

# Intrinsic Physical Concepts Discovery with Object-Centric Predictive Models

Qu Tang<sup>1,2\*</sup> Xiangyu Zhu<sup>1,2\*</sup> Zhen Lei<sup>1,2,3</sup> Zhaoxiang Zhang<sup>1,2,3\*</sup>

<sup>1</sup>School of Artificial Intelligence, University of Chinese Academy of Sciences

<sup>2</sup>MAIS, Institute of Automation, Chinese Academy of Sciences

<sup>3</sup>Centre for Artificial Intelligence and Robotics, Hong Kong Institute of Science & Innovation,  
Chinese Academy of Sciences

{tangqu2020, zhaoxiang.zhang}@ia.ac.cn, {xiangyu.zhu, zlei}@nlpr.ia.ac.cn

## Abstract

The ability to discover abstract physical concepts and understand how they work in the world through observing lies at the core of human intelligence. The acquisition of this ability is based on compositionally perceiving the environment in terms of objects and relations in an unsupervised manner. Recent approaches learn object-centric representations and capture visually observable concepts of objects, e.g., shape, size, and location. In this paper, we take a step forward and try to discover and represent intrinsic physical concepts such as mass and charge. We introduce the *PHYSical Concepts Inference NETWORK (PHYCINE)*, a system that infers physical concepts in different abstraction levels without supervision. The key insights underlining *PHYCINE* are two-fold, commonsense knowledge emerges with prediction, and physical concepts of different abstraction levels should be reasoned in a bottom-up fashion. Empirical evaluation demonstrates that variables inferred by our system work in accordance with the properties of the corresponding physical concepts. We also show that object representations containing the discovered physical concepts variables could help achieve better performance in causal reasoning tasks, i.e., *ComPhy*.

## 1. Introduction

Why do objects bounce off after the collision? Why do magnets attract or repel each other? Objects cover many complex physical concepts which define how they interact with the world [5]. Humans have the ability to discover abstract concepts about how the world works through just observation. In a preserved video, some concepts are obvious in the visual appearances of objects, like location, size, velocity, etc, while some concepts are hidden in the behaviors of objects. For example, intrinsic physical concepts

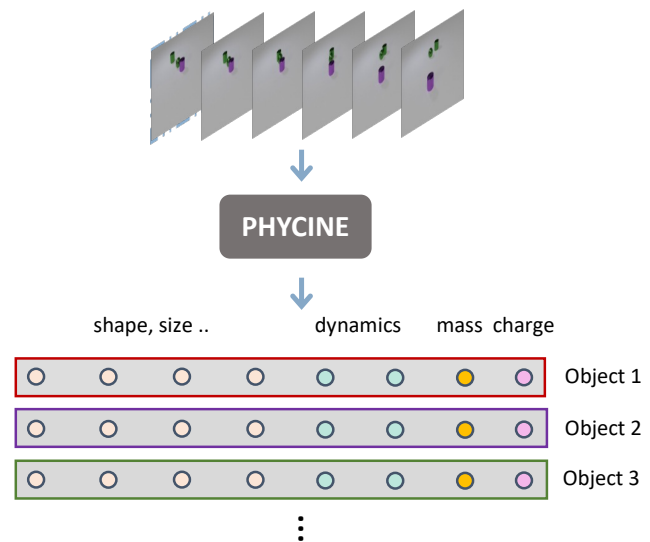


Figure 1. Given a video from the ComPhy [5] dataset, PHYCINE decomposes the scene into multi-object representations that contain physical concepts of different abstraction levels, including visual attributes, dynamics, mass, and charge. (shown in different colors).

like mass and charge, are unobservable from static scenes and can only be discovered from the object dynamics. Objects carrying the same or opposite charge will exert a repulsive or attractive force on each other. After a collision, the lighter object will undergo larger changes in its motion compared with the massive one. How can machines learn to reveal and represent such common sense knowledge? We believe the answer lies in two key steps: decomposing the world into object-centric representations, and building a predictive model with abstract physical concepts as latent variables to handle uncertainty.

Object-centric representation learning [3,9,18,22,23,29] aims at perceiving the world in a structured manner to improve the generalization ability of intelligent systems and achieve higher-level cognition. VAE [13]-based models,

\*Corresponding author. \*Equal contribution

like IODINE [9] learn disentangled features that separate interpretable object properties (e.g., shape, color, location) in a common format. However, abstract concepts like mass and charge, can not be distilled by generative models since building object-level representations does not necessarily model these higher-level-abstracted physics. CPL [5] successfully learns high-level concepts with graph networks, but it relies on supervision signals from the ground-truth object-level concept labels. There are also several studies investigating the effectiveness of object-centric representations in learning predictive models to solve predicting and planning tasks [11, 14, 23, 24]. Nevertheless, to our knowledge, there is no work yet trying to discover and represent object-level intrinsic physical concepts in an unsupervised manner.

The idea that abstract concepts can be learned through prediction has been formulated in various ways in cognitive science, neuroscience, and AI over several decades [19, 20]. Intuitively, physics concepts such as velocity, mass, and charge, may emerge gradually by training the system to perform long-term predictions at the object representation level [16]. Through predictions at increasingly long-time scales, more and more complex concepts about how the world works may be acquired in a bottom-up fashion. In this paper, we focus on the main challenge: enabling the model to represent and disentangle the unfolded concepts. We follow common sense: with a neural physics engine, if the prediction of an object trajectory fails, there must be physical concepts that have not been captured. Therefore, a latent variable that successfully models the uncertainty of prediction defines a new physical concept, and a better physical engine can be built. Following this idea, we categorize physical concepts into three levels of abstraction: extrinsic concepts, dynamic concepts, and intrinsic concepts. Firstly, the extrinsic properties (e.g., color, shape, material, size, depth, location) can be referred to as object contexts belonging to the lowest level of abstraction, and a perception module can directly encode the contexts. Secondly, the dynamic properties (e.g., velocity) in the middle level are hidden in the temporal and spatial relationships of visual features and should be inferred from short-term prediction. Thirdly, intrinsic properties like mass and charge can neither be directly observed nor inferred from short-term prediction. They can only be inferred by analyzing the way how objects exert force on each other. For example, inferring mass needs the incorporation of a collision event, and inferring charge needs to observe the change of object dynamics, which depends on a long-term observation or prediction.

