Explainability of neural networks through the use of GENEOs

Patrizio Frosini

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

Explainable, Trustworthy and Robust Artificial Intelligence
September 8 - 11, 2025, Lake Como School of Advanced Studies

Outline

Some epistemological assumptions

Assumption 1: Data are processed by observers

Data have no meaning without an observer to interpret them.

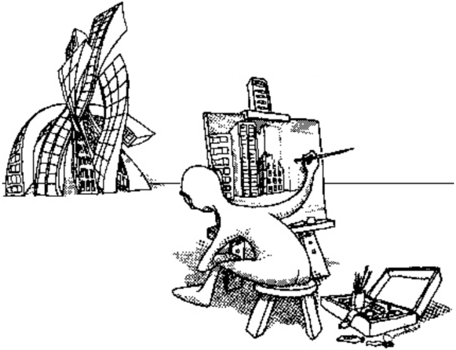

An observer is an agent that transforms data while preserving their symmetries.

Assumption 2: Observers are variables

Data interpretation strongly depends on the chosen observer.

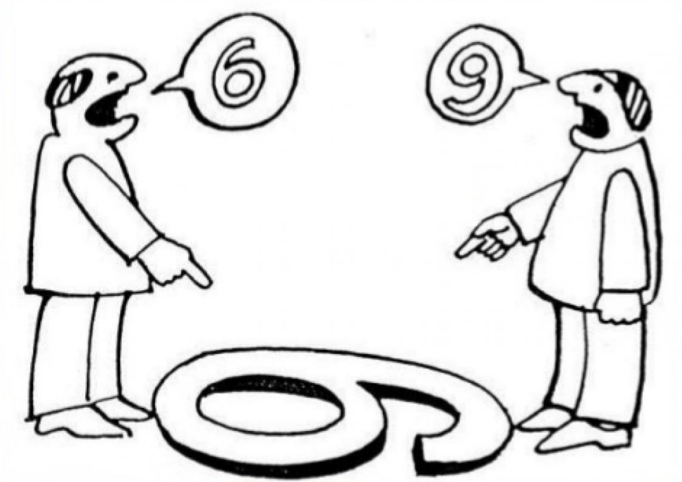

Assumption 3: Observers are important

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.

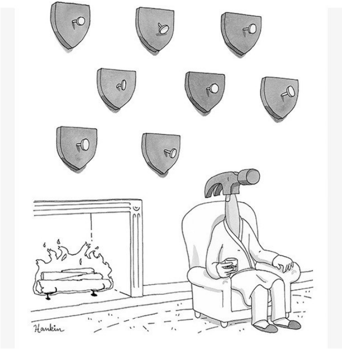

The shape is not in the data but in the eyes of the observer.

How can we translate these ideas into mathematics?

?

?

Perception spaces and GENEOs

PERCEPTION
SPACE

GROUP
EQUIVARIANT
NON-EXPANSIVE
OPERATOR
(GENEO)

Some epistemological assumptions

What is a GENEO?

Some basics on the theory of GENEOs

Building linear and nonlinear GENEOs

GENEOs and XAI

What is a GENEO?

- A Group Equivariant Non-Expansive Operator (GENEO) is a mathematical tool used to approximate observers that act on data.
- The theory of GENEOs is based on the idea that the geometric characteristics of observers significantly influence the interpretation of data.
- In these lectures, we will explore the core properties of GENEOs, examine their role in machine learning, and discuss their promising applications in explainable artificial intelligence.

Some epistemological assumptions

What is a GENEO?

Some basics on the theory of GENEOs

Building linear and nonlinear GENEOs

GENEOs and XAI

Let us begin by recalling the concept of a group action

A *group* $\mathbf{G} = (G, \circ, \mathrm{id}_G)$ consists of a set G, an associative operation $\circ\colon G \times G \to G$ having a unit element $\mathrm{id}_G \in G$ such that, for all $g \in G$, there exists $g^{-1} \in G$ satisfying $g \circ g^{-1} = g^{-1} \circ g = \mathrm{id}_G$.

