

Hierarchical Relational Inference

Aleksandar Stanić, Sjoerd van Steenkiste, Jürgen Schmidhuber

Swiss AI Lab IDSIA, USI, SUPSI
Lugano, Switzerland
{aleksandar, sjoerd, juergen}@idsia.ch

Abstract

Common-sense physical reasoning in the real world requires learning about the interactions of objects and their dynamics. The notion of an abstract object, however, encompasses a wide variety of physical objects that differ greatly in terms of the complex behaviors they support. To address this, we propose a novel approach to physical reasoning that models objects as hierarchies of parts that may locally behave separately, but also act more globally as a single whole. Unlike prior approaches, our method learns in an unsupervised fashion directly from raw visual images to discover objects, parts, and their relations. It explicitly distinguishes multiple levels of abstraction and improves over a strong baseline at modeling synthetic and real-world videos.

1 Introduction

Common-sense physical reasoning in the real world involves making predictions from complex high-dimensional observations. Humans somehow discover and represent abstract objects to compactly describe complex visual scenes in terms of ‘building blocks’ that can be processed separately (Spelke and Kinzler 2007). They model the world by reasoning about dynamics of high-level objects such as footballs and football players and the consequences of their interactions. It is natural to expect that artificial agents operating in the real world will benefit from a similar approach (Lake, Salakhutdinov, and Tenenbaum 2015; Greff, van Steenkiste, and Schmidhuber 2020).

Real world objects vary greatly in terms of their properties. This complicates modelling their dynamics. Some have deformable shapes, e.g., clothes, or consist of parts that support a variety of complex behaviors, e.g., arms and fingers of a human body. Many objects can be viewed as a *hierarchy* of parts that locally behave somewhat independently, but also act more globally as a single whole (Mrowca et al. 2018; Lingelbach et al. 2020). This suggests to simplify models of object dynamics by explicitly distinguishing multiple levels of abstraction, separating hierarchical sources of influence.

Prior approaches to common-sense physical reasoning explicitly consider objects and relations at a representational level, e.g., (Chang et al. 2017; Battaglia et al. 2016; van Steenkiste et al. 2018; Kipf et al. 2018). They decompose

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

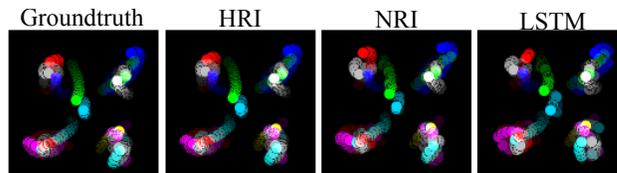


Figure 1: HRI outperforms baselines at modeling interacting objects that are coupled via hierarchically organized springs.

complex physical interactions in the environment into pairwise interactions between objects, modelled efficiently by Graph Networks (Battaglia et al. 2018). Here the representation of each object is updated at each time step by propagating ‘messages’ through the corresponding interaction graph. While recent approaches (specifically) address the challenge of learning object representations from raw visual data (Greff, van Steenkiste, and Schmidhuber 2017; Kosiorek et al. 2018; van Steenkiste et al. 2018; Burgess et al. 2019; Greff et al. 2019) and of dynamically inferring relationships between objects (van Steenkiste et al. 2018; Kipf et al. 2018; Goyal et al. 2019; Veerapaneni et al. 2019), reasoning about the dynamics and interactions of complex objects remains difficult without incorporating additional structure. On the other hand, approaches that consider part-based representations of objects and hierarchical interaction graphs lack the capacity to learn from raw images and dynamically infer relationships (Mrowca et al. 2018; Lingelbach et al. 2020).

Here we propose *Hierarchical Relational Inference* (HRI), a novel approach to common-sense physical reasoning capable of learning to discover objects, parts, and their relations, directly from raw visual images in an unsupervised fashion. HRI extends *Neural Relational Inference* (NRI) (Kipf et al. 2018), which infers relations between objects and models their dynamics while assuming access to their state (e.g., obtained from a physics simulator). HRI improves upon NRI in two regards. Firstly, it considers part-based representations of objects and infers *hierarchical* interaction graphs to simplify modeling the dynamics (and interactions) of more complex objects. This necessitates a more efficient message-passing approach that leverages the hierarchical structure, which we will also introduce. Secondly, it provides a mechanism for applying NRI (and

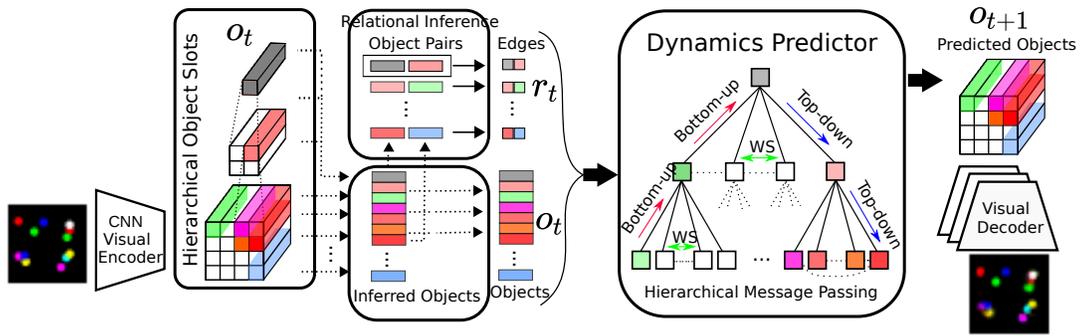


Figure 2: The proposed HRI model. An encoder infers part-based object representations, which are fed to a relational inference module to obtain a hierarchical interaction graph. A dynamics predictor uses hierarchical message-passing to make predictions about future object states. Their ‘rendering’, produced by a decoder, is compared to the next frame to train the system.

thereby HRI) to raw visual images that infers part-based object representations spanning multiple levels of abstraction.

Our main contributions are as follows: (i) We introduce HRI, an end-to-end approach for learning hierarchical object representations and their relations directly from raw visual input. (ii) It includes novel modules for extracting a hierarchical part-based object representations and for hierarchical message passing. The latter can operate on a (hierarchical) interaction graph more efficiently by propagating effects between all nodes in the graph in a single message-passing phase. (iii) On a trajectory prediction task from object states, we demonstrate that the hierarchical message passing module is able to discover the latent hierarchical graph and greatly outperforms strong baselines (Figure 1). (iv) We also demonstrate how HRI is able to infer objects and relations directly from raw images. (v) We apply HRI to synthetic and real-world physical prediction tasks, including real-world videos of moving human bodies, and demonstrate improvements over strong baselines.

