On the use of Group Equivariant Non-Expansive Operators for Topological Data Analysis and Machine Learning

Patrizio Frosini

Department of Computer Science, University of Pisa
patrizio.frosini@unipi.it

Logic, Categories, and Applications Seminar
May 16, 2025, Bologna

Outline

Some epistemological assumptions

Some basics on the theory of GENEOs

GENEOs and XAI

Assumption 1: Data are processed by observers

Data have no meaning without an observer to interpret them.

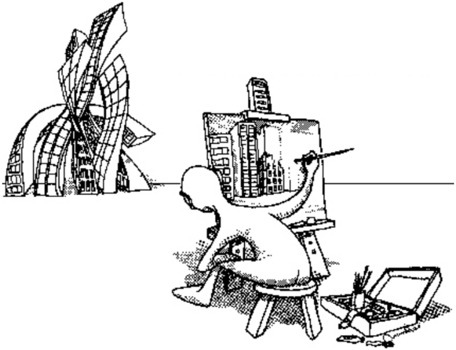

An observer is an agent that transforms data while preserving their symmetries.

Assumption 2: Observers are variables

Data interpretation strongly depends on the chosen observer.

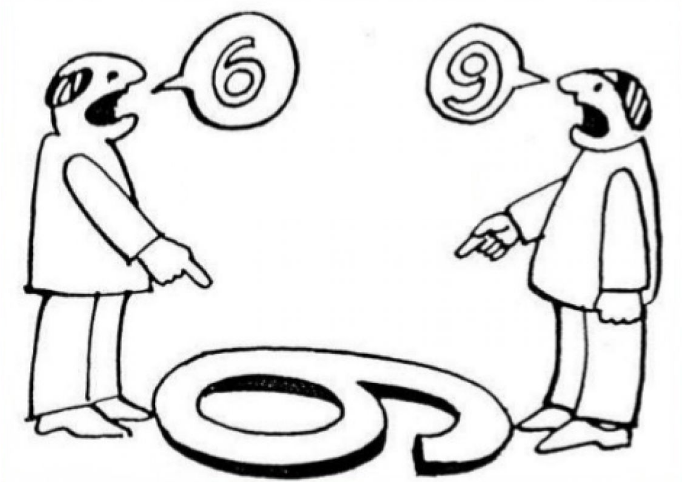

Assumption 3: Observers are important

We are rarely directly interested in the data, but rather in how observers react to their presence.

Consequently, we should focus more on the properties of the observers than on the properties of the data.

Assumption 4: There is no structure in the data

Generally speaking, data lack inherent structure. Instead, the structure of data reflects the observer's own structure.

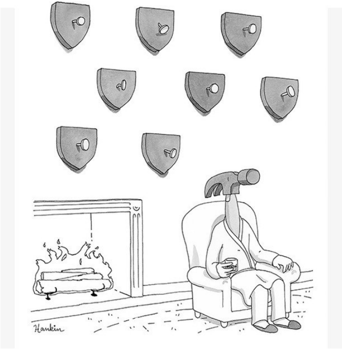

The shape is not in the data but in the eyes of the observer.

How can we translate these ideas into mathematics?

? ?

Perception spaces and GENEOs

PERCEPTION
SPACE

GROUP
EQUIVARIANT
NON-EXPANSIVE
OPERATOR
(GENEO)

Let's start by defining perception spaces

We recall that a pseudo-metric is just a metric d without the property $d(x_1, x_2) = 0 \implies x_1 = x_2$.

Definition

Let us consider

1. A nonempty set Φ endowed with a pseudo-metric D_Φ.
2. Let us denote by the symbol $*$ the left action of a group $(G, \circ)$ on Φ, and endow G with the pseudo-metric D_G defined by setting $D_G(g_1, g_2) := \sup_{\varphi \in \Phi} D_\Phi(g_1 * \varphi, g_2 * \varphi)$ for any $g_1, g_2 \in G$. We will also assume that the action of the group G on the metric space (Φ, D_Φ) is isometric, i.e., for every $\varphi_1, \varphi_2 \in \Phi$ and every $g \in G$, $D_\Phi(g * \varphi_1, g * \varphi_2) = D_\Phi(\varphi_1, \varphi_2)$.

We say that (Φ, G) is a perception space.