In this work, we build a system called PHYSical Concepts Inference NETWORK (PHYCINE). In the system, there are features arranged in a bottom-up pyramid that represents physical concepts in different abstraction levels. These fea-

tures cooperatively perform reconstruction and prediction to interpret the observed world. Firstly, the object context features reconstruct the observed image with a generative model. Secondly, object dynamics features predict the next-step object contexts by learning a state transition function. Finally, the mass and charge features model the interaction between objects. PHYCINE uses a relation module to calculate pair-wise forces for all entities, and adaptively learn the variables that represent object mass and charge with proper regularization. During training, all representations are randomly initialized and iteratively refined using gradient information about the evidence lower bound (ELBO) obtained by the current estimate of the parameters. The model can be trained using only raw input video data in an end-to-end fashion. As shown in Figure 1, taking a raw video as input, PHYCINE not only extracts extrinsic object contexts (i.e., size, shape, color, location, material), but also infers more abstract concepts (i.e., object dynamics, mass, and charge). In our experiments, we demonstrate the model’s ability to discover and disentangle physical concepts, which can be used to solve downstream tasks. We evaluate the learned representation on ComPhy, a causal reasoning benchmark.

Our main contributions are as follows: (i) We challenge a problem of physical concept discovery by observing raw videos and successfully discovering intrinsic concepts of velocity, mass, and charge. (ii) We introduce a framework PHYCINE, a hierarchical object-centric predictive model that infers physical concepts from low (e.g., color, shape) to high (e.g., mass, charge) abstract levels, leading to disentangled object-level physical representations. (iii) We demonstrate the effectiveness of the representation learned by PHYCINE on ComPhy.

## 2. Related Work

**Object-centric representation learning.** Several recent lines of work learn object-centric representations from visual inputs without explicit supervision. A common framework for static scenes is returning a set of object representation vectors or slot vectors via a form of auto-encoders. MoNet [3] firstly segment the image with an attention module, then iteratively explain away and reconstruct the scene. IODINE [9], which is closely related to our work, use amortized variational inference to refine the posterior parameters. Slot-Attention [18] learns a grouping strategy to decompose input features into a set of slot representations. Further, many researchers have used architectures shown to be effective for static images to extract object-level representations from videos. ViMON [25] expands MoNet to videos by adding a gated recurrent unit to aggregate information over time. Similarly, Steve [22] applies slot attention to more complex and naturalistic videos by using a transformer-based image decoder conditioned on slots.