2 Method

Motivated by how humans learn to perform common-sense physical reasoning, we propose Hierarchical Relational Inference (HRI). It consists of a *visual encoder*, a *relational inference module*, a *dynamics predictor*, and a *visual decoder*. All are trained end-to-end in an unsupervised manner. First, the visual encoder produces hierarchical (i.e. part-based) representations of objects that are grounded in the input image. This representation serves as input to the relational inference module, which infers pairwise relationships between objects (and parts) given by the edges in the corresponding interaction graph. The dynamics predictor performs hierarchical message-passing on this graph, using the learned representations of parts and objects for the nodes. The resulting predictions (based on the updated representations) are then decoded back to image space using the visual decoder, to facilitate an unsupervised training objective. An overview is shown in Figure 2. We note that HRI consists of standard building blocks (CNNs, RNNs, and GNNs) that are well understood. In this way, we add only a minimal inductive bias, which helps facilitate scaling to more complex real-world visual settings, as we will demonstrate.

2.1 Inferring Objects, Parts, and their Relations

To make physical predictions about a stream of complex visual observations, we will focus on the underlying *interaction graph*. It distinguishes objects or parts (corresponding to nodes) and the relations that determine interactions between them (corresponding to the edges), which must be inferred. Using this more abstract (neuro-symbolic) representation of a visual scene allows us to explicitly consider certain invariances when making predictions (e.g., the number of objects present) and reasoning about complex interactions in terms of simpler pair-wise interactions between objects.

Inferring Object/Part Representations The task of the visual encoder is to infer separate representations for each object from the input image. Intuitively, these representations contain information about its state, i.e., its position, behavior and appearance. In order to relate and compare these representations efficiently, it is important that they are described in a common format. Moreover, since we are concerned with a hierarchical (i.e. part-based) representation of objects, we also require a mechanism to relate the part representations to the corresponding object representation.

Here we address these challenges by partitioning the feature maps learned by a CNN according to their spatial coordinates to obtain object representations. This is a natural choice, since CNNs are known to excel at representation learning for images (Ciregan, Meier, and Schmidhuber 2012; Krizhevsky, Sutskever, and Hinton 2012) and because weight-sharing across spatial locations ensures that the resulting object representations are described in a common format. Indeed, several others have proposed to learn object representations in this way (Santoro et al. 2017; Zambaldi et al. 2019). Here, we take this insight a step further and propose to learn hierarchical object representations in a similar way. In particular, we leverage the insight that the parts belonging to real-world objects tend to be spatially close, to apply a sequence of convolutions followed by down-sampling operations to extract object-level representations from part-level representations (left side of Figure 2). While this leaves the network with ample freedom to develop its own internal notion of an object, we find that representational slots learn

to describe physical objects (Figure 4a).

The choice of kernel-size and degree of down-sampling allow us to adjust how representations at one level of abstraction are combined at the next level. Similarly, the parameters of the CNN layers that produce the initial set of feature maps determine the size of the input region captured by these ‘spatial slots’ (Greff, van Steenkiste, and Schmidhuber 2020) at the lowest level of abstraction. Note the distinction between *parts*, *objects* and *slots*. Parts refer to objects at the lowest level of the visual hierarchy, while the more general notion of an object applies to nodes at all levels. Slots are variable placeholders (of a function) at a representational level, which at each point in time are expected to contain information about a particular object/part. Therefore, an architectural hierarchy of slots reflects a hierarchy of objects. In general, we construct a 3-level part-based hierarchy, which is then fed into the *relational module*.

Neural Relational Inference To infer relations between object representations, we will make use of NRI (Kipf et al. 2018), which learns *explicit* relations. This is advantageous, since it allows one to incorporate prior beliefs about the overall connectivity of the interaction graph (e.g., a degree of sparsity) and associate a representation with each relation to distinguish between multiple different relation types. By default, NRI takes as input a set of object trajectories (states) and infers their pairwise relations (edges) using a Graph Neural Network (GNN) (Scarselli et al. 2009; Gilmer et al. 2017; Battaglia et al. 2018). It assumes a static interaction graph, and performs relational inference by processing the entire input sequence at once, i.e., the whole sequence (length 50) of a particular object is concatenated, and only a single, static, “node embedding” is created via an MLP. In contrast, we will consider a dynamic interaction graph, since objects move across the image and may end up in different spatial slots throughout the sequence. This is achieved by inferring edges at each time step based on the ten most recent object states, concatenating the latent vectors of a particular object and using an MLP to obtain a “node embedding”.

More formally, given a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with nodes (objects) $o \in \mathcal{V}$ and edges (relations) $r_{i,j} = (o_i, o_j) \in \mathcal{E}$, NRI defines a single node-to-node message passing operation in a GNN similar to (Gilmer et al. 2017):

$$\mathbf{e}_{i,j} = f_e([\mathbf{o}_i, \mathbf{o}_j, \mathbf{r}_{i,j}]), \mathbf{o}'_j = f_o([\sum_{i \in \mathcal{N}_{\mathbf{o}_j}} \mathbf{e}_{i,j}, \mathbf{o}_j]) \quad (1)$$

where $\mathbf{e}_{i,j}$ is an embedding (effect) of the relation $\mathbf{r}_{i,j}$ between objects \mathbf{o}_i and \mathbf{o}_j , \mathbf{o}'_j is the updated object embedding, \mathcal{N}_j the set of indices of nodes connected by an incoming edge to object \mathbf{o}_j and $[\cdot, \cdot]$ indicates concatenation. Functions f_o and f_e are node- and edge-specific neural networks (MLPs in practice). By repeatedly applying (1), multiple rounds of message passing can be performed.

The NRI ‘encoder’ receives as input a sequence of object state trajectories $\mathbf{o} = (\mathbf{o}^1, \dots, \mathbf{o}^T)$, which in our case are inferred. It consists of a GNN f_ϕ that defines a probability distribution over edges $q_\phi(\mathbf{r}_{ij}^t | \mathbf{o}^{t-k:t}) = \text{softmax}(f_\phi(\mathbf{o}^{t-k:t})_{ij})$, where k is the window size, and

relations are one-hot encoded. The GNN performs the following message passing operations, where the initial node representations \mathbf{o}_i are obtained by concatenating the corresponding object states across the window:

$$\begin{aligned} \mathbf{o}'_j &= f_o^1(\mathbf{o}_j), \mathbf{e}'_{i,j} = f_e^1([\mathbf{o}'_i, \mathbf{o}'_j]), \mathbf{o}''_j = f_o^2(\sum_{i \neq j} \mathbf{e}'_{i,j}), \\ \mathbf{e}''_{i,j} &= f_e^2([\mathbf{o}''_i, \mathbf{o}''_j]), f_\phi(\mathbf{o}^{t-k:t})_{ij} = \mathbf{e}''_{i,j} \end{aligned}$$

where ϕ contains the parameters of the message-passing functions, which are simple MLPs, and \mathbf{o}' , \mathbf{e}' and \mathbf{o}'' , \mathbf{e}'' are node- and edge-embeddings after first and second message passing operations respectively. To backpropagate through the sampling from $q_\phi(\mathbf{r}_{ij} | \mathbf{o})$, NRI uses a continuous approximation of the discrete distribution to obtain gradients via the reparameterization trick (Maddison, Mnih, and Teh 2017; Jang, Gu, and Poole 2017).