A *group homomorphism* $T\colon (G, \circ_G, \mathrm{id}_G) \to (K, \circ_K, \mathrm{id}_K)$ is a function $T\colon G \to K$ such that, for all $g_1, g_2 \in G$,
$T(g_2 \circ_G g_1) = T(g_2) \circ_K T(g_1)$.

Given a group $(G, \circ, \mathrm{id}_G)$ and a set X, a *group left action* is a function $*\colon G \times X \to X$ such that, for all $x \in X$ and $g_1, g_2 \in G$,

$$\mathrm{id}_G * x = x \qquad \text{and} \qquad (g_2 \circ g_1) * x = g_2 * (g_1 * x).$$

An example of a group action

Let R be the group of rigid motions (isometries) of the Euclidean plane $\mathbb{R}^2$. The group operation is composition $\circ$:

- Composition is associative: $f_3 \circ (f_2 \circ f_1) = (f_3 \circ f_2) \circ f_1$.
- The *identity* rigid motion id is the unit: $\mathrm{id} \circ f = f \circ \mathrm{id} = f$.
- Every $f \in R$ has an inverse $f^{-1} \in R$ with $f \circ f^{-1} = f^{-1} \circ f = \mathrm{id}$.

The group R acts on $\mathbb{R}^2$ on the left by evaluation:

$$f * P := f(P), \qquad f \in R,\ P \in \mathbb{R}^2.$$

This is a well-defined left action since

$$\mathrm{id} * P = P \quad \text{and} \quad (f_2 \circ f_1) * P = (f_2 \circ f_1)(P) = f_2\big(f_1(P)\big) = f_2 * (f_1 * P).$$

Perception spaces

Recall that a pseudo-metric is a distance function d satisfying nonnegativity, symmetry, and the triangle inequality, but not necessarily 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. A group $(G, \circ)$ acting on Φ on the left, denoted by $*$. We assume the action is by isometries, 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 call (Φ, G) an (extended) perception space.

Moreover, the action induces a pseudo-metric D_G on G defined by

$$D_G(g_1, g_2) := \sup_{\varphi \in \Phi} D_\Phi(g_1 * \varphi, g_2 * \varphi), \qquad g_1, g_2 \in G.$$

Perception spaces

The set Φ represents the data we may obtain from our measuring tools (functions, graphs, point clouds, ...). The group G represents the possible transformations of the data that 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 is given by the set of electrocardiograms, represented as functions of the time variable, while G can be the group of time translations.

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).

Proposition

If $(F, T) : (\Phi, G) \to (\Psi, K)$ *is a GEO, with* F *non-expansive and surjective, then* T *is non-expansive.*

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 (T is the identity taking the **group of isometries** to itself).

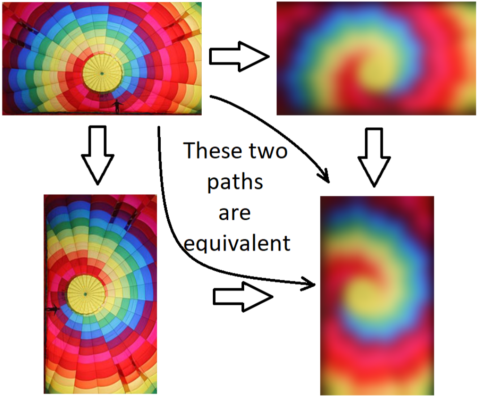

These two paths are equivalent

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, T is the identity taking the **group of isometries** to itself).

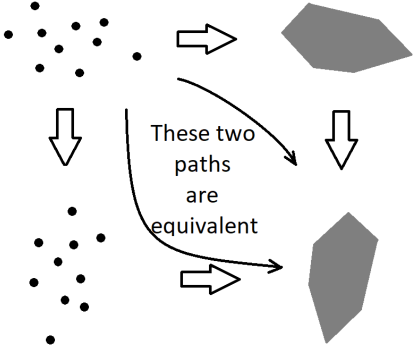

Here, the Hausdorff distance between compact sets is used.