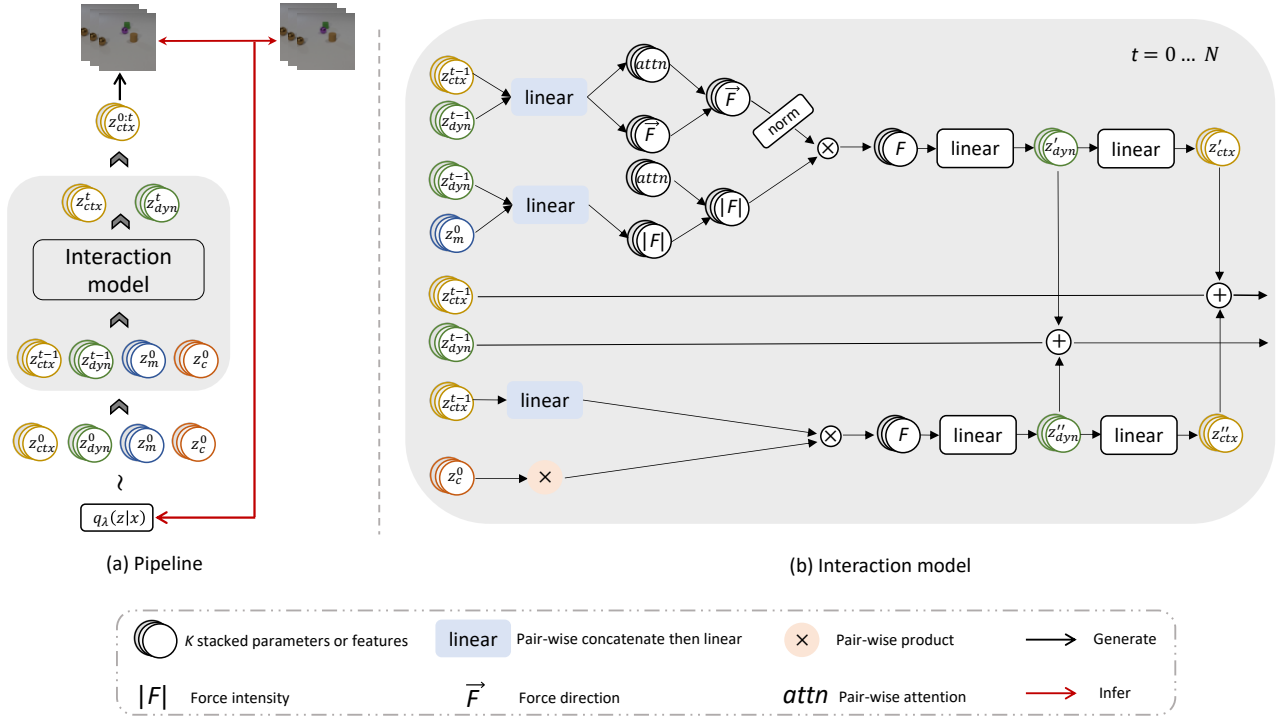


Figure 2. (a). Schematic of PHYCINE. Given a video with images of  $x^0 \cdot x^t$ , we aim to get four physical embeddings of each object, including context  $z_{ctx}$ , dynamics  $z_{dyn}$ , mass  $z_m$  and charge  $z_c$ . Firstly, representations of the time-step 0 are sampled from an arbitrary posterior distribution. Then,  $z_{ctx}^0$  is required to reconstruct image  $x^0$ , and the interaction model takes as inputs  $z_{ctx}^0, z_{dyn}^0, z_m^0$ , and  $z_c^0$  to iteratively output  $z_{ctx}^t$ . In this way, a video could be rendered. The gradient information from image generation is used to iteratively refine the posterior, which could finally accurately represent the physical concepts of all present objects. (b). Interaction model that models collision force and charge force with all the identical nodes pointing to one shared tensor.

**Object-centric predictive models.** Object-centric models are used to reach multi-object 3D scene prediction with or without action conditioning. PROVIDE [29] fits IODINE to videos leveraging a 2D-LSTM and can be used to predict the trajectories of each object separately. ODDN [23] achieves prediction with interaction by explicitly distilling object dynamics and modeling pair-wise object relations. In C-SWM [14], contrastive learning is proposed to learn a transition model in latent space. OP3 [24] uses IODINE as the inference model with action conditioned, achieving object-centric prediction and planning.

**Physical reasoning.** Our work is also closely related to research on learning scene dynamics for physical and causal reasoning. The earlier work focus on image or environment with symbolic representation. CoPhy [2] studies physical dynamics prediction in a counterfactual setting. CATER [8] introduces a synthetic video dataset for temporal reasoning associated with collocational actions. CLEVRER [28] incorporating dynamics modeling with compositional causal reasoning, and grounding the reasoning tasks to the language domain. ComPhy [5] studies objects' intrinsic physical properties from objects' interactions and how these properties affect their motions in

future and counterfactual scenes to answer corresponding questions.

### 3. Method

Our goal is to learn an object-oriented and disentangled representation of physical concepts from observing the environment state. In addition, we would like to learn a predictive system that models interactions between objects under the adjustment of intrinsic physical concepts. We explicitly factorize the object representations by segmenting the latent variables  $z$  into contexts  $z_{ctx}$ , dynamics  $z_{dyn}$ , mass  $z_m$ , and charge  $z_c$  as shown in Figure 2.  $z_{ctx}$  models all the extrinsic properties, including color, shape, material, location, size, etc.  $z_{ctx}$  can be inferred by reconstructing the current time step image.  $z_{dyn}$  mainly encodes velocity which directly determines the transition of  $z_{ctx}$  to the next time step, and could be inferred by one-step prediction with a state transition function. In simple cases,  $z_{ctx}$  and  $z_{dyn}$  have been enough to simulate the scenes with only uniform linear motions. However, when objects exert forces on each other through different forms of interactions, such as attraction, repulsion, or collision, the object velocities would be affected by the forces, and  $z_m$  and  $z_c$  are required to

model these interactions and update  $z_{dyn}$ . During training, extrinsic physical concepts can be inferred from the reconstruction of the image, and intrinsic physical concepts may emerge by enforcing the system to perform longer-term predictions. Besides, increasingly abstract levels of concepts are acquired in a hierarchical fashion, and the complete object representation is composed by the cascade of concepts of different abstract levels.

We start by introducing a generative model of multi-object representations that infers object contexts in Section 3.1. Then, in Section 3.2, we explain how PHYCINE infers object dynamics through short-term prediction. Finally, we introduce the interaction model and the way it helps infer mass and charge with long-term prediction in Section 3.3.

### 3.1. Learning Object Contexts with Reconstruction

To achieve a systematic generalization like humans, the multi-slot framework introduced by IODINE decomposes a static image into  $K$  slots, and each ideally describes an independent object or background of the input scene. The slots can capture the object’s unique contexts such as color, shape, location, and size in a disentangled style. The image  $x$  is modeled with a spatial Gaussian mixture model where each slot corresponds to a single object. With a generative model, each latent vector  $z_{ctx}^{(k)}$  could be decoded into a pixel-wise mean  $\mu_i^{(k)}$  and a pixel-wise mask  $m_i^{(k)}$  describing the assignment probability and appearance of pixel  $i$  for object  $k$ . The generative mechanism is shared such that any ordering of the  $K$  slots produces the same image. The  $K$  mask logits are normalized with a softmax function to ensure that the sum of  $K$  mask logits for each pixel is 1:

$$p(x|z_{ctx}) = \prod_{i=1}^D \sum_{k=1}^K m_i^{(k)} \mathcal{N}(x_i; \mu_i^{(k)}, \sigma^2), \quad (1)$$

where  $\mu$  and  $m$  parameterize the final spatial Gaussian mixture distribution, and  $\sigma$  is the same and fixed for all  $i$  and  $k$ .

