
Inductive Biases for Object-Centric Representations in the Presence of Complex Textures

Samuele Papa¹ Ole Winther^{2,3,4} Andrea Dittadi^{2,5}

Abstract

Understanding which inductive biases could be helpful for the unsupervised learning of object-centric representations of natural scenes is challenging. In this paper, we systematically investigate the performance of two models on datasets where neural style transfer was used to obtain objects with complex textures while still retaining ground-truth annotations. We find that by using a single module to reconstruct both the shape and visual appearance of each object, the model learns more useful representations and achieves better object separation. In addition, we observe that adjusting the latent space size is insufficient to improve segmentation performance. Finally, the downstream usefulness of the representations is significantly more strongly correlated with segmentation quality than with reconstruction accuracy.

1. Introduction

A core motivation for object-centric learning is that humans interpret complex environments such as natural scenes as the composition of distinct interacting objects. Evidence for this claim can be found in cognitive psychology and neuroscience (Spelke, 1990; Téglás et al., 2011; Wagemans, 2015; Dehaene, 2020). Current object-centric learning approaches try to merge the advantages of connectionist and symbolic methods by representing each object with a distinct vector (Greff et al., 2020). The problem of object separation becomes central for unsupervised methods that can only use the data itself to learn how to isolate objects. Several methods have been proposed to provide inductive biases to

¹POP-AART Lab, University of Amsterdam and The Netherlands Cancer Institute ²Technical University of Denmark ³University of Copenhagen ⁴Rigshospitalet, Copenhagen University Hospital ⁵Max Planck Institute for Intelligent Systems, Tübingen. Correspondence to: Samuele Papa <s.papa@uva.nl>, Andrea Dittadi <adit@dtu.dk>.



Figure 1: Left: samples from the original datasets. Right: samples from the same datasets with neural style transfer.

achieve this objective (e.g., Burgess et al., 2019; Engelcke et al., 2020b; Locatello et al., 2020; Engelcke et al., 2020a; Kipf et al., 2021). However, they are typically tested on simple datasets where objects show little variability in their texture, often being monochromatic. This characteristic may allow models to successfully separate objects by relying on low-level characteristics, such as color (Greff et al., 2019), over more desirable high-level ones, such as shape.

Research in the direction of natural objects is still scarce (Karazija et al., 2021; Engelcke et al., 2021; Kipf et al., 2021), as such datasets often do not provide exhaustive knowledge of the factors of variation, which are very rich in natural scenes. In this context, unsupervised methods struggle to learn object-centric representations, and the reason for this remains unexplained (Greff et al., 2019, Section 5).

In this paper, we conduct a systematic experimental study on the inductive biases necessary to learn object-centric representations when objects have complex textures. To obtain significant and interpretable results, we focus on static images and use *neural style transfer* (Gatys et al., 2016) to apply complex textures to the objects of the Multi-Sprites (Kabra et al., 2019) and CLEVR (Johnson et al., 2017) datasets, as shown in Fig. 1. This makes the task significantly more challenging while still preserving all the advantages of a procedurally generated dataset, and avoiding the above-mentioned pitfalls of natural datasets.

We investigate MONet (Burgess et al., 2019) and Slot Attention (Locatello et al., 2020), two popular slot-based autoencoder models that learn to represent objects separately and in a common format. The latter obtains object representations by applying Slot Attention to a convolutional embedding of the input, and then decodes each representation into shape and visual appearance via a *single* component. In contrast, MONet reconstructs them with *two* components: a recurrent attention network that

segments the input, and a variational autoencoder (VAE) (Kingma & Welling, 2014; Rezende et al., 2014) that separately learns a representation for each object by learning to reconstruct its shape and visual appearance. Unlike in Slot Attention, the shape reconstructed by the VAE is *not* used to reconstruct the final image; instead the shape from the recurrent attention network is used. To still have shape information in the latent representation of the VAE, the training loss includes a KL divergence between the mask predicted by the VAE and the one predicted by the attention network. Therefore shape and visual appearance are, in effect, *separate* unless the KL divergence provides a strong enough signal. For this study, we posit two desiderata for object-centric models, adapted from Dittadi et al. (2021b):

Desideratum 1. *Object separation and reconstruction.* The models should have the ability to accurately separate and reconstruct the objects in the input, even those with complex textures. For the models considered here, this means that they should correctly segment the objects and reproduce their properties in the reconstruction, including their texture.

Desideratum 2. *Object representation.* The models should capture and represent the fundamental properties of each object present in the input. When ground-truth properties are available for the objects, this can be evaluated via a downstream prediction task.

We summarize our **main findings** on the two models analyzed as follows: **(1)** Models that better balance the importance of both shape and visual appearance of the objects seem to be less prone to what we call *hyper-segmentation* (see Section 3). We show how this can be achieved with an architecture that uses a single module to obtain both shape and visual appearance of each object. When this is not the case, it becomes significantly more challenging for a model to correctly separate objects and learn useful representations. **(2)** Hyper-segmentation of the input leads to the inability of the model to obtain useful representations. Separation is a strong indicator of representation quality. **(3)** The *representation bottleneck* is not sufficient to regulate a model’s ability to segment the input. Tuning other hyperparameters such as encoder and decoder capacity appears to be necessary.

2. Methods

Datasets and models. Similarly to Dittadi et al. (2021b), we use neural style transfer (Gatys et al., 2016) to increase the complexity of the texture of the objects in the Multi-dSprites and CLEVR datasets (see Appendix B for details). This allows for textures that have high variability but are still correlated with the shape of the object, as opposed to preset patterns as done in Greff et al. (2019) and Karazija et al. (2021) or completely random ones. We apply neural style transfer to the entire image and then select the

objects using the ground-truth segmentation masks (see Fig. 1). Keeping the background simple allows for a more straightforward interpretation of the models’ performance. As object-centric models, we consider MONet (Burgess et al., 2019) and Slot Attention (Locatello et al., 2020), chosen because they implement very distinct mechanisms to solve the the problem of separation (see Section 1).

Evaluation. Following the two desiderata in Section 1, we separately focus on the *separation*, *reconstruction*, and *representation* performance of the models. *Separation* is measured by the Adjusted Rand Index (ARI) (Hubert & Arabi, 1985), which quantifies the similarity between two partitions of a set. *Reconstruction* is measured using the Mean Squared Error (MSE) between input and reconstructed images. *Representation* is measured by the performance of a simple downstream model trained to predict the properties of each object using only the object representations as inputs. Following previous literature (Locatello et al., 2020; Dittadi et al., 2021b), we match ground-truth objects with object representations such that the overall loss is minimized.

Performance studies. The *baseline* performance of the models on the style transfer datasets is established using the hyperparameters from the literature. We then vary parameters and architectures to improve performance. In MONet, we reduce the number of skip-connections of the U-Net in the attention module, we change the latent space size, the number of channels in the encoder and decoder of the VAE, and the β and σ parameters in the training objective. In Slot Attention, we increase the number of layers and channels in both encoder and decoder and increase the size of the latent space. For both models, we investigate how the latent space size alone affects performance. We use multiple random seeds to account for variability in performance when feasible (see Appendix E for further details).