Another example (for those familiar with TDA)

The operator taking filtering functions to persistence diagrams is another example of GENEO.

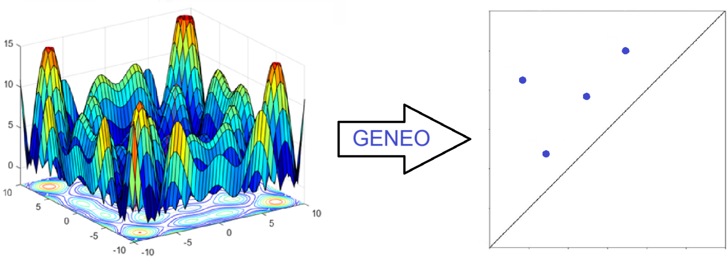

Equivariance ⇒ invariance of persistence diagrams
under reparameterization of the domain.

Nonexpansiveness ⇒ stability of persistence diagrams.

Why are GENEOs interesting?

- GENEOs rest on a rigorous topological/geometric framework (in what follows we outline several results).
- GENEOs encode prior knowledge about the chosen observer.
- The non-expansiveness property of GENEOs imposes a strong constraint that enables meaningful data simplification.
- GENEOs enable a compositional approach to deep learning.
- Analyzing the geometry of the observer space (as represented by GENEOs) is often more informative than analyzing the geometry of the data space.

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 some applications in deep learning more transparent and interpretable and speed up the learning process.

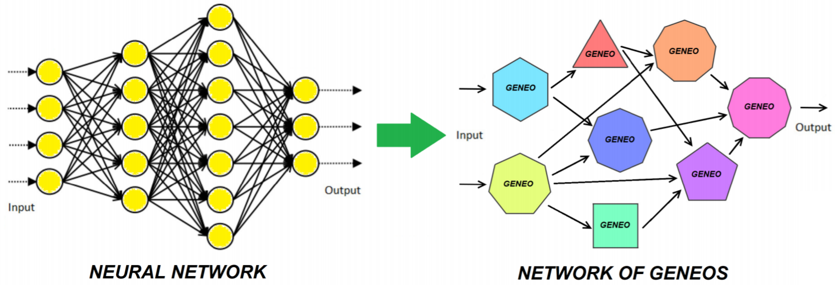

Some research projects concerning GENEOs (I)

CNIT / WiLab - Huawei Joint Innovation Center (JIC)

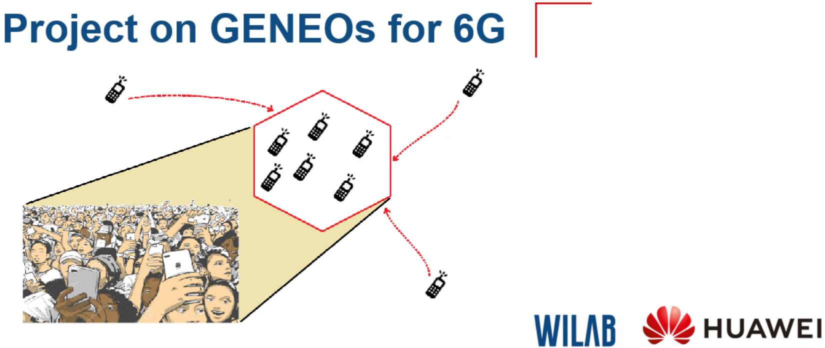

Project on GENEOs for 6G

Some research projects concerning GENEOs (II)