The representation is inferred by amortized variational inference:

$$z_{ctx} \sim q_\lambda(z_{ctx}|x); \lambda = \{\mu, \sigma\}. \quad (2)$$

$$\lambda \leftarrow \lambda + f(z_{ctx}, x, \mu, m, \nabla_\mu p(x|z_{ctx}), \nabla_m p(x|z_{ctx})). \quad (3)$$

$q(z_{ctx}|x)$  is the posterior probability of  $z_{ctx}$  given the image  $x$ , with the parameters of  $\lambda$ ,  $f$  is a network for parameter updating. The posterior parameters start with an arbitrary guess, and the latent representations  $z_{ctx}$  are firstly sampled from  $q_\lambda$ . In each refinement step, input image  $x$  together with a set of variables (means  $\mu$ , masks  $m$ , mean gradient  $\nabla_\mu p(x|z_{ctx})$ , mask gradient  $\nabla_m p(x|z_{ctx})$ ), which are computed from  $z$  and  $x$  are fed into a refinement network  $f$  to obtain the additive updates for the posterior.

### 3.2. Learning Object Dynamics with Short-term Prediction

Since object appearance is sufficient to conduct state image reconstruction, the object-context model can not infer dynamics. Given  $z_{ctx}$ , PHYCINE represents object dynamics  $z_{dyn}$  by extending the idea of amortized variational inference. From the perspective of common sense, object dynamics (e.g., velocity) transit object contexts (location for rigid objects in 3D scene) in each time step. PHYCINE learns a transition function that outputs next-step contexts  $z_{ctx}^{t+1}$  with  $z_{ctx}^t$  and  $z_{dyn}^t$  as inputs. Mathematically:

$$z_{ctx}^{t+1} = FC(z_{dyn}^t) + z_{ctx}^t. \quad (4)$$

In the beginning, the predicted  $z_{ctx}^{t+1}$  could be arbitrary values, and as the training goes on, the transition function is enforced to apply dynamics to location and depth other than other irrelevant concepts, e.g., shape and color. Specifically,  $z_{ctx}^{t+1}$  is required to generate  $x^{t+1}$ , with the shared decoder as used in the reconstruction task. In order to predict the future frame image,  $z_{dyn}^t$  need to be well inferred. This process could be done in one-step prediction, and could be described as:

$$z_{dyn}^t \sim q_\lambda(z_{dyn}^t|x^{t+1}, z_{ctx}^t). \quad (5)$$

We can see that  $z_{ctx}$  could be directly inferred by observing the image, and the discovery of  $z_{dyn}$  lies in the correctness of predicting the transition of  $z_{ctx}$ , which makes  $z_{dyn}$  of a higher abstraction level than  $z_{ctx}$ . However, object dynamics could be constantly changed because of intrinsic physical properties. Therefore, when higher level physical concepts get involved,  $z_{dyn}^t$  could only predict  $z_{ctx}^{t+1}$  correctly, since  $z_{dyn}^{t+1}$  may be updated by interactions between objects, which are defined by some intrinsic concepts.

### 3.3. Learning Mass and Charge with Long-term Prediction

In order to model the change in object dynamics, the system should consider the intrinsic concepts and disentangle the interactions from different sources. In turn, carrying out longer-term predictions endows the system with the capability of inferring the intrinsic physical concepts. To achieve this, PHYCINE explicitly designs collision force, as well as the attraction and repulsion force (i.e., the charge and force).

The interaction model (Figure 2 (b)) incorporates collision force and charge force to transit  $z_{ctx}$  and  $z_{dyn}$  from time step  $t = 0$  to  $t = N$ . By enforcing  $z_m$  and  $z_c$  to interpret the two kinds of forces, and keeping  $z_m$  and  $z_c$  constant throughout the prediction process, the model can adaptively learn the concepts of mass and charge in these two variables.

**Collision.** Objects exert force on each other and change their velocities at the moment of collision. The system should capture the moment when a collision event occurs, and how the velocities of the participants change. In the model,  $z_{ctx}$  and  $z_{dyn}$  are used to estimate whether objects are getting close and the direction of the force.  $z_{dyn}$  and  $z_m$  are used to estimate the intensity of the force.

Specifically, we first concatenate the contexts and dynamics of each object-pair and perform self-attentions to get pair-wise attentions  $attn$  and force directions  $\vec{F}$ :

$$attn^{k_i, k_j} = W([z_{ctx}^{(k_i)}; z_{dyn}^{(k_i)}; z_{ctx}^{(k_j)}; z_{dyn}^{(k_j)}]), \quad (6)$$

$$\vec{F}^{k_i, k_j} = W([z_{ctx}^{(k_i)}; z_{dyn}^{(k_i)}; z_{ctx}^{(k_j)}; z_{dyn}^{(k_j)}]), \quad (7)$$

where attention  $attn^{k_i, k_j}$  represents whether or not objects  $k_i$  and  $k_j$  are close enough and decide whether a force is exerted, and  $\vec{F}^{k_i, k_j}$  indicates the directions of the pair-wise forces. Then, we concatenate the mass and dynamics of each object-pair and estimate the force intensities:

$$|F|^{k_i, k_j} = W([z_m^{(k_i)}; z_{dyn}^{(k_i)}; z_m^{(k_j)}; z_{dyn}^{(k_j)}]). \quad (8)$$

With the interaction attention, force directions, and force intensities, we can get the complete force representation for each object:

$$F^{k_i} = \sum_{j \neq i} l2norm(\vec{F} * attn) * \sum_{j \neq i} (|F| * attn), \quad (9)$$

where l2 normalization ensures that  $\vec{F}$  only carries the direction information.

**Attraction and Repulsion.** The forces of attraction and repulsion are also pair-wise interactions. Unlike the collision force, repulsion and attraction do not happen in an instant but exist all the time.