2.2 Physical Reasoning

Physical reasoning is performed by the *dynamics predictor*, which leverages the inferred object representations and edges to predict object states at the next time step. To distinguish between representations at different levels of abstractions, HRI makes use of *hierarchical message passing*. We will also use recurrent units in the non-Markovian setting.

Hierarchical message passing The default approach to physical reasoning based on message-passing employed in NRI can only propagate effects between directly connected nodes. This is costly, as it requires several iterations for information to propagate across the whole graph, especially when the number of nodes increases (a consequence of modeling objects as hierarchies of parts). Alternatively, we can leverage the hierarchical structure of the interaction graph to propagate all effects across the entire graph in a *single* step, i.e. evaluating each relation only once. To achieve this, we introduce a hierarchical message passing module which propagates effects between objects using a three-phase sequential mechanism, loosely inspired by Mrowca et al..

Starting from the leaf nodes, the *bottom-up phase* computes the effect on parent nodes \mathbf{o}_p based on messages from its children, $\mathbf{e}_p^1 = \mathbf{e}_p^0 + f_{MP}^{bu}(\{\mathbf{e}_c^0\}_{c \in \mathcal{C}_p}, \mathbf{e}_p^0, \{\mathbf{r}_{cp}\}_{c \in \mathcal{C}_p})$ where \mathcal{C}_p is the set of children indices of object \mathbf{o}_p and the initial effects \mathbf{e}^0 are simply the object embeddings. In this way, global information is propagated from every node in the hierarchy to the root node. Afterwards, the bottom-up effect \mathbf{e}_i^1 on node \mathbf{o}_i is combined with effects from its siblings (*within-sibling* phase) $\mathbf{e}_i^2 = \mathbf{e}_i^1 + f_{MP}^{ws}(\{\mathbf{e}_s^1\}_{s \in \mathcal{S}_i}, \mathbf{e}_i^1, \{\mathbf{r}_{si}\}_{s \in \mathcal{S}_i})$, where \mathcal{S}_i is the set of sibling indices of object \mathbf{o}_i . Starting from the root node, the *top-down* phase then propagates top-down effects that are incorporated by computing $\mathbf{e}_c^3 = \mathbf{e}_c^2 + f_{MP}^{td}(\mathbf{e}_p^2, \mathbf{e}_c^2, \mathbf{r}_{pc})$ for all children \mathbf{o}_c based on its parent \mathbf{o}_p . Functions f_{MP}^{bu} , f_{MP}^{ws} , and f_{MP}^{td} perform a single node-to-edge and edge-to-node message passing operation as in (1) and have shared weights. Note that this mechanism is independent of the choice of object and relational inference module and can act on any hierarchical interaction graph.

Dynamics predictor Physical reasoning is performed by the dynamics predictor, which predicts future object states $p_\theta(\mathbf{o}^{t+1}|\mathbf{o}^{1:t}, \mathbf{r}^{1:t})$ from the sequence of object states and interactions. We implement this as in the NRI ‘decoder’ (Kipf et al. 2018), i.e. using a GNN that passes messages between objects, but with two notable differences. Firstly, we will pass messages *only* if an edge is inferred between two nodes, as opposed to also considering a separate dynamics predictor for the “non-edge” relation that causes information exchange between unconnected nodes¹. Secondly, we will leverage the hierarchical structure of the inferred interaction graph to perform *hierarchical message-passing*.

If we assume Markovian dynamics, then we have $p_\theta(\mathbf{o}^{t+1}|\mathbf{o}^{1:t}, \mathbf{r}^{1:t}) = p_\theta(\mathbf{o}^{t+1}|\mathbf{o}^t, \mathbf{r}^t)$ and can use hierarchical message passing to predict object states at the next step:

$$p_\theta(\mathbf{o}^{t+1}|\mathbf{o}^t, \mathbf{r}^t) = \mathcal{N}(\mathbf{o}^t + \Delta\mathbf{o}^t, \sigma^2\mathbf{I}), \quad (2)$$

where σ^2 is a fixed variance, $\Delta\mathbf{o}^t = f_O([\mathbf{o}^t, \mathbf{e}^t])$, $\mathbf{e}^t = f_H(\mathbf{o}^t, \mathbf{r}^t)$ is the effect computed by the hierarchical message passing module f_H , and f_O is an output MLP. Notice how we learn to predict the *change* in the state of an object, which is generally expected to be easier. When encoding individual images, no velocity information can be inferred to form the object state. In this non-Markovian case we adapt (2) to include an LSTM (Hochreiter and Schmidhuber 1997) that models $p_\theta(\mathbf{o}^{t+1}|\mathbf{o}^{1:t}, \mathbf{r}^{1:t})$ directly:

$$\begin{aligned} \mathbf{h}^{t+1}, \mathbf{c}^{t+1} &= f_{LSTM}(\mathbf{o}^t, \mathbf{e}^t, \mathbf{c}^t), \quad \mathbf{o}^{t+1} = f_O(\mathbf{h}^{t+1}), \\ p(\mathbf{o}_j^{t+1}|\mathbf{o}^{1:t}, \mathbf{r}^{1:t}) &= \mathcal{N}(\mathbf{o}^{t+1}, \sigma^2\mathbf{I}), \end{aligned}$$

where \mathbf{c} and \mathbf{h} are LSTM’s cell and hidden state respectively, and $\mathbf{e}^t = f_H(\mathbf{h}^t, \mathbf{r}^t)$.