PANDORA

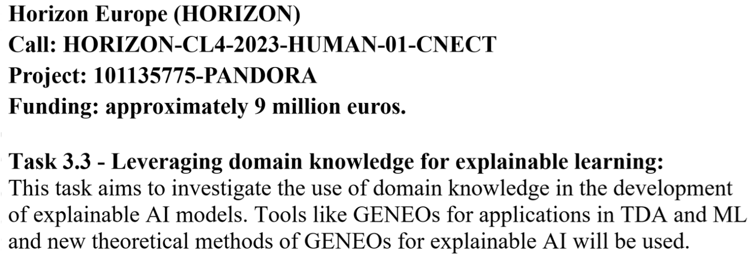

Horizon Europe (HORIZON)
Call: HORIZON-CL4-2023-HUMAN-01-CNECT
Project: 101135775-PANDORA
Funding: approximately 9 million euros.

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

https://pandora-heu.eu/consortium/

Some research projects concerning GENEOs (III)

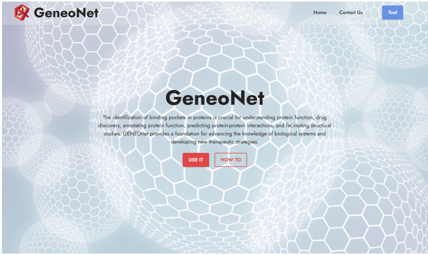

GeneoNet

The identification of binding pockets in proteins is crucial for understanding protein function, drug discovery, annotating protein function, predicting protein-protein interactions, and facilitating structural studies. GENEOnet provides a foundation for advancing the knowledge of biological systems and developing new therapeutic strategies.

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

⟹ Attend Giovanni Bocchi's lecture on Wednesday!

Some basic theoretical results

The rest of this morning lecture will be devoted to illustrating, without proofs, some theoretical results about GENEOs.

QUESTION: Why spend time on the mathematical details of GENEOs?

ANSWER: For the same reason that, e.g., in linear algebra one introduces notions such as linear dependence, generating sets, and bases: to equip oneself with general tools that enable the efficient, controlled, and uniform application of common methods to a wide range of problems and applications.

A pseudo-metric on G

We recall that G is endowed 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 $g_1, g_2 \in G$.

Proposition

Let (Φ, G) be a perception space. The followings hold.

(a) *G is a topological group.*

(b) *The action of G on Φ is continuous.*

Theorem

If Φ is compact and G is complete, then G is compact.

Recall that a pseudometric space is said to be compact if every sequence in the space has a convergent subsequence.

Compactness of the space of GENEOs

From now on, let $\mathscr{F}_T^{\mathrm{all}}$ denote the set of all GENEOs (F, T) from a perception space (Φ, G) to a perception space (Ψ, K), with fixed homomorphism T. We equip $\mathscr{F}_T^{\mathrm{all}}$ with the following distance:

$$D_{\mathrm{GENEO}}((F,T),(F',T)) := \sup_{\varphi\in\Phi} D_\Psi\big(F(\varphi), F'(\varphi)\big),$$

where D_Ψ is the metric chosen on Ψ.

Theorem

If (Φ, D_Φ), (Ψ, D_Ψ) are compact, then the space $(\mathscr{F}_T^{\mathrm{all}}, D_{\mathrm{GENEO}})$ is compact.

Compactness of the space of GENEOs

The compactness of the space of GENEOs is also an important property from an applied standpoint. It implies that the space of GENEOs can be approximated arbitrarily well by a finite set of GENEOs, and that every continuous loss function defined on this space attains an absolute minimum (i.e., there exists an optimal GENEO with respect to the loss function).