Perception spaces

The set Φ represents the data we may get from our measuring tools (functions, graphs, cloud of points,...). The group G represents the possible invariances of data the observer may be interested in.
For example, Φ can be a set of grey-level images represented as functions from $\mathbb{R}^2$ to $[0,1]$, while G can be the group of isometries of the real plane.
Another simple example can be given by the set of electrocardiograms represented as functions of the time variable, while G can be the group of time translations.
In any case, the following statement holds.

Proposition

$(G, \circ)$ *is a topological group and the action of* G *on* Φ *is continuous.*

GEOs and GENEOs

Definition

- Let (Φ, G), (Ψ, K) be two perception spaces. If a map $F : \Phi \to \Psi$ and a group homomorphism $T : G \to K$ are given, such that $F(g * \varphi) = T(g) * F(\varphi)$ for every $\varphi \in \Phi$, $g \in G$, we say that (F, T) is an (extended) *group equivariant operator* (GEO).
- If (F, T) is non-expansive (i.e. $D_\Psi(F(\varphi_1), F(\varphi_2)) \leq D_\Phi(\varphi_1, \varphi_2)$ for every $\varphi_1, \varphi_2 \in \Phi$, and $D_K(T(g_1), T(g_2)) \leq D_G(g_1, g_2)$ for every $g_1, g_2 \in G$), we say that (F, T) is an (extended) *group equivariant non-expansive operator* (GENEO).

An example of GENEO

When we blur an image by applying a **convolution** with a rotationally symmetric kernel whose mass is less than 1 in L^1, we are applying a GENEO (here, we are considering the **group of isometries**).

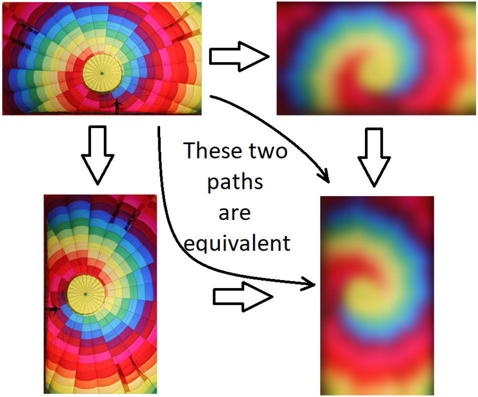

Here, the maximum distance between functions is used.

Another example of GENEO

When we compute the **convex hull** of a cloud of points, we are applying a GENEO (here, we are considering the **group of isometries**).

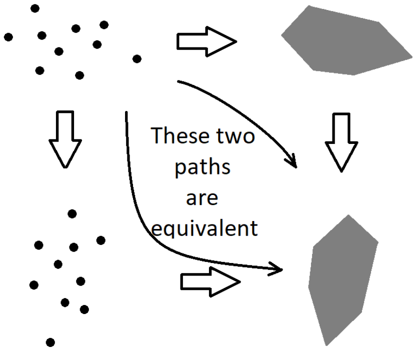

Here, the Hausdorff distance between compact sets is used.

Two key observations (1)

Our main goal is to build a robust geometric and compositional theory for approximating an ideal observer through GENEOs and GEOs.

$$\approx \boxed{\begin{array}{c} (\Psi, K) \\ \uparrow \\ (F, T) \\ | \\ (\Phi, G) \end{array}}$$

Two key observations (2)

GENEOs can be taken as inputs of higher-level GENEOs. Data obtained through measuring instruments can be seen as GENEOs of level 0. Therefore, hierarchies of GENEOs can be considered.

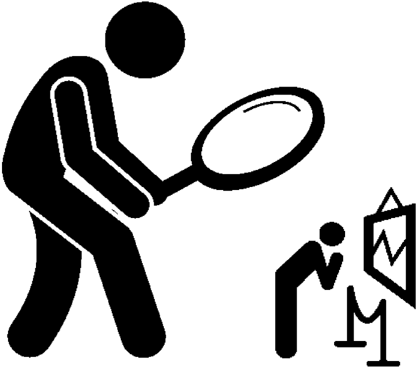

Construction of GENEOs

How can we build GENEOs?

When data are represented by real-valued functions, the space of GENEOs is closed under composition, computation of minimum and maximum, translation, direct product, and convex combination. (However there is much more than this...)

GENEOs are like LEGO bricks that can be combined together to form more complex GENEOs.

