

Guiding Video Prediction with Explicit Procedural Knowledge

Patrick Takenaka^{1,2}, Johannes Maucher¹, Marco F. Huber^{2,3}

Institute for Applied AI, Hochschule der Medien Stuttgart, Germany¹

Institute of Industrial Manufacturing and Management IFF, University of Stuttgart, Germany²

Fraunhofer Institute for Manufacturing Engineering and Automation IPA, Stuttgart, Germany³

{takenaka,maucher}@hdm-stuttgart.de, marco.huber@ieee.org

Abstract

We propose a general way to integrate procedural knowledge of a domain into deep learning models. We apply it to the case of video prediction, building on top of object-centric deep models and show that this leads to a better performance than using data-driven models alone. We develop an architecture that facilitates latent space disentanglement in order to use the integrated procedural knowledge, and establish a setup that allows the model to learn the procedural interface in the latent space using the downstream task of video prediction. We contrast the performance to a state-of-the-art data-driven approach and show that problems where purely data-driven approaches struggle can be handled by using knowledge about the domain, providing an alternative to simply collecting more data.

1. Introduction

The integration of expert knowledge in deep learning systems reduces the complexity of the overall learning problem, while offering domain experts an avenue to add their knowledge into the system, potentially leading to improved data efficiency, controllability, and interpretability. [29] showed in detail the various types of knowledge that are currently being integrated in deep learning, ranging from logic rules to regularize the learning process [38], to modelling the underlying graph structure in the architecture [33]. Especially in the physical sciences, where exact and rigid performance of the model is of importance, data-driven systems have shown to struggle on their own, and many complex problems cannot be described through numerical solvers alone [30]. This highlights the need for hybrid modelling approaches that make use of both theoretical domain knowledge, and of collected data. When viewing this approach from the perspective of deep learning, if the model is able to understand and work with integrated domain knowledge, it could potentially render many data samples redundant w.r.t. information gain.

In addition to the recognized knowledge integration categories [29], we propose to view procedural knowledge described through programmatic functions as its own category, as it is equally able to convey domain information in a structured manner as other types, while bringing with it an already established ecosystem of definitions, frameworks, and tools. Such inductive domain biases in general can help models to obtain a more structured view of the environment [7] and lead them towards more desirable predictions by either restricting the model hypothesis space, or by guiding the optimization process [3].

We argue that by incorporating procedural knowledge we can give neural networks powerful learning shortcuts where data-driven approaches struggle, and as a result reduce the demand for data, allow better out-of-distribution performance, and enable domain experts to control and better interpret the predictions. In summary, our contributions are:

- Specification of a general architectural scheme for procedural knowledge integration.
- Application of this scheme to video prediction, involving a novel latent space separation scheme to facilitate learning of the procedural interface.
- Performance analysis of our proposed method in contrast to a purely data-driven approach.

The paper is structured as follows: First, our proposed procedural knowledge integration scheme is introduced in Sec. 2, followed by its specification for the video prediction use case in Sec. 2.1. We show relevant related work in Sec. 3 and continue by describing the concrete model and overall setup that we used in Sec. 4, after which several experiments regarding the model performance and feasibility are made in Sec. 5.

2. Proposed Architecture

We view the integrated procedural knowledge as an individual module in the overall architecture, and the learn-




Figure 1. Abstract structure of our proposed procedural knowledge integration interface. Features of X are extracted in model M , resulting in intermediate feature maps (shown in grey). Out of these, a selected feature map z is then used to decode it into the input space of the integrated procedural module f through $M_{f_{in}}$, and the output of f is encoded back into the latent space of z using $M_{f_{out}}$. M continues with this updated latent state to obtain prediction \hat{y} .

ing objective corresponds to the correct utilization of this module, i.e., the learning of the program interface, to solve the task at hand. More specifically, we consider the case where the integrated knowledge is only solving an intermediate part of the overall task, i.e., it neither directly operates on the input data, nor are its outputs used as a prediction target.

More formally, given data sample X and procedural module f , the model latent state z is decoded into and encoded from the function input space through learned modules $M_{f_{in}}$ and $M_{f_{out}}$, respectively. Here, z corresponds to an intermediate feature map of an arbitrary deep learning model M whose target domain at least partially involves processes that are described in f . The output of $M_{f_{out}}$ is then fused with z using an arbitrary operator \oplus . This structure is shown in Fig. 1.

Procedural knowledge in general and programmatic functions in particular operate on a discrete set of input and output parameters. The aforementioned interface thus needs to disentangle the relevant parameters in the distributed representation and bind them to the correct inputs, and perform the reverse operation on the output side, tasks that are still challenging in many cases [9].

We show that in our setup we are able to learn this interface implicitly by focusing on a downstream task instead.

2.1. Case Study: Video Prediction

Video Prediction is an important objective in the current deep learning landscape. With it, many visual downstream tasks can be enhanced or even enabled that utilize temporal information. Example tasks are model predictive control (MPC) [14], visual question answering (VQA) [37], system

identification [14], or even content generation [11].

