
Generalization properties of contrastive world models

Kandan Ramakrishnan*
Kandan.Ramakrishnan@bcm.edu

R. James Cotton †
rcotton@sralab.org

Xaq Pitkow *§
xaq@cmu.edu

Andreas S. Tolias *
astolias@bcm.edu

Abstract

Recent work on object-centric world models aim to factorize representations in terms of objects in a completely unsupervised or self-supervised manner. Such world models are hypothesized to be a key component to address the generalization problem. While self-supervision has shown improved performance however, OOD generalization has not been systematically and explicitly tested. In this paper, we conduct an extensive study on the generalization properties of contrastive world model. We systematically test the model under a number of different OOD generalization scenarios such as extrapolation to new object attributes, introducing new conjunctions or new attributes. Our experiments show that the contrastive world model fails to generalize under the different OOD tests and the drop in performance depends on the extent to which the samples are OOD. When visualizing the transition updates and convolutional feature maps, we observe that any changes in object attributes (such as previously unseen colors, shapes, or conjunctions of color and shape) breaks down the factorization of object representations. Overall, our work highlights the importance of object-centric representations for generalization and current models are limited in their capacity to learn such representations required for human-level generalization.

1 Introduction

One of the main challenges in AI is to learn high-level causal variables from low-level variables like pixels in images that enable generalization. Learning world models [5, 7, 9, 19] in a self-supervised manner might be a key component for generalization [10]. Such world models also form a core component of human cognition [3]. A number of studies in cognitive psychology and neuroscience [13, 16, 20] show that world models aid in generalization and robust learning. For example, infants learn about the physical properties of objects as entities that behave consistently over time and are able to re-apply their knowledge to new scenarios involving previously unseen objects [14]. Thus, learning object-centric world models in a self-supervised manner seem to be a key component of cognition that allow humans to predict and generalize to novel situations.

A desirable property of world models is their potential ability to generalize to novel objects as observed in infants [15]. While contrastive world models [8] have been previously evaluated for their prediction performance, very few studies have tested their effectiveness for generalization to novel data samples. In a recent study [1], world models are shown to have generalization abilities beyond their training environment or number of objects. This study evaluates generalization in terms of single

*Department of Neuroscience, Baylor College of Medicine.

†Shirley Ryan AbilityLab, Northwestern University

‡Neuroscience Institute, Department of Machine Learning, Carnegie Mellon University

§Department of ECE, Rice University

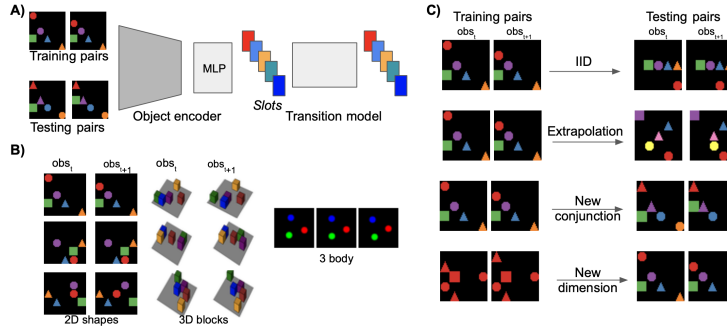


Figure 1: A) Model architecture of object-centric world model used in our experiments. The world model consists of an object encoder with slot architecture and followed by a Graph Neural Network as the transition model. B) Datasets used to evaluate models for out-of-distribution generalization. The grid based 2D shapes, 3D blocks and 3 body dataset are visualized. C) Illustration of the generalization tests on 2D shapes dataset - i) IID: training and testing data are the same, ii) New conjunction : Testing on novel color-shape combinations not seen during training, iii) Extrapolation : Testing for a new shape or new color different from the training dataset and iv) New dimension - Either variation in shape or color is seen during training while the testing contain variation in both shape and color.

scenario such as novel task and also under different environmental conditions. However in this study it is assumed that the input image is factored into constituent objects by an object detection module. This, however is not a thorough test of generalization for world models. Our aim is to conduct an exhaustive evaluation of the generalization abilities of contrastive world models under a number of different OOD conditions.

In our experiments, we train a contrastive structured world model (CSWM) on a next step prediction task given an input observation and an associated action. The models are trained on different datasets - 2D shapes, 3D blocks and 3-body dataset [6]. The trained models are then tested for different types of generalization. We find that: (1) Object-centric world models are unable to factorize representations under OOD; and (2) the drop in generalization performance depends on the number of time steps and number of objects that are OOD. Overall, our findings challenge the notion that contrastive learning of object-centric world models potentially help with OOD generalization and this requires design of novel learning paradigm to preserve factorization of representations critical for generalization.