3. Experiments

Architectural biases. Qualitatively, the MONet baseline (Fig. 2a) segments primarily according to color, resulting in each slot encoding fragments of multiple objects that share the same color. We call this behavior *hyper-segmentation*. On the other hand, the Slot Attention baseline (Fig. 2c) produces blurred reconstructions and avoids hyper-segmentation. Here, some objects are still split across more than one slot but, unlike in MONet, we do not observe multiple objects that are far apart in the scene being (partially) modeled by the same slot. We observe this quantitatively in Fig. 3 (top): compared to the Slot Attention baseline, the MONet baseline has a significantly worse ARI score but a considerably better MSE. Fig. 12 in Appendix G shows similar results on CLEVR.

These observations can guide our search for better model

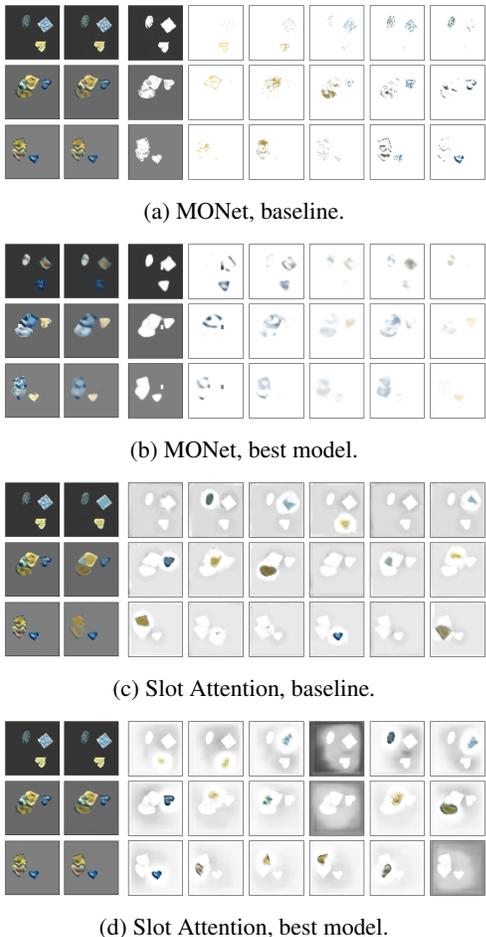


Figure 2: Qualitative results on Multi-dSprites (validation set). From left to right in each subfigure: original input, final reconstruction, and product of the reconstruction and mask for each of the six slots. MONet still fails to separate objects correctly although it blurs the reconstructions. See Fig. 12 for results on CLEVR.

parameters. Slot Attention is blurring away the small details of the texture and focuses on the shape to separate them. MONet, instead, achieves good reconstructions but does so by using the attention module to select pixels that share the same color, as opposed to entire objects, while the VAE simply reconstructs plain colors (see Appendix G for more details). Therefore, for MONet we attempt to sacrifice some reconstruction quality in exchange for better object separation. For Slot Attention, we investigate whether improving the reconstructions negatively affects object separation. We refer to Appendix E for further details and results on the hyperparameter search for MONet and Slot Attention.

We now discuss the results obtained using the combination of hyperparameters that achieves the best performance, called *best model*. In Fig. 2b, we see that MONet still hyper-segments even though the reconstructions are now blurred. For Slot Attention (Fig. 2d), we observe that the quality of reconstructions has improved, and it more often represents an entire object in a single slot. Although the ARI for

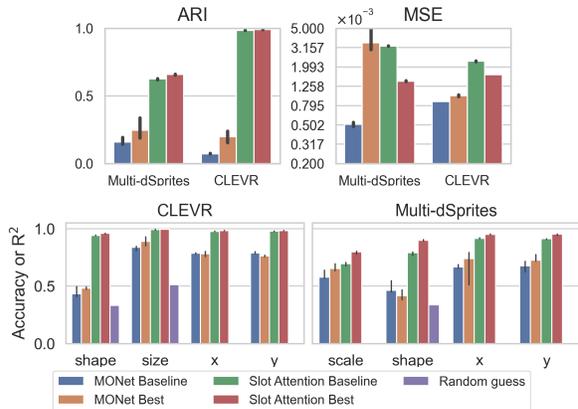


Figure 3: Median performance of the different seeds trained for each of the indicated models (error bars denote 95% confidence intervals). Top: ARI (\uparrow) and MSE (\downarrow) for each dataset and model. Bottom: performance of the downstream model on each feature of the objects. Accuracy is used for categorical features and R^2 for numerical features. A random guess baseline is shown in purple.

MONet has also improved, the separation problem is still far from solved, while Slot Attention shows a significant improvement both in terms of ARI and MSE (see two upper-most plots in Fig. 3). Note how, for Slot Attention, the ARI is significantly lower in Multi-dSprites, when compared to CLEVR. The likely reason is that, when a significant portion of an object is occluded by another, the visible shape is being altered significantly and the edges of objects are not clear. Therefore, two explanations of the same scene can still be reasonable while not corresponding to the ground truth. This extreme overlap never occurs in CLEVR.

Overall, even when MONet sacrifices reconstruction quality and blurs away the details, hyper-segmentation is still present as evidenced by our qualitative and quantitative analyses. This suggests that the separation problem in MONet may not simply be caused by the training objective, but rather by its architectural biases. Indeed, improving the reconstruction performance of Slot Attention has, instead, yielded both better separation and more detailed reconstructions, suggesting that generating shape and appearance using a single module is a more favorable inductive bias for learning representations of objects with complex textures.

Representation performance. To understand the interplay between separation and learned representations, we explore the performance on a downstream property prediction task that was trained using the object representations as inputs (see Appendix C.1 for details). From Fig. 3, we observe how MONet fails to capture some of the properties in the representations and consistently performs worse than Slot Attention, for both the baseline and the improved versions. This suggests that, as highlighted in Dittadi et al. (2021b), a model that is not capable of correctly separating objects will also fail to accurately represent them. The trend is also clear from Fig. 4, which shows that a higher ARI score

		CLEVR						Multi-dSprites			
ARI		92	92	93	85	ARI		79	89	81	88
MSE		-29	-20	-15	-1	MSE		-21	-24	-2	-11
	shape		size	x	y		scale	shape	x	y	

Figure 4: Rank correlation of the ARI and MSE scores with downstream property prediction performance. Correlations are computed over *all* the models trained with that dataset in our study.

strongly correlates with an increased performance of the downstream model on all object properties. The correlation with MSE is much weaker, which highlights how *strong visual reconstruction performance is not the ultimate indicator for good object representations*. This result does not contradict previous findings (Dittadi et al., 2021b) as here the properties we expect the downstream model to predict have little to do with the texture of the object and, therefore, the model can have poorer reconstructions while still obtaining useful representations.

Representation bottleneck. Here, we investigate how the size of the *representation bottleneck* (see Appendix A) affects performance. In Fig. 7, we observe that the MSE improves slightly and quickly plateaus for both MONet and Slot Attention when the latent space size increases. However, the ARI does not significantly change even with increased latent size. The increase in latent space size arguably increases the model’s capacity, but it does not prove to be enough to improve the separation and reconstruction performance significantly.

