out fires randomly, and the fires also

move with some randomness. The agent randomly explores with actions *up*, *down*, *left*, *right*, *jump* and *noop*. All these dynamic objects interact with the environment and other objects according to the underlying physics engine. Moreover, gravity and jump model has a long-term dynamics effects, leading to a partial observation problem. To test whether our model can truly learn the underlying physical mechanism behind the visual observations and perform relational reasoning, we set the k -to- m zero-shot generalization experiment (Figure 3), where we use k different environments for training and m different unseen environments for testing.

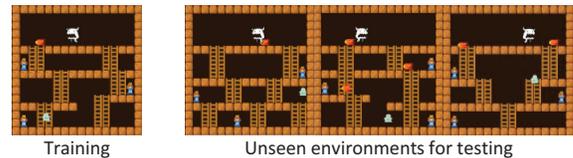


Figure 3: An Example of 1-to-3 zero-shot generalization.

To make a sufficient comparison with previous methods on object dynamics learning and video prediction, we conduct 1-5, 2-5 and 3-5 generalization experiments with a variety of evaluation indices. We use n -error accuracy to measure the performance of object dynamics prediction, which is defined as the proportion that the difference between the predicted and ground-true agent locations is less than n pixel. We also add an extra pixel-based measurement (denoted by object RMSE), which compares the pixel difference near dynamic objects between the predicted and ground-truth images.

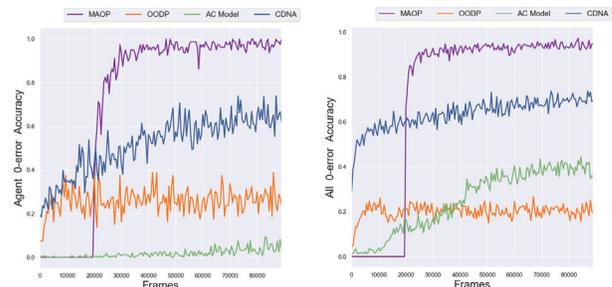


Figure 4: The performance of object dynamics prediction in unseen environments as training progresses on 3-to-5 generalization problem of *Monster Kong*. Since we use the first 20k samples to train the level of dynamic instance segmentation, the curve of MAOP starts at iteration 20001.

As shown in Table 1, MAOP significantly outperforms other methods in all experiment settings in terms of generalization ability and sample efficiency of object dynamics learning. It can achieve 90% 0-error accuracy in unseen environments even trained with 3k samples from a single environment, while other methods have a much lower accuracy (less than 45%). In addition, MAOP with only 3k training samples outperforms CDNA using 300k samples. Although AC Model achieves high accuracy in training environments, its performance in unseen scenes is much worse, which is

Table 1: Prediction performance on *Monster Kong*. k - m means the k -to- m generalization problem. § indicates training with only 3000 samples. ALL represents all dynamic objects. The first column shows the number of samples used for training the models.

Models		Training environments				Unseen environments											
		1-5§		1-5		2-5		3-5		1-5§		1-5		2-5		3-5	
		Agent	All	Agent	All	Agent	All	Agent	All	Agent	All	Agent	All	Agent	All	Agent	All
Training Samples	MAOP	3k		100k		100k		100k		100k		-		-		-	
	OODP	3k		200k		200k		200k		200k		-		-		-	
	AC Model	3k		500k		500k		500k		500k		-		-		-	
	CDNA	3k		300k		300k		300k		300k		-		-		-	
0-error accuracy	MAOP	0.95	0.92	0.98	0.95	0.99	0.96	0.99	0.95	0.94	0.90	0.97	0.92	0.98	0.93	0.99	0.94
	OODP	0.15	0.15	0.18	0.16	0.22	0.17	0.26	0.20	0.14	0.15	0.20	0.15	0.18	0.15	0.26	0.18
	AC Model	0.01	0.36	0.87	0.94	0.80	0.93	0.77	0.92	0.01	0.20	0.08	0.16	0.30	0.48	0.45	0.66
	CDNA	0.28	0.62	0.77	0.84	0.78	0.82	0.78	0.84	0.26	0.44	0.79	0.80	0.78	0.78	0.81	0.83
Object RMSE	MAOP	24.58		21.96		21.97		23.04		29.67		27.22		25.55		24.30	
	OODP	65.63		66.44		66.66		64.73		65.46		67.41		67.78		64.95	
	AC Model	71.02		18.88		22.39		21.30		77.24		57.41		55.45		43.48	
	CDNA	40.92		24.52		24.37		24.18		51.08		37.15		27.67		25.33	

