"Re-defining disentanglement in Representation Learning for artificial agents."

Hugo Caselles-Dupré, Michael Garcia-Ortiz, David Filliat

Flowers Laboratory, INRIA & ENSTA ParisTech

AI Lab, Softbank Robotics Europe

Plan

1) Introduction about my research

2) Dive into one paper

3) Discussion & future work

My PhD

CIFRE PhD thesis shared between SBRE and Flowers@ENSTA

Former SBRE supervisor: Michael Garcia-Ortiz

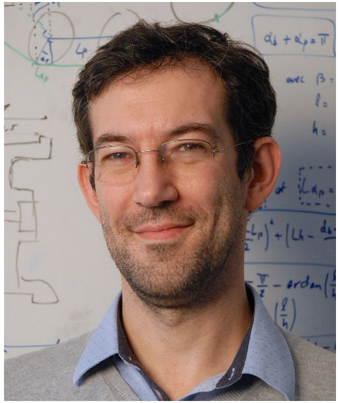

ENSTA Paris supervisor: David Filliat

Big picture of the research area

Motivation: Create intelligent systems to solve tasks automatically.

Problem: How do you construct such systems?

Problem division

Problem: How do you construct such systems?

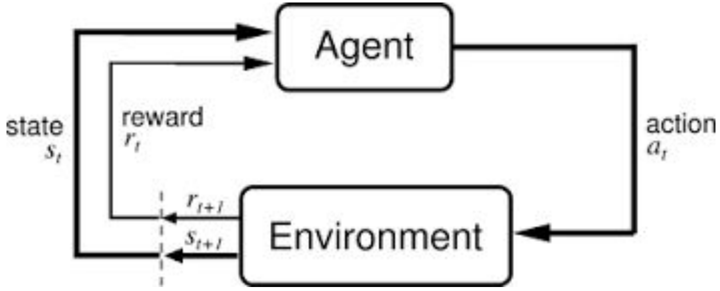

Sub-problem: How do you construct agents that can solve a wide range of Reinforcement Learning (RL) tasks ?

PhD goal

> "Construct a continual perception system that can be used to solve Reinforcement Learning tasks."

PhD goal

"Construct a continual perception system that can be used to solve Reinforcement Learning tasks."

Lifelong, continual, incremental, never-ending …

PhD goal

“Construct a continual perception system that can be used to solve Reinforcement Learning tasks.”

Analyzing, and extracting information from the sensorimotor flux

Approach: State Representation Learning

Goal: build an agent that understand and abstract how the world works.

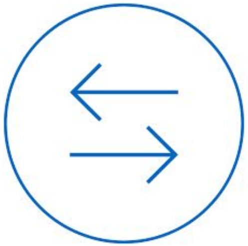

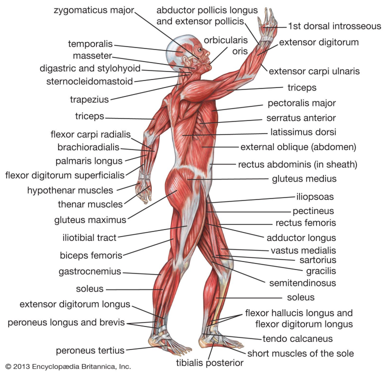

Sensors

Motors

How can humans process the multi-modal high-dimensional sensorimotor flux and abstract simple concepts like space, objects and physics?

We then use it to perform tasks.