4. Related work

Learning representations that reflect the underlying structure of data is believed to be useful for downstream learning and systematic generalization (Bengio et al., 2013; Lake et al., 2017; Greff et al., 2020). While many recent empirical studies have investigated the usefulness of disentangled representations and the inductive biases involved in learning them (Locatello et al., 2019; Van Steenkiste et al., 2019; Träuble et al., 2021; Montero et al., 2021; Dittadi et al., 2021c; 2022), analogous experimental studies in the context of object-centric representations are scarce. The study by Engelcke et al. (2020a) presents an investigation into inductive biases for object separation, focusing on one model and traditional synthetic datasets. In this work, we study hyper-segmentation on datasets where objects have complex textures. Karazija et al. (2021) recently proposed ClevrTex, a dataset that introduces challenging textures to scenes from CLEVR (Johnson et al., 2017). The experiments reported show that, without any tuning, some models fail to segment complex scenes by focusing on colors. The authors, similarly to what we highlighted in this work, state that “ignoring confounding aspects of the scene rather than representing them might aid in the overall task [of segmentation].” In

our work, we investigate the mechanism behind the ability of some models to ignore superfluous details and more successfully segment the image, proposing a useful inductive bias to achieve better object representations.

Recently, works that propose new object-centric learning methods also include evaluations on more complex datasets. Greff et al. (2019) train IODINE on Textured MNIST and ImageNet, and observe that the model separates the image primarily according to color when the input is complex. GENESIS-V2 (Engelcke et al., 2021) was trained on Sketchy and APC, two real-world robot manipulation datasets. However, the authors do not explore the mechanism behind the performance of the models they tested, and do not attempt to optimize the architectures. In the video domain, Kipf et al. (2021) include evaluations on a dataset with complex textures, training their model to predict optical flow rather than a reconstruction of the input.

5. Conclusions

In this paper, we have investigated which inductive biases may be most useful for slot-based unsupervised models to obtain good object-centric representations of scenes where objects have complex textures. We found that using a single module to reconstruct both shape and visual appearance of objects naturally balances the importance of these two aspects in the generation process, thereby avoiding hyper-segmentation and achieving a better compromise between precise texture reconstructions and correct object segmentation. Therefore, our recommendation is that models should have separation as an integral part of the representation process. Additionally, we showed that separation strongly correlates with the quality of the representations, while reconstruction accuracy does not: this justifies sacrificing some reconstruction quality. Finally, we observed that the representation bottleneck is not a sufficient inductive bias, as increasing the latent space size can be counterproductive unless the model is already separating the input correctly.

Although the models considered in our study have been shown to be among the most successful ones on this type of data, it would be interesting to explore if the same conclusions hold for other models. Another interesting avenue for future work is to extend our study to more complex downstream tasks involving abstract reasoning, e.g., in a neuro-symbolic system, where symbol manipulation can be performed either within a connectionist framework (Evans & Grefenstette, 2018; Smolensky, 1990; Battaglia et al., 2018) or by purely symbolic methods (Asai & Fukunaga, 2018; Mao et al., 2019; Dittadi et al., 2021a). Finally, it would be relevant to validate our conclusions on additional datasets, and to introduce objects with varying texture complexity, as this could require different model capacities to achieve separation (Engelcke et al., 2020a).

References

- Asai, M. and Fukunaga, A. Classical planning in deep latent space: Bridging the subsymbolic-symbolic boundary. In *Proceedings of the aaai conference on artificial intelligence*, volume 32, 2018.
- Bachlechner, T., Majumder, B. P., Mao, H., Cottrell, G., and McAuley, J. ReZero is all you need: Fast convergence at large depth. In *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence*, pp. 1352–1361. PMLR, December 2021.
- Battaglia, P. W., Hamrick, J. B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V., Malinowski, M., Tacchetti, A., Raposo, D., Santoro, A., Faulkner, R., Gulcehre, C., Song, F., Ballard, A., Gilmer, J., Dahl, G., Vaswani, A., Allen, K., Nash, C., Langston, V., Dyer, C., Heess, N., Wierstra, D., Kohli, P., Botvinick, M., Vinyals, O., Li, Y., and Pascanu, R. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018.
- Bengio, Y., Courville, A., and Vincent, P. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, 2013.
- Burgess, C. P., Matthey, L., Watters, N., Kabra, R., Higgins, I., Botvinick, M., and Lerchner, A. Monet: Unsupervised scene decomposition and representation. *arXiv preprint arXiv:1901.11390*, 2019.
- Dehaene, S. *How We Learn: Why Brains Learn Better than Any Machine ... for Now*. Viking, New York, 2020.
- Dittadi, A., Drachmann, F. K., and Bolander, T. Planning from pixels in atari with learned symbolic representations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 4941–4949, 2021a.
- Dittadi, A., Papa, S., De Vita, M., Schölkopf, B., Winther, O., and Locatello, F. Generalization and robustness implications in object-centric learning. *arXiv preprint arXiv:2107.00637*, 2021b.
- Dittadi, A., Träuble, F., Locatello, F., Wüthrich, M., Agrawal, V., Winther, O., Bauer, S., and Schölkopf, B. On the transfer of disentangled representations in realistic settings. In *International Conference on Learning Representations*, 2021c.
- Dittadi, A., Träuble, F., Wüthrich, M., Widmaier, F., Gehler, P. V., Winther, O., Locatello, F., Bachem, O., Schölkopf, B., and Bauer, S. The role of pretrained representations for the OOD generalization of RL agents. In *International Conference on Learning Representations*, 2022.
- Engelcke, M., Jones, O. P., and Posner, I. Reconstruction bottlenecks in object-centric generative models. *arXiv preprint arXiv:2007.06245*, 2020a.
- Engelcke, M., Kosiorek, A. R., Jones, O. P., and Posner, I. GENESIS: Generative Scene Inference and Sampling with Object-Centric Latent Representations. In *ICLR 2020*, 2020b.
- Engelcke, M., Parker Jones, O., and Posner, I. GENESIS-V2: Inferring unordered object representations without iterative refinement. *arXiv preprint arXiv:2104.09958*, 2021.
- Evans, R. and Grefenstette, E. Learning explanatory rules from noisy data. *Journal of Artificial Intelligence Research*, 61:1–64, 2018.
- Gatys, L., Ecker, A., and Bethge, M. A Neural Algorithm of Artistic Style. *Journal of Vision*, 16(12):326–326, September 2016. ISSN 1534-7362. doi: 10.1167/16.12.326.
- Greff, K., Kaufman, R. L., Kabra, R., Watters, N., Burgess, C., Zoran, D., Matthey, L., Botvinick, M., and Lerchner, A. Multi-object representation learning with iterative variational inference. In *International Conference on Machine Learning*, pp. 2424–2433. PMLR, 2019.
- Greff, K., van Steenkiste, S., and Schmidhuber, J. On the binding problem in artificial neural networks. *arXiv preprint arXiv:2012.05208*, 2020.
- Hubert, L. and Arabi, P. Comparing partitions. *Journal of Classification*, 2(1):193–218, December 1985. ISSN 1432-1343. doi: 10.1007/BF01908075.
- Jacq, A. Neural Transfer Using PyTorch — PyTorch Tutorials 1.10.1+cu102 documentation. https://pytorch.org/tutorials/advanced/neural_style_tutorial.html, 2021.
- Johnson, J., Hariharan, B., Van Der Maaten, L., Fei-Fei, L., Lawrence Zitnick, C., and Girshick, R. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2901–2910, 2017.
- Kabra, R., Burgess, C., Matthey, L., Kaufman, R. L., Greff, K., Reynolds, M., and Lerchner, A. Multi-object datasets, 2019.
- Karazija, L., Laina, I., and Ruppel, C. Clevrtex: A texture-rich benchmark for unsupervised multi-object segmentation. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.