The main point in the approach based on GENEOs

In perspective, we are looking for a good compositional theory for building efficient and transparent networks of GENEOs.
Some preliminary experiments suggest that replacing neurons with GENEOs could make deep learning more transparent and interpretable and speed up the learning process.

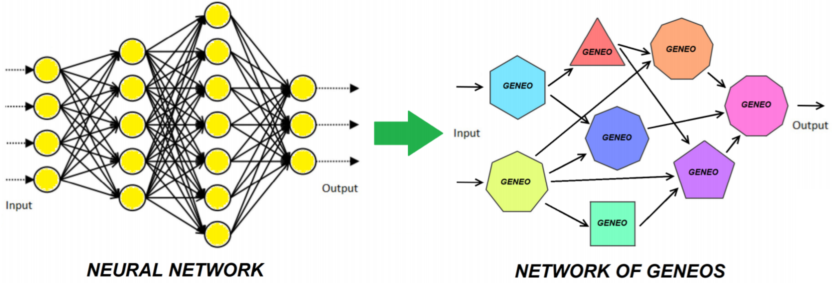

GENEOs and Machine Learning

If interested, you can find more details about the theory of GENEOs in these papers:

- M. G. Bergomi, P. Frosini, D. Giorgi, N. Quercioli,
  Towards a topological-geometrical theory of group equivariant non-expansive operators for data analysis and machine learning,
  Nature Machine Intelligence, vol. 1(9) (2019), 423–433.
  https://www.nature.com/articles/s42256-019-0087-3
- G. Bocchi, P. Frosini, M. Ferri,
  A novel approach to graph distinction through GENEOs and permutants,
  Scientific Reports, 15 (2025), 6259.
  https://www.nature.com/articles/s41598-025-90152-7

GENEOs and Machine Learning

For more details about the use of GENEOs in Machine Learning, you can have a look at this paper:

https://ems.press/journals/mag/articles/10389352

Research projects (I)

CNIT / WiLab - Huawei Joint Innovation Center (JIC)

Project on GENEOs for 6G

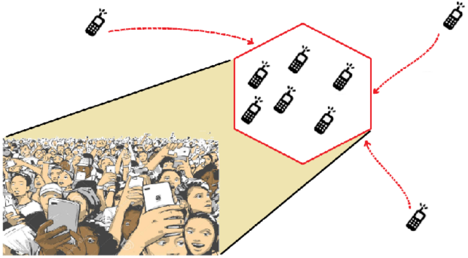

Research projects (II)

PANDORA

Horizon Europe (HORIZON)
Call: HORIZON-CL4-2023-HUMAN-01-CNECT
Project: 101135775-PANDORA
Funding: approximately 9 million euros.

Task 3.3 - Leveraging domain knowledge for explainable learning:
This task aims to investigate the use of domain knowledge in the development of explainable AI models. Tools like GENEOs for applications in TDA and ML and new theoretical methods of GENEOs for explainable AI will be used.

The project has received funding form the European Union's Horizon Europe Framework Programme (Horizon) under grant agreement No 101135775

Research projects (III)

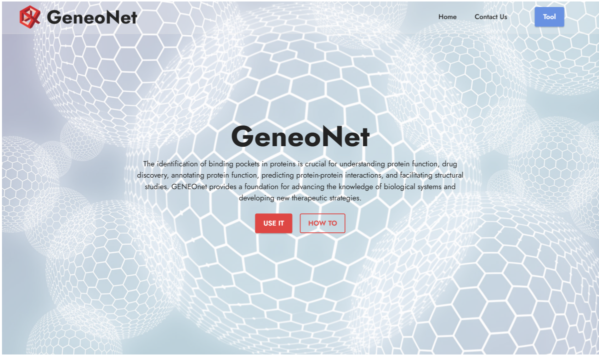

The GeneoNet webservice represents the outcome of our partnership with Italian Pharmaceutical Company Dompé Farmaceutici S.p.A.: https://geneonet.exscalate.eu/

Research projects (III)

GENEOs can be used for the detection of druggable protein pockets.