Some of these benefit even more if the system is controllable and thus, allows the integration of human intention into the inference process. This is typically done by conditioning the model on additional modalities such as natural language [11, 40] or by disentangling the latent space [28, 31].

More recently, researchers have shown [37] that object-centric learning offers a suitable basis for video prediction, as learning object interactions is difficult without suitable representations. We propose to build on top of such models, since knowledge about objects in the environment is an integral aspect of many domain processes and as such facilitates our approach.

We proceed by reducing these distributed object-centric representations further to individual object properties, which are then usable by our procedural module—a simple differentiable physics engine modeling the underlying scene dynamics.

We follow the approach of SlotFormer [37] and utilize a frozen pretrained Slot Attention for Video (SAVi) [16] model trained on video object segmentation to encode and decode object latent states for each frame. Also similarly, our proposed model predicts future frames in an autoregressive manner using a specialized rollout module, with the assumption that the first N frames of a video are given in order to allow the model to observe the initial dynamics of the scene.

Within the rollout module, our first goal for each object latent state is to disentangle object factors that are relevant as function input from those that are not. In our case, these are the dynamics and appearance—or Gestalt [27]—factors, respectively. However, as the upstream SAVi model is frozen and did not assume such disentanglement, we first have to apply a non-linear transformation on its latent space to enable the model to learn to separate the latent state into dynamics and Gestalt parts based on the inductive biases of the architecture that follows. We then use these—still distributed—latent states to obtain discrete physical state representations—i.e., in our case 3D vectors representing position and velocity—that can be processed by our explicit dynamics module in order to predict the state of the next time step. In order to avoid bottlenecks in the information flow, we introduce a parallel model that predicts both a dynamics correction and the future Gestalt state. The reasoning here is that in many cases both are dependent of each other and thus, need to be modelled jointly. Both dynamics predictions are then averaged over to produce the final dynamics state. The fused dynamics state and the predicted Gestalt state are finally concatenated to obtain the latent state of the next time step. This latent state is finally transformed non-linearly back into the latent space of the pretrained SAVi model, before it is decoded into pixel space.




Figure 2. Overview of the prediction of latent state z at time step t , given the previous latent states \mathbf{z} of time steps $t - N \dots t - 1$, where N is the number of context frames. The encoder model S_{enc} transforms the fixed latent space of \mathbf{z} into a separable latent space composed of dynamics state \mathbf{z}_d and Gestalt state \mathbf{z}_g . Both are fed through the joint Gestalt and dynamics prediction model G to obtain dynamics correction $z_{d_{\text{cor}}}$ and future Gestalt state z_g , whereas only \mathbf{z}_d is given to the explicit dynamics model D to get the explicit dynamics prediction $z_{d_{\text{exp}}}$. Both $z_{d_{\text{exp}}}$ and $z_{d_{\text{cor}}}$ are fused with fusion method F , resulting in the future dynamics state z_d . Finally, z_d and z_g are concatenated and fed through the state decoder S_{dec} to obtain future latent state z . In terms of Fig. 1, the integrated physics engine within D corresponds to f , with $M_{f_{\text{in}}}$ being the computational graph starting from S_{enc} up until the physics engine, and $M_{f_{\text{out}}}$ the subsequent dynamics computations until after S_{dec} . The dynamics correction and Gestalt computations are not shown explicitly and are part of M .

The rollout module can be seen in detail in Fig. 2. We verify in our experiments that even without additional auxiliary loss terms to regularize the latent state our model is able to correctly utilize the integrated dynamics module, indicating that the inductive bias of a correctly predicted and decoded physics state is sufficient for better visual predictions.

3. Related Work

Physics-Guided Deep Learning for Videos. The explicit representation of dynamics prevalent in a video within a deep learning model is a popular shortcut to learning the underlying concepts in the scene, and oftentimes necessary due to the inherent difficulty and ambiguity of many tasks [30]. The main objectives are usually the estimation of underlying system parameters and rules [35, 39, 20], or the adherence of the model output to certain environmental constraints [4, 42, 34, 43], leading to more accurate predictions. With that—as is the case for neuro-symbolic approaches [36, 35, 44] in most cases—the idea is to also inherently benefit from an improvement in interpretability and data efficiency.

A long-standing approach is to represent the dynamics by an individual module, i.e., a physics engine, and use different means to join it with a learnable model. Early work [36, 35] utilized this to predict physical outcomes, while simultaneously learning underlying physical parameters. Later work extended this towards video prediction, in which the output of the physics engine is used for rendering through a learnable decoder. Some used custom decoder networks for the given task [13, 41], or integrated a

complete differentiable renderer in addition [23]. However, these were limited to specialized use cases for the first, and required perfect knowledge of the visual composition of the environment for the latter. Another common direction is the use of Spatial Transformers (STs) [12], since they allow easy integration of spatial concepts such as position and rotation in the decoding process. However, these approaches [17, 14, 15]—albeit similar to our approach—assumed that (1) no data-driven correction of physics state is necessary and (2) the visuals of the scene outside of the dynamics properties remain static and can be encoded in the network weights, limiting their applicability to more complex settings. With our proposed approach we can model such properties.