- Kingma, D. P. and Welling, M. Auto-encoding variational Bayes. In *International Conference on Learning Representations*, 2014.
- Kipf, T., Elsayed, G. F., Mahendran, A., Stone, A., Sabour, S., Heigold, G., Jonschkowski, R., Dosovitskiy, A., and Greff, K. Conditional object-centric learning from video. *arXiv preprint arXiv:2111.12594*, 2021.
- Lake, B. M., Ullman, T. D., Tenenbaum, J. B., and Gershman, S. J. Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40, 2017.
- Locatello, F., Bauer, S., Lucic, M., Gelly, S., Schölkopf, B., and Bachem, O. Challenging common assumptions in the unsupervised learning of disentangled representations. In *International Conference on Machine Learning*, 2019.
- Locatello, F., Weissenborn, D., Unterthiner, T., Mahendran, A., Heigold, G., Uszkoreit, J., Dosovitskiy, A., and Kipf, T. Object-centric learning with slot attention. *arXiv preprint arXiv:2006.15055*, 2020.
- Mao, J., Gan, C., Kohli, P., Tenenbaum, J. B., and Wu, J. The neuro-symbolic concept learner: Interpreting scenes, words, and sentences from natural supervision. In *International Conference on Learning Representations*, 2019.
- Montero, M. L., Ludwig, C. J., Costa, R. P., Malhotra, G., and Bowers, J. The role of disentanglement in generalisation. In *International Conference on Learning Representations*, 2021.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. Pytorch: An imperative style, high-performance deep learning library. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 32*, pp. 8024–8035. Curran Associates, Inc., 2019.
- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082*, 2014.
- Smolensky, P. Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial intelligence*, 46(1-2):159–216, 1990.
- Spelke, E. S. Principles of object perception. *Cognitive Science*, 14:29–56, 1990. ISSN 0364-0213. doi: 10.1016/0364-0213(90)90025-R.
- Téglás, E., Vul, E., Giroto, V., Gonzalez, M., Tenenbaum, J. B., and Bonatti, L. L. Pure Reasoning in 12-Month-Old Infants as Probabilistic Inference. *Science*, 332:1054–1059, May 2011. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.1196404.
- Träuble, F., Creager, E., Kilbertus, N., Locatello, F., Dittadi, A., Goyal, A., Schölkopf, B., and Bauer, S. On disentangled representations learned from correlated data. In *International Conference on Machine Learning*, pp. 10401–10412. PMLR, 2021.
- Van Steenkiste, S., Locatello, F., Schmidhuber, J., and Bachem, O. Are disentangled representations helpful for abstract visual reasoning? *arXiv preprint arXiv:1905.12506*, 2019.
- Wagemans, J. *The Oxford handbook of perceptual organization*. Oxford Library of Psychology, 2015.

A. Remarks on notation

The term *representation bottleneck* should not be confused with the *reconstruction bottleneck* introduced by Engelcke et al. (2020a). The representation bottleneck refers to the small size of the latent space, while the reconstruction bottleneck refers to how easy it is for the model to reconstruct the data. In Engelcke et al. (2020a), the reconstruction bottleneck is posited to be the reason behind the models' inability to separate objects into different slots.

Often, in the paper, we refer to *object-centric representations* and *slots* as synonyms, although this is only true for slot-based models.

The term *hyper-segmentation* refers to when a model splits the input into different slots with little to no regard to high-level characteristics of the input, such as the shapes of objects, and instead just uses low-level characteristics, primarily color. This often results in slots that consist of small clusters of pixels with similar color, which means that several objects can be partially represented in the same slot and at the same time each object may be partially represented in multiple slots. This phenomenon is distinct from *over-segmentation* (Engelcke et al., 2020b), where multiple slots may reconstruct a single object but no slot reconstructs (parts of) multiple objects. Examples are shown in the main text of the paper (see Fig. 2b and Fig. 2b), where MONet is hyper-segmenting the input, while Slot Attention is sometimes over-segmenting it.

B. Datasets

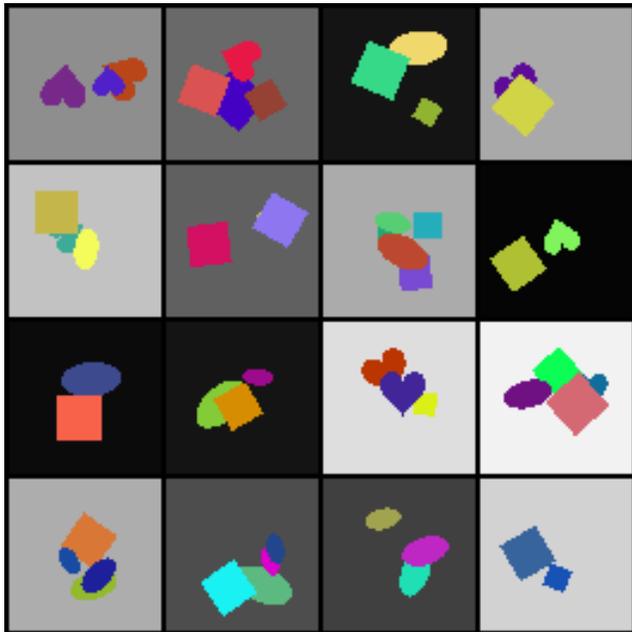
The original versions of both datasets are taken from Kabra et al. (2019).

CLEVR. The CLEVR dataset consists of 3D objects placed on a gray background at different distances from the camera. Overlap between objects is kept to a minimum. There are spheres, cylinders, and cubes of eight different colors. The objects can be metallic or opaque. There is a big and small variant of each object and they can be placed in several different orientations. We use the variant of the dataset that has no more than 6 objects in it, as was done in previous object-centric learning research. The total number of samples in the training dataset is 49483, and we leave 2000 samples for validation and 2000 samples for testing.

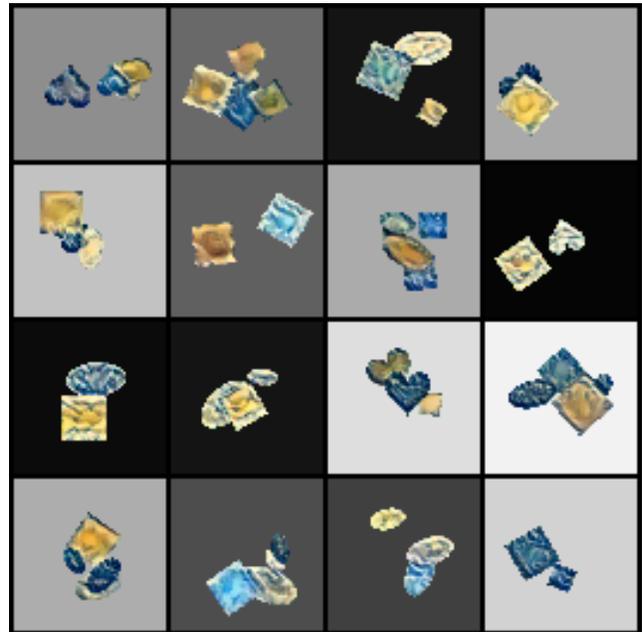
Multi-dSprites. The Multi-dSprites dataset places several 2D objects on a grayscale background. The objects can be a square, an ellipse, or a heart. They can have any RGB color, orientation, and different levels of overlap. Here, we use 90000 samples for the training, 5000 for validation, and 5000 for testing.