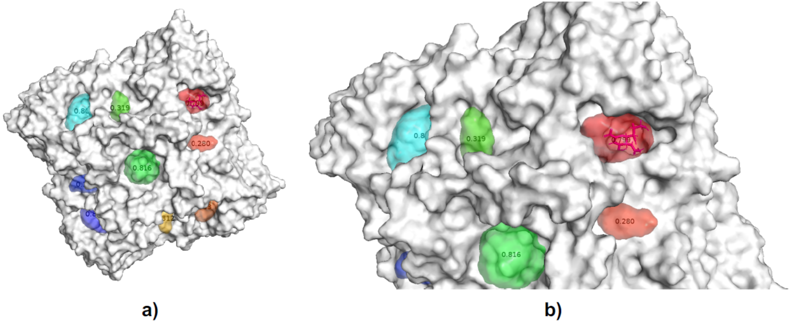

Model predictions for protein 2QWE. In Figure a) the global view of the prediction is shown, where different pockets are depicted in different colors and are labelled with their scores. In Figure b) the zoomed of the pocket containing the ligand is shown.

Research projects (III)

More information about GeneoNet is available in this paper:

G. Bocchi, P. Frosini, A. Micheletti, A. Pedretti, G. Palermo, D. Gadioli, C. Gratteri, F. Lunghini, A. R. Beccari, A. Fava, C. Talarico, *A geometric XAI approach to protein pocket detection*, The 2nd World Conference on eXplainable Artificial Intelligence, Valletta, Malta, July 17-19, 2024, , vol. 3793, 217-224 (2024). https://ceur-ws.org/Vol-3793/paper_28.pdf

Collaborators on this research

- Filippo Bonchi (University of Pisa)
- Jacopo Joy Colombini (Scuola Normale Superiore, Pisa)
- Francesco Giannini (Scuola Normale Superiore, Pisa)
- Fosca Giannotti (Scuola Normale Superiore, Pisa)
- Roberto Pellungrini (Scuola Normale Superiore, Pisa)

Reference:

- Jacopo Joy Colombini, Filippo Bonchi, Francesco Giannini, Fosca Giannotti, Roberto Pellungrini and Patrizio Frosini, *Mathematical Foundation of Interpretable Equivariant Surrogate Models*, World Conference on Explainable Artificial Intelligence (XAI-2025), Novel Post-hoc & Ante-hoc XAI Approaches, 09-11 July, 2025 - Istanbul, Turkey.

What is an explanation?

How can we mathematically and generally formalize the concept of an explanation provided by an agent, viewed as an operator?

Basic idea

Informal idea: We could say that the action of an agent A is explained by another agent B from the perspective of an agent C if:

1. C perceives A and B as similar to each other;
2. C perceives B as less complex than A.

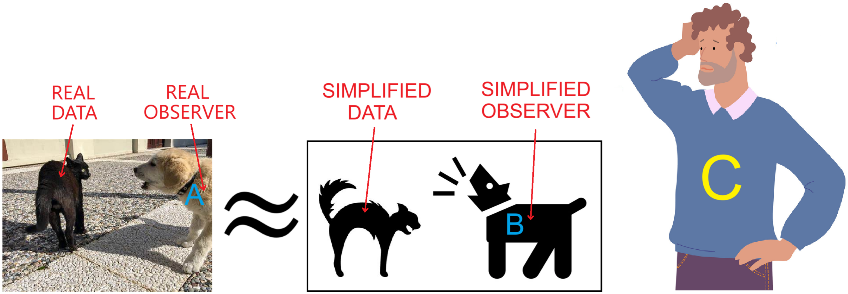

Basic idea

How can we transform our informal idea into a precise mathematical model?

Let us begin by formalizing property 1.

Informal idea: We could say that the action of an agent A is explained by another agent B from the perspective of an agent C if:

1. C perceives A and B as similar to each other;
2. C perceives B as less complex than A.

An extended pseudo-metric for *ALL* GEOs

We must introduce a pseudo-metric between GEOs that remains well-defined even when the GEOs operate on **different domains** and produce outputs in **distinct codomains**. This poses a non-trivial challenge.

$$\begin{array}{c} (\Psi_\alpha, K_\alpha) \\ \uparrow \\ (F_\alpha, T_\alpha) \\ \mid \\ (\Phi_\alpha, G_\alpha) \end{array}$$

What's the distance between these two GEOs?

$$\begin{array}{c} (\Psi_\beta, K_\beta) \\ \uparrow \\ (F_\beta, T_\beta) \\ \mid \\ (\Phi_\beta, G_\beta) \end{array}$$

In other words, what does it mean for two GEOs to behave approximately the same way?