2.3 Learning

Standard approaches to modelling physical interactions that do not assume access to states uses a prediction objective in pixel space (van Steenkiste et al. 2018; Veerapaneni et al. 2019). This necessitates a mechanism to ‘render’ the updated object representations. In this case, HRI can be viewed as a type of Variational Auto-Encoder (Kingma and Welling 2013; Rezende, Mohamed, and Wierstra 2014), where the inferred edges and objects are treated as latent variables, and the ELBO can be maximized for the predicted frames:

$$\begin{aligned} \mathcal{L} = & \mathbb{E}_{q_\phi(\mathbf{r}|\mathbf{x})} [\log p_\theta(\mathbf{x}|\mathbf{r}, \mathbf{o})] - D_{KL}[q_{\phi_o}(\mathbf{o}|\mathbf{x})||p_{\theta_o}(\mathbf{o})] \\ & - D_{KL}[q_{\phi_r}(\mathbf{r}|\mathbf{x}, \mathbf{o})||p_{\theta_r}(\mathbf{r})]. \end{aligned} \quad (3)$$

The relational module $q_{\phi_r}(\mathbf{r}|\mathbf{x}, \mathbf{o})$ outputs a factorized distribution over \mathbf{r}_{ij} , which in our case is a categorical variable that can take on two values (one-hot encoded) that indicate the presence of an edge between \mathbf{o}_i and \mathbf{o}_j . The edge prior $p_{\theta_r}(\mathbf{r}) = \prod_{i \neq j} p_{\theta_r}(\mathbf{r}_{ij})$ is a factorized uniform distribution, which controls the sparsity of the learned graph. The object inference module $q_{\phi_o}(\mathbf{o}|\mathbf{x})$ outputs a factorized distribution over \mathbf{o}_i , and the object prior $p_{\theta_o}(\mathbf{o})$ is Gaussian, as in a standard VAE. Given the inferred interaction graph, the dynamics predictor and visual decoder define $p_\theta(\mathbf{x}|\mathbf{r}, \mathbf{o})$.

¹Although messages are passed only if an edge between two nodes exists, a “non-edge” categorical variable is used to allow the model to infer that there is no edge between two nodes.

Visual Decoder The visual decoder renders the updated object states and we will consider two different implementations of this mapping. The first variant, which we call *SlotDec*, ensures compositionality in pixel space by decoding objects separately followed by a summation to produce the final image. In Figure 2 it is depicted as a set since individual object slots that are decoded separately, and decoders share weights. This implements a stronger inductive bias that encourages each slot to correspond to a particular object (since images are composed of objects) and also makes it easy to inspect the representational content of each slot. On the other hand, summation in pixel space is problematic when objects in scenes are more cluttered and occlude one another. For this reason we implement a second variant, which we call *ParDec*, where all states are decoded together as in a standard convolutional decoder. As a result of decoding all object slots together *ParDec* is less interpretable than *SlotDec*, but computationally more efficient and potentially more scalable to real-world datasets since it does not make strong assumptions on how information about objects should be combined. This may also make it easier to handle background, although this is not explored.

3 Related Work

More generic approaches to future frame prediction are typically based on RNNs, which are either optimized for next-step prediction directly (Srivastava, Mansimov, and Salakhudinov 2015; Lotter, Kreiman, and Cox 2017), e.g. using a variational approach (Babaeizadeh et al. 2018; Denton and Fergus 2018; Kumar et al. 2020), or (additionally) require adversarial training (Lee et al. 2019; Vondrick, Pirsivash, and Torralba 2016). Several such approaches were also proposed in the context of physical reasoning (Finn, Goodfellow, and Levine 2016; Lerer, Gross, and Fergus 2016; Li et al. 2016). However, unlike our approach, they do not explicitly distinguish between objects and relations. This is known to affect their ability to accurately model physical interactions and extrapolation (van Steenkiste et al. 2018; Garnelo and Shanahan 2019), despite their remarkable capacity for modeling more complex visually scenes.

More closely related approaches to physical prediction explicitly distinguish object representations and incorporate a corresponding relational inductive bias, typically in the form of a graph network (Battaglia et al. 2018). In this case, reasoning is based on message-passing between object states and relations are either inferred heuristically (Chang et al. 2017; Mrowca et al. 2018; Lingelbach et al. 2020), implicitly by the message passing functions (Sukhbaatar et al. 2015; Battaglia et al. 2016; Santoro et al. 2017; Watters et al. 2017; Sanchez-Gonzalez et al. 2018, 2020) e.g. via attention (Hoshen 2017; van Steenkiste et al. 2018), or explicitly as in NRI (Kipf et al. 2018). Typically, these approaches assume access to supervised state descriptions, and only few works also infer object representations from raw visual images (van Steenkiste et al. 2018; Veerapaneni et al. 2019; Watters et al. 2019). RIMs (Goyal et al. 2019) impose a modular structure into an RNN, whose sparse communication can be seen as a kind of relational inference. However, neither of these additionally incorporate a hierarchical in-

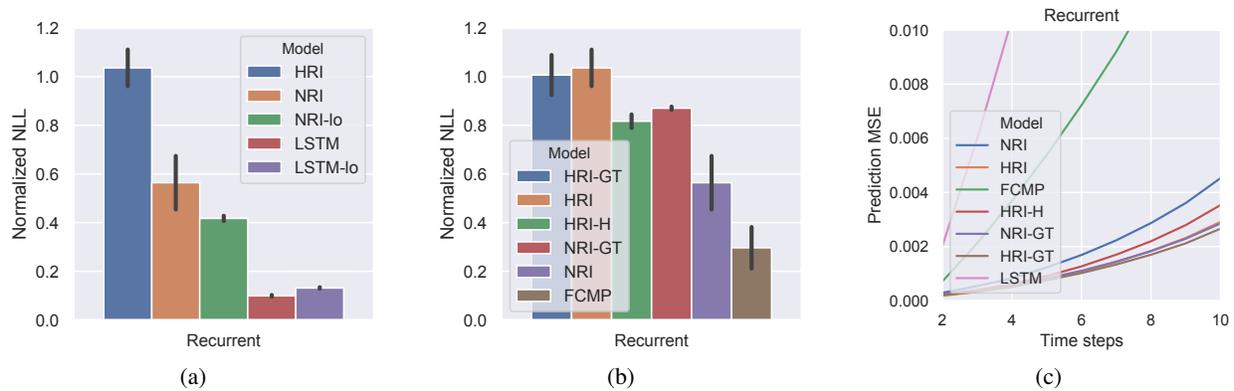


Figure 3: Performance on 4-3-state-springs. We compare HRI to (a) baselines and (b) ablations in terms of the “normalized” negative log likelihood (higher is better). (c) MSE for future prediction (prediction rollouts).

ductive bias to cope with more complex physical interactions (and corresponding objects). Arguably, this is due to the difficulty of inferring corresponding part-based object representations, and indeed prior approaches that do incorporate a hierarchical inductive bias rely on supervised state descriptions (Mrowca et al. 2018; Lingelbach et al. 2020).