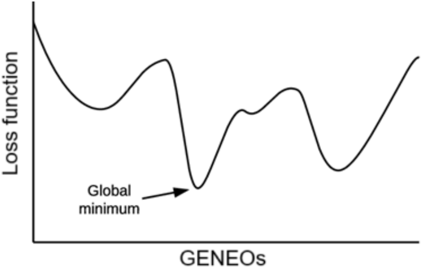

Convexity of the space of GENEOs

Now assume that

- Φ, Ψ are normed real vector spaces, with distances D_Φ, D_Ψ induced by the norms $\|\cdot\|_\Phi$, $\|\cdot\|_\Psi$ (i.e., $D_\Phi(\varphi, \varphi') = \|\varphi - \varphi'\|_\Phi$ and $D_\Psi(\psi, \psi') = \|\psi - \psi'\|_\Phi$);
- the group actions are linear, i.e.,
  $g * (a\varphi + b\varphi') = a(g * \varphi) + b(g * \varphi')$,
  $k * (a\psi + b\psi') = a(k * \psi) + b(k * \psi')$, for all $a, b \in \mathbb{R}$, $g \in G$, $k \in K$, $\varphi, \varphi' \in \Phi$, $\psi, \psi' \in \Psi$.

Let $(F_1, T), (F_2, T), \ldots, (F_n, T)$ be GENEOs from (Φ, G) to (Ψ, K). If $(a_1, \ldots, a_n) \in \mathbb{R}^n$ and $\sum_{i=1}^n |a_i| \leq 1$, define

$$F_\Sigma(\varphi) := \sum_{i=1}^{n} a_i F_i(\varphi).$$

Convexity of the space of GENEOs

Proposition

(F_{Σ}, T) is a GENEO from (Φ, G) to (Ψ, K).

Corollary

The space $\mathscr{F}_T^{\mathrm{all}}$ is convex.

Two key observations (1)

- While the space of data is often non-convex (and hence averaging data does not make sense), the assumption of convexity of Ψ implies the convexity of the space of observers and allows us to consider the "average of observers".

Two key observations (2)

- Our main goal is to develop a good geometric and compositional theory to approximate an ideal observer. In our model, "to approximate an observer" means to look for a GENEO F that minimizes a suitable "cost function" $c(F)$. The cost function quantifies the error that is committed by taking the GENEO F instead of the ideal observer. Since the space of GENEOs is compact and convex (under the assumption that the data spaces are compact and convex), if the cost function $c(F)$ is strictly convex we have that there is one and only one GENEO that best approximates the ideal observer.

The functional case of GENEOs

A particularly important case of GENEO is the one in which the perception spaces involved consist of sets of data expressed as functions with values in $\mathbb{R}$ or $\mathbb{R}^n$.

In the remainder of this lecture we shall restrict our attention to the case of data expressed as real-valued functions.

In every perception space (Φ, G) under consideration, Φ will be a set of $\mathbb{R}$-valued functions defined on a domain X (denoted by $\mathrm{Dom}(\Phi)$) and endowed with the metric $D_\Phi(\varphi, \varphi') = \|\varphi - \varphi'\|_\infty$).

The group G will be a group of permutations of X such that, if $\varphi \in \Phi$ and $g \in G$, then $\varphi \circ g^{-1} \in \Phi$ as well, and the left action of G on Φ will be given by $g * \varphi = \varphi \circ g^{-1}$.

The functional case of GENEOs

This is indeed a left action of G on Φ:

$$\begin{aligned}(g_2 \circ g_1) * \varphi &= \varphi \circ (g_2 \circ g_1)^{-1} \\ &= \varphi \circ (g_1^{-1} \circ g_2^{-1}) \\ &= (\varphi \circ g_1^{-1}) \circ g_2^{-1} \\ &= g_2 * (\varphi \circ g_1^{-1}) \\ &= g_2 * (g_1 * \varphi).\end{aligned}$$

Furthermore, this action is linear:

$$\begin{aligned}g * (a\varphi_1 + b\varphi_2) &= (a\varphi_1 + b\varphi_2) \circ g^{-1} \\ &= a(\varphi_1 \circ g^{-1}) + b(\varphi_2 \circ g^{-1}) \\ &= a(g * \varphi_1) + b(g * \varphi_2).\end{aligned}$$

Some epistemological assumptions

What is a GENEO?