Our main goal: observer approximation

The previous pseudo-metric is necessary to build a geometric theory for approximating an ideal observer through GENEOs and GEOs.

$$\approx \boxed{\begin{array}{c}(\Psi, K)\\ \uparrow \\ (F, T) \\ | \\ (\Phi, G)\end{array}}$$

An extended pseudo-metric for *ALL* GEOs

Informally speaking, two GEOs are considered similar if there exist two horizontal GENEOs that make this diagram "nearly commutative", with the same holding true in the opposite direction:

$$\begin{tikzcd}
(\Psi_\alpha, K_\alpha) & (\Psi_\beta, K_\beta) \arrow[l] \\
(\Phi_\alpha, G_\alpha) \arrow[u, "{(F_\alpha, T_\alpha)}"] \arrow[r] & (\Phi_\beta, G_\beta) \arrow[u, "{(F_\beta, T_\beta)}"']
\end{tikzcd}$$

We can measure the non-commutativity of each diagram by a **cost function**.

An example

Suppose we have two neural networks for edge detection in images, represented as GEOs.

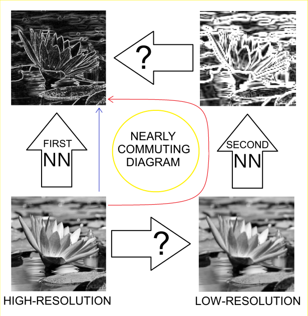

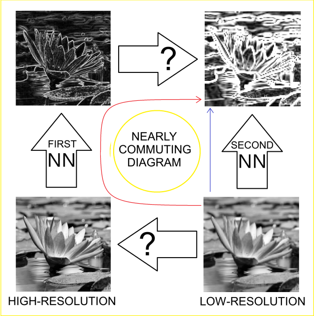

The two neural networks are considered close if there exist two pairs of horizontal GENEOs that make these diagrams "nearly commutative".

An extended pseudo-metric for *ALL* GEOs

To formalize our new pseudo-metric d_E between GEOs, let us consider the category $\mathbf{S}_{all}$ whose objects are all perception spaces, and whose morphisms $(F, T) : (\Phi, G) \to (\Phi', G')$ are GENEOs.
The morphisms in $\mathbf{S}_{all}$ are called *translation GENEOs*. These morphisms describe the possible "logical correspondences" between data represented by different perception spaces.

For example, a translation GENEO might transform high-resolution images into low-resolution images.

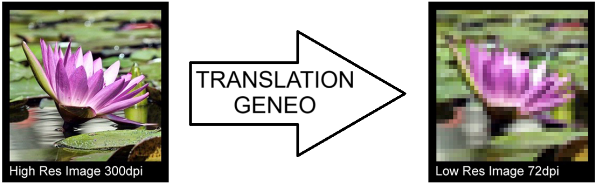

An extended pseudo-metric for *ALL* GEOs

Let us choose a set $\mathscr{G}$ of GEOs. Therefore,

$$\mathscr{G} = \{(F_\alpha, T_\alpha) : (\Phi_\alpha, G_\alpha) \to (\Psi_\alpha, K_\alpha)\}_{\alpha \in A}.$$

To proceed with the definition of our pseudo-metric on $\mathscr{G}$, we need to specify which logical correspondences between data we consider admissible. To this end, let us consider a small subcategory $\mathbf{S}$ of the category $\mathbf{S}_{all}$.

$\mathscr{G}$ will be the set of GEOs where we will define our pseudo-metric, while the morphisms in $\mathbf{S}$ will be the translation GENEOs considered admissible.

Definition of the explainability distance

Let

$$(F_\alpha, T_\alpha) : (\Phi_\alpha, G_\alpha) \to (\Psi_\alpha, K_\alpha)$$
$$(F_\beta, T_\beta) : (\Phi_\beta, G_\beta) \to (\Psi_\beta, K_\beta)$$

be two GEOs in the given set of GEOs $\mathscr{G}$.
Let us consider a pair

$$\pi = \Big((L_{\alpha,\beta}, P_{\alpha,\beta}), (M_{\beta,\alpha}, Q_{\beta,\alpha})\Big)$$

of morphisms in $\mathbf{S}$, with

- $(L_{\alpha,\beta}, P_{\alpha,\beta})$ a morphism from (Φ_α, G_α) to (Φ_β, G_β),
- $(M_{\beta,\alpha}, Q_{\beta,\alpha})$ a morphism from (Ψ_β, K_β) to (Ψ_α, K_α),

Note that the two GENEOs have opposite directions. We say that π is a crossed pair of translation GENEOs from (F_α, T_α) to (F_β, T_β).

