

The role of object-centric representations, guided attention, and external memory on generalizing visual relations

Guillermo Puebla (guillermo.puebla@cenia.cl)

National Center for Artificial Intelligence, Vicuña Mackenna 4860
Macul, Chile

Jeffrey S. Bowers (j.bowers@bristol.ac.uk)

School of Psychological Science, University of Bristol, 12a Priory Road
Bristol, United Kingdom

Abstract

Visual reasoning is a long-term goal of vision research. In the last decade, several works have attempted to apply deep neural networks (DNNs) to the task of learning visual relations from images, with modest results in terms of the generalization of the relations learned. In recent years, several innovations in DNNs have been developed in order to enable learning abstract relation from images. In this work, we systematically evaluate a series of DNNs that integrate mechanism such as slot attention, recurrently guided attention, and external memory, in the simplest possible visual reasoning task: deciding whether two objects are the same or different. We found that, although some models performed better than others in generalizing the same-different relation to specific types of images, no model was able to generalize this relation across the board. We conclude that abstract visual reasoning remains largely an unresolved challenge for DNNs.

Keywords: visual reasoning; relational generalization; deep neural networks; guided attention; external memory

Introduction

Detecting relations is one of the fundamental operations of the visual system. This allows us to form a coherent representation of the environment as sets of relations between objects (Vo, Boettcher, & Draschkow, 2019). It is also the basis of robust object recognition, because representing an object as a set of relations between parts frees us from recognizing it solely on the basis of its superficial features (Biederman, 1987). Furthermore, representing relations between entities forms the basis of the kind of reasoning abilities that set us apart from other species (Gentner, Shao, Simms, & Hespos, 2021). Given this predominant role across different forms of visual processing, several researchers have attempted to apply deep neural networks to visual reasoning, in particular to the same-different task (i.e., classifying an image with two objects as an example of the categories "same" or "different"). This previous research found that, in contrast with earlier machine learning models, convolutional neural networks (CNNs) can learn to classify images with abstract shapes as same or different (e.g., Messina, Amato, Carrara, Gennaro, & Falchi, 2021; Funke et al., 2021). However, Puebla and Bowers

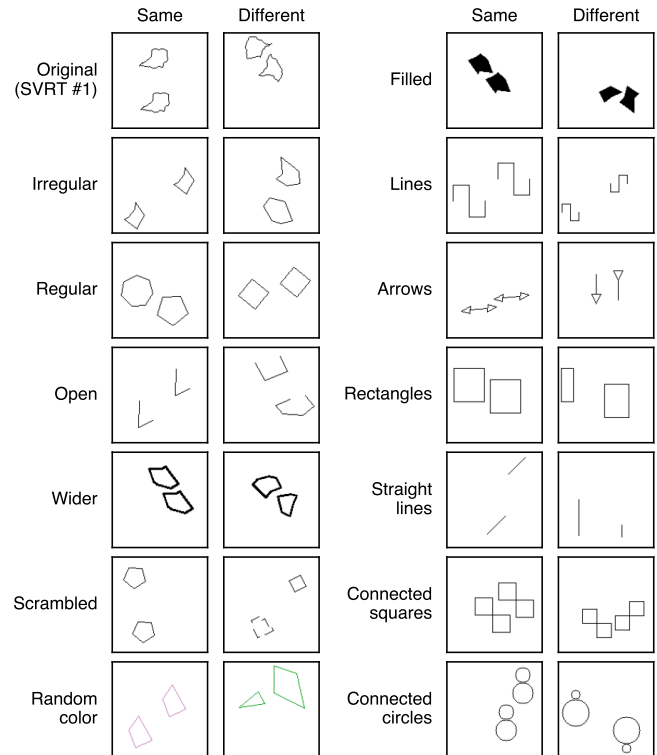


Figure 1: Example images of all datasets tested.

(2022) showed that the representations learned by these models are highly specific: when trained in task #1 of the synthetic visual reasoning test (SVRT, Fleuret et al., 2011, see Original condition Figure 1), CNNs tended to classify correctly images that were superficially similar to the ones they were trained on (e.g., Irregular or Regular conditions in Figure 1) and misclassify images that illustrated the same relation but were more superficially dissimilar (e.g., Lines or Arrows conditions). In the meantime, a number of new DNNs have introduced architectural innovations targeted at achieving relational visual reasoning. In this work, we test relational generalization of the same-different task on these models.

Models

ResNet50 (He, Zhang, Ren, & Sun, 2016) We included this model as a baseline deep CNN.



Slot Attention (Locatello et al., 2020) This model segregates objects in an image through a key-value attention mechanism that assigns different parts of the image to competing slots.

Recurrent vision transformer (RViT, Messina, Amato, Carrara, Gennaro, & Falchi, 2022) This model applies a standard vision transformer encoder recurrently, that is the model takes as input its own output, for a number of 4 processing steps.