probably because its pure pixel-level inference easily leads to overfitting. CDNA performs better than AC Model, but still cannot efficiently generalize with limited training samples. Since OODP has structural limitation and optimization difficulty, it has innate difficulty on frames with multiple dynamic objects. In Figure 4 and Supplementary Figure S3 (<https://arxiv.org/abs/1904.07482>), we also plot the learning curve of these models. Compared to other models, MAOP achieves higher prediction accuracy for unseen environments at a faster rate during the training process. We further add a video (see <https://arxiv.org/abs/1904.07482>) for better perceptual understanding of prediction performance in unseen environments.

Table 2: Accuracy of dynamics prediction on *Flappy Bird* and *Freeway*. Since only the agent’s ground-true location is accessible in ALE, we just show the dynamics prediction of the agent. Actually, we observe that predictions of other dynamic objects are also accurate by comparing predicted with ground-true images (see Supplementary Figure S5 at <https://arxiv.org/abs/1904.07482>).

Models	Flappy Bird (0-acc)		Freeway (Agent)		
	100 samples		300 samples		100 samples
	Agent	All	Agent	All	0-acc 1-acc 2-acc
MAOP	0.83	0.89	0.83	0.92	0.78 0.84 0.89
OODP	0.01	0.18	0.02	0.15	0.26 0.33 0.42
AC Model	0.03	0.18	0.04	0.23	0.31 0.38 0.42
CDNA	0.13	0.77	0.30	0.81	0.42 0.43 0.47

We also evaluate MAOP on *Flappy Bird* and *Freeway*. *Flappy Bird* is a side-scroller game with a moving camera. *Freeway* is an Atari game, which has a large number of dynamic objects. Since the testing environments will be similar with the training ones without limitation of samples, we limit the training samples to form a sufficiently challenging generalization task. MAOP still outperforms existing baseline methods (Table 2), which demonstrates that MAOP is effec-

tive for the concurrent dynamics prediction of a large number of objects. In addition, we conduct a modular test to better understand the contribution of each learning level (see Supplementary Section 3 at <https://arxiv.org/abs/1904.07482>). The results show that each level of MAOP can independently perform well and has a good robustness to the proposals generated by the more abstracted level. Taken together, the above results demonstrates that MAOP has superiority of sample efficiency and generalization ability, which suggests MAOP is good at relational reasoning and learns the object-level dynamics, rather than learn some patterns from mass data to recover the dynamics as the conventional neural networks do.

To further test the priority and limitation of MAOP, we have applied MAOP on a diverse set of games on Atari. Testing results with 3k training samples on *Skiing*, *MsPacman*, *Krull*, *Pong*, *MontezumaRevenge*, and *Breakout* are shown in Table 3. We observe a superior performance of MAOP over baseline models on *Skiing*, *MsPacman*, *Krull* and *Pong*, and a slightly worse performance on *Breakout*. Because our model is designed for scenes with multiple dynamic objects, its performance may be lower than some simpler baseline methods on environments with only one or two dynamic objects, such as *MontezumaRevenge* and *Breakout*.

Model-Based Planning in Unseen Environments

Although RL has achieved considerable successes, most RL researches tend to “train on the test set” (Nichol et al. 2018; Pineau 2018). It is critical yet challenging to develop model-based RL approaches that support generalization over unseen environments. Monte Carlo tree search (MCTS) (Browne et al. 2012) is developed to leverage the environment models to conduct efficient lookahead search, which has shown remarkable effectiveness on long-term planning, such as AlphaGo (Silver et al. 2016). Considering that our learned dynamics model can efficiently generalize to unseen environments, we can directly use our learned model to perform MCTS in unseen environments. To perform long-range planning, we first test our performance of long-range prediction, as shown in

Table 3: Accuracy of dynamics prediction on six Atari games.