PhD goal

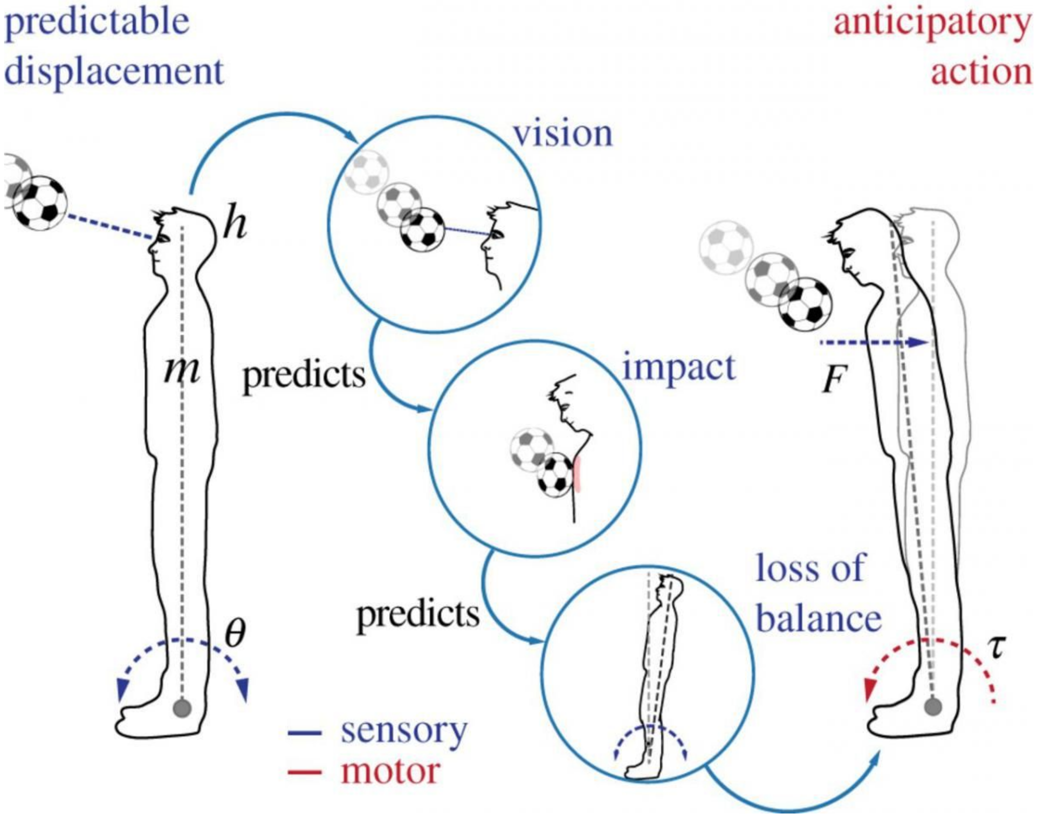

predictable displacement
anticipatory action
vision
h
m
predicts
impact
F
predicts
loss of balance
θ
τ
— sensory
— motor

Common approaches

Coupled

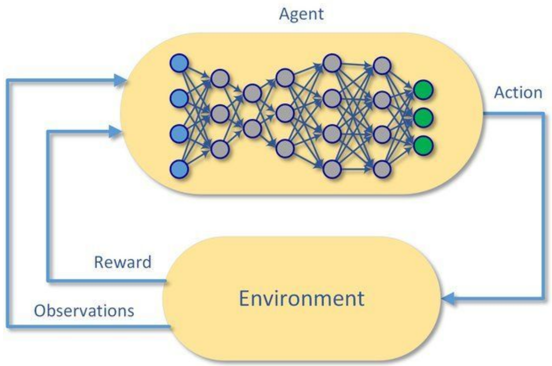

or

Decoupled

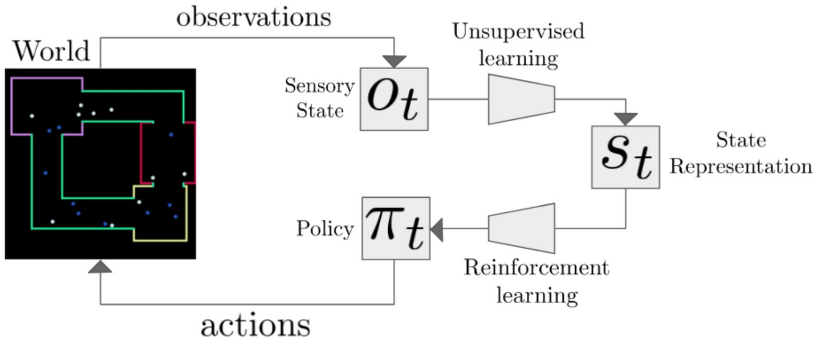

Problems with common approaches

Is current SRL abstracting any useful concept?