Assuming that all object pairs have a universal charge force, we estimate it from  $z_{ctx}$  and  $z_c$ . The pair-wise product of  $z_c$  can be positive or negative, or zero:

$$F^{k_i} = \sum_{j \neq i} W([z_{ctx}^{(k_i)}; z_{ctx}^{(k_j)}]) * z_c^{(k_i)} * z_c^{(k_j)}. \quad (10)$$

Forces from both modules update  $z_{dyn}$  and further update  $z_{ctx}$  to the next time step, and the inference of mass and charge could be formulated as:

$$z_m^t, z_c^t \sim q\lambda(z_m^t, z_c^t | x^{\leq t+N}, z_{ctx}^t, z_{dyn}^t). \quad (11)$$

### 3.4. Objective

In our system, Besides the initial posterior  $\lambda$ , the parameters of three models need to be optimized: parameters of

Table 1. Dimensions of physics latent codes, and their abstraction levels and inference conditions among different concepts.

Concept	dim	abstraction	inference
Contexts	12	low	image reconstruction
Dynamics	2	middle	short-term prediction
Mass	1	high	long-term prediction
Charge	1	high	long-term prediction

the generative model  $p_\theta$  that generate images from  $z_{ctx}$ ; parameters of the inference model  $q_\phi$  that updates the posterior parameters using gradient information about the ELBO obtained by the current estimate of the parameters; parameters of the interaction model  $r_\delta$  which roll out a sequence of  $z_{ctx}^{1:N}$ :

$$z_{ctx}^0, z_{dyn}^0, z_m^0, z_c^0 \sim \lambda. \quad (12)$$

$$z_{ctx}^{1:N}, z_{dyn}^{1:N} = r_\delta(z_{ctx}^0, z_{dyn}^0, z_m^0, z_c^0). \quad (13)$$

We train all the parameters by gradient descent through the reconstruction and prediction objective. The training loss  $\mathcal{L}$  is the negative ELBO which could be described as:

$$\mathcal{L} = \sum_{t=0}^N (-\log(p_\theta(x^t | z_{ctx}^t)) + \beta KL(q_\phi(z^t | x^t, z^{<t}) || p(z))) \quad (14)$$

The first term is the decoder negative log-likelihood, given our mixture of decoder distribution, as discussed above. The second term is the Kullback–Leibler divergence (KL) divergence of the latent posterior with the latent prior, weighted with a hyperparameter  $\beta$ .

### 3.5. Preventing Collapse by Regularization

The way of predicting  $z_{ctx}^{t:t+N}$  based on  $z_{ctx}^t$  with a series of latent variables (i.e.,  $z_{dyn}^t, z_m^t$ , and  $z_c^t$ ) is easy to collapse when the latent variable has excessive information capacity. For example, given  $z_{ctx}^t$ , for all  $z_{ctx}^{t+1}$  there exists a trivial solution that produces zero prediction loss where  $z_{dyn}^t$  being identity with  $z_{ctx}^{t+1}$ .

One promising way to prevent collapse is minimizing the information content of the latent variable used in the prediction. We conduct this solution by pre-defining  $z_{dyn}, z_m$ , and  $z_c$  as low-dimensional variables. Another purpose of this is that the low-dimensional concepts are more unambiguous and thus conducive to the discovery of the upper-level concepts. For example, 2-dimensional  $z_{dyn}$  represents two orthogonal velocity directions, which remain unchanged after normalization, and this is the basis for PHYCINE to learn the concept of object mass.

In addition to limiting the capacity of latent variables, we also limit their functions by making them work in a regularized way. As demonstrated in Figure 2:

Table 2. Question-answering accuracy of visual reasoning models on COMPHY. We report per-option and per-question accuracies for each sub-task. CPL uses a supervised object detector to obtain object contexts, such as Faster/Mask R-CNN [10,21], then learns representations of mass and charge with object-centric labels as supervision. We omit the self-supervised learning strategy of ALOE for implementation simplicity.

Model	learned physics	Factual	Predictive		Counterfactual	
			per opt.	per ques.	per opt.	per ques.
CNN-LSTM [1]	-	46.6	59.5	29.8	58.6	14.6
HCRN [15]	-	47.3	62.7	32.7	58.6	14.2
MAC [12]	-	<b>68.6</b>	60.2	32.2	60.2	16.0
$CPL^{phy-label}$ [5]	ctx+m+c	80.5	75.3	56.4	68.3	29.1
ALOE-IODINE [6,9]	ctx	54.4	60.7	29.1	63.8	20.9
ALOE-PHYCINE	ctx+dyn+m+c	59.1 (+4.7)	<b>63.3</b> (+2.6)	<b>33.4</b> (+4.3)	<b>68.2</b> (+4.4)	<b>25.2</b> (+4.3)

Table 3. Ablation experiments. We study the contribution of each concept representation and the importance of the bottom-to-up training strategy and interaction modeling. We adjust the contexts dimension to ensure that all compared features are the same dimension.

	Factual	Predictive		Counterfactual	
		per opt.	per ques.	per opt.	per ques.
ctx	54.4	60.7	29.1	63.8	20.9
ctx + m	55.2	62.8	31.5	65.5	22.6
ctx + c	55.7	61.8	31.3	66.0	22.1
ctx + dyn	55.5	63.1	33.2	64.4	21.5
ctx+dyn+m+c	<b>59.1</b>	<b>63.8</b>	<b>34.2</b>	<b>68.2</b>	<b>25.2</b>
w/o b2u training	58.0	63.4	33.6	67.0	22.8
w/o interaction	47.3	57.8	25.5	63.9	20.1

- $z_m$  together with  $z_{ctx}$  and  $z_{dyn}$  model the collision.
- $z_c$  and  $z_{ctx}$  model the charge force.
- $z_{dyn}^t$  (updated after interaction) outputs the deltas between  $z_{ctx}^t$  and  $z_{ctx}^{t+1}$ .