Related approaches that specifically focus on the issue of learning object representations from raw visual observations typically make use of mixture models (Greff, van Steenkiste, and Schmidhuber 2017; Greff et al. 2019) or attention (Eslami et al. 2016; Kosiosek et al. 2018; Stanić and Schmidhuber 2019; Burgess et al. 2019; Crawford and Pineau 2019; Jiang et al. 2020). In our case, we adopt an approach based on spatial slots (van Steenkiste, Greff, and Schmidhuber 2019; Greff, van Steenkiste, and Schmidhuber 2020) that puts less emphasis on the exact correspondence to objects. Unlike prior work (Santoro et al. 2017; Zambaldi et al. 2019), we demonstrate how spatial slots can be extended to a hierarchical setting.

Hierarchical structure learning has previously been explored in vision. In this case, hierarchies are either learned from 3D volumetric primitives (Tulsiani et al. 2017), or using ground truth pixel-wise flow fields to guide object or object parts segmentation (Liu et al. 2019). In concurrent work (Mittal et al. 2020), RIMs were extended to consider a hierarchy of modules, although their correspondence to perceptual objects (and parts) is unclear. A more relevant approach was recently proposed by (Deng, Zhi, and Ahn 2019), which is able to infer a hierarchy of objects and parts, directly from raw visual images. However, it was not explored how this approach can be adapted to video, and how these abstractions can be used for physics prediction.

4 Experiments

In this section we evaluate HRI on four different dynamics modelling tasks: state trajectories of objects connected via finite-length springs in a hierarchical structure (*state-springs*); corresponding rendered videos (*visual-springs*); rendered joints of human moving bodies (*Human3.6M*); and raw videos of moving humans (*KTH*). We compare HRI to NRI (Kipf et al. 2018), which performs relational in-

ference but lacks a hierarchical inductive bias, and to an LSTM (Hochreiter and Schmidhuber 1997) that concatenates representations from all objects and predicts them jointly, but lacks a relational inference mechanism altogether. Appendix A contains all experimental details. Reported results are mean and standard deviations over 5 seeds.

4.1 State Springs Dataset

We consider synthetic physical systems containing simple objects connected via finite-length springs that can be organized according to a hierarchical interaction graph (Figure 5, middle row). Here, an approach that attempts to model the underlying system dynamics (which are highly complex) would clearly benefit from a hierarchical inductive bias, which allows us to validate our design choices. In all configurations we consider hierarchies with 3 levels (containing a root node, intermediate nodes, and leaf nodes), whose connectivity is decided randomly while ensuring that all nodes are connected. Corresponding physical objects are initialized at random locations with the root node biased towards the center and the intermediate and leaf nodes towards a particular quadrant to reduce clutter (see also Appendix A).

We experiment with hierarchies containing 4 intermediate nodes, each having 3 or 4 leaf nodes, denoted as *4-3-state-springs* and *3-3-state-springs*, respectively. Each model receives as input the 4-dimensional state trajectories: $x(t), y(t), \Delta x(t), \Delta y(t)$.

Comparison to baselines We compare HRI to NRI and LSTM on *4-3-state-springs* (Figure 3a) and *3-3-state-springs* (Figure 7a in Appendix B), in terms of the negative log likelihood inversely proportional to a version of HRI that operates on the ground-truth interaction graph (HRI-GT)². In this case, values closer to 1.0 are better, although we also provide raw negative log likelihoods in Figure 8. It can be observed that HRI markedly outperforms NRI on this task, and that both significantly improve over the LSTM (which was

²This allows us to factor out the complexity of the task and make it easier to compare results *between* tasks.

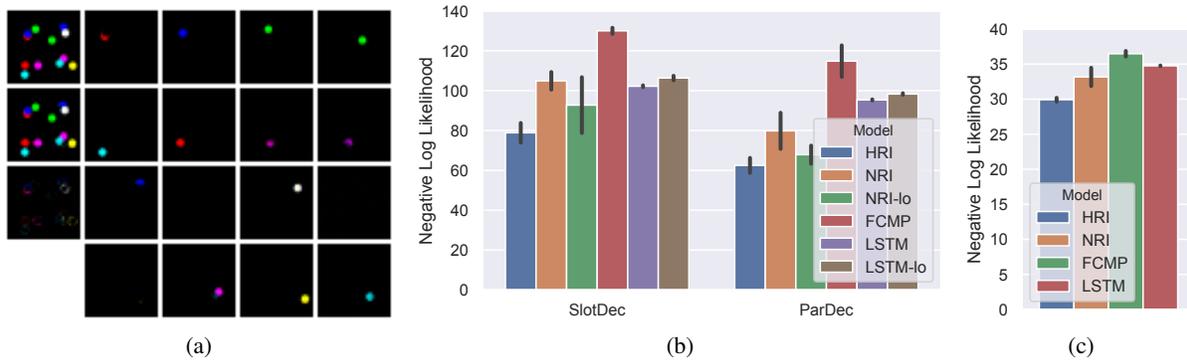


Figure 4: (a) HRI introspection: first column contains ground-truth, prediction and their difference. The other columns show the 16 object slots decoded separately. (b) Negative log likelihood for all models on the 4-3-visual-springs and (c) Human3.6M.

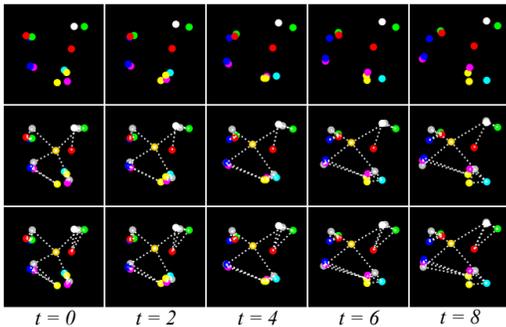


Figure 5: Rendered inputs for 4-3-state-springs (leaf objects only) (top); full interaction graph, unobserved by model (middle); Predictions and inferred edges by HRI (bottom).

expected). These findings indicate that the hierarchical inductive bias in HRI is indeed highly beneficial for this task.

We also compare to baselines that focus only on the leaf objects (*NRI-lo* and *LSTM-lo*) and are therefore not encouraged to separate hierarchical sources of influence (i.e. we impose less structure). It can be observed that this does not result in better predictions and the gap between *NRI-lo* and *NRI* indicates that explicitly considering multiple levels of abstraction, as in HRI, is indeed desirable.

Ablation & Analysis We conduct an ablation of HRI, where we modify the relational inference module. We consider FCMP, a variation of NRI that does not perform relational inference but assumes a fixed fully connected graph, and HRI-H, which is given knowledge about the ‘valid’ edges in the ground-truth hierarchical graph to be inferred and performs relational inference only on those.