Neural Style Transfer. Neural Style Transfer was applied in its most basic form (Jacq, 2021) except for a few additions to make running it on several datasets easier. We opted to use *The Starry Night* by Dutch painter *Vincent Van Gogh* as a reference style image (we used the photo from Wikimedia Commons, which is in the public domain). We experimented with several parameters, and we noticed a lot of variability between runs and a more pleasant result from the most basic implementation of the algorithm.

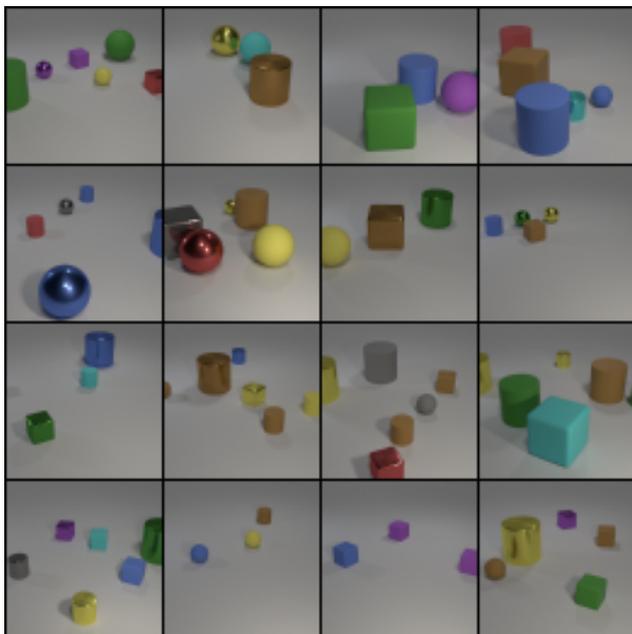
The final version of the datasets was obtained by first applying neural style transfer to each image (optimization happens on an image-to-image basis). This results in the entire scene having the style of the reference image. After obtaining the neural style transfer version of the image, we applied the original segmentation masks of the objects to obtain an image where only the foreground objects have a complex texture, while the background remains the original one.



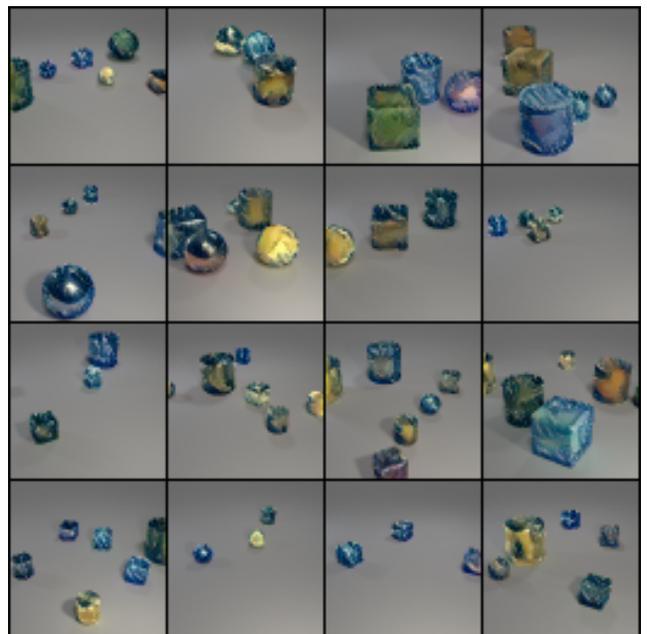
(a) Samples from original Multi-dSprites.



(b) Samples from style transfer Multi-dSprites.



(c) Samples from original CLEVR.



(d) Samples from style transfer CLEVR.

Figure 5: Samples from the original and neural style transfer datasets.

C. Evaluation

C.1. Downstream feature prediction task

The setup for the feature prediction task is the same as the one used in [Dittadi et al. \(2021b\)](#). The models used are a simple linear model and an MLP with one hidden layer having 256 neurons and enough outputs to predict all of the features of an object. The input to the model is the object representation of a single object and the output is the predicted features for that object. Let r be the representation of an object, M the model, $\hat{y} = M(r)$ is the output of the model, and y is the target vector such that $y_{i_k:i_{k+1}}$ is the k th feature of the object, a vector of dimension $(i_{k+1} - i_k + 1)$. We use the MSE loss for numerical features and the cross-entropy loss for categorical ones.

Now, it is important to note that, in order to correctly train the model, the representation r needs to be matched with the target vector y of the object that r is representing. However, this is very challenging, as the models can represent any object in any of the slots. Therefore, following [Dittadi et al. \(2021b\)](#), we adopted two different strategies to match slots with objects. The first is called *loss matching*: The loss for each pair of slot and object is computed, resulting in a loss matrix L , where $L_{i,j}$ is the loss between the predicted features from the j th slot and the target features from the i th object in the scene. Then, the Hungarian matching algorithm is used to find the pairs of slots-objects that minimize the sum of the loss. The second approach is called *mask matching*: The masks predicted by the models and the ground masks are matched, to find the pairs that have the smallest difference. By using loss matching, the assumption is that the initial errors that are inevitable (because the downstream model has not been trained yet) will eventually disappear. When using mask matching, this problem disappears, however, we rely on the ability of the models to generate masks that closely match the ground truth, which is not the case for models that are hyper-segmenting the input, as is often the case in our study.

C.2. ARI and MSE

We use the standard definitions of Adjusted Rand Index (ARI) and Mean Square Error (MSE).

The ARI, as the name suggests, is the Adjusted version of the Rand Index. The Rand Index is defined as follows. Given a set of n elements S and two partitions A and B of this set, the Rand Index looks at the number of pairs of elements from S and what set they belong to in the two partition. The definition is the following:

$$R := \frac{m_{11} + m_{00}}{m_{11} + m_{00} + m_{10} + m_{01}},$$

where m_{11} is the number of pairs that are in the same set in both A and B , m_{00} the number of pairs that are in different sets in both partitions, m_{10} the number of pairs that are in the same set in A but different sets in B , and m_{01} counts how many pairs are in different sets in A but the same set in B .

Then, the ARI normalizes and corrects the bias from the Rand Index based on a given null hypothesis ([Hubert & Arabi, 1985](#)), such that an ARI of 0 results from completely random partitions of the set, while 1 when the two partitions coincide.

It is important to highlight that the ARI is dependent on the size of the partitions and the number of elements present in each set of the partition, which are both assumed to not change across the entire dataset, which is not always true for the current application.

The MSE is computed between each channel in each pixel of the image, following the traditional formula.

D. Implementation of the models

The models were re-implemented in PyTorch ([Paszke et al., 2019](#)) and run on NVIDIA A100 or NVIDIA TitanRTX GPUs. The total approximate training time to reproduce this study is 300 GPU days.

E. Hyperparameter searches

E.1. Baselines

The baselines were obtained by training the models on the two datasets with 3 different seeds. The parameters are taken from the original papers, but for MONet we use different values for the foreground and background sigma, as suggested by [Greff et al. \(2019\)](#). We stopped the training for all runs in our study, even the ones described later, to 500k steps.

E.2. Improving MONet

Starting from the baseline results, we first explored the hyperparameter space manually, to develop an intuition regarding the effect of each hyperparameter on the performance.