Task specific

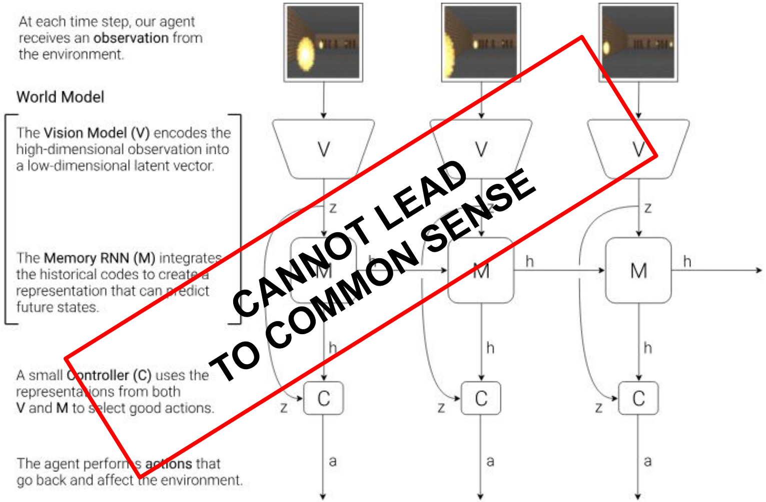

Where are the motors in the SRL?

> This does not help mastering the sensorimotor flux of data.

PhD goal

“Construct a continual sensorimotor control system that can be used to solve Reinforcement Learning tasks.”

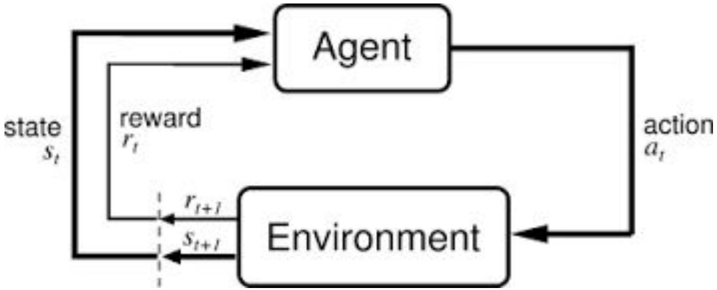

Related work

Mostly SRL, RL and Robotics:

- Reinforcement Learning
- Developmental Robotics
- Continual Learning
- Curiosity
- Deep Learning
- Model-based RL
- Transfer learning in RL
- Lifelong learning
- ...

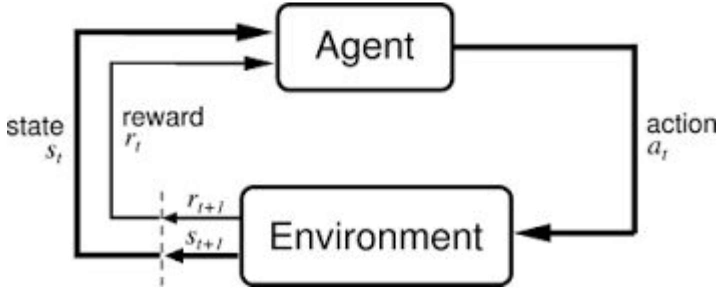

Works

- ## State Representation Learning

H. Caselles-Dupré, M. Garcia-Ortiz, D. Filliat, “On the Sensory Commutativity of Action Sequences for Embodied Agents”, 2020 arXiv preprint arXiv:2002.05630.

H. Caselles-Dupré, M. Garcia-Ortiz, D. Filliat, "Symmetry-Based Disentangled Representation Learning requires Interaction with Environments", Neural Information Processing Systems (NeurIPS) 2019, December 8-14, 2019, Vancouver, Canada.

H. Caselles-Dupré, M. Garcia-Ortiz, D. Filliat, "S-TRIGGER: Continual State Representation Learning via Self-Triggered Generative Replay", Workshop on Continual Learning, Neural Information Processing Systems (NIPS) 2018, December 2–9, 2018, Montréal, Canada.