Definition of the explainability distance

$$
\begin{tikzcd}
(\Psi_\alpha, K_\alpha) \arrow[drr, phantom, "\pi"] & & (\Psi_\beta, K_\beta) \arrow[ll, "(M_{\beta,\alpha}, Q_{\beta,\alpha})"'] \\
(\Phi_\alpha, G_\alpha) \arrow[u, "(F_\alpha, T_\alpha)"] \arrow[rr, "(L_{\alpha,\beta}, P_{\alpha,\beta})"] & & (\Phi_\beta, G_\beta) \arrow[u, "(F_\beta, T_\beta)"']
\end{tikzcd}
$$

Figure: A crossed pair of translation GENEOs.

Definition of the explainability distance

To proceed, we need to equip each metric space Φ_α with a Borel probability measure μ_α. In simple terms, the measure μ_α represents the probability of the data points in Φ_α appearing in our experiments.

We will assume that all GENEOs in **S** are not just distance-decreasing (i.e., non-expansive) but also measure-decreasing, i.e., if $(L_{\alpha,\beta}, P_{\alpha,\beta}) : (\Phi_\alpha, G_\alpha) \to (\Phi_\beta, G_\beta)$ belongs to **S** and the set $A \subseteq \Phi_\alpha$ is measurable for μ_α, then $L_{\alpha,\beta}(A)$ is measurable for μ_β, and $\mu_\beta(L_{\alpha,\beta}(A)) \leq \mu_\alpha(A)$ (We recall that GENEOs are not surjective, in general).

Definition of the explainability distance

We also assume that the function that takes each $\varphi \in \Phi_\alpha$ to $f_{\alpha,\beta}(\varphi) := D_\Psi\Big((M_{\beta,\alpha} \circ F_\beta \circ L_{\alpha,\beta})(\varphi), F_\alpha(\varphi)\Big)$ is integrable with respect to the probability measure μ_α defined on the dataset Φ_α. The functional cost of π is defined by setting

$$\mathrm{cost}(\pi) := \int_{\Phi_\alpha} D_\Psi\Big((M_{\beta,\alpha} \circ F_\beta \circ L_{\alpha,\beta})(\varphi), F_\alpha(\varphi)\Big)\, d\mu_\alpha.$$

The value $\mathrm{cost}(\pi)$ quantifies how far the two paths in the next figure are from being equivalent, on average, when φ is randomly selected in Φ_α according to the probability measure μ_α.

Definition of the explainability distance

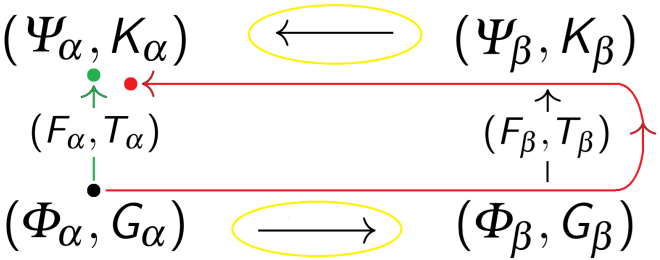

Figure: The explainability distance we are going to define measures how far the green path and the red path are from being equivalent, on average.

Definition of the explainability distance

We can formalize the new pseudo-metric d_E on $\mathscr{G}$ by defining $d_E(GEO1, GEO2)$ as the infimum of the maximum between the cost of π_1 and the cost of π_2, over all crossed pairs π_1 of admissible translation GENEOs from GEO1 to GEO2 and all crossed pairs π_2 of admissible translation GENEOs from GEO2 to GEO1.

Formally, $d_E(GEO1, GEO2)$ is equal to

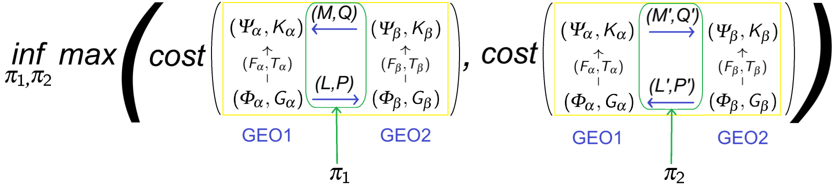

Definition of the explainability distance

Proposition

d_E is an extended pseudo-distance.