2 Related work

In [12] it has been shown that learning disentangled representations requires an inductive bias in the model architecture and the data. Other studies have shown its relation to fairness [11], its usefulness for downstream tasks that enable quicker learning from fewer samples [18]. Few studies also study its relation to generalization. [4] demonstrates that learning disentangled representations is a good predictor for out-of-distribution task performance. However, in the same study, they also show that VAE-based approaches [12] do not learn to disentangle complex datasets, which instead requires increased model capacity. In another study, [17] investigate the ability of disentangled representations to generalize under OOD conditions. However none of these studies have investigated generalization performance of contrastive world models.

3 Results

3.1 Model evaluation on IID data

The evaluation of CSWM model on the 2D shapes IID data shows perfect factorization of object-level representations as shown in Figure 2. Given the perfect factorization, CSWM model achieves perfect performance on IID data.

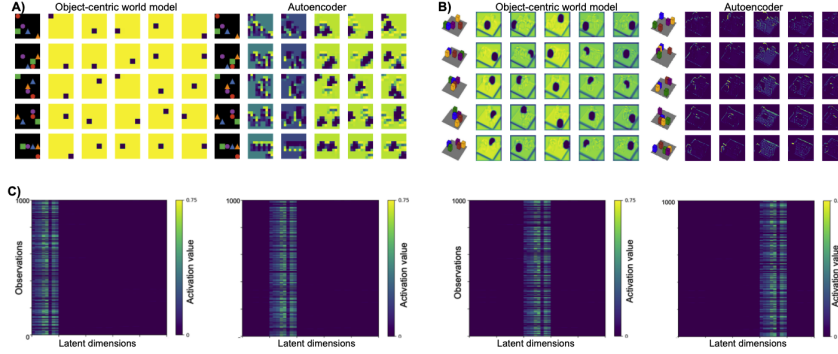


Figure 2: Understanding the factorization of representations: A) visualization of the activation maps from the convolutional backbone on 2D shapes dataset of both CSWM and AE. Each map corresponds to an object from the input image which indicates to what extent the encoder is able to factorize the representation space as per objects. B) Visualization of convolutional maps on 3D blocks dataset of both CSWM and AE models. C) Each plot is the state transitions when only one object is moved in the environment.

3.2 Out-of-distribution evaluation

To test the hypothesis whether object-centric representations generalize to novel data, we evaluate the models prediction performance under OOD settings by changing the attributes (shape or color) of objects in the dataset. Given that the CSWM outperforms other slot-based models, we evaluate CSWM for OOD prediction performance. For each generalization test, we also vary the number of objects with attributes changed in the test set.

Figure 3 shows the CSWM performance under different types of generalization on the 2D shapes dataset. We observe that for all types of generalization, the performance deteriorates. The prediction performance depends to what extent the test samples are OOD: as more the number of objects are changed, the performance worsened. Additionally, we also notice that the performance degrades over longer time-steps. Single-step prediction performance on OOD remains close to IID, especially with only one object changed. However, for higher time steps (5 and 10 time steps) there is a considerable drop in performance. Similarly the CSWM model fails to generalize to 3D blocks and 3 body dataset (Figure 4).

The drop in OOD performance of the model indicates that either the convolution encoder doesn't factorize the representations or that the transition model fails to update the state representation accurately. To identify we run a qualitative analysis on both the encoder representations and the transition updates. As in the IID analysis, we visualize the convolution filters to see if the objects are factored under the OOD conditions. Figure 3B shows the output from the convolution layers of the CSWM object encoder trained on 2D shapes dataset. Unlike in IID, each convolution map does not correspond to a distinct object in the input image. We observe that the convolution map activations mixes mask corresponding to multiple objects. The CSWM model is thus unable to factorize object-level representation under all the different OOD settings. Given that the factorization breaks down, the transition model can no longer correctly identify the right object to update its state. This is indeed the case, as shown in Figure 2C. We visualize the transition updates such that for each subplot of observations the same object is acted upon. In OOD conditions, we observe that the transition model updates the state for multiple objects in each observation, even though the true prediction is for the transition model to update the state of just one object in an observation. This is in contrast to IID scenario where we observed that one object state was updated for every observation. The failure in accurate factorization of representations results in inaccurate transition updates as seen in the low prediction performance.

Overall, the slot mechanism of CSWM model is unable to factorize the representations when attributes of the object are changed during test time, resulting in poor prediction performance.