- ## Continual Learning

T. Lesort*, H. Caselles-Dupré*, M. Garcia-Ortiz, J-F. Goudou, D. Filliat, "Generative Models from the perspective of Continual Learning", Workshop on Continual Learning, Neural Information Processing Systems (NIPS) 2018, December 2–9, 2018, Montréal, Canada.

R. Traoré*, H. Caselles-Dupré*, T Lesort*, T. Sun, N. Díaz-Rodríguez, D. Filliat, "Continual Reinforcement Learning deployed in Real-life using Policy Distillation and Sim2Real Transfer", Workshop on “Multi-Task and Lifelong Reinforcement Learning”, International Conference on Machine Learning (ICML) 2019, June 9–15, 2019, Long Beach, USA.

- ## Simulator

H. Caselles-Dupré, L. Annabi, O. Hagen, M. Garcia-Ortiz, D. Filliat, "Flatland: a Lightweight First-Person 2-D Environment for Reinforcement Learning", Workshop on Continual Unsupervised Sensorimotor Learning, ICDL-EpiRob 2018, September 16–20, 2018, Tokyo, Japan.

“Symmetry-Based Disentangled Representation Learning requires Interaction with Environments”

Extension of : I. Higgins et al., “Towards a Definition of Disentangled Representations”, 2018.

Neural Information Processing Systems 2019

Motivation of Higgins et al.

- Representation learning: vectorial representations of data
- Disentanglement: isolate factors of variation (latent variable <-> high-level features)

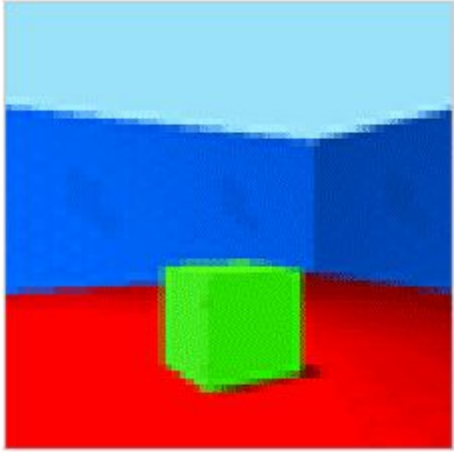

- Problem: disentanglement needs a proper definition, otherwise confusing

I. Higgins et al., [“Towards a Definition of Disentangled Representations”](), 2018.

Solution: Symmetries

- Which transformations change some properties of the underlying world state, while leaving all other properties invariant?

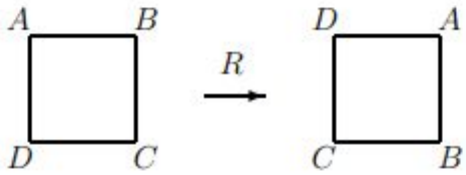

- Successful approach in physics: symmetries have revolutionized the understanding of the world structure.

I. Higgins et al., [“Towards a Definition of Disentangled Representations”](), 2018.

Why group theory?

- Group theory: the symmetry group of an object (image, signal, etc.) is the group of all transformations under which the object is invariant.

- Ex, scene understanding: transformations include translations, rotations and changes in object colour.

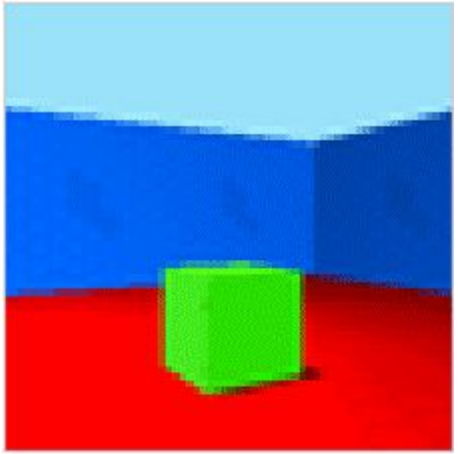

[3Dshapes dataset]

I. Higgins et al., "Towards a Definition of Disentangled Representations", 2018.

Disentangled representation definition I

- World: observations O of world states W.
- Symmetries: translations.

