

The role of disentanglement in generalisation

Milton L. Montero^{1, 2} Casimir J.H. Ludwig¹ Rui Ponte Costa²
Gaurav Malhotra¹ Jeffrey S. Bowers¹

¹School of Psychological Science

²Computational Neuroscience Unit, Department of Computer Science
University of Bristol, United Kingdom




May 3, 2021

- ▶ **Disentangled representations** encode separate factors of variation of the data in separate variables of latent representation.



- ▶ **Disentangled representations** encode separate factors of variation of the data in separate variables of latent representation.
- ▶ These “should capture the compositional structure of the world” (Duan et al., 2020).

- ▶ **Disentangled representations** encode separate factors of variation of the data in separate variables of latent representation.
- ▶ These “should capture the compositional structure of the world” (Duan et al., 2020).

Learning disentangled representations





Learning disentangled representations





Can disentangled representations support out-of-training-distribution generalisation?


Testing generalisation




Testing generalisation




Testing generalisation





Testing generalisation





Testing generalisation





Testing generalisation





Testing generalisation





Testing generalisation





Testing generalisation




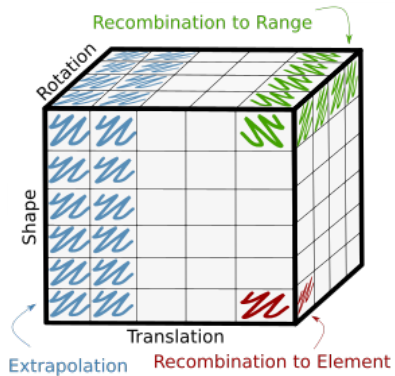
Testing generalisation





Testing generalisation



Testing generalisation



Results: dSprites



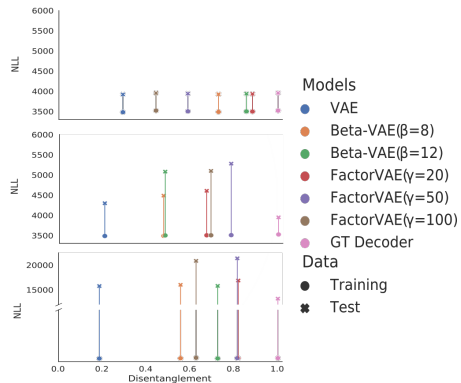

Models

- VAE
- Beta-VAE($\beta=8$)
- Beta-VAE($\beta=12$)
- FactorVAE($\gamma=20$)
- FactorVAE($\gamma=50$)
- FactorVAE($\gamma=100$)
- GT Decoder

Data


- Training
- ✱ Test

Results: 3DShapes

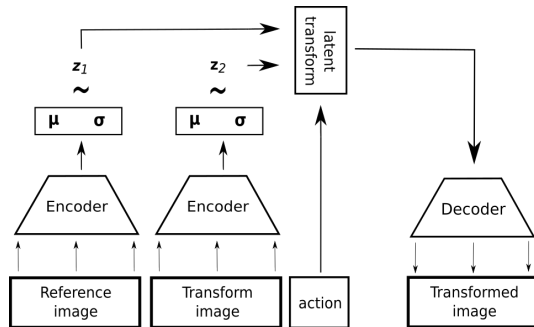



Can we push the models to learn compositional representations by changing the task?

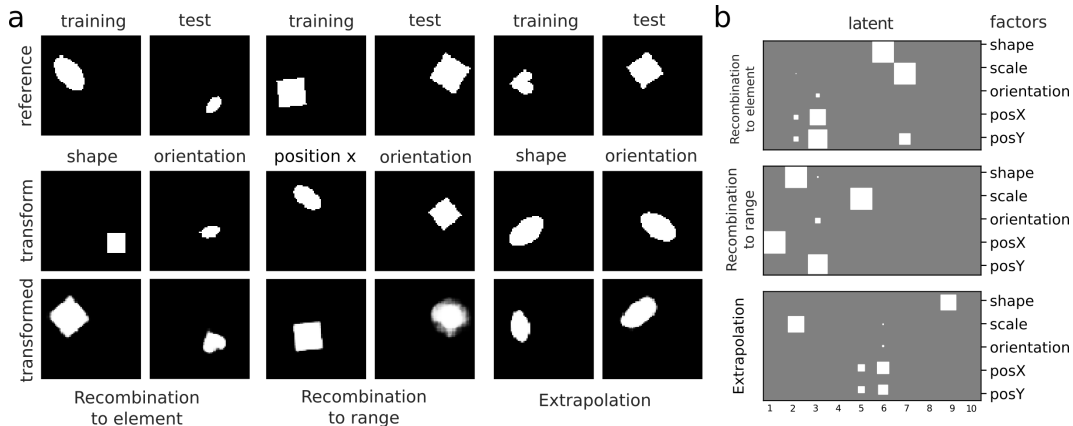
Composition task and model



Composition task and model



Compositon task: results



1. There are several ways that we can define generalisation outside the training space.

Conclusions

1. There are several ways that we can define generalisation outside the training space.
2. Disentanglement, does not seem to allow the models to generalise to these settings any more than non-entangled ones.

Conclusions

1. There are several ways that we can define generalisation outside the training space.
2. Disentanglement, does not seem to allow the models to generalise to these settings any more than non-entangled ones.
3. We can obtain high disentanglement by training a model to manipulate the world.

1. There are several ways that we can define generalisation outside the training space.
2. Disentanglement, does not seem to allow the models to generalise to these settings any more than non-entangled ones.
3. We can obtain high disentanglement by training a model to manipulate the world.
4. This works adds to other contributions that highlight the need for better definitions, metrics, models and training settings .

Many thanks to my co-authors:

Jeff Bowers, Gaurav Malhotra, Rui Ponte Costa & Casimir Ludwig



European Research Council
Established by the European Commission