In Figure 3b and Figure 7b (Appendix B) it can be observed how the lack of a relational inference module further harms performance. Indeed, in this case it is more complex for FCMP to implement a similar solution, since it requires ‘ignoring’ edges between objects that do not influence each other. It can also be observed how HRI outperforms HRI-H (but see Figure 7b where they perform the same), which is surprising since the latter essentially solves a simpler task.

We speculate that training interactions could be responsible for this gap. Switching to static graph inference yielded a significant drop in performance for NRI ($\sim 74\%$ worse), while HRI remained mostly unaffected ($\sim 5\%$ worse).

We also consider the benefit of hierarchical message-passing in isolation by comparing HRI to NRI-GT, which receives the ground-truth interaction graph. In Figure 3b and Figure 7b (Appendix B) it can be seen how the lack of hierarchical message-passing explains part of the previously observed gap between HRI and NRI, but not all of it. It suggests that by explicitly considering multiple levels of abstraction (as in HRI), conducting relational inference generally becomes easier on these tasks. Finally, we consider the effect of using a feedforward dynamics predictor (Figures 8 and 9 in Appendix B). It can be seen that in the feedforward case, position and velocity are not enough to derive interactions between objects. For this reason we exclusively focus on the recurrent predictor from now on.

Long-term Prediction and Graph Inference We investigate if HRI is able to perform long-term physical predictions by increasing the number of rollout steps at test-time. In Figure 3c we report the MSE between the predicted and ground truth trajectories and compare to previous baselines and variations. It can be observed that HRI outperforms all other models, sometimes even HRI-GT, and this gap increases as we predict deeper into the future. To qualitatively evaluate the plausibility of HRI’s rollout predictions, we generated videos by rendering them. As can be seen in Figure 5 (bottom row) both the predicted trajectories and the inferred graph connectivity closely matches the ground-truth. Similar results are obtained for 3-3-state-springs in Appendix B.

4.2 Visual Datasets

To generate visual data for springs we rendered 4-3-state-springs and 3-3-state-springs with all balls having the same radius and one of the 7 default colors, assigned at random³. Additionally, we consider *Human3.6M* (Ionescu et al. 2013)

³Similar results were obtained when using homogeneous colors (Figures 11 and 12) and when varying shape (Figures 17 to 21).

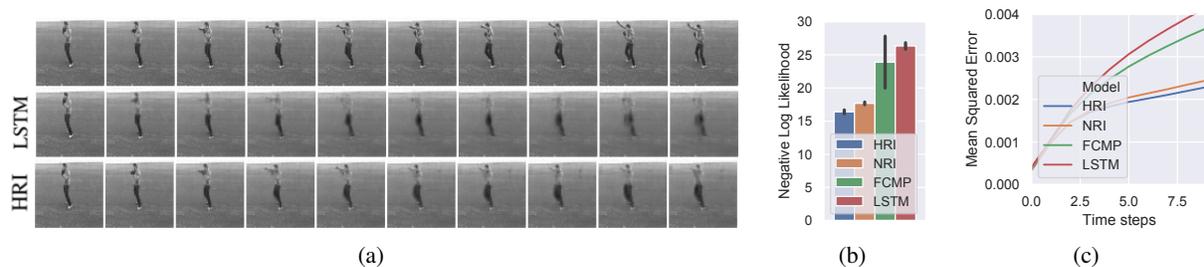


Figure 6: (a) Ground truth (top) and predicted 10 time steps rollout for LSTM (middle) and HRI (bottom) row on KTH. (b) Negative log likelihood for all models on KTH. (c) Rollout predictions accuracy for all models on KTH.

(using rendered joints) and *KTH* (Schuldt, Laptev, and Caputo 2004) (using raw videos). We train each model in two stages, which acts as a curriculum: first we train the visual encoder and decoder on a reconstruction task, afterwards we optimize the dynamics parameters on the prediction task.

Visual Springs To create the visual springs dataset, we render only the *leaf nodes*. Hence the visible balls themselves are the “parts” that must be grouped in line with more abstract “objects” corresponding to the intermediate and root nodes of the interaction graph (e.g. the gold and silver balls in the bottom row of Figure 5) which are not observed.

In Figure 4b we compare HRI to several baselines and previously explored variations (the results for *3-3-visual-springs* are available in Figure 10 in Appendix B). We are able to observe similar results, in that HRI is the best performing model, although the added complexity of inferring object states results in smaller differences. This is better understood by comparing the difference between HRI and NRI⁴ to the difference between HRI and NRI-lo. It can be seen that NRI-lo performs slightly worse than HRI (although it is less stable) and is much better than NRI, which suggests that inferring the intermediate and root node objects is one of the main challenges in this setting. Notice, however, that HRI continues to markedly improve over the LSTM and visual NRI in this setting. We note that comparing to NRI-GT or HRI-GT is not possible, since the mapping from learned object representations to those in the graph is unknown.

Comparing SlotDec and ParDec in Figure 4b, we observe that better results are obtained using the latter. We suspect that this is due to the visual occlusions that occur, which makes summation in pixel space more difficult. Nonetheless, in most cases we observe a clear correspondence between spatial slots and information about individual objects (see Figure 4a). Finally, Figure 13 (Appendix B) demonstrates how future predictions by HRI match the ground-truth quite well. Similarly, we observe that the inferred interaction graph by HRI resembles the ground-truth much more closely compared to NRI (Figure 15 in Appendix B).

When evaluating HRI trained on *4-3-visual-springs* on *3-3-visual-springs* (i.e. when *extrapolating*) the relative ordering is preserved (Figure 22 in Appendix B).

⁴Note that NRI requires object states. Here, we consider an adaptation using the encoder and decoder proposed for HRI.

Real-world Data We consider *Human3.6M*, which consists of 3.6 million 3D human poses composed of 32 joints and corresponding images taken from professional actors in various scenarios. Here we use the provided 2D pose projections to render 12 joints in total (3 per limb) as input to the model. This task is significantly more complex, since the underlying system dynamics are expected to vary over time (i.e. they are non-stationary). Figure 4c demonstrates the performance of HRI and several baselines on this task. Note that HRI is the best performing model, although the gap to NRI and LSTM is smaller. This can be explained by the fact that many motions, such as when sitting, eating, and waiting involve relatively little motion or only motion in a single joint (and thereby lack hierarchical interactions). Example future predictions by HRI can be seen in Figure 14.

Finally, we apply HRI directly on the *raw* videos of *KTH*, which consists of people performing one of six actions (walking, jogging, running, etc.). Figure 6a shows a comparison between the LSTM baseline and HRI model. In the video generated by the LSTM, the shape of the human is not preserved, while HRI is able to clearly depict human limbs. We prescribe this to the ability of HRI to reason about the objects in the scene in terms of its parts and their interactions, which in turn simplifies the physics prediction. This observation is also reflected in the quantitative analysis in Figures 6b and 6c, where large differences can be observed, especially when predicting deeper into the future.