O = images, W = (x,y), G = translations

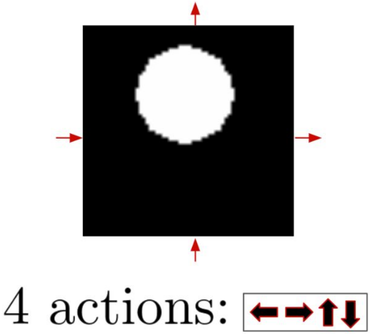

- Goal: disentangled encoder f (projects W on Z)

$$\begin{tikzcd} W \arrow[r] \arrow[rr, bend left, "f"] & O \arrow[r, "h"] & Z \end{tikzcd}$$

I. Higgins et al., *"Towards a Definition of Disentangled Representations"*, 2018.

Disentangled representation definition II

- G applies an action on W: $g_x \cdot_W w = ((x+1) \mod N, y)$

  → This action should be the same on Z. $g \cdot_Z f(w) = f(g \cdot_W w)$

Defines Symmetry-Based representations

$$\begin{tikzcd} G \times W \arrow[r, "\cdot_W"] \arrow[d, "id_G \times f"'] & W \arrow[d, "f"] \\ G \times Z \arrow[r, "\cdot_Z"'] & Z \end{tikzcd}$$

I. Higgins et al., ["Towards a Definition of Disentangled Representations"](), 2018.

Disentangled representation definition III

- Hypothesis: G can be decomposed into sub-groups that do not affect each other.

  → A sub-group only acts on a subspace of the representation.

Defines disentanglement

$$G = G1 \times G2$$

G = translations
G1 = translations along x axis
G2 = translations along y axis

I. Higgins et al., "Towards a Definition of Disentangled Representations", 2018.

In practice?

- I. Higgins et al. use CCI-VAE with still samples.

- Manage to learn: f(w)=f((x,y))=(λ1*x, λ2*y).

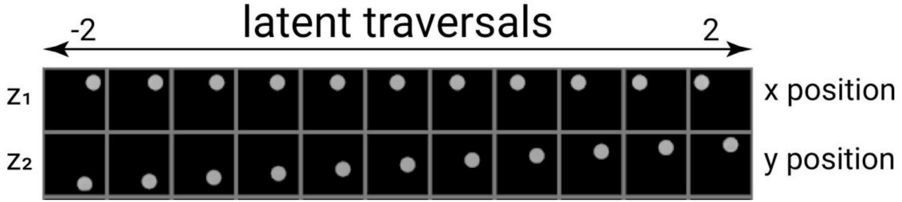

- But what about the group action? $g \cdot_Z f(w) = f(g \cdot_W w)$

Same problem as with common approaches

Task specific

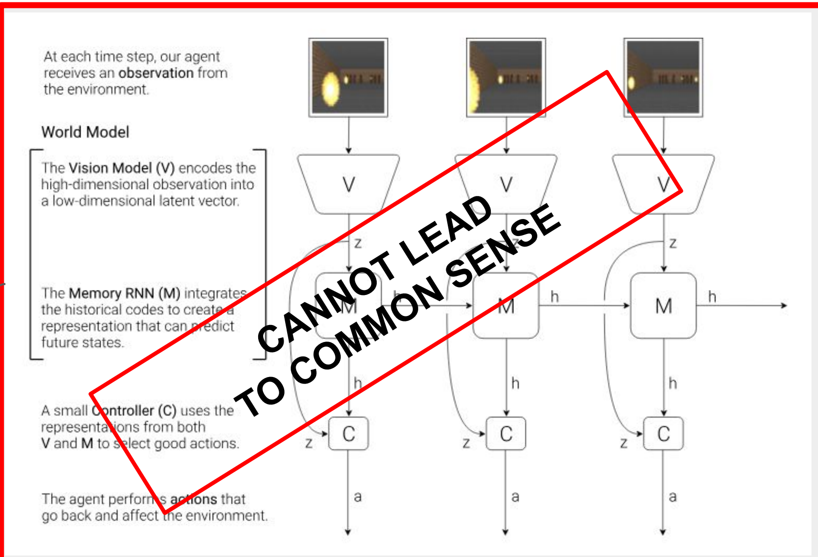

Where are the motors in the SRL?