We then performed a hyperparameter search for MONet. We ran a full search, resulting in 36 runs. Because of the high number of runs, we decided to use a single seed. The parameters are listed in Table 1. Those that are not listed were kept unchanged. All combinations of parameters are tested, but `foreground sigma` and `background sigma` have been changed in pairs, so that when the foreground sigma is 0.05, the background sigma is 0.03 and when foreground sigma is 0.5, background sigma is 0.3 to keep consistent weights of the reconstruction loss.

Some analysis on the results of these models can be seen in Fig. 6, where we can see how the parameters have little to no impact on the overall performance of the model. What proved to be most effective was reducing the number of skip connections in the U-Net and using a small sigma for the loss function. However, these results are not very conclusive, as a small number of skip connections is actually just increasing the ARI slightly by reconstructing bigger patches of objects in the slots and not actually separating them correctly.

Parameter	Value(s)
<code>foreground sigma</code>	0.05, 0.5
<code>background sigma</code>	0.03, 0.3
<code>gamma</code>	0.05, 1, 5
<code>latent size</code>	64
<code>latent space MLP size</code>	128
<code>decoder input channels</code>	66
<code>number of skip connections in U-Net</code>	0, 3, 5
<code>dataset</code>	CLEVR, Multi-dSprites

Table 1: Hyperparameter search for MONet.

E.3. Representation bottleneck study

The representation bottleneck study was done by changing the latent space of both MONet and Slot Attention with 2 seeds and without changing any other parameter, resulting in 32 runs. The latent sizes tested are shown in Table 2. The findings are summarized in the main text of the paper.

MONet	Slot Attention
8	32
16	64
32	128
64	256
128	512

Table 2: Latent space sizes tested in the study for each of the two models.

E.4. Improving Slot Attention

We tried to increase the size of the encoder and decoder architecture as much as possible, while being reasonable regarding training time and GPU memory. We tested several architectures, with the objective of improving the overall reconstruction quality by reducing the blurriness. We quickly realised that we needed a very deep architecture, therefore, we opted to use residual layers. The final architecture managed to achieve the best results when averaged over 3 different seeds. Each layer is a stack of two convolutional layers, with Leaky ReLU activation functions, a skip connection and we also employ the re-zero strategy (Bachlechner et al., 2021). We increased the latent size to 512, used upscaling in the encoder and downscaling in the decoder. We fixed the broadcast size of the decoder to 32. We used a stack of 16 residual blocks. The architecture of

the encoder is described in Table 3, and the decoder is symmetrical (starting from 256 channels going down and instead of downscaling we have upscaling). To map from the input number of channels to the desired ones we use an additional convolutional layer, the same for the output channels. We did not experiment with the number of iterations that the Slot Attention Module performs, but it would be interesting to understand whether this parameter is very influential in natural scenes.

Name	Number of channels	Activation/ Comment
Residual Block	64	Leaky ReLU
Residual Block	64	Leaky ReLU
Residual Block	64	Leaky ReLU
Residual Block	64	Leaky ReLU
Downscaling		Only for CLEVR
Residual Block	64	Leaky ReLU
Residual Block	64	Leaky ReLU
Residual Block	64	Leaky ReLU
Residual Block	64	Leaky ReLU
Downscaling		
Residual Block	128	Leaky ReLU
Residual Block	128	Leaky ReLU
Residual Block	128	Leaky ReLU
Residual Block	128	Leaky ReLU
Residual Block	256	Leaky ReLU
Residual Block	256	Leaky ReLU
Residual Block	256	Leaky ReLU
Residual Block	256	Leaky ReLU

Table 3: Final encoder used for the Slot Attention model that obtained the best results in terms of both ARI and MSE. Residual Blocks always have 2 convolutions and use ReZero (Bachlechner et al., 2021), two downscaling operations are used for the CLEVR dataset, while one for Multi-dSprites. The decoder is perfectly symmetrical to this structure.