Emergent symbol binding network (ESBN, Webb, Sinha, & Cohen, 2021) This model consist of a recurrent neural network augmented with a key-value external memory. This model aims to bind its memory values (direct representation of the input) and keys (inferred representation of the input's role in the sequence).

Guided Attention Model for (visual) Reasoning (GAMR, Vaishnav & Serre, 2023) This model is composed of three modules. An encoder builds a representation of the input. A recurrent controller guides an attention mechanism to select relevant object representations and write them into an external memory. A graph neural network module computes relations between the objects stored in the external memory.

Object-Centric Recurrent Attention (OCRA, Adeli, Ahn, & Zelinsky, 2022) This model consist of a recurrent encoder that controls an attention window that trades of the area it covers by its resolution. At the same time, a recurrent decoder control a write window that modifies a reconstruction output at each time step. The encoder feeds a two-layer capsule neural network that predicts a class label.

Methods

We trained 10 runs of all the models in task #1 of SVRT until reaching a validation accuracy of approximately 99%. This dataset consist of 28,000 128×128 RGB images. We tested all the models in all 14 datasets illustrated in Figure 1 (5,600 images per dataset), with the exception that OCRA was not tested on the Random colors dataset since this model takes only grey scale images as an input.

Results and discussion

As can be seen in Figure 1, all models achieved high accuracy in the test split of task #1 of SVRT except for the ESBN model, which performed at chance in all datasets. Further analysis showed that this model can learn the same-different task only when the objects (presented individually in a sequence of two images) are centered in the image, which severely questions Webb et al. (2021)'s conclusions regarding the relational reasoning capabilities of the model.

Furthermore, Slot Attention, GAMR and OCRA tended to show better generalization on the datasets that were harder for ResNet50. However, no single model achieved high levels of accuracy across all the test datasets, which is what would be expected if a model learned an abstract representation of the relations "same" and "different".




Figure 2: Mean accuracy on 10 runs of each model per condition. Error bars are standard errors of the mean.

Acknowledgments

The first author has received funding for this project from the National Center for Artificial Intelligence CENIA FB210017, Basal ANID.

The second author has received funding for this project from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 741134) for the second author.

References

- Adeli, H., Ahn, S., & Zelinsky, G. J. (2022). A brain-inspired object-based attention network for multi-object recognition and visual reasoning. *bioRxiv*, 2022–04.
- Biederman, I. (1987). Recognition-by-components: a theory of human image understanding. *Psychological review*, 94(2), 115.
- Fleuret, F., Li, T., Dubout, C., Wampller, E. K., Yantis, S., & Geman, D. (2011). Comparing machines and humans on a visual categorization test. *Proceedings of the National Academy of Sciences*, 108(43), 17621–17625.
- Funke, C. M., Borowski, J., Stosio, K., Brendel, W., Wallis, T. S., & Bethge, M. (2021). Five points to check when comparing visual perception in humans and machines. *Journal of Vision*, 21(3), 16–16.
- Gentner, D., Shao, R., Simms, N., & Hespos, S. (2021). Learning same and different relations: cross-species comparisons. *Current Opinion in Behavioral Sciences*, 37, 84–89.
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 770–778).
- Locatello, F., Weissenborn, D., Unterthiner, T., Mahendran, A., Heigold, G., Uszkoreit, J., . . . Kipf, T. (2020). Object-centric learning with slot attention. *Advances in Neural Information Processing Systems*, 33, 11525–11538.
- Messina, N., Amato, G., Carrara, F., Gennaro, C., & Falchi, F. (2021). Solving the same-different task with convolutional neural networks. *Pattern Recognition Letters*, 143, 75–80.
- Messina, N., Amato, G., Carrara, F., Gennaro, C., & Falchi, F. (2022). Recurrent vision transformer for solving visual reasoning problems. In *Image analysis and processing—iciap 2022: 21st international conference, lecce, italy, may 23–27, 2022, proceedings, part iii* (pp. 50–61).
- Puebla, G., & Bowers, J. S. (2022). Can deep convolutional neural networks support relational reasoning in the same-different task? *Journal of Vision*, 22(10), 11–11.
- Vaishnav, M., & Serre, T. (2023). GAMR: A guided attention model for (visual) reasoning. In *The eleventh international conference on learning representations*. Retrieved from <https://openreview.net/forum?id=iLMgk2IGNyv>
- Vo, M. L.-H., Boettcher, S. E., & Draschkow, D. (2019). Reading scenes: How scene grammar guides attention and aids perception in real-world environments. *Current opinion in psychology*, 29, 205–210.
- Webb, T. W., Sinha, I., & Cohen, J. (2021). Emergent symbols through binding in external memory. In *International conference on learning representations*. Retrieved from <https://openreview.net/forum?id=LSFCEb3GYU7>