For object-centric scenarios it is common to also take into account the relational structure of dynamical scenes in order to model object interactions by utilizing graph-based methods in the architecture [1, 2, 18, 25].

Disentangled Video Dynamics. Latent factor disentanglement in general assumes that the data is composed of a set of—sometimes independent—latent factors. Once the target factors can be disentangled, control over the environment becomes possible, and as such these approaches are of special interest in generative models. Early work heavily built on top of Variational Autoencoders (VAEs) [10]. However, later on it was proven that inductive biases are necessary to achieve disentanglement, and earlier work instead only exploited biases in the data [21]. Typically, these inductive biases are in the form of factor labels [22]. Such models were also used for disentanglement of physical properties and dynamics [45, 26]. In this domain, instead of only providing labels to achieve disentanglement, it is also common to help the model discover underlying dynamics by modeling them as Partial Differential Equations (PDEs)[19, 6, 42]. For video data that does not necessarily follow certain physical rules, some use a more general approach and focus on the disentanglement of position and Gestalt factors, with the idea that many object factors are independent of their position in the frame [27]. Having explicit encoding or decoding processes also helps in obtaining disentangled dynamics [17, 14, 23, 15].

4. Setup

As is done in the original SAVi paper [16], we condition the SAVi slots on the first frame object bounding boxes and pre-train on sequences of six video frames, optimizing the reconstruction of the optical flow map for each frame. Experiments have shown that optical flow reconstruction leads to better object segmentations, which we find is a better proxy for evaluating correct object dynamics than video reconstruction itself. After convergence we freeze the SAVi model.

For the video prediction task, we encode the initial six

frames using this frozen model, and use these as initial context information for the video prediction model. We then let the model auto-regressively predict the next 12 frames during training—or 24 frames during validation—always keeping the most recent six frames as reference. While more than a single reference frame would not be necessary for the integrated dynamics knowledge, the six frames are instead used in the transformer-based joint dynamics and Gestalt predictor model. In order to give the model a hint about the magnitude of the dynamics state values, we condition the dynamics state of the first frame on the ground-truth state.

4.1. Implementation Details

For the SAVi model we mainly follow the implementations of SlotFormer [37] and the original work [16]. The encoder consists of a standard Convolutional Neural Network (CNN) with a subsequent positional embedding. To obtain slot representations for a given frame we perform two iterations of slot attention, followed by a transformer model with multi-head self attention for modelling slot interactions and a final Long Short-Term Memory (LSTM) model in order to transition the representation into the next time step. We set the number of slots to six, each with size 128. The representations obtained after the slot attention rounds are decoded into the target frames using a Spatial Broadcast Decoder [32] with a broadcast size of 8.

For the video prediction model we denote the most recent N context frame representations of time steps $t - N \dots t - 1$ in bold as \mathbf{z} and the latent representation prediction for time step t as z in order to improve readability. Both the latent state encoder S_{enc} and decoder of the video prediction model S_{dec} are MLPs, each with a single ReLU activated hidden layer of size 128. They have shown to introduce sufficient non-linearity to allow state disentanglement. The latent state obtained from S_{enc} is kept the same size as the slot size and is split into two equally sized parts z_d and z_g for the subsequent dynamics and Gestalt models.

The dynamics model—i.e., the explicit physics engine—takes a physical state representation consisting of a 3D position and 3D velocity of a single frame as input, which is obtained from a linear readout layer of the most recent context frame of latent state z_d , or directly from groundtruth for the very first predicted frame. The physics engine itself is fully differentiable and consists of no learnable parameters. It calculates the dynamics taking place as in the original data simulation using a regular semi-implicit Euler integration scheme. Pseudo code of this engine can be seen in Listing 1. Its output—consisting of again a 3D position and 3D velocity of the next timestep—is then transformed back into the latent state $z_{d_{\text{exp}}}$ with another linear layer. For the Gestalt properties we utilize a prediction setup and configuration as in the original SlotFormer model: First, the latent state \mathbf{z}_g is enriched with temporal positional encodings after

which a multi head self attention transformer is used for obtaining future latent representations $z_{d_{\text{cor}}}$ and z_g . Both $z_{d_{\text{exp}}}$ and $z_{d_{\text{cor}}}$ are merged by taking their mean, and the resulting vector z_d is concatenated with z_g in order to obtain the latent representation of the future frame. S_{dec} is finally used to transform this vector back into the latent representation z of SAVi, where it can be decoded into pixel space by the pretrained frozen SAVi decoder.

