On representations, symmetries, groups, variational principles

... and all that



Overview

- Problem of **disentangled representations**.
- How **symmetries** have been used to solve it.
- Formalised as **Groups**.
- Hint to its potential relation with **variational principles** (Free Energy Principle).
- Hint to richer representations using **Groupoids**.

Disentangled Representation Learning Caselles-Dupré et al.



- The world and its representations are too complex.
- We need to find low-dimensional representations of the world for which the underlying structure is separated into disjoint parts (i.e., disentangled) reflecting its compositional nature.
- Two issues:
 - reduce computational complexity (as in the input size) and
 - transfer knowledge by generalizing over "similar" representations, that is, representations that are composed of and preserved (up to isomorphism?) over the same parts and that, in turn, support prediction.

Disentangled Representation Learning Caselles-Dupré et al.

- Symmetry transformations change some properties of the underlying world state, while leaving all other properties invariant --gives exploitable structure.
- Formalised using **group theory**. *Groups* are composed of such transformations. *Group actions* are the effect of the transformations on the state of the world and representation.



What is a group (algebraic definition)

- A group is a set, G, together with a binary operation that combines any two elements a and b to form another element, denoted a • b.
- (*G*, •) must satisfy four *group axioms*:
 - Closure: For all a, b in G, the result of the operation, a b, is also in G.
 - Associativity: For all a, b and c in $G, (a \bullet b) \bullet c = a \bullet (b \bullet c)$.
 - Identity element: There exists an element *e* in *G* such that, for every element *a* in *G*, the equation *e a* = *a e* = *a* holds.
 - Inverse element: For each a in G, there exists an element b in G, such that a b = b a = e.
- Interestingly, $a \bullet b = b \bullet a$ may not always be true.
- Example: The set of integers together with the addition operation. But not the set {0,1,2} under addition or the set of integers under subtraction.

Symmetry group

 The symmetry group (of a geometric object) is the group of all transformations under which the object is invariant, endowed with the group operation of composition. Such a transformation is an invertible mapping which takes the object to itself, and which preserves all the relevant structure of the object. Examples: rotations, reflections, translations (in space and time).



Some history ...



- The study of groups flourished in the XIX century, originally to solve algebraic equations (E. Galois).
- In geometry, F. Klein proposed the *Erlangen Program* to classify various geometries (Euclidean, affine, and projective) with respect to geometrical properties that are left invariant under rotations and reflections.
- It was also in Göttingen where E. Noether proved the connection between symmetries and conservation laws (e.g., total energy is conserved under translation in time).
- In Physics, the special theory of relativity unified seemingly contradictory mechanical and electromagnetic phenomena of the hand of Lorentz groups; and the general theory of relativity explained gravity under the group of all diffeomorphisms of a space-time. The Standard Model classifies all elementary particles and their interactions according to their flavour, charge and colour symmetries, and, in so doing, unifies electromagnetism, QED and QCD and explains electroweak interactions (H. Weyl).

Why symmetries?

- We attribute symmetry properties to theories and laws (symmetry principles). It is useful to derive the laws of nature and to test their validity by means of the laws of invariance, rather than to derive the laws of invariance from what we believe to be the laws of nature.
- We may derive specific consequences with regard to particular phenomena on the basis of their symmetry properties (symmetry arguments). P. Curie postulated a necessary condition for a given phenomenon to happen, namely, that it is compatible with the symmetry conditions established by a principle.



Back to representations

- Let W be a set of **world-states** $W = (w_1, ..., w_m)$.
- There is a generative process b : W → O leading from world-states to observations (these could be pixel, retinal, or any other potentially multi-sensory observations), and an inference process h : O → Z leading from observations to an agent's representations.
- We consider the composition $f: W \rightarrow Z$ ($f = h \bullet b$).

Interlude: Free Energy Principle (K. Friston)



- Biological systems resist tendency to disorder by maintaining their states in the face of a changing environment (homeostasis).
- Tradeoff between value (expected reward, expected utility) and its complement, surprise (prediction error, expected cost).
- States must have low entropy (high probability that the system is in any of a small number of states). Entropy is the long-term average of surprise.
- Biological agents must minimize surprise.
- Free energy is an upper bound on surprise. Thus, agents must minimize free energy: Free Energy Principle.

Principle of Least Action





- Of all possible paths, why a parabola?
- It *minimizes* the "action", i.e., the integral of kinetic minus potential.
- Gives equations of motion.
- Uses calculus of variations (that is, small changes to find the minimum).
- Encapsulates conservation laws and symmetries.