### 3.6. Training Strategy.

In PHYCINE, the concept of mass is used to estimate the intensity of the force of a collision event, which requires the system to first learn the concept of collision. Besides, both collision and charge forces can change objects’ velocity, it is difficult to disentangle them when joint training these two interaction mechanisms. Thus, we train PHYCINE in a bottom-up fashion with three stages. In the first stage, the model learns the concepts of context, dynamics, and collision. Specifically, we train the model on the scenes where the objects have the same mass and are uncharged. We randomly activate the collision module and enables the model to learn small attention value and dynamic change when a collision is not happening. In the second stage, we learn how mass affects the collision by training the model on scenes where objects have different masses and are uncharged with a longer-term prediction objective. In the third stage, we learn the concept of charge in the complete dataset. It is worth noting that, none of the physics labels are employed to supervise the model.

## 4. Experiments

In experiments, we evaluate PHYCINE on the ComPhy dataset to illustrate the effectiveness of the learned representation. We also detaily analyze the interpretability of the discovered physical concepts.

### 4.1. Dataset

Compositional Physical Reasoning (ComPhy) studies objects’ intrinsic physical properties from objects’ interactions and how these properties affect their motions in future and counterfactual scenes to answer corresponding questions. Objects in ComPhy contain compositional appearance attributes like color, shape, and material. ComPhy also studies intrinsic physical concepts, mass, and charge, which can not be directly captured from objects’ static appearance. Thus, in addition to collision, there are also attraction and repulsion events in ComPhy. Accordingly, there are questions related to object concepts of mass and charge, e.g., ‘Is the cube heavier than the cylinder’ in Factual questions and ‘If the sphere were oppositely charged, what would happen?’ in counterfactual questions. ComPhy has 8,000 sets for training, 2,000 sets for validation, and 2,000 for testing.

### 4.2. Video Understanding and Causal Reasoning

**Baselines.** Approaches on causal reasoning benchmarks [5,28] can be divided into two paradigms: neural networks [6,12,23,26] and neuro-symbolic models [4,5,7,28]. Neuro-symbolic models leverage various independently-learned modules and perform better than neural network baselines. ALOE is a state-of-the-art cross-modal transformer architecture that takes in object-centric visual features and word embeddings as inputs, then reasons the answers. ALOE performs even better than the neural-symbolic models in CLEVRER. In this work, we focus on object-centric learning and verify the learned representation with ALOE.

**Results on ComPhy.** We conduct experiments on ComPhy against several counterparts: CNN-LSTM, HCRN, MAC, CPL, and ALOE. Among them, CPL is a




Figure 3. Example result of PHYCINE inferring  $z^0$  given 6 frames, then regenerating the video and imagining the future (shown with 4 more frames).

high-performance interpretable symbolic model. MAC decomposes visual question answering into multiple attention-based reasoning steps and predicts the answer based on the last step. ALOE is a state-of-the-art end-to-end method on CLEVRER, which serves as our baseline.

We can observe from Table 2 that, video question-answering models CNN-LSTM and HCRN perform worse than the others. The reason is that models like HCRN are mostly designed for human action video, which mainly focuses on temporal modeling rather than spatial-temporal modeling of compositional physical events. CPL performs the best overall. However, explicit signals of mass and charge are used as supervision to learn the representation, pre-trained object detector is also required. MAC outperforms ALOE-IODINE on factual questions, showing the effectiveness of its compositional attention architecture. ALOE-IODINE gains the competitive edge on predictive and counterfactual questions compared with MAC. When feeding ALOE with PHYCINE features instead of IODINE, performances are boosted in all types of questions, which indicates the effectiveness of the learned representation of the intrinsic physical concepts.

We also study how a single type of physical concept (i.e., dynamics, mass, and charge) affects the results where feature dimensions are kept equal for all experiments (Table 3). For factual questions, three types of concept representation are beneficial, and we believe it is quite reasonable. Concretely, dynamics could help answer questions like ‘Are there metal cubes that enter the scene?’. mass and charge could be helpful to reason the relative weight of two objects or how they are charged. dynamics is the most effective concept in reasoning predictive questions, compared with mass or charge. This is easy to explain: It is more straightforward to estimate what will happen with an explicit representation of object velocity. In counterfactual questions, the object’s weight or charge is assumed to be a counterfactual value, then the model is required to measure what subsequent events will happen. In this setting, concepts of mass and charge are crucial, which is validated in the experiments (Table 2: mass and charge improve the per-option accuracy of counterfactual questions from 63.8 to 65.5 and 66.0 respectively. More importantly, not only the representation of every single concept contributes to the results, their combi-

nation further boosts the performance.

### 4.3. Video Regeneration and Prediction

PHYCINE infers multi-object representations with intrinsic physical concepts of one image by fitting  $N$  time-steps of future frames. In other words,  $z^0$  is inferred from  $x^0$  to  $x^N$ . If the obtained representation is correct,  $z^0$  should be capable of regenerating the watched video clip by firstly predicting in the latent space to get  $z^1$  to  $z^N$ , then rendering the contexts representation to video.

Evaluating how accurately the system can regenerate the given video with one set of latent representations is not enough to prove that the physical concepts are well understood. One reason is that  $z^0$  may only overfit a given video without involving any physical laws (e.g.,  $z_{dyn}$ ,  $z_m$ , and  $z_c$  encode the object location transition rather than actual intrinsic physical concepts). In order to verify whether PHYCINE fully understands how objects interact with the world with various physical concepts, we test if the system could further predict images from  $x^N$  to  $x^T$ . Note that We only train PHYCINE to watch and predict (re-generate from the first time-step representation) images for  $N$  time-steps.

Figure 3 shows qualitative results for the regenerated and predicted trajectories. There are two events in this dynamic scene, a collision between the sphere and the cube, and a repulsion between the sphere and the cylinder. In the beginning, the sphere and the cylinder with the same charge move toward each other. Then, the sphere collides with the stationary cube and changes its velocity. At last, the sphere and the cylinder accelerate away from each other because of the repulsion force, and the cube moves with a uniform velocity. We can observe from the result that, PHYCINE could accurately capture interactions from collision force and charge force. Besides, PHYCINE can also imagine the future without further training, which shows the interpretability of our system.