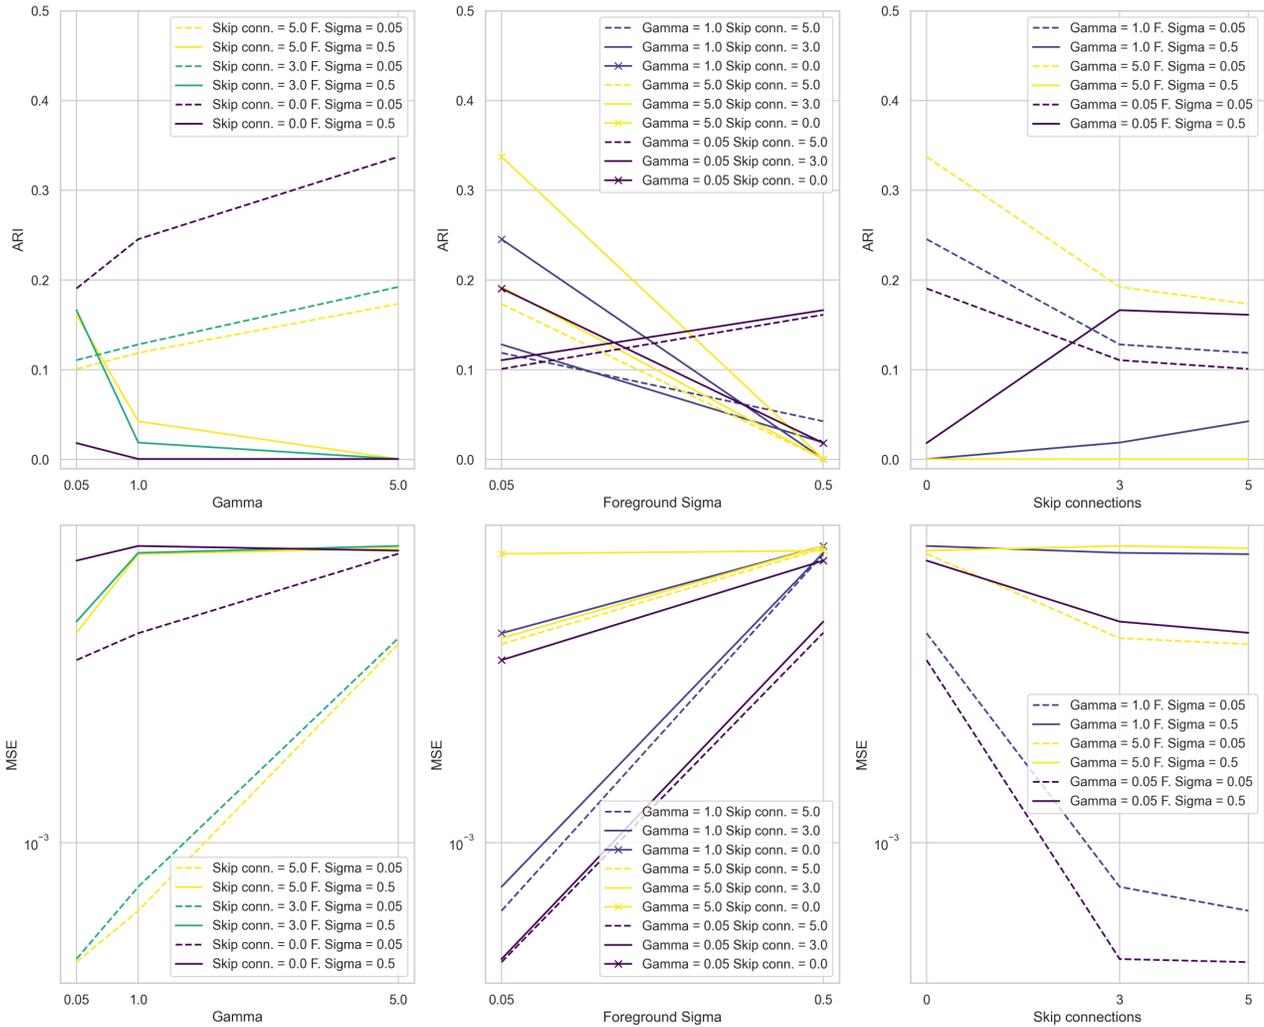


Figure 6: Results (top row: ARI; bottom row: MSE) from the hyperparameter search for MONet. Although increasing gamma or foreground sigma in the loss function successfully deteriorates reconstruction performance (MSE), they are not sufficient to improve the ARI (in fact, the increase of foreground sigma actually decreases the ARI). A smaller number of skip connections also achieves the desired higher MSE, which corresponds to higher ARI only for small values of sigma. Often, having big values for gamma and sigma results in trends opposing the desired ones in terms of ARI.

F. Additional results

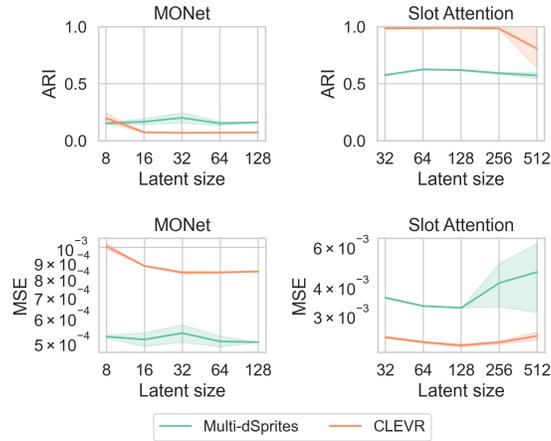


Figure 7: The ARI (\uparrow) and MSE (\downarrow) show no significant change across latent space sizes, supporting that the *representation bottleneck* (see main text) is not a sufficient inductive bias for object separation. Training instability can be seen in Slot Attention with latent size above 128. Two seeds for each latent size are used. Shown are mean (line) and 95% confidence interval (shaded area).

	CLEVR				Multi-dSprites				
ARI	78	81	92	90	80	85	82	84	Model = linear Matching = loss
MSE	-15	-8	-23	-26	-16	-25	-17	-14	
ARI	92	92	93	85	79	89	81	88	Model = MLP1 Matching = loss
MSE	-29	-20	-15	-1	-21	-24	-2	-11	
ARI	83	88	94	93	75	74	88	88	Model = linear Matching = mask
MSE	-21	-11	-16	-13	17	15	-19	-17	
ARI	89	94	93	95	83	85	90	87	Model = MLP1 Matching = mask
MSE	-22	-16	-13	-19	4	7	-24	-21	
	shape	size	x	y	scale	shape	x	y	

Figure 8: Pearson’s correlation coefficient for all runs, grouped by dataset and showing the different combinations of matching and downstream model. The correlation between downstream model’s performance and the ARI (\uparrow) and MSE (\downarrow) metrics shows that ARI is a strongest indicator of good representation quality when the object-centric models are being trained on data with complex texture. Difference in correlation between different matching methods and downstream models is again to be attributed to the poor mask generation quality, which makes mask matching very challenging.

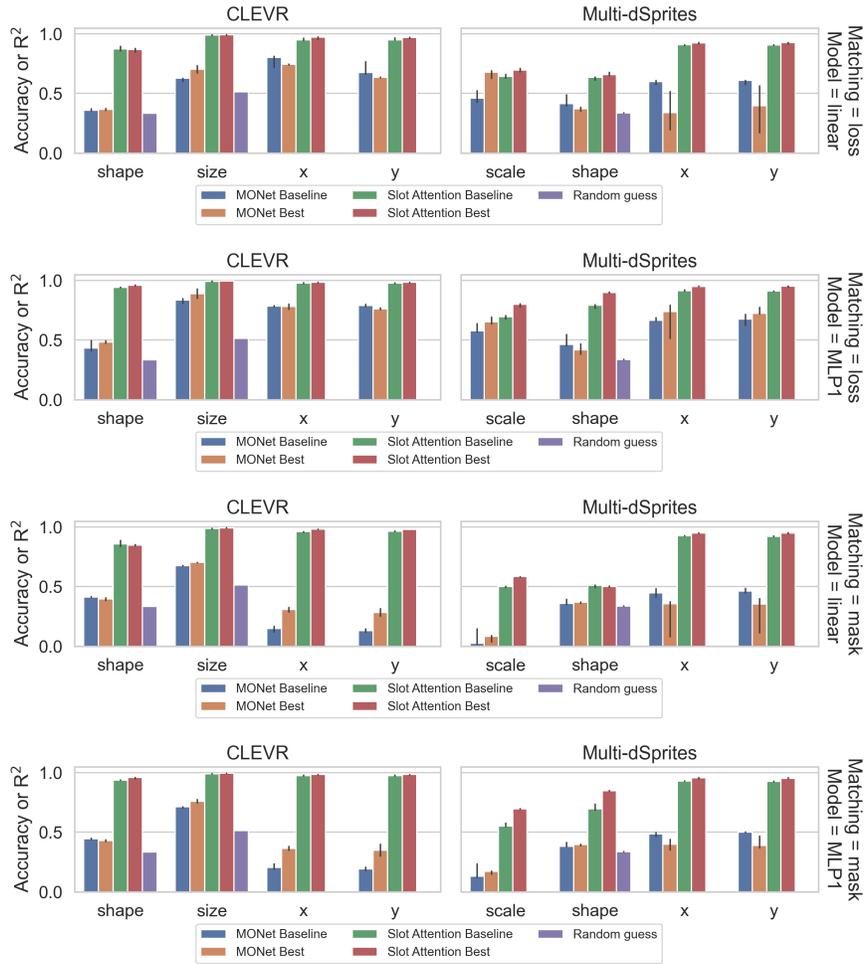


Figure 9: Comparison of the downstream performance for all combinations of slot-object matching and model type (results on from the test set, downstream models trained on the validation set). We notice how accuracy (\uparrow) and R^2 (\uparrow) both increase significantly for both MONet models when using loss matching compared to mask matching (especially for numerical features). This is expected, as the masks generated by MONet suffer substantially from hypersegmentation, which makes mask-matching a very unstable way to pair slots with the correct object. Instead, Slot Attention manages to generate more accurate masks, which results in more consistent performance between mask and loss matching methods.