10

... free-energy is basically prediction error



action and perception to suppress prediction errors and minimise surprise

How do we minimize Free Energy? The Bayesian Brain

- A probabilistic model that can generate predictions, against which sensory samples are tested to update beliefs about their causes. This generative model is decomposed into a likelihood (the probability of sensory data, given their causes) and a prior (the a priori probability of those causes).
- Perception then becomes the inference process of inverting the likelihood model (mapping from causes to sensations) to access the posterior probability of the causes, given sensory data (mapping from sensations to causes).
- Rings a bell?

Back to disentanglement

- A group G of symmetries acting on W via a group action $\varphi : G \times W \rightarrow W$.
- We would like to find a corresponding group action $\psi : G \ge Z \Rightarrow Z$ so that the symmetry structure of W is reflected in Z (Condition 1).
- We also want the group action ψ to be **disentangled**, which means that applying G_i to Z leaves all sub-spaces of Z unchanged but the one corresponding to the transformation G_i (Condition 2).

- Formally, the representation Z is disentangled with respect to the decomposition $G = G_1 \times ... \times G_n$ if:
- 1. There is a group action $\psi : G \times Z \rightarrow Z$.
- 2. The map $f: W \rightarrow Z$ is equivariant between the group actions on W and Z:



3. There is a decomposition $Z = Z_1 \times ... \times Z_n$ such that each Z_i is fixed by the action of all G_j , $j \neq i$ and affected only by G_i .

14

Dísentanglement

Problem in paradise

- How can one learn a disentangled representation? This task involves knowledge about how the group action affects *Z*. The group action is defined to be the effect of symmetries on the representation.
- These symmetries can be **translations**, **rotations**, **time translations**, etc. In AI, we would design an algorithm that learns from examples. We thus need, in practice, a way to **apply these transformations** on observations of the world $(o_t)_{t=1...n}$ and observe the result $(g \psi o_t = o_{t+1})_{t=1...n}$.

Actions ...

- Analogy between the effect of a symmetry g (by the group action φ) on the environment (o_1 , g, $g \varphi o_1 = o_2$), and a transition that uses the dynamics f of the environment (o_t , a_t , $f(o_t$, a_t) = o_{t+1}). It allows us to consider a more realistic scenario where we have an agent in an environment, and we can apply the group actions to this agent. In our analogy we simply say that $o_1 = o_t$, $o_2 = o_{t+1}$ and $a_t = g$ and $\varphi = f$.
- So, how to discover the symmetries of the world? Use active perception or causal manipulations of the world to empirically determine them.

Back to FEP for a second

- Agents can suppress free energy by changing the two things it depends on: they can change sensory input by acting on the world or they can change their recognition density by changing their internal states.
- Action can reduce free energy only by increasing accuracy (that is, selectively sampling data that are predicted).
- Conversely, optimizing brain states makes the representation an approximate conditional density on the causes of sensory input. This enables action to avoid surprising sensory input.

Crisis ... what crisis?

• Choose actions that minimize surprise (FEP)!

- That is, follow a variational principle.
- Such principle will, in turn, be formalized as symmetry groups.
- Importantly, it will constraint the the agent's policy (sequence of actions, aka trajectory or path), that is, the dynamics of the system (its "law of motion").





Ma non troppo

• FEP is not the only variational principle of cognition (e.g., Betti and Gori).



- At the end of the day, FEP also relies in priors.
- There is no evidence of FEP **nor** that the brain optimizes a given function **nor** that the brain operates as a Bayesian inference engine.
- Actions are not strictly necessary to learn. Associative learning, Pavlovian (aka classical) conditioning in particular, occurs any time two events (stimuli) are presented "together".
- There is plenty of evidence for associative learning.

Stop the press!

- Quessard, Barrett and Clements claim that they have a come up with a framework in which agents learn the underlying group structure of environments and disentangled representations without any prior knowledge of the symmetry group.
 - They assume a "large" symmetry group, SO(n), and define a metric on disentanglement --acting on a minimum dimension.
 - They assume a certain structure (compositional rotation) for the group actions.
 - Does this qualify as "no prior knowledge"?
- Nice thing: they extend previous work to **continuous** symmetries. Pfau et al., have also recently proposed to work with Lie Groups, but can only be learned **if the true metric of the manifold is known**.