### 4.4. Interpretability and Disentanglement

Disentanglement is a desirable property of representations that measures how the learned features correspond to individual, interpretable attributes of objects in the scene. IODINE, as our baseline framework, encourages the emergence of object context concepts in the slots with the de-

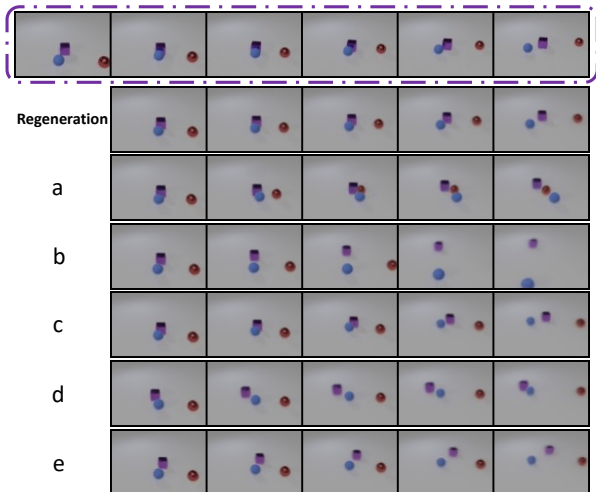


Figure 4. Disentanglement of the discovered intrinsic physical concepts, i.e., dynamics, mass, and charge, with counterfactual prediction.

coder implementing an inductive bias about the scene composition. In the design of PHYCINE, we expect the disentanglement also be maintained for the intrinsic concepts, i.e., dynamics, mass, and charge. In this section, we demonstrate the disentanglement with counterfactual prediction, which means manually adjusting the value of a certain concept and observing the corresponding effects.

Figure 4 shows the comparison between 5 counterfactual situations ( $a \rightarrow e$ ) and the original inferred result, where two spheres have the same charge and are initially getting close to each other and then repelled by the charge force. There is also a collision between the blue sphere and the uncharged and stationary cube:

- (a) When giving the blue sphere an opposite charge, the repulsion force turns into an attraction force and is followed by two collisions, the cube, and the red sphere, and the two spheres, which proves the correctness and interpretability of the concepts of charge and collision.
- (b) When giving the uncharged cube the same charge as spheres, three objects repel and away from each other.
- (c) When turning down the mass value of the cube, the collision brings more changes to its velocity. On the contrary, it gives less velocity change to the opponent. This proves the correctness and interpretability of the concept of mass.
- (d) When giving the cube an initial velocity in the first dimension of  $z_{dyn}$ , the cube heads in one direction and avoids the collision.
- (e) When giving the cube an initial velocity in the second dimension of  $z_{dyn}$ , the cube heads in another direction and also avoids the collision.

The counterfactual experiments reveal the mechanisms of each dimension of the learned variables, which on the

one hand consistent with the desired physical concepts, and on the other hand, are disentangled.

#### 4.5. Bottom-up Training and Regularization

We conduct experiments to verify the importance of the bottom-up training strategy and regularization, with quantity results shown in Table 3.

Though regularized with the pre-defined interaction model, both mass and charge have access to directly monitor and adjust contexts, instead of through dynamics. When training step by step, the model gains a better and more accurate interpretation of the world with the accumulation of knowledge. Results indicate that the representation quality without bottom-up training decreased in the case of answering questions that related to concepts of mass and charge, i.e., Factual and Counterfactual questions.

We get rid of regularization completely by omitting the interaction model, which means that the model needs to compress the input video into one latent variable and a simple state transition function. Results show that not only no meaningful conceptual features are learned, but also the finite-dimensional features are even worse at describing the current frame contexts than the baseline model (ctx *vs* w/o interaction).

### 5. Discussion and Conclusion

Relations and interactions among objects are governed by physical laws which helps us interpret the configurations of objects accurately [17]. In this work, we studied the configuration of intrinsic physical concepts (mass and charge) for objects, by observing their interactions. We proposed PHYCINE, a framework that infers disentangled physical concepts of different abstraction levels by fitting future physical events with amortized inference, trained using only raw input data without supervision. Based on the visualization of counterfactual predictions, it is confirmed that the inferred concepts are interpretable. The discovery of the intrinsic physical concepts helps to understand and predict physical events, and richer representations of video could be obtained, which leads to better performance in solving downstream reasoning tasks. Lastly, PHYCINE has shown the practicability of discovering and configuring intrinsic physical concepts with object-centric predictive models, it seems worth investigating further in this direction. [27]

#### Acknowledgments

This work was supported in part by the National Key Research & Development Program (No. 2020YFC2003901), Chinese National Natural Science Foundation Projects #62176256, #62276254, #61976229, the Youth Innovation Promotion Association CAS (#Y2021131), and the InnoHK program.