G. Qualitative results

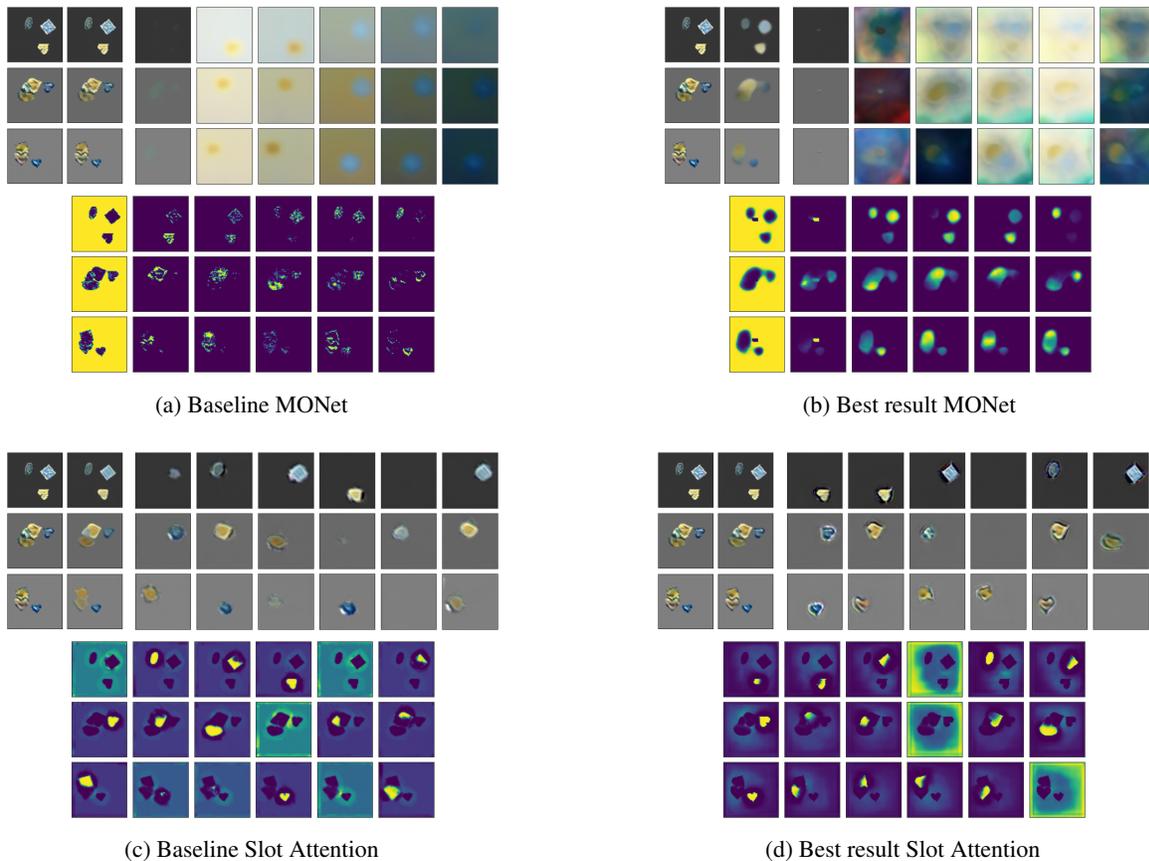


Figure 11: Qualitative results for the separation performance of the models in the comparative study on *Multi-dSprites*. From left to right in all subfigures: (top) input, final reconstruction, reconstruction for each of the six slots (no predicted mask is applied here, only the visual appearance part of the reconstruction is shown), (bottom) mask for each of the six slots. Here the difference between the two versions of Slot Attention is even more noticeable, and we can see how MONet is blurring the masks. However, MONet never manages to reconstruct the correct visual appearance, even when a more accurate shape of the objects is being predicted by the attention module. Balancing visual appearance and shape is much more challenging in MONet.

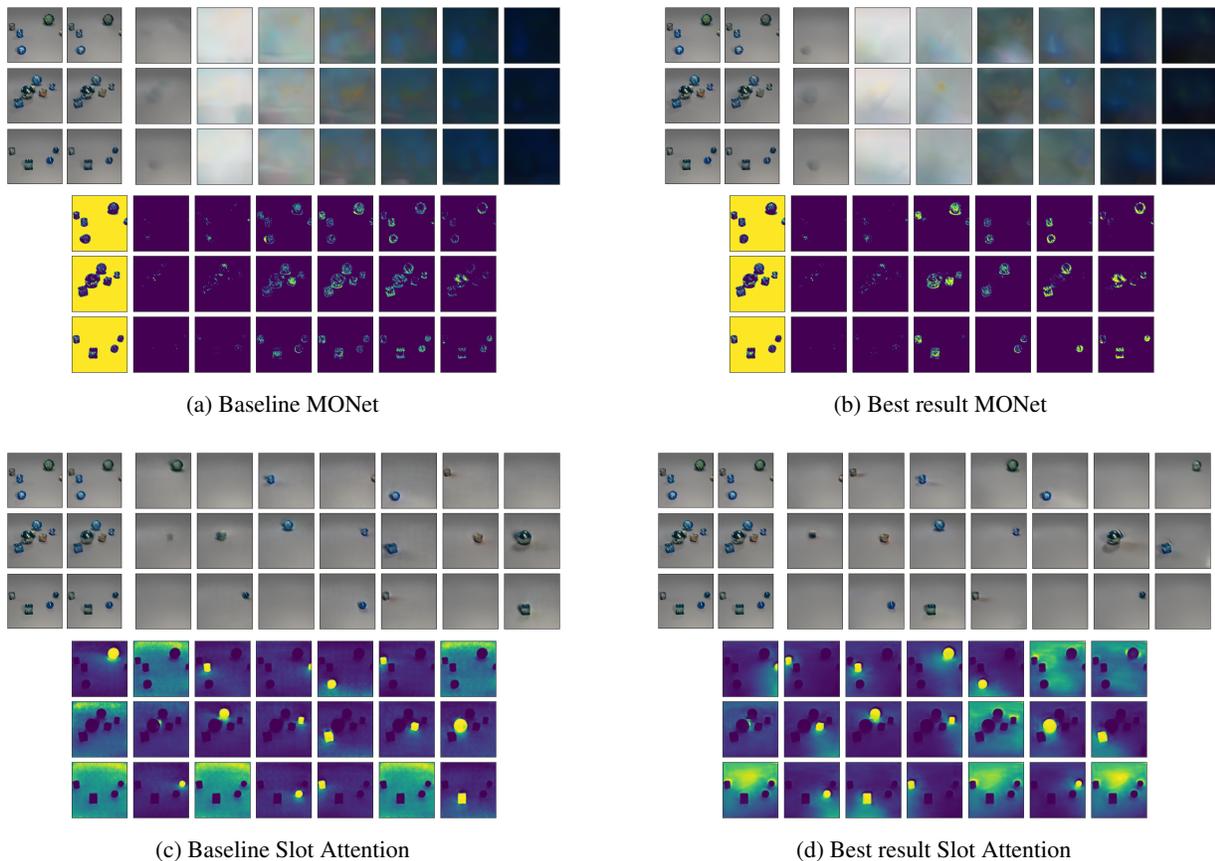


Figure 12: Qualitative results for the separation performance of the models in the comparative study on *CLEVR*. From left to right in all subfigures: (top) input, final reconstruction, reconstruction for each of the six slots (no predicted mask is applied here, only the visual appearance part of the reconstruction is shown), (bottom) mask for each of the six slots. The masks on the improved Slot Attention start to include more of the background for each object. In both baseline and best result, Slot Attention isolates each object in a distinct slot, rarely over-segmenting the input, a stark difference when comparing to Multi-dSprites. For MONet, it manages to perform better in *CLEVR* than Multi-dSprites, however, the best result is still hypersegmenting the input and not blurring it. Overall, MONet cannot reconstruct the visual appearance using the VAE, and leaves all the heavy lifting to the attention module.