In summary

- Symmetries have been proposed as a tool to establish disentangled representations. They reflect structure (reduce dimensionality) and may be instrumental for generalization and transfer learning.
- The world shows symmetries.
- Which,, allegedly, we learn by interacting with it.
- Useful "isomorphism" between "action" (as an agent) and group actions (transformations).
- Groups formalize the notion of symmetry.
- Symmetries (and groups) embed variational principles.
- The Free Energy Principle is one of such principles which proposes to minimize surprise by action and prediction.





1	2	3	4
5	6	7	8
9	10	11	12
13	14	15	

- Mathematically, the fifteen puzzle is similar to Rubik's cube: the goal is to arrive to a certain position by performing the correct sequence of moves. However, unlike the Rubik's cube, the available moves depend on the current position.
 For example, if the blank square is in the lower-right corner, it can only be moved left or up.
- Unlike the transformations of the Rubik's cube, transformations of the fifteen puzzle cannot always be composed. Specifically, you can only compose two transformations if the ending position of the first is the same as the starting position of the second.
- The former is a group, the latter a groupoid.

Groupoid

Algebraically ...

- A group is a special case of a groupoid.
- Groups are groupoids that have only one object. A groupoid is a group with many objects (positions in the example).



Good or bad?







- A sphere has (lots of) symmetry.
- A bowling ball may not show total symmetries, yet it has obvious **partial** symmetries.
- Without the notion of groupoid, such symmetries are lost.
- Besides, in real systems, symmetries are formed and broken (new symmetries emerge) –which cannot be explained using groups.
- This also affects the evolution in the representations of objects.

Let's put it another way





Yves Klein, 1KB 191, 1962

Raffaello Sanzío, Scuola dí Atene, 1509-11

One more intuition

• Groupoids describe reversible processes which

may traverse a number of states.



- One approach to capturing the topology of the European road system is to list the journeys one can make beginning and ending in Bilbao.
- However, it might appear a little strange to privilege Bilbao and the act of staying put there. Each city might be thought to deserve equal treatment.
- Moreover, if you want to know about trips from Paris to Rome, it would seem perverse to have to sift through the set of round trips from Bilbao which pass through Paris and then Rome.
- More reasonable, then, to list all trips between any pair of cities, where ordered pairs of trips can be composed, if and only if the destination of the first trip matches the starting point of the second.

Groupoid (algebraic structure and category)

- A groupoid is a set, G, together with and a partial function
 (it is not necessarily defined for all elements of G, that is, is not a binary operation).
- A groupoid is a **category** in which every morphism is an isomorphism, i.e. invertible.
- The collapse of a groupoid into a mere collection of groups loses some information because it is not (categorically) natural.





- Categories abstract away from objects and focus on relations (maps or morphisms). Objects are only understood in terms of their relations to other objects (categories give an interpretative context, semantics or type).
- They are hierarchical, in that relations between categories are formalized as structure-preserving functors, and between functors as natural transformations, etc. forming *n*-categories.
- Groupoids are, in essence, 2-categoríes –that is, groupoids form hierarchies (whereas groups don't).
- Maps in categories can also represent **processes** as dictated by dynamic laws (e.g., using string diagrams). Processes need also to preserve structure, that is, symmetries that embed variational principles.

29

Summary on groups and groupoids

- We may need richer structures that show partial symmetries among multiple objects, and that form hierarchies naturally.
- Groupoids seem to be a promising area to investigate (and their relation to group actions).
- They also encode symmetry breaking, which refers to creativity.

References

- Symmetry-Based Disentangled Representation Learning requires Interaction with Environments, H. Caselles-Dupré, M. Garcia-Ortiz and D. Filliat, <u>https://arxiv.org/pdf/1904.00243.pdf</u>
- On the Sensory Commutativity of Action Sequences for Embodied Agents, H. Caselles-Dupré, M. Garcia-Ortiz and D. Filliat, <u>https://arxiv.org/pdf/2002.05630.pdf</u>
- Learning Group Structure and Disentangled Representations of Dynamical Environments, R. Quessard, T.D. Barrett and W.R. Clements, <u>https://arxiv.org/abs/2002.06991</u>
- Disentangling by Subspace Diffusion, D. Pfau, I. Higgings, A. Botev and S. Racaniere, <u>https://arxiv.org/pdf/2006.12982.pdf</u>
- The free-energy principle: a rough guide to the brain?, K. Friston, https://www.fil.ion.ucl.ac.uk/~karl/The%20free-energy%20principle%20-%20a%20rough%20guide%20to%20the%20brain.pdf
- Least Action Principles and Well-Posed Learning Problems, A. Betti and M. Gori, https://arxiv.org/pdf/1907.02517.pdf