Listing 1. Python pseudo code of the integrated function for our data domain which calculates a future physical state consisting of position and velocity of each object. G in the code corresponds to the gravitational constant. As is done in the original simulation each predicted frame is subdivided into smaller simulation steps—a standard approach for numerical-based physics simulations.

```
def dynamics_step(pos, vel):
    for sim_idx in range(simulation_steps):
        # Position delta between objects
        pos_delta = get_pos_delta(pos)
        # Squared distances between objects
        r2 = sum(pow(pos_delta, 2))
        # Calculate force direction vector
        F_dir = pos_delta / sqrt(r2)
        # Calculate force
        F = F_dir * (G * (mass / r2))
        # F = ma
        a = F / mass
        # Semi-implicit euler
        vel = vel + simulation_dt * a
        pos = pos + simulation_dt * vel
    return pos, vel
```

4.2. Data

Our dataset consists of a simulated environment of multiple interacting objects resulting in complex nonlinear dynamics. The idea was to generate an object-centric dataset for which current state-of-the-art video prediction models struggle and where the integration of knowledge about the environment is possible and sensitive. Datasets used in existing object-centric video prediction literature either did not feature complex nonlinear dynamics, or involved non-differentiable dynamics (*e.g.*, collisions) that are out of scope for now. However for the latter we note that non-differentiable dynamics such as collisions could still be integrated with our approach by building a computational graph that covers all conditional pathways. Although this approach is computationally more inefficient and does not directly convey collision event information to the learning algorithm, work exists [5] that show that this can still be exploited well enough and is simultaneously easy to implement in current deep learning frameworks with dynamic computational graphs.

The future states are predicted using a simple physics engine that simulates gravitational pull between differently sized spherical objects without collisions, as in the three

body problem [24]. In order to keep objects in the scene, we add an invisible gravitational pull towards the camera focus point and limit the movement in x and y direction. Objects are then rendered in 3D space using slight illumination and no background. Each object can have different material properties, which change their visuals slightly.

We create 10k RGB video samples consisting of 32 frames and spatial size 64×64 each with their corresponding optical flow and segmentation masks using kubric [8], which combines a physical simulator with a 3D rendering engine, allowing the generation of arbitrary physical scenes. We render four frames per second, and subdivide each frame into 60 physical simulation steps. Each sample uses the same underlying dynamics but with different starting conditions for the objects. The number of objects randomly varies per sample from 3-5 objects. For each object, we also store its physical state at each frame consisting of the 3D world position and velocity. All objects have the same fixed mass.

5. Experiments

In all experiments, we compare our proposed architecture with a SlotFormer model, representing a purely data-driven approach. To improve comparability, the transformer architectures of both our joint dynamics and Gestalt predictor G and the SlotFormer rollout module are the same. Also, both use the same underlying frozen SAVi model as object-centric encoder and decoder.

We train SAVi and the video prediction models for at maximum 100k steps each or until convergence is observed by early stopping, using a batch size of 64. We clip gradients to a maximum norm of 0.05 and train using Adam with an initial learning rate of 0.0001.

For evaluation purposes, we report the aggregated object segmentation performance over three seeds using the Adjusted Rand Index (ARI) and mean Intersection-Over-Union (mIoU) scores, in addition to their foreground (FG) variants ARI-FG and mIoU-FG which disregard background predictions.

We first analyze the baseline performance of our proposed approach in Sec. 5.1, followed by an experiment focusing on the completeness of the integrated function in Sec. 5.2. We then consider the model performance for very limited data availability in Sec.5.3 and conclude with an ablation experiment regarding the latent state separation in Sec.5.4.

5.1. Baseline

Here, we integrate the complete underlying dynamics of the environment in our model. As such, we also verify the utility of still keeping a parallel auto-regressive joint Gestalt and dynamics model by replacing it with an identity function and observing the performance, since with perfect




Figure 3. The mIoU performance w.r.t. each auto-regressive frame prediction. While the data-driven model exponentially becomes more inaccurate over time, the integration of the dynamics knowledge helps to keep the prediction performance stable. The pure variant of our architecture without a data-driven Gestalt and dynamics predictor follows the slope of our main architecture, albeit at a lower magnitude. Their difference indicates the missing handling of Gestalt and dynamics interdependencies.

knowledge about the dynamics and the initial frame appearance the model should have all necessary information for an accurate prediction.