The non-expansiveness of GENEOs is a key component in the definition of d_E.

In simple terms, the value $d_E((F_\alpha, T_\alpha),(F_\beta, T_\beta))$ measures the *cost* of changing (F_α, T_α) into (F_β, T_β).

When $d_E((F_\alpha, T_\alpha),(F_\beta, T_\beta))$ is small, it indicates that the GEOs (F_α, T_α) and (F_β, T_β) act approximately in the same way on the data they process, on average.

Back to the basic idea of explanation

Let us recall our informal idea.

Informal idea: We could say that the action of an agent A is explained by another agent B from the perspective of an agent C if:

1. C perceives A and B as similar to each other; $\boxed{\checkmark}$
2. C perceives B as less complex than A.

The formalization of 1 is completed using the pseudo-metric d_E.
How about the formalization of 2?

Complexity of GEOs

Let us assume a set $\Gamma = \{(F_i, T_i) : (\Phi_i, G_i) \to (\Psi_i, K_i)\}$ of GEOs is given. We will say that Γ is our internal library. For each GEO $(F_i, T_i) \in \Gamma$ we arbitrarily choose a value c_i representing the complexity $\mathrm{comp}((F_i, T_i))$ of (F_i, T_i).

Let us now consider the closure of Γ, i.e., the minimal set $\bar{\Gamma}$ such that

- $\bar{\Gamma} \supseteq \Gamma$;
- $\bar{\Gamma}$ is closed under composition (i.e., if $(F, T), (F', T') \in \bar{\Gamma}$ are composable, then $(F', T') \circ (F, T) \in \bar{\Gamma}$);
- $\bar{\Gamma}$ is closed under direct product (i.e., if the GEOs $(F, T), (F', T') \in \bar{\Gamma}$, then $(F, T) \otimes (F', T') \in \bar{\Gamma}$).

Each composition and direct product is associated with a complexity.

The complexity of each GEO $(F, T) \in \bar{\Gamma}$ is obtained by minimizing the sum of the complexities of the GEOs (F_i, T_i) that we use and the complexities of the compositions and direct products that we apply to build (F, T).

Back to the basic idea of explanation

Informal idea: We could say that the action of an agent A is explained by another agent B from the perspective of an agent C if:

1. C perceives A and B as similar to each other; $\boxed{\checkmark}$
2. C perceives B as less complex than A. $\boxed{\checkmark}$

Our theoretical construction is now complete.

A mathematical concept of explanation

Now we can formalize our mathematical concept of **explanation**. Specifically, we can define it as follows: The action of an agent represented by a GEO (F_α, T_α) is $\boxed{\text{explained } \textbf{at a level } \varepsilon}$ by the action of another agent of $\boxed{\textbf{complexity less than } k}$ represented by a GEO $(F_\beta, T_\beta) \in \bar{\Gamma}$ when $d_E((F_\alpha, T_\alpha), (F_\beta, T_\beta)) \leq \varepsilon$.

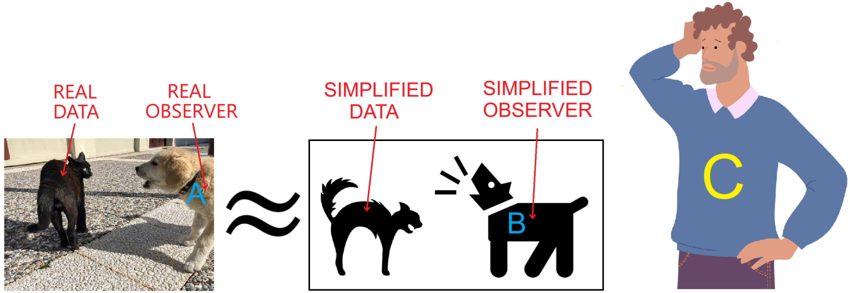

Summary

To sum up, GENEOs are novel mathematical tools designed to approximate agents acting on data—particularly equivariant neural networks—through a compositional approach. They are generally interpretable, which makes them potentially valuable for explainable artificial intelligence (XAI).


Thanks for
your attention