We are in this case !

Main result: intuition

SBDRL requires transitions instead of still samples.

The specific order in which the transitions happen, i.e. the transition function of the world or “physics” is not learnable using only still samples.

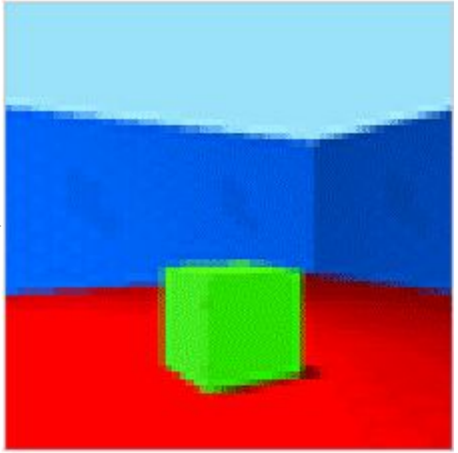

$$\begin{tikzcd} G \times W \arrow[r, "\cdot W"] \arrow[d, "id_G \times f"'] & W \arrow[d, "f"] \\ G \times Z \arrow[r, "\cdot Z"'] & Z \end{tikzcd}$$

Main result: formulation

SBDRL requires transitions instead of still samples.

Theorem 1. *Suppose we have a SB representation* $(f, \cdot_Z)$ *of a world* $\mathcal{W}_0 = (W = (w_1, .., w_m) \in \mathbb{R}^{m \times d}, \cdot_{\mathcal{W}_0})$ *w.r.t to* $G = G_1 \times ... \times G_n$ *using a training set* $\mathcal{T}$ *of unordered observations of* $\mathcal{W}_0$*. Let* W_k *be the set of possible values for the* k^{th} *dimension of* $w \in W$*.*
Then:

1. *There exists at least* $k_{W,G} = n[(\min_k(card(W_k))!] - 1$ *worlds* $(\mathcal{W}_1, .., \mathcal{W}_{k_{W,G}})$ *equipped with the same world states* $\mathcal{W}_i = (w_1, .., w_m)$ *and symmetries* G*, but different group actions* $\cdot_{\mathcal{W}_i}$*.*
2. *For these worlds,* $(f, \cdot_Z)$ *is not a SB representation.*
3. *These worlds can produce exactly the same training set* $\mathcal{T}$ *of still images.*

Practical options

How to learn SB-disentangled representations in practice?

2 options arise

Option 1: a la World Models

Definition requires:

$$g \cdot_Z f(w) = f(g \cdot_W w)$$

$$\begin{tikzcd}
G \times W \arrow[r, "\cdot_W"] \arrow[d, "id_G \times f"'] & W \arrow[d, "f"] \\
G \times Z \arrow[r, "\cdot_Z"'] & Z
\end{tikzcd}$$

Two-steps approach:

Learn state representation f → Learn group action of G on f

Option 2: End-to-end learning

Learn state representation f and group action at the same time

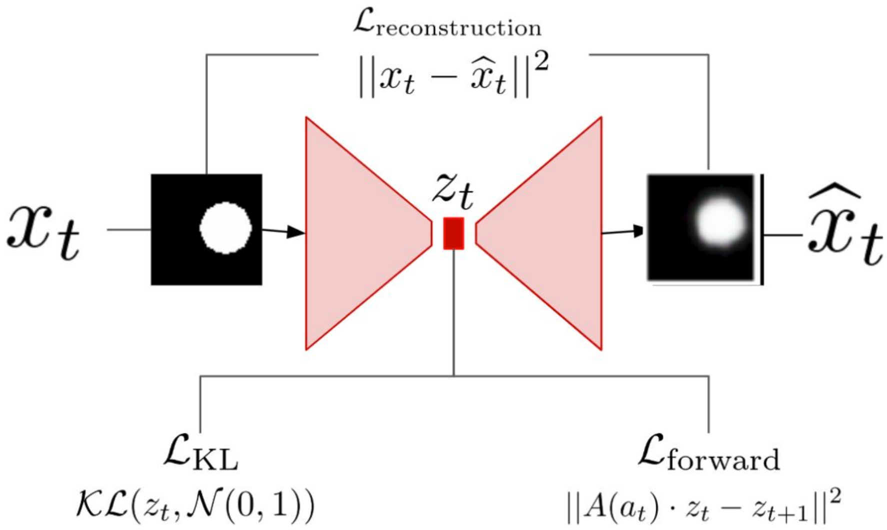

Results

Both approaches are successful empirically.

Option 2 makes more sense as latent space has to organize specifically.

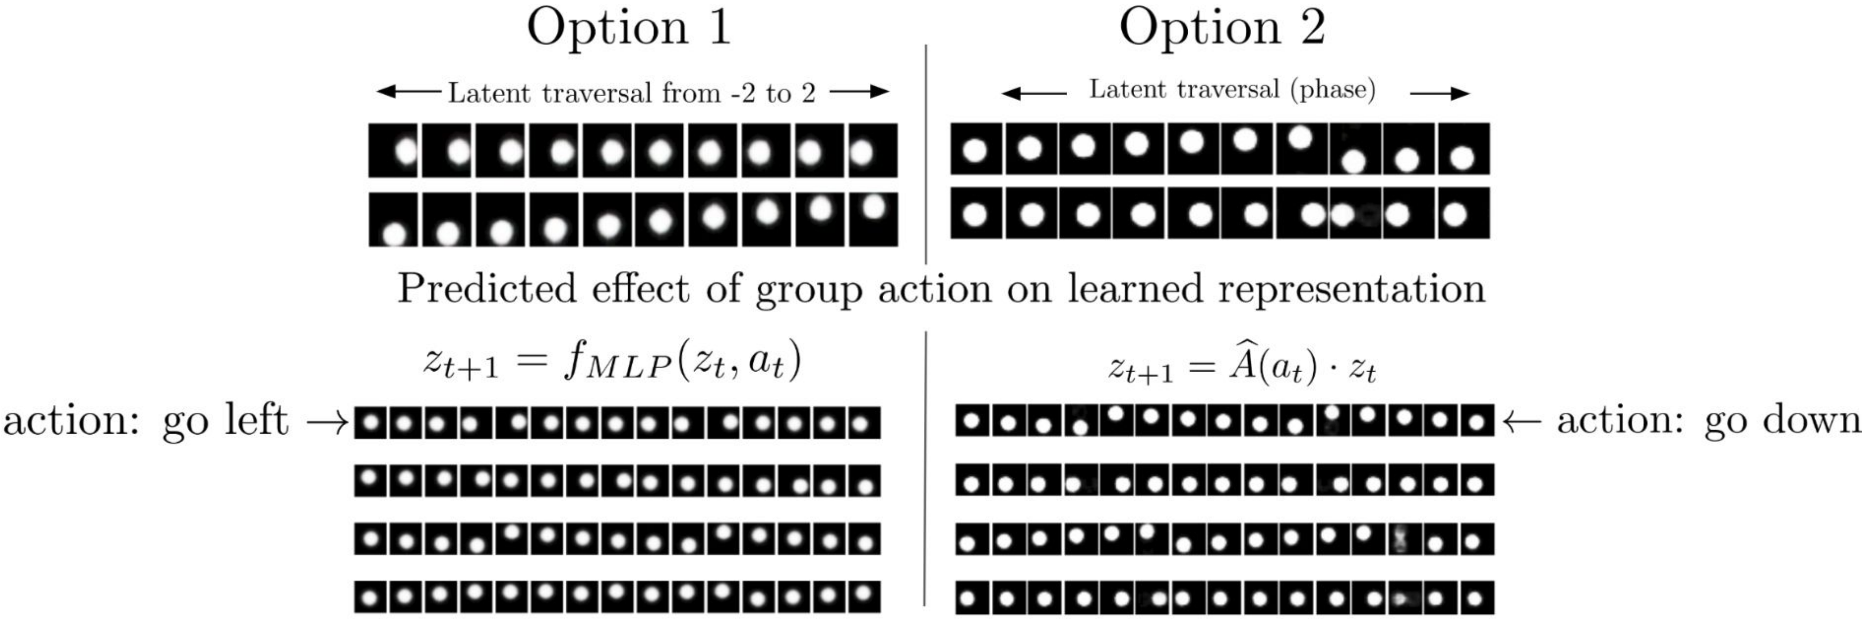

Usefulness of learned representations

Learning downstream tasks should be easier.