Model	Skiing			MsPacman			Krull			Pong			MontezumaRevenge			Breakout		
	0acc	1acc	2acc	0acc	1acc	2acc	0acc	1acc	2acc									
MAOP	0.97	0.99	1.00	0.65	0.88	0.94	0.14	0.48	0.73	0.63	0.73	0.83	0.95	1.00	1.00	0.52	0.66	0.77
OODP	0.62	0.80	0.90	0.30	0.35	0.46	0.02	0.08	0.16	0.46	0.64	0.66	0.66	0.92	0.99	0.73	0.66	0.80
AC Model	0.27	0.47	0.56	0.44	0.52	0.54	0.01	0.05	0.13	0.37	0.40	0.42	0.63	0.79	0.95	0.45	0.57	0.66
CDNA	0.76	0.95	0.99	0.52	0.68	0.74	0.27	0.41	0.51	0.65	0.66	0.79	0.96	1.00	1.00	0.63	0.71	0.77

Supplementary Table S1 (<https://arxiv.org/abs/1904.07482>). MAOP only trained for 1-step prediction can achieve 90% 2-error accuracy in unseen environments when predicting 3 steps of the future, while the accuracy is 73% when predicting 6 steps of the future, which is also a satisfactory performance for lookahead search. Supplementary Figure S7 (<https://arxiv.org/abs/1904.07482>) illustrates a case visualizing the 6-step prediction of MAOP in unseen environments.

We evaluate our performance of model-based planning on *Monster Kong*. In this game, the goal of the agent is to approach the princess and a reward will be given when the straight-line distance from agent to princess gets smaller than that in the agent’s history. The value of such a reward is proportional to the shrinking distance. The agent will win with an extra reward +5 when touching the princess, and lose with an extra reward -5 when hitting the fires. To gain a better understanding of the contribution of MAOP to the MCTS agent, we compare MCTS in conjunction with MAOP to DQN (Mnih et al. 2015) and to an ablation (i.e., using the real simulator of the unseen environments in MCTS). We provide the same ground-true reward functions for all dynamics model during MCTS. We conduct random experiments in 5 unseen environments, where the agent and the princess randomly generate. We train all models in training environments with 5k samples, and test zero-shot generalization of the model-free behavior policy (i.e., DQN) and model-based planning policy (i.e., MAOP-based, CDNA-based, OODP-based and AC-based) in unseen environments.

As shown in Table 4, MAOP achieves almost the same performance with the true environment model for model-based planning in unseen environments and significantly outperforms other baseline models and DQN. The model-free approach DQN tends to overfit the training environments and cannot learn to plan in unseen environments, leading to a much higher death rate and a much lower score. The learning curves in Supplementary Figure S8 (<https://arxiv.org/abs/1904.07482>) also verify this. In addition, we observe that MCTS in conjunction with MAOP acquires intriguing forward-looking skills, such as jumping over the fires and jumping across the big gap that are critical for survival and reaching the goal (we provide videos for the learned policies at <https://arxiv.org/abs/1904.07482>).

Interpretable Representations and Knowledge

MAOP takes a step towards interpretable dynamics model learning. Through interacting with environments, it learns visually and semantically interpretable knowledge in a self-supervised manner, which contributes to unlocking the “black box” of the dynamics prediction and potentially opens the

Table 4: The performance of using MCTS with different dynamics models, and DQN in unseen environments. REAL indicates the real simulator. Time Out indicates exceeding 100 steps. Reward is averaged over 21 runs.

Methods	Reward	Win	Lose	Time Out
MCTS + MAOP	38.19	47.62%	9.52%	42.86%
MCTS + REAL	38.41	52.38%	9.52%	38.10%
MCTS + CDNA	6.83	0%	33.33%	66.67%
MCTS + OODP	13.95	0%	52.38%	47.62%
MCTS + AC	7.50	0%	47.62%	52.38%
DQN	13.67	26.7%	23.8%	49.5%

avenue for further researches on object-oriented RL, model-based RL, and hierarchical RL.