5 Conclusion

Hierarchical Relational Inference (HRI) is a novel approach to common-sense physical reasoning capable of learning to discover objects, parts, and their relations, directly from raw visual images in an unsupervised fashion. It builds on the idea that the dynamics of complex objects are best modeled as hierarchies of parts that separate different sources of influence. We compare HRI’s predictions to those of a strong baseline on synthetic and real-world physics prediction tasks where it consistently outperforms. Several ablation studies validate our design choices. A potential limitation of HRI is that it groups parts into objects based on spatial proximity, which may be suboptimal in case of severe occlusions. Further, it would be interesting to provide for inter-object interactions, and unknown hierarchy depths. Both might be handled by increasing the hierarchy depth, and relying on the relational inference module to ignore non-existent edges.

Acknowledgments

We thank Anand Gopalakrishnan, Róbert Csordás, Imanol Schlag, Louis Kirsch and Francesco Faccio for helpful comments and fruitful discussions. This research was supported by the Swiss National Science Foundation grant 407540_167278 EVAC - Employing Video Analytics for Crisis Management, and the Swiss National Science Foundation grant 200021_165675/1 (successor project: no: 200021_192356). We are also grateful to NVIDIA Corporation for donating several DGX machines to our lab and to IBM for donating a Minsky machine.

References

- Babaeizadeh, M.; Finn, C.; Erhan, D.; Campbell, R. H.; and Levine, S. 2018. Stochastic Variational Video Prediction. In *International Conference on Learning Representations*.
- Battaglia, P.; Pascanu, R.; Lai, M.; Rezende, D. J.; et al. 2016. Interaction networks for learning about objects, relations and physics. In *Advances in neural information processing systems*, 4502–4510.
- Battaglia, P. W.; Hamrick, J. B.; Bapst, V.; Sanchez-Gonzalez, A.; Zambaldi, V.; Malinowski, M.; Tacchetti, A.; Raposo, D.; Santoro, A.; Faulkner, R.; et al. 2018. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*.
- Burgess, C. P.; Matthey, L.; Watters, N.; Kabra, R.; Higgins, I.; Botvinick, M.; and Lerchner, A. 2019. Monet: Unsupervised scene decomposition and representation. *arXiv preprint arXiv:1901.11390*.
- Chang, M.; Ullman, T.; Torralba, A.; and Tenenbaum, J. 2017. A Compositional Object-Based Approach to Learning Physical Dynamics. In *ICLR*.
- Ciregan, D.; Meier, U.; and Schmidhuber, J. 2012. Multi-column deep neural networks for image classification. In *2012 IEEE conference on computer vision and pattern recognition*, 3642–3649. IEEE.
- Crawford, E.; and Pineau, J. 2019. Spatially invariant unsupervised object detection with convolutional neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 3412–3420.
- Deng, F.; Zhi, Z.; and Ahn, S. 2019. Generative Hierarchical Models for Parts, Objects, and Scenes. *arXiv preprint arXiv:1910.09119*.
- Denton, E.; and Fergus, R. 2018. Stochastic Video Generation with a Learned Prior. In *International Conference on Machine Learning*, 1174–1183.
- Eslami, S. A.; Heess, N.; Weber, T.; Tassa, Y.; Szepesvari, D.; Hinton, G. E.; et al. 2016. Attend, infer, repeat: Fast scene understanding with generative models. In *Advances in Neural Information Processing Systems*, 3225–3233.
- Finn, C.; Goodfellow, I.; and Levine, S. 2016. Unsupervised learning for physical interaction through video prediction. In *Advances in neural information processing systems*, 64–72.
- Garnelo, M.; and Shanahan, M. 2019. Reconciling deep learning with symbolic artificial intelligence: representing objects and relations. *Current Opinion in Behavioral Sciences* 29: 17 – 23. ISSN 2352-1546.
- Gilmer, J.; Schoenholz, S. S.; Riley, P. F.; Vinyals, O.; and Dahl, G. E. 2017. Neural message passing for quantum chemistry. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 1263–1272. JMLR.
- Goyal, A.; Lamb, A.; Hoffmann, J.; Sodhani, S.; Levine, S.; Bengio, Y.; and Schölkopf, B. 2019. Recurrent independent mechanisms. *arXiv preprint arXiv:1909.10893*.
- Greff, K.; Kaufman, R. L.; Kabra, R.; Watters, N.; Burgess, C.; Zoran, D.; Matthey, L.; Botvinick, M.; and Lerchner, A. 2019. Multi-Object Representation Learning with Iterative Variational Inference. In *International Conference on Machine Learning*, 2424–2433.
- Greff, K.; van Steenkiste, S.; and Schmidhuber, J. 2017. Neural expectation maximization. In *Advances in Neural Information Processing Systems*, 6691–6701.
- Greff, K.; van Steenkiste, S.; and Schmidhuber, J. 2020. On the Binding Problem in Artificial Neural Networks. *arXiv preprint arXiv:2012.05208*.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long short-term memory. *Neural computation* 9(8): 1735–1780.
- Hoshen, Y. 2017. Vain: Attentional multi-agent predictive modeling. In *Advances in Neural Information Processing Systems*, 2701–2711.
- Ionescu, C.; Papava, D.; Olaru, V.; and Sminchisescu, C. 2013. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE transactions on pattern analysis and machine intelligence* 36(7): 1325–1339.
- Jang, E.; Gu, S.; and Poole, B. 2017. Categorical reparameterization with gumbel-softmax. In *International Conference on Learning Representations*.
- Jiang, J.; Janghorbani, S.; Melo, G. D.; and Ahn, S. 2020. SCALOR: Generative World Models with Scalable Object Representations. In *International Conference on Learning Representations*.
- Kingma, D. P.; and Welling, M. 2013. Auto-Encoding Variational Bayes. In *International Conference on Learning Representations*.
- Kipf, T. N.; Fetaya, E.; Wang, K.; Welling, M.; and Zemel, R. S. 2018. Neural Relational Inference for Interacting Systems. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, 2693–2702.
- Kosiorok, A. R.; Kim, H.; Posner, I.; and Teh, Y. W. 2018. Sequential Attend, Infer, Repeat: Generative Modelling of Moving Objects. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS’18*, 8615–8625. USA: Curran Associates Inc.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, 1097–1105.
- Kumar, M.; Babaeizadeh, M.; Erhan, D.; Finn, C.; Levine, S.; Dinh, L.; and Kingma, D. 2020. Videoflow: A flow-based generative model for video. In *ICLR*.
- Lake, B. M.; Salakhutdinov, R.; and Tenenbaum, J. B. 2015. Human-level concept learning through probabilistic program induction. *Science* 350(6266): 1332–1338.
- Lee, A. X.; Zhang, R.; Ebert, F.; Abbeel, P.; Finn, C.; and Levine, S. 2019. Stochastic Adversarial Video Prediction. In *International Conference on Learning Representations*.
- Lerer, A.; Gross, S.; and Fergus, R. 2016. Learning Physical Intuition of Block Towers by Example. In *International Conference on Machine Learning*, 430–438.