Preliminary experiments: Inverse model $a_t = f_{inverse}(s_t, s_{t+1})$

1K transitions samples

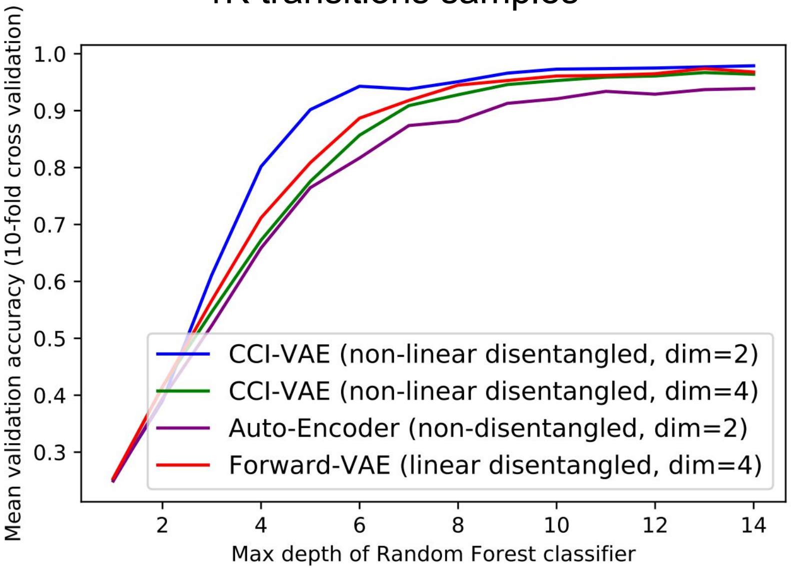

10K transitions samples

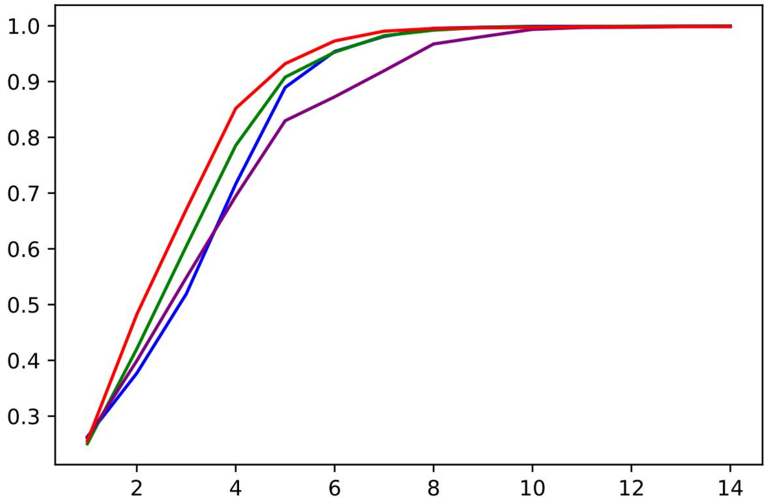

Conclusion

- Formal definition of disentanglement is needed.

- SBDRL: Disentanglement is defined w.r.t a decomposition of the symmetry group of the world.

Learning a SB-disentangled representation requires transitions instead of still samples.

Code and paper

Code (Colab)

arXiv.org Arxiv

Open question

How to learn SBD representations in more complex environments?

- Actions might be associated to a symmetry.

- Local symmetries.

Thanks!

[Code (Colab)]

arXiv.org [Arxiv]

Contact :

casellesdupre.hugo@gmail.com @Caselles Site