Table 1. Performance comparison of a purely data-driven model (SlotFormer), our proposed model (Ours), and a variant of our architecture with an identity function as the joint Gestalt and dynamics predictor (Ours-Pure). For reference the performance of the underlying SAVi model is also reported, describing the upper bound performance that any downstream video prediction model can achieve.

	mIoU \uparrow	mIoU-FG \uparrow	ARI \uparrow	ARI-FG \uparrow
Ours	32.1\pm0.7	29.2\pm0.8	48.2\pm1.0	82.9\pm2.8
Ours-Pure	29.5 \pm 0.1	26.1 \pm 0.2	43.8 \pm 0.2	71.6 \pm 0.6
SlotFormer	15.1 \pm 0.1	9.4 \pm 0.1	16.9 \pm 0.3	6.1 \pm 0.4
SAVi	36.1	34.0	55.1	93.0

As we can see in Tab. 1, our proposed architecture outperforms a purely data-driven approach such as SlotFormer by a large margin, and comes close to the performance of the underlying SAVi model, which in contrast to video prediction methods has access to every video frame and simply needs to segment them. However, even when integrating perfect dynamics knowledge it is still beneficial to keep a parallel data-driven Gestalt and dynamics predictor, highlighting the need to model the dependency between appearance and dynamics in the scene. Both our models are also able to predict the future object positions and velocities in the physics state space accurately, with a Mean Absolute Error (MAE) close to 0 across all predicted frames when compared to the groundtruth.

Regarding the unroll performance, i.e., the frame-by-frame prediction performance, the SlotFormer model’s performance quickly deteriorates, while both variants of our architecture keep the performance more stable over time, as

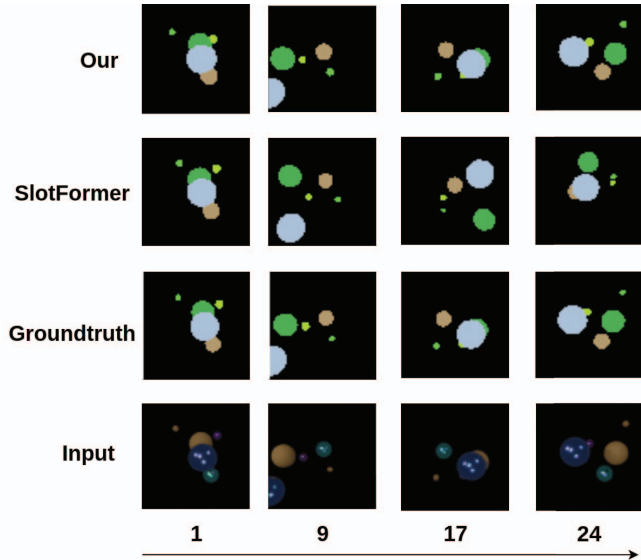


Figure 4. Sample prediction comparisons for different unroll steps. While both models are able to keep object shapes intact, the dynamics of the SlotFormer model are diverging quickly, while our model can keep up with the complex dynamics.

Table 2. Performance of the model using inaccurate dynamics information (Ours-Inaccurate) in contrast to the base models. It can be observed that although the performance decreases, it still stays above that of the purely data-driven SlotFormer model.

	mIoU \uparrow	mIoU-FG \uparrow	ARI \uparrow	ARI-FG \uparrow
Ours-Inaccurate	21.7 \pm 0.6	17.1 \pm 0.7	29.0 \pm 1.4	27.0 \pm 2.6
Ours	32.1 \pm 0.7	29.2 \pm 0.8	48.2 \pm 1.0	82.9 \pm 2.8
SlotFormer	15.1 \pm 0.1	9.4 \pm 0.1	16.9 \pm 0.3	6.1 \pm 0.4

seen in Fig. 3. As seen in Fig. 4, the performance decrease stems mainly from wrong dynamics, as the object shapes are kept intact even for the SlotFormer model.

5.2. Inaccurate Dynamics Knowledge

In the previous setup, the integrated function described the underlying dynamics perfectly and as such might allow the model to learn undesirable shortcuts. Here, we therefore evaluate whether inaccuracies in the integrated dynamics knowledge hinder the utilization of the integrated dynamics. We introduce these inaccuracies by using wrong simulation time steps, which results in wrong state predictions, albeit with the same underlying dynamics. We report the results in Tab. 2.

While the performance has clearly deteriorated, it is still above the purely data-driven approach. As such we can see that just the information about the dynamics process in itself carries valuable information for the final predictions, not only the concrete dynamics state.

Table 3. Performance comparison of a purely data-driven model (SlotFormer) and our proposed architecture using only 300 training samples. While the performance of both models deteriorates, the SlotFormer model predictions are now close to random predictions. On the other hand, our model is still performing better than the SlotFormer model with the full dataset available.

	mIoU \uparrow	mIoU-FG \uparrow	ARI \uparrow	ARI-FG \uparrow
Ours-300	23.4 \pm 1.7	19.0 \pm 2.0	33.0 \pm 2.8	35.5 \pm 5.3
Slotformer-300	12.1 \pm 0.3	5.9 \pm 0.3	11.1 \pm 0.8	2.6 \pm 0.2

Ours	32.1 \pm 0.7	29.2 \pm 0.8	48.2 \pm 1.0	82.9 \pm 2.8
SlotFormer	15.1 \pm 0.1	9.4 \pm 0.1	16.9 \pm 0.3	6.1 \pm 0.4