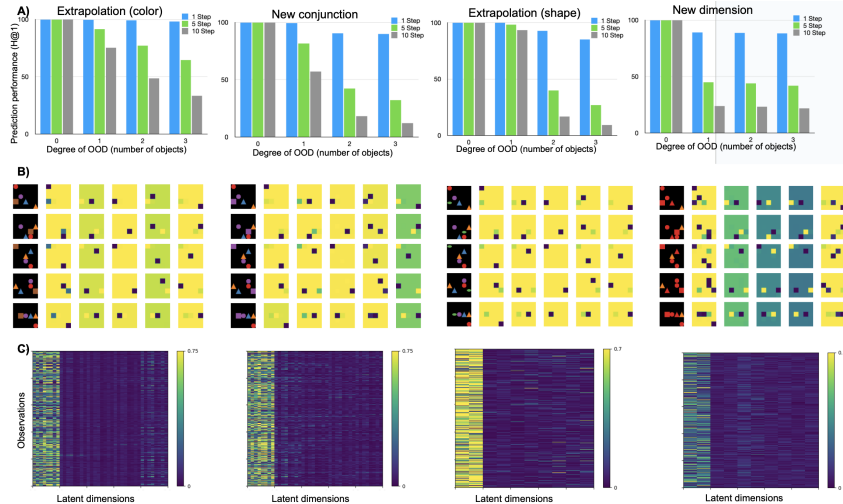


Figure 3: Evaluation of CSWM model under OOD generalization. A) H@1 prediction performance of the model under different types and extent of OOD. B) Visualization of convolutional maps of the model corresponding to each generalization test. C) Transition updates corresponding to each generalization test.




Figure 4: OOD generalization performance of CSWM on 3D blocks and 3 body dataset.

4 Discussion

In this paper, we find that the tested world model is unable to generalize to novel data which can be attributed to the breakdown in factorization of representations. While we aim to conduct an exhaustive evaluation to understand the generalization properties of slot-based models, there are limitations regarding datasets and models. The datasets used in the study are based on synthetic images and have relatively simple object dynamics. As regards to the model, there are other approaches to building object-centric models that we did not test. A number of approaches use generative models [2] to build object-centric representations.

Acknowledgements

We would like to thank Zhe Li for helpful discussions. This research has been funded by the NSF NeuroNex program through grant DBI-1707400 and UF1 NS126566 awarded to AST and XP. This work was also supported in part upon work supported by the Air Force Office of Scientific Research (AFOSR) under award number FA9550-21RT0750 to XP.

References

- [1] Biza, O., Kipf, T., Klee, D., Platt, R., van de Meent, J.W., Wong, L.L., 2022. Factored world models for zero-shot generalization in robotic manipulation. arXiv preprint arXiv:2202.05333 .
- [2] Burgess, C.P., Matthey, L., Watters, N., Kabra, R., Higgins, I., Botvinick, M., Lerchner, A., 2019. Monet: Unsupervised scene decomposition and representation. arXiv preprint arXiv:1901.11390 .
- [3] Craik, K.J.W., 1967. The nature of explanation. volume 445. CUP Archive.
- [4] Dittadi, A., Träuble, F., Locatello, F., Wüthrich, M., Agrawal, V., Winther, O., Bauer, S., Schölkopf, B., 2020. On the transfer of disentangled representations in realistic settings. arXiv preprint arXiv:2010.14407 .
- [5] Huang, Q., He, H., Singh, A., Zhang, Y., Lim, S.N., Benson, A.R., 2020. Better set representations for relational reasoning. Advances in Neural Information Processing Systems 33, 895–905.
- [6] Jaques, M., Burke, M., Hospedales, T., 2019. Physics-as-inverse-graphics: Joint unsupervised learning of objects and physics from video. arXiv preprint arXiv:1905.11169 .
- [7] Karl, M., Soelch, M., Bayer, J., Van der Smagt, P., 2016. Deep variational bayes filters: Unsupervised learning of state space models from raw data. arXiv preprint arXiv:1605.06432 .
- [8] Kipf, T., Van der Pol, E., Welling, M., 2019. Contrastive learning of structured world models. arXiv preprint arXiv:1911.12247 .
- [9] Kossen, J., Stelzner, K., Hussing, M., Voelcker, C., Kersting, K., 2019. Structured object-aware physics prediction for video modeling and planning. arXiv preprint arXiv:1910.02425 .
- [10] LeCun, Y., 2022. A path towards autonomous machine intelligence version 0.9. 2, 2022-06-27. Open Review 62.
- [11] Locatello, F., Abbati, G., Rainforth, T., Bauer, S., Schölkopf, B., Bachem, O., 2019a. On the fairness of disentangled representations. Advances in neural information processing systems 32.
- [12] Locatello, F., Bauer, S., Lucic, M., Raetsch, G., Gelly, S., Schölkopf, B., Bachem, O., 2019b. Challenging common assumptions in the unsupervised learning of disentangled representations, in: international conference on machine learning, PMLR. pp. 4114–4124.
- [13] Spelke, E.S., 1990. Principles of object perception. Cognitive science 14, 29–56.
- [14] Spelke, E.S., Gutheil, G., Van de Walle, G., 1995. The development of object perception. Visual cognition: An invitation to cognitive science 2, 297–330.
- [15] Spelke, E.S., Kinzler, K.D., 2007. Core knowledge. Developmental science 10, 89–96.
- [16] Téglás, E., Vul, E., Giroto, V., Gonzalez, M., Tenenbaum, J.B., Bonatti, L.L., 2011. Pure reasoning in 12-month-old infants as probabilistic inference. science 332, 1054–1059.
- [17] Träuble, F., Creager, E., Kilbertus, N., Locatello, F., Dittadi, A., Goyal, A., Schölkopf, B., Bauer, S., 2021. On disentangled representations learned from correlated data, in: International Conference on Machine Learning, PMLR. pp. 10401–10412.
- [18] Van Steenkiste, S., Locatello, F., Schmidhuber, J., Bachem, O., 2019. Are disentangled representations helpful for abstract visual reasoning? Advances in Neural Information Processing Systems 32.
- [19] Veerapaneni, R., Co-Reyes, J.D., Chang, M., Janner, M., Finn, C., Wu, J., Tenenbaum, J., Levine, S., 2020. Entity abstraction in visual model-based reinforcement learning, in: Conference on Robot Learning, PMLR. pp. 1439–1456.
- [20] Wagemans, J., 2015. The Oxford handbook of perceptual organization. OUP Oxford.