- Li, W.; Azimi, S.; Leonardis, A.; and Fritz, M. 2016. To fall or not to fall: A visual approach to physical stability prediction. *arXiv preprint arXiv:1604.00066* .
- Lingelbach, M. J.; Mrowca, D.; Haber, N.; Fei-Fei, L.; and Yamins, D. L. K. 2020. Towards Curiosity-Driven Learning of Physical Dynamics. *ICLR 2020 Bridging AI and Cognitive Science Workshop* 6.
- Liu, Z.; Wu, J.; Xu, Z.; Sun, C.; Murphy, K.; Freeman, W. T.; and Tenenbaum, J. B. 2019. Modeling Parts, Structure, and System Dynamics via Predictive Learning. In *International Conference on Learning Representations*.
- Lotter, W.; Kreiman, G.; and Cox, D. 2017. Deep predictive coding networks for video prediction and unsupervised learning. In *ICLR*.
- Maddison, C. J.; Mnih, A.; and Teh, Y. W. 2017. The concrete distribution: A continuous relaxation of discrete random variables. In *International Conference on Learning Representations*.
- Mittal, S.; Lamb, A.; Goyal, A.; Voleti, V.; Shanahan, M.; Lajoie, G.; Mozer, M.; and Bengio, Y. 2020. Learning to combine top-down and bottom-up signals in recurrent neural networks with attention over modules. In *International Conference on Machine Learning*, 6972–6986. PMLR.
- Mrowca, D.; Zhuang, C.; Wang, E.; Haber, N.; Fei-Fei, L. F.; Tenenbaum, J.; and Yamins, D. L. 2018. Flexible neural representation for physics prediction. In *Advances in neural information processing systems*, 8799–8810.
- Rezende, D. J.; Mohamed, S.; and Wierstra, D. 2014. Stochastic Backpropagation and Approximate Inference in Deep Generative Models. In *International Conference on Machine Learning*, 1278–1286.
- Sanchez-Gonzalez, A.; Godwin, J.; Pfaff, T.; Ying, R.; Leskovec, J.; and Battaglia, P. W. 2020. Learning to simulate complex physics with graph networks. *arXiv preprint arXiv:2002.09405* .
- Sanchez-Gonzalez, A.; Heess, N.; Springenberg, J. T.; Merel, J.; Riedmiller, M. A.; Hadsell, R.; and Battaglia, P. 2018. Graph Networks as Learnable Physics Engines for Inference and Control. In *ICML*, 4467–4476.
- Santoro, A.; Raposo, D.; Barrett, D. G.; Malinowski, M.; Pascanu, R.; Battaglia, P.; and Lillicrap, T. 2017. A simple neural network module for relational reasoning. In *Advances in neural information processing systems*, 4967–4976.
- Scarselli, F.; Gori, M.; Tsoi, A. C.; Hagenbuchner, M.; and Monfardini, G. 2009. The Graph Neural Network Model. *IEEE Transactions on Neural Networks* 20(1): 61–80. ISSN 1045-9227. doi:10.1109/TNN.2008.2005605.
- Schuldte, C.; Laptev, I.; and Caputo, B. 2004. Recognizing human actions: a local SVM approach. In *Proceedings of the 17th International Conference on Pattern Recognition, 2004. ICPR 2004.*, volume 3, 32–36. IEEE.
- Spelke, E. S.; and Kinzler, K. D. 2007. Core knowledge. *Developmental science* 10(1): 89–96.
- Srivastava, N.; Mansimov, E.; and Salakhudinov, R. 2015. Unsupervised learning of video representations using lstms. In *International conference on machine learning*, 843–852.
- Stanić, A.; and Schmidhuber, J. 2019. R-SQAIR: Relational Sequential Attend, Infer, Repeat. *NeurIPS PGR Workshop* .
- Sukhbaatar, S.; Weston, J.; Fergus, R.; et al. 2015. End-to-end memory networks. In *Advances in neural information processing systems*, 2440–2448.
- Tulsiani, S.; Su, H.; Guibas, L. J.; Efros, A. A.; and Malik, J. 2017. Learning shape abstractions by assembling volumetric primitives. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2635–2643.
- van Steenkiste, S.; Chang, M.; Greff, K.; and Schmidhuber, J. 2018. Relational Neural Expectation Maximization: Unsupervised Discovery of Objects and their Interactions. In *International Conference on Learning Representations*.
- van Steenkiste, S.; Greff, K.; and Schmidhuber, J. 2019. A Perspective on Objects and Systematic Generalization in Model-Based RL. In *ICML workshop on Generative Modeling and Model-Based Reasoning for Robotics and AI*.
- Veerapaneni, R.; Co-Reyes, J. D.; Chang, M.; Janner, M.; Finn, C.; Wu, J.; Tenenbaum, J. B.; and Levine, S. 2019. Entity Abstraction in Visual Model-Based Reinforcement Learning. In *Conference on Robot Learning*.
- Vondrick, C.; Pirsivash, H.; and Torralba, A. 2016. Generating videos with scene dynamics. In *Advances in neural information processing systems*, 613–621.
- Watters, N.; Matthey, L.; Bosnjak, M.; Burgess, C. P.; and Lerchner, A. 2019. COBRA: Data-Efficient Model-Based RL through Unsupervised Object Discovery and Curiosity-Driven Exploration. *ICML workshop on Generative Modeling and Model-Based Reasoning for Robotics and AI* .
- Watters, N.; Zoran, D.; Weber, T.; Battaglia, P.; Pascanu, R.; and Tacchetti, A. 2017. Visual interaction networks: Learning a physics simulator from video. In *Advances in neural information processing systems*, 4539–4547.
- Zambaldi, V.; Raposo, D.; Santoro, A.; Bapst, V.; Li, Y.; Babuschkin, I.; Tuyls, K.; Reichert, D.; Lillicrap, T.; Lockhart, E.; Shanahan, M.; Langston, V.; Pascanu, R.; Botvinick, M.; Vinyals, O.; and Battaglia, P. 2019. Deep reinforcement learning with relational inductive biases. In *International Conference on Learning Representations*.