Table 4. Performance comparison of our proposed architecture (Ours) and a variant that does not separate the latent state into Gestalt and dynamics factors (Ours-Single). For reference the performance of the SlotFormer model is also shown.

	mIoU \uparrow	mIoU-FG \uparrow	ARI \uparrow	ARI-FG \uparrow
Ours-Single	21.6 \pm 1.1	16.7 \pm 1.3	28.1 \pm 1.9	26.1 \pm 4.6
Ours	32.1 \pm 0.7	29.2 \pm 0.8	48.2 \pm 1.0	82.9 \pm 2.8
SlotFormer	15.1 \pm 0.1	9.4 \pm 0.1	16.9 \pm 0.3	6.1 \pm 0.4

5.3. Data Efficiency

Next, we analyze the prediction performance when using only 300 data samples, amounting to 3% of the original data. We report the results in Tab. 3. As expected, the performance of both models drops, however the SlotFormer predictions are now close to random predictions, indicated by the very low foreground scores. In contrast, our proposed model still achieves a better overall performance than the SlotFormer model using the complete dataset.

5.4. Joint Latent State

Here we analyze whether the separation of the latent state into Gestalt and dynamics factors is necessary by working on only a single latent state without separation, without both the latent state encoder and decoder. As can be seen in Tab. 4, the performance decreases significantly when not performing latent state separation. However, the performance was still above that of the SlotFormer model, indicating that even poor dynamics integration can be beneficial.

6. Conclusion

We have introduced a scheme to integrate procedural knowledge into deep learning models and specialized this approach for a video prediction case. We have shown that the prediction performance can be significantly improved if one uses knowledge about underlying dynamics as opposed to learning in a data-driven fashion alone. However, we also highlighted the benefit of (1) a sensible latent state separation in order to facilitate the use of the procedural knowledge, and (2) the use of a parallel prediction model that corrects the dynamics prediction and models Gestalt

and dynamics interdependencies. Future work is focused on increasing the benefit further for inaccurate or incomplete knowledge integration, as this enables the use in more complex settings. Also, the current need for ground truth conditioning in the first frame limits applicability in some settings, and as such semi-supervised or even completely unsupervised state discovery increase the utility of our approach. Last, the application to video prediction downstream tasks such as MPC, VQA, or more complex system parameter estimation are all potential extensions of this work.

References

- [1] Peter Battaglia, Razvan Pascanu, Matthew Lai, Danilo Jimenez Rezende, and Koray kavukcuoglu. Interaction networks for learning about objects, relations and physics. In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS'16*, pages 4509–4517, Red Hook, NY, USA, Dec. 2016. Curran Associates Inc.
- [2] Daniel M. Bear, Chaofei Fan, Damian Mrowca, Yunzhu Li, Seth Alter, Aran Nayebi, Jeremy Schwartz, Li Fei-Fei, Jiajun Wu, Joshua B. Tenenbaum, and Daniel L.K. Yamins. Learning physical graph representations from visual scenes. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS'20*, pages 6027–6039, Red Hook, NY, USA, Dec. 2020. Curran Associates Inc.
- [3] Andrea Borghesi, Federico Baldo, and Michela Milano. Improving Deep Learning Models via Constraint-Based Domain Knowledge: A Brief Survey. *arXiv:2005.10691 [cs, stat]*, May 2020.
- [4] Emmanuel de Bézenac, Arthur Pajot, and Patrick Gallinari. Deep learning for physical processes: Incorporating prior scientific knowledge*. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(12):124009, Dec. 2019.