A Experimental details

In this section, we provide an overview of the setup and experiments we conduct. We first introduce the datasets, model architecture, and evaluation metrics used in the study. Finally, we outline the task and experiments on which the models’ predictive performance is evaluated.

A.1 Data

Our experiments use three different datasets – 2D Shape, 3D blocks, and 3 body physics dataset. 2D shapes is a 5x5 grid world consisting of 5 different randomly placed different objects. The overall size of the grid is 50x50. Each object of size 10x10 can only occupy an empty location in the grid and has a unique combination of shape and color. When an action takes place, one object is moved by one position (Fig 1B) and the object cannot be moved if the location is occupied by another object or outside the grid. The 3D blocks dataset is also a block pushing environment, similar to the one used in [8]. The rendering component is changed with different perspective and partial occlusions, making it a slightly challenging task. The 3-body environment is an interacting system based on classical gravitational dynamics [6] without any actions. The input to the model is two consecutive images of 50x50 that implicitly contains the velocity information.

To generate an experience buffer for training, we initialize the environment by uniformly sampling objects at random locations. At every time step, we provide state observations as $50 \times 50 \times 3$ tensors with RGB color channels, normalized to [0, 1]. We randomly sample an object-specific action. Actions are provided as a 4-dimensional one-hot vector (if an action is applied) that encodes the directional movement of a particular object. Else it is represented as a vector of zeros if no action is applied to a particular object. Note that only a single object receives an action per time step.

A.2 Model architecture

The world model as illustrated in Figure 1A comprises an object encoder and a transition model. For the contrastive structured world model (CSWM), we follow the architectural details as given in [8]. The object encoder is a convolutional neural network and the transition model is a Graph Neural Network. The transition model accepts a factored latent state and the action vector to predict the next latent state by outputting the transition update as a residual. The network models pairwise interactions between latent state factors (corresponding to objects in the environment) using a fully-connected node network and an edge network. Both are implemented as MLPs with one hidden layer. The edge network outputs an embedding for each directed edge and then aggregated using the node network in order to update the state of each node. Finally, the graph neural network outputs a vector of updated factors. Finally we also use a standard Autoencoder (AE) in our experiments.

A.3 Task evaluation

We use standard metrics used for evaluating world models directly in latent space. We ask the model to predict the representation of state given an observation and action. The true state representation (obtained from the true observation by taking action in the environment) is compared against the predicted state representation against reference states stored in a buffer. This measure is the Hits at Rank 1 (H@1) and Mean reciprocal rank (MRR). We only report Hits at Rank 1 (H@1) in our experiments.

A.4 Experiments

In our experiments we use the contrastive structured world model (CSWM) [8] as illustrated in Figure 1A. Our experiments use three different datasets – 2D Shape, 3D blocks, and 3 body physics dataset. For the OOD evaluation, we test the model under three types of generalization conditions (Figure 1C) i) Extrapolation - object with either a new a color or shape previously unseen during training is introduced ii) New conjunction - object with a unique shape-color combination unseen during training but either the shape or color seen during training iii) New dimension - either shape or color variation is introduced in testing.