Visual Interpretability. To demonstrate the model interpretability of MAOP in unseen environments, we visualize the learned masks of dynamic and static objects. We highlight the attentions of object masks by multiplying the raw images by the binarized masks. Note that MAOP does not require the actual number of objects but a maximum number and some learned object masks may be redundant. Thus, we only show the informative object masks. As shown in Figure 5, our model captures all the key objects in the environments including the controllable agents (cowboy, bird, and chicken), the uncontrollable dynamic objects (monster, fires, pipes and cars), and the static objects that have effects on the motions of dynamic objects (ladders, walls and the free space), which demonstrates that model can learn disentangled object representations and distinguish the objects by both appearance and dynamic property.

Dynamical Interpretability. To show the dynamical interpretability behind image prediction, we test our predicted motions by comparing RMSEs between the predicted and ground-truth motions in unseen environments (Supplementary Table S2 at <https://arxiv.org/abs/1904.07482>). Intriguingly, most predicted motions are quite accurate, with the RMSEs less than 1 pixel. Such a visually indistinguishable error also verifies the accuracy of our dynamics learning.

Discovery of the Controllable Agent. With the learned knowledge in MAOP, we can easily uncover the action-controlled agent from all the dynamic objects, which is useful semantic information that can be used in heuristic algorithms. For example, it allows agents to efficiently explore (e.g., contingency awareness (Choi et al. 2019), empowerment (Karl et al. 2017), megalomania-drivenness (Song et al. 2019), and distance-based rewards (Srinivas et al. 2018)).

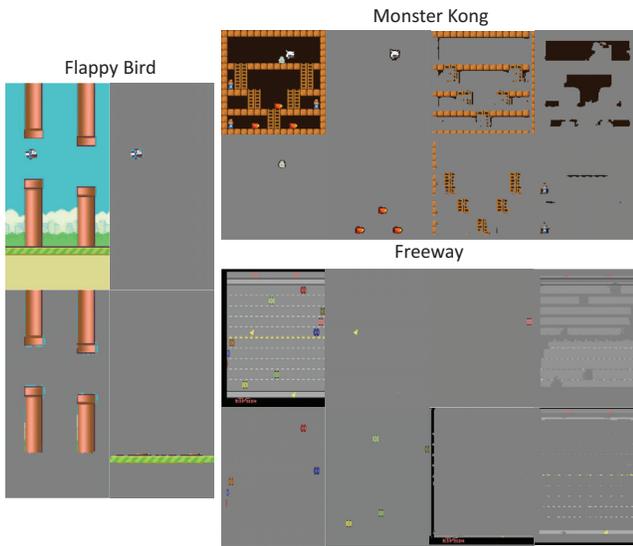


Figure 5: Visualization of the masked images in unseen environments. Top left corner is the raw image.

Specifically, the object that has the maximal variance of total effects over actions is the action-controlled agent. Denote the total effects as $E_i^t = (\sum_j E^t(c_i, j)) + E_{\text{self}}^t(c_i)$, the label of the action-controlled agent is calculated as, $\arg \max_i \sum_t \text{Var}_{a^t}(E_i^t)$. We observe that our discovery of the controllable agent achieves right or near 100% accuracy in unseen environments (see Supplementary Table S3 at <https://arxiv.org/abs/1904.07482>).

Conclusion and Discussion

This paper presents a self-supervised multi-level learning framework for learning action-conditioned object-based dynamics. It enables sample-efficient and interpretable model learning, and achieves zero-shot generalization over novel environments with multiple dynamic objects and different static object layouts. The learned dynamics model enables an agent to directly plan in unseen environments. MAOP can easily generalize the learned knowledge over environments with similar objects but may not work well with those with totally new objects, which is an important direction for future work.