- [5] Mingyu Ding, Zhenfang Chen, Tao Du, Ping Luo, Joshua B. Tenenbaum, and Chuang Gan. Dynamic Visual Reasoning by Learning Differentiable Physics Models from Video and Language. In *Advances in Neural Information Processing Systems*, Nov. 2021.
- [6] Jérémie Donà, Jean-Yves Franceschi, Sylvain Lamprier, and Patrick Gallinari. PDE-Driven Spatiotemporal Disentanglement. In *International Conference on Learning Representations*, Jan. 2021.
- [7] Anirudh Goyal and Yoshua Bengio. Inductive biases for deep learning of higher-level cognition. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 478(2266):20210068, Oct. 2022.
- [8] Klaus Greff, Francoisk Belletti, Lucas Beyer, Carl Doersch, Yilun Du, Daniel Duckworth, David J Fleet, Dan Gnanaprasam, Florian Golemo, Charles Herrmann, Thomas Kipf, Abhijit Kundu, Dmitry Lagun, Issam Laradji, Hsueh-Ti Liu, Henning Meyer, Yishu Miao, Derek Nowrouzezahrai, Cengiz Oztireli, Etienne Pot, Noha Radwan, Daniel Rebain, Sara Sabour, Mehdi S. M. Sajjadi, Matan Sela, Vincent Sitzmann, Austin Stone, Deqing Sun, Suhani Vora, Ziyu Wang, Tianhao Wu, Kwang Moo Yi, Fangcheng Zhong, and Andrea Tagliasacchi. Kubric: A scalable dataset generator. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3739–3751, June 2022.
- [9] Klaus Greff, Sjoerd van Steenkiste, and Jürgen Schmidhuber. On the Binding Problem in Artificial Neural Networks. *arXiv:2012.05208 [cs]*, Dec. 2020.
- [10] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. Beta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework. In *International Conference on Learning Representations*, July 2022.
- [11] Yaosi Hu, Chong Luo, and Zhenzhong Chen. Make It Move: Controllable Image-to-Video Generation with Text Descriptions. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18198–18207, June 2022.
- [12] Max Jaderberg, Karen Simonyan, Andrew Zisserman, and koray kavukcuoglu. Spatial transformer networks. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015.
- [13] Michael Janner, Sergey Levine, William T. Freeman, Joshua B. Tenenbaum, Chelsea Finn, and Jiajun Wu. Reasoning About Physical Interactions with Object-Oriented Prediction and Planning. In *International Conference on Learning Representations*, Sept. 2018.
- [14] Miguel Jaques, Michael Burke, and Timothy Hospedales. Physics-as-Inverse-Graphics: Unsupervised Physical Parameter Estimation from Video. In *International Conference on Learning Representations*, Sept. 2019.
- [15] Rama Krishna Kandukuri, Jan Achterhold, Michael Moeller, and Joerg Stueckler. Physical Representation Learning and Parameter Identification from Video Using Differentiable Physics. *International Journal of Computer Vision*, 130(1):3–16, Jan. 2022.
- [16] Thomas Kipf, Gamaleldin Fathy Elsayed, Aravindh Mahendran, Austin Stone, Sara Sabour, Georg Heigold, Rico Jonschkowski, Alexey Dosovitskiy, and Klaus Greff. Conditional Object-Centric Learning from Video. In *International Conference on Learning Representations*, Jan. 2022.
- [17] Adam Kosiorek, Hyunjik Kim, Yee Whye Teh, and Ingmar Posner. Sequential Attend, Infer, Repeat: Generative Modelling of Moving Objects. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- [18] Jannik Kossen, Karl Stelzner, Marcel Hussing, Claas Voelcker, and Kristian Kersting. Structured Object-Aware Physics Prediction for Video Modeling and Planning. In *International Conference on Learning Representations*, Sept. 2019.
- [19] Vincent Le Guen and Nicolas Thome. Disentangling Physical Dynamics From Unknown Factors for Unsupervised Video Prediction. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11471–11481, Seattle, WA, USA, June 2020. IEEE.