## References

- [1] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433, 2015. 6
- [2] Fabien Baradel, Natalia Neverova, Julien Mille, Greg Mori, and Christian Wolf. Cophy: Counterfactual learning of physical dynamics. *arXiv preprint arXiv:1909.12000*, 2019. 3
- [3] Christopher P Burgess, Loic Matthey, Nicholas Watters, Rishabh Kabra, Irina Higgins, Matt Botvinick, and Alexander Lerchner. Monet: Unsupervised scene decomposition and representation. *arXiv preprint arXiv:1901.11390*, 2019. 1, 2
- [4] Zhenfang Chen, Jiayuan Mao, Jiajun Wu, Kwan-Yee Kenneth Wong, Joshua B Tenenbaum, and Chuang Gan. Grounding physical concepts of objects and events through dynamic visual reasoning. *arXiv preprint arXiv:2103.16564*, 2021. 6
- [5] Zhenfang Chen, Kexin Yi, Yunzhu Li, Mingyu Ding, Antonio Torralba, Joshua B Tenenbaum, and Chuang Gan. Comphy: Compositional physical reasoning of objects and events from videos. *arXiv preprint arXiv:2205.01089*, 2022. 1, 2, 3, 6
- [6] David Ding, Felix Hill, Adam Santoro, Malcolm Reynolds, and Matt Botvinick. Attention over learned object embeddings enables complex visual reasoning. *Advances in neural information processing systems*, 34:9112–9124, 2021. 6
- [7] Mingyu Ding, Zhenfang Chen, Tao Du, Ping Luo, Josh Tenenbaum, and Chuang Gan. Dynamic visual reasoning by learning differentiable physics models from video and language. *Advances In Neural Information Processing Systems*, 34:887–899, 2021. 6
- [8] Rohit Girdhar and Deva Ramanan. Cater: A diagnostic dataset for compositional actions and temporal reasoning. *arXiv preprint arXiv:1910.04744*, 2019. 3
- [9] Klaus Greff, Raphaël Lopez Kaufman, Rishabh Kabra, Nick Watters, Christopher Burgess, Daniel Zoran, Loic Matthey, Matthew Botvinick, and Alexander Lerchner. Multi-object representation learning with iterative variational inference. In *International Conference on Machine Learning*, pages 2424–2433. PMLR, 2019. 1, 2, 6
- [10] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 6
- [11] Jun-Ting Hsieh, Bingbin Liu, De-An Huang, Li F Fei-Fei, and Juan Carlos Niebles. Learning to decompose and disentangle representations for video prediction. *Advances in neural information processing systems*, 31, 2018. 2
- [12] Drew A Hudson and Christopher D Manning. Compositional attention networks for machine reasoning. *arXiv preprint arXiv:1803.03067*, 2018. 6
- [13] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 1
- [14] Thomas Kipf, Elise Van der Pol, and Max Welling. Contrastive learning of structured world models. *arXiv preprint arXiv:1911.12247*, 2019. 2, 3
- [15] Thao Minh Le, Vuong Le, Svetha Venkatesh, and Truyen Tran. Hierarchical conditional relation networks for video question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9972–9981, 2020. 6
- [16] Yann LeCun. A path towards autonomous machine intelligence. *preprint posted on openreview*, 2022. 2
- [17] Wenbin Li, Ales Leonardis, and Mario Fritz. Visual stability prediction and its application to manipulation. In *2017 AAAI Spring symposium series*, 2017. 8
- [18] Francesco Locatello, Dirk Weissenborn, Thomas Unterthiner, Aravindh Mahendran, Georg Heigold, Jakob Uszkoreit, Alexey Dosovitskiy, and Thomas Kipf. Object-centric learning with slot attention. *Advances in Neural Information Processing Systems*, 33:11525–11538, 2020. 1, 2
- [19] Gerrit W Maus, Jason Fischer, and David Whitney. Motion-dependent representation of space in area mt+. *Neuron*, 78(3):554–562, 2013. 2
- [20] Nora Nortmann, Sascha Rekauzke, Selim Onat, Peter König, and Dirk Jancke. Primary visual cortex represents the difference between past and present. *Cerebral Cortex*, 25(6):1427–1440, 2015. 2
- [21] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015. 6
- [22] Gautam Singh, Yi-Fu Wu, and Sungjin Ahn. Simple unsupervised object-centric learning for complex and naturalistic videos. *arXiv preprint arXiv:2205.14065*, 2022. 1, 2
- [23] Qu Tang, Xiangyu Zhu, Zhen Lei, and Zhaoxiang Zhang. Object dynamics distillation for scene decomposition and representation. In *International Conference on Learning Representations*, 2021. 1, 2, 3, 6
- [24] Rishi Veerapaneni, John D Co-Reyes, Michael Chang, Michael Janner, Chelsea Finn, Jiajun Wu, Joshua Tenenbaum, and Sergey Levine. Entity abstraction in visual model-based reinforcement learning. In *Conference on Robot Learning*, pages 1439–1456. PMLR, 2020. 2, 3
- [25] Marissa A Weis, Kashyap Chitta, Yash Sharma, Wieland Brendel, Matthias Bethge, Andreas Geiger, and Alexander S Ecker. Benchmarking unsupervised object representations for video sequences. *J. Mach. Learn. Res.*, 22:183–1, 2021. 2
- [26] Ziyi Wu, Nikita Dvornik, Klaus Greff, Thomas Kipf, and Animesh Garg. Slotformer: Unsupervised visual dynamics simulation with object-centric models. *arXiv preprint arXiv:2210.05861*, 2022. 6
- [27] Sirui Xie, Ari S Morcos, Song-Chun Zhu, and Ramakrishna Vedantam. Coat: Measuring object compositionality in emergent representations. In *International Conference on Machine Learning*, pages 24388–24413. PMLR, 2022. 8
- [28] Kexin Yi, Chuang Gan, Yunzhu Li, Pushmeet Kohli, Jiajun Wu, Antonio Torralba, and Joshua B Tenenbaum. Clevrer: Collision events for video representation and reasoning. *arXiv preprint arXiv:1910.01442*, 2019. 3, 6
- [29] Polina Zablotskaia, Edoardo A Dominici, Leonid Sigal, and Andreas M Lehrmann. Unsupervised video decomposition

using spatio-temporal iterative inference. *arXiv preprint*  
*arXiv:2006.14727*, 2020. [1](#), [3](#)