As abrupt changes (e.g., colors) are often predictable from a long-term view or memory, our model can be extended to more domains by incorporating memory networks (e.g., LSTM). In addition, our future work includes extending our model for deformation prediction (e.g., object appearing, disappearing and non-rigid deformation) and incorporating a camera motion prediction network module introduced by (Vijayanarasimhan et al. 2017) for applications such as FPS games and autonomous driving. Learning 3D dynamics from 2D video is extremely challenging. Conventional neural networks try to learn such 3D dynamics by remembering some patterns from 2D data as they do for the non-rigid deformation, such as AC Model (Oh et al. 2015) and CDNA (Finn, Goodfellow, and Levine 2016). This approach achieves good

performance in training environments, but it requires a large number of data and does not really recover the true 3D dynamics model. To learn generalized 3D dynamics model, object-oriented learning paradigm in conjunction with 3D CNN (3D data input) is necessary, which is an important direction for future work.

References

- Battaglia, P.; Pascanu, R.; Lai, M.; et al. 2016. Interaction networks for learning about objects, relations and physics. In *Advances in Neural Information Processing Systems*.
- Battaglia, P. W.; Hamrick, J. B.; Bapst, V.; et al. 2018. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*.
- Browne, C. B.; Powley, E.; Whitehouse, D.; et al. 2012. A survey of monte carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in games* 4(1).
- Chang, M. B.; Ullman, T.; Torralba, A.; and Tenenbaum, J. B. 2016. A compositional object-based approach to learning physical dynamics. *arXiv preprint arXiv:1612.00341*.
- Chiappa, S.; Racaniere, S.; Wierstra, D.; and Mohamed, S. 2017. Recurrent environment simulators. *International Conference on Learning Representations*.
- Choi, J.; Guo, Y.; Moczulski, M. L.; Oh, J.; Wu, N.; Norouzi, M.; and Lee, H. 2019. Contingency-aware exploration in reinforcement learning.
- Cobo, L. C.; Isbell, C. L.; and Thomaz, A. L. 2013. Object focused q-learning for autonomous agents. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*.
- Diuk, C.; Cohen, A.; Littman, M. L.; et al. 2008. An object-oriented representation for efficient reinforcement learning. In *Proceedings of the 25th international conference on Machine learning*.
- Finn, C., and Levine, S. 2017. Deep visual foresight for planning robot motion. In *Robotics and Automation (ICRA), 2017 IEEE International Conference on*.
- Finn, C.; Goodfellow, I.; and Levine, S. 2016. Unsupervised learning for physical interaction through video prediction. In *Advances in Neural Information Processing Systems*.
- Fragkiadaki, K.; Agrawal, P.; Levine, S.; and Malik, J. 2016. Learning visual predictive models of physics for playing billiards. *International Conference on Learning Representations*.
- Gu, S.; Lillicrap, T.; Sutskever, I.; and Levine, S. 2016. Continuous deep q-learning with model-based acceleration. In *International Conference on Machine Learning*.
- Guestrin, C.; Koller, D.; Gearhart, C.; et al. 2003. Generalizing plans to new environments in relational mdps. In *Proceedings of the 18th international joint conference on Artificial intelligence*.
- Guo, X.; Wang, X.; Yang, L.; Cao, X.; and Ma, Y. 2014. Robust foreground detection using smoothness and arbitrariness constraints. In *European Conference on Computer Vision*.

- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*.
- Jaderberg, M.; Simonyan, K.; Zisserman, A.; et al. 2015. Spatial transformer networks. In *Advances in Neural Information Processing Systems*.
- Johansson, G. 1975. Visual motion perception. *Scientific American* 232(6).
- Kansky, K.; Silver, T.; Mély, D. A.; et al. 2017. Schema networks: Zero-shot transfer with a generative causal model of intuitive physics. *arXiv preprint arXiv:1706.04317*.
- Karl, M.; Soelch, M.; Becker-Ehmck, P.; Benbouzid, D.; van der Smagt, P.; and Bayer, J. 2017. Unsupervised real-time control through variational empowerment. *arXiv preprint arXiv:1710.05101*.
- Lee, D.-S. 2005. Effective gaussian mixture learning for video background subtraction. *IEEE Transactions on Pattern Analysis & Machine Intelligence* (5).
- Li, Y.; Qi, H.; Dai, J.; et al. 2017. Fully convolutional instance-aware semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- Liu, L.; Ouyang, W.; Wang, X.; et al. 2018. Deep learning for generic object detection: A survey. *arXiv preprint arXiv:1809.02165*.
- Liu, B.; He, X.; and Gould, S. 2015. Multi-class semantic video segmentation with exemplar-based object reasoning. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*.
- Lo, B., and Velastin, S. 2001. Automatic congestion detection system for underground platforms. In *Proceedings of International Symposium on Intelligent Multimedia, Video and Speech Processing*.
- Lu, Z.-L., and Sperling, G. 1995. The functional architecture of human visual motion perception. *Vision research* 35(19).
- Maddalena, L., and Petrosino, A. 2008. A self-organizing approach to background subtraction for visual surveillance applications. *IEEE Transactions on Image Processing* 17(7).
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; et al. 2015. Human-level control through deep reinforcement learning. *Nature* 518(7540).
- Nichol, A.; Pfau, V.; Hesse, C.; Klimov, O.; and Schulman, J. 2018. Gotta learn fast: A new benchmark for generalization in rl. *arXiv preprint arXiv:1804.03720*.
- Oh, J.; Guo, X.; Lee, H.; Lewis, R. L.; and Singh, S. 2015. Action-conditional video prediction using deep networks in atari games. In *Advances in Neural Information Processing Systems*.
- Pineau, J. 2018. Reproducible, reusable, and robust reinforcement learning (Invited Talk). *Advances in Neural Information Processing Systems*.
- Pinheiro, P. O.; Collobert, R.; and Dollár, P. 2015. Learning to segment object candidates. In *Advances in Neural Information Processing Systems*.
- Racanière, S.; Weber, T.; Reichert, D.; et al. 2017. Imagination-augmented agents for deep reinforcement learning. In *Advances in Neural Information Processing Systems*.
- Santoro, A.; Raposo, D.; Barrett, D. G.; et al. 2017. A simple neural network module for relational reasoning. In *Advances in neural information processing systems*.
- Silver, D.; Huang, A.; Maddison, C. J.; et al. 2016. Mastering the game of go with deep neural networks and tree search. *nature* 529(7587).
- Song, Y.; Wang, J.; Lukasiewicz, T.; Xu, Z.; Zhang, S.; and Xu, M. 2019. Mega-reward: Achieving human-level play without extrinsic rewards. *arXiv preprint arXiv:1905.04640*.
- Srinivas, A.; Jabri, A.; Abbeel, P.; Levine, S.; and Finn, C. 2018. Universal planning networks. *arXiv preprint arXiv:1804.00645*.
- Tasfi, N. 2016. Pygame learning environment. <https://github.com/ntasfi/PyGame-Learning-Environment>.
- van Steenkiste, S.; Chang, M.; Greff, K.; and Schmidhuber, J. 2018. Relational neural expectation maximization: Unsupervised discovery of objects and their interactions. *arXiv preprint arXiv:1802.10353*.
- Vijayanarasimhan, S.; Ricco, S.; Schmid, C.; Sukthankar, R.; and Fragkiadaki, K. 2017. Sfm-net: Learning of structure and motion from video. *arXiv preprint arXiv:1704.07804*.
- Watters, N.; Zoran, D.; Weber, T.; et al. 2017. Visual interaction networks. In *Advances in Neural Information Processing Systems*.
- Wu, J.; Lu, E.; Kohli, P.; Freeman, B.; and Tenenbaum, J. 2017. Learning to see physics via visual de-animation. In *Advances in Neural Information Processing Systems*.
- Xu, Z.; Liu, Z.; Sun, C.; et al. 2019. Unsupervised discovery of parts, structure, and dynamics. *arXiv preprint arXiv:1903.05136*.
- Zambaldi, V.; Raposo, D.; Santoro, A.; et al. 2018. Relational deep reinforcement learning. *arXiv preprint arXiv:1806.01830*.
- Zhang, C., and Shah, J. A. 2016. Co-optimizing multi-agent placement with task assignment and scheduling. In *IJCAI*.
- Zhou, X.; Yang, C.; and Yu, W. 2013. Moving object detection by detecting contiguous outliers in the low-rank representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35(3).
- Zhu, G.; Huang, Z.; and Zhang, C. 2018. Object-oriented dynamics predictor. *Advances in Neural Information Processing Systems*.