- [20] Martin Link, Max Schwarz, and Sven Behnke. Predicting Physical Object Properties from Video. In *International Joint Conference on Neural Networks*, June 2022.
- [21] Francesco Locatello, Stefan Bauer, Mario Lucic, Gunnar Rätsch, Sylvain Gelly, Bernhard Schölkopf, and Olivier Bachem. A sober look at the unsupervised learning of disentangled representations and their evaluation. *The Journal of Machine Learning Research*, 21(1):209:8629–209:8690, Jan. 2020.
- [22] Francesco Locatello, Michael Tschannen, Stefan Bauer, Gunnar Rätsch, Bernhard Schölkopf, and Olivier Bachem. Disentangling Factors of Variations Using Few Labels. In *International Conference on Learning Representations*, Dec. 2019.
- [23] J. Krishna Murthy, Miles Macklin, Florian Golemo, Vikram Voleti, Linda Petrini, Martin Weiss, Breandan Considine, Jérôme Parent-Lévesque, Kevin Xie, Kenny Erleben, Liam Paull, Florian Shkurti, Derek Nowrouzezahrai, and Sanja Fidler. gradSim: Differentiable simulation for system identification and visuomotor control. In *International Conference on Learning Representations*, Oct. 2020.
- [24] Z. E. Musielak and B. Quarles. The three-body problem. *Reports on Progress in Physics*, 77(6):065901, June 2014.
- [25] Qu Tang, XiangYu Zhu, Zhen Lei, and ZhaoXiang Zhang. Intrinsic Physical Concepts Discovery with Object-Centric Predictive Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 23252–23261, June 2023.
- [26] Romain Thoreau, Laurent Risser, Véronique Achard, Béatrice Berthelot, and Xavier Briottet. p3VAE: A physics-integrated generative model. Application to the semantic segmentation of optical remote sensing images, Apr. 2023.
- [27] Manuel Traub, Sebastian Otte, Tobias Menge, Matthias Karlbauer, Jannik Thuemmel, and Martin V. Butz. Learning What and Where: Disentangling Location and Identity Tracking Without Supervision. In *The Eleventh International Conference on Learning Representations*, Feb. 2023.
- [28] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. MoCoGAN: Decomposing Motion and Content for Video Generation. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1526–1535, June 2018.
- [29] Laura von Rueden, Sebastian Mayer, Katharina Beckh, Bogdan Georgiev, Sven Giesselbach, Raoul Heese, Birgit Kirsch, Julius Pfrommer, Annika Pick, Rajkumar Ramamurthy, Michal Walczak, Jochen Garcke, Christian Bauckhage, and Jannis Schuecker. Informed Machine Learning – A Taxonomy and Survey of Integrating Prior Knowledge into Learning Systems. *IEEE Transactions on Knowledge and Data Engineering*, 35(1):614–633, Jan. 2023.
- [30] Rui Wang and Rose Yu. Physics-Guided Deep Learning for Dynamical Systems: A Survey, Feb. 2023.
- [31] Yaohui Wang, Piotr Bilinski, Francois Bremond, and Antitza Dantcheva. G3AN: Disentangling Appearance and Motion for Video Generation. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5263–5272, Seattle, WA, USA, June 2020. IEEE.
- [32] Nicholas Watters, Loic Matthey, Christopher P. Burgess, and Alexander Lerchner. Spatial Broadcast Decoder: A Simple Architecture for Learning Disentangled Representations in VAEs, Aug. 2019.
- [33] Nicholas Watters, Daniel Zoran, Theophane Weber, Peter Battaglia, Razvan Pascanu, and Andrea Tacchetti. Visual Interaction Networks: Learning a Physics Simulator from Video. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [34] Hao Wu, Wei Xiong, Fan Xu, Xiao Luo, Chong Chen, Xian-Sheng Hua, and Haixin Wang. PastNet: Introducing Physical Inductive Biases for Spatio-temporal Video Prediction, May 2023.
- [35] Jiajun Wu, Joseph Lim, Hongyi Zhang, Joshua Tenenbaum, and William Freeman. Physics 101: Learning Physical Object Properties from Unlabeled Videos. In *Proceedings of the British Machine Vision Conference 2016*, pages 39.1–39.12, York, UK, 2016. British Machine Vision Association.
- [36] Jiajun Wu, Ilker Yildirim, Joseph J Lim, Bill Freeman, and Josh Tenenbaum. Galileo: Perceiving Physical Object Properties by Integrating a Physics Engine with Deep Learning. In *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015.
- [37] Ziyi Wu, Nikita Dvornik, Klaus Greff, Thomas Kipf, and Animesh Garg. SlotFormer: Unsupervised Visual Dynamics Simulation with Object-Centric Models. In *The Eleventh International Conference on Learning Representations*, Feb. 2023.
- [38] Jingyi Xu, Zilu Zhang, Tal Friedman, Yitao Liang, and Guy Broeck. A Semantic Loss Function for Deep Learning with Symbolic Knowledge. In *Proceedings of the 35th International Conference on Machine Learning*, pages 5502–5511. PMLR, July 2018.
- [39] Kai Xu, Akash Srivastava, Dan Gutfreund, Felix Sosa, Tomer Ullman, Josh Tenenbaum, and Charles Sutton. A Bayesian-Symbolic Approach to Reasoning and Learning in Intuitive Physics. In *Advances in Neural Information Processing Systems*, volume 34, pages 2478–2490. Curran Associates, Inc., 2021.
- [40] Yucheng Xu, Li Nanbo, Arushi Goel, Zijian Guo, Zonghai Yao, Hamidreza Kasaei, Mohammadreza Kasaei, and Zhibin Li. Controllable Video Generation by Learning the Underlying Dynamical System with Neural ODE, Apr. 2023.
- [41] Tsung-Yen Yang, Justinian P. Rosca, Karthik R. Narasimhan, and Peter Ramadge. Learning Physics Constrained Dynamics Using Autoencoders. In *Advances in Neural Information Processing Systems*, Oct. 2022.
- [42] Yuan Yin, Vincent Le Guen, Jérémie Dona, Emmanuel de Bézenac, Ibrahim Ayed, Nicolas Thome, and Patrick Gallinari. Augmenting Physical Models with Deep Networks for Complex Dynamics Forecasting. *Journal of Statistical Mechanics: Theory and Experiment*, 2021(12):124012, Dec. 2021.
- [43] Yuan Yin, Matthieu Kirchmeyer, Jean-Yves Franceschi, Alain Rakotomamonjy, and Patrick Gallinari. Continuous PDE Dynamics Forecasting with Implicit Neural Representations. In *The Eleventh International Conference on Learning Representations*, Sept. 2022.

- [44] Dongran Yu, Bo Yang, Dayou Liu, and Hui Wang. A Survey on Neural-symbolic Systems. *arXiv:2111.08164 [cs]*, Nov. 2021.
- [45] Deyao Zhu, Marco Munderloh, Bodo Rosenhahn, and Jörg Stückler. Learning to Disentangle Latent Physical Factors for Video Prediction. In Gernot A. Fink, Simone Frintrop, and Xiaoyi Jiang, editors, *Pattern Recognition*, volume 11824, pages 595–608. Springer International Publishing, Cham, 2019.