Some basics on the theory of GENEOs

Building linear and nonlinear GENEOs

GENEOs and XAI

Elementary methods to build GENEOs

Proposition (Composition)

If $(F_1, T_1) \in \mathrm{GENEO}((\Phi, G), (\Psi, H))$*, and* $(F_2, T_2) \in \mathrm{GENEO}((\Psi, H), (\chi, K))$*, then* $(F_2 \circ F_1, T_2 \circ T_1) \in \mathrm{GENEO}((\Phi, G), (\chi, K))$*.*

Proposition (Image by a 1-Lipschitz function)

If $(F_1, T), \ldots, (F_n, T) \in \mathrm{GENEO}((\Phi, G), (\Psi, H))$*,* L *is a 1-Lipschitz map from* $\mathbb{R}^n$ *to* $\mathbb{R}$*, and* $L^*(F_1, \ldots, F_n)(\Phi) \subseteq \Psi$ *(where* L^* *is the map induced by* L*), then* $(L^*(F_1, \ldots, F_n), T) \in \mathrm{GENEO}((\Phi, G), (\Psi, H))$*.*

The next three statements follow from the last proposition.

Elementary methods to build GENEOs

Proposition (LATTICE OF GENEOS)

If $(F_1, T), \dots, (F_n, T) \in \mathrm{GENEO}((\Phi, G), (\Psi, H))$, *and* $\max(F_1, \dots, F_n)(\Phi), \min(F_1, \dots, F_n)(\Phi) \subseteq \Psi$, *then* $(\max(F_1, \dots, F_n), T), (\min(F_1, \dots, F_n), T) \in \mathrm{GENEO}((\Phi, G), (\Psi, H))$.

Proposition (Translation)

If $(F, T) \in \mathrm{GENEO}((\Phi, G), (\Psi, H))$, *and* $F_b(\Phi) \subseteq \Psi$ *for* $F_b(\varphi) := F(\varphi) - b$, *then* $(F_b, T) \in \mathrm{GENEO}((\Phi, G), (\Psi, H))$.

Proposition (Convex combination)

If $(F_1, T), \dots, (F_n, T) \in \mathrm{GENEO}((\Phi, G), (\Psi, H))$, $(a_1, \dots, a_n) \in \mathbb{R}^n$ *with* $\sum_{i=1}^n |a_i| \leq 1$, *and* $F_\Sigma(\Phi) \subseteq \Psi$ *for* $F_\Sigma(\varphi) := \sum_{i=1}^n a_i F_i(\varphi)$, *then* $(F_\Sigma, T) \in \mathrm{GENEO}((\Phi, G), (\Psi, H))$.

Permutant measures

Let us consider the set $\Phi = \mathbb{R}^X \cong \mathbb{R}^n$ of all functions from a finite set $X = \{x_1, \ldots, x_n\}$ to $\mathbb{R}$, and a subgroup G of the group $\mathrm{Bij}(X)$ of all bijections from X to X.

Definition

A finite (signed) measure μ on $\mathrm{Bij}(X)$ is called a *permutant measure* with respect to G if every subset H of $\mathrm{Bij}(X)$ is measurable and μ is invariant under the conjugation action of G (i.e., $\mu(H) = \mu(gHg^{-1})$ for every $g \in G$).

An example of permutant measure

Let us consider the set X of the vertices of a cube in $\mathbb{R}^3$, and the group G of the orientation-preserving isometries of $\mathbb{R}^3$ that take X to X. Let π_1, π_2, π_3 be the three planes that contain the center of mass of X and are parallel to a face of the cube. Let $h_i : X \to X$ be the orthogonal symmetry with respect to π_i, for $i \in \{1, 2, 3\}$.

We can now define a permutant measure μ on the group $\mathrm{Bij}(X)$ by setting $\mu(h_1) = \mu(h_2) = \mu(h_3) = c$, where c is a positive real number, and $\mu(h) = 0$ for any $h \in \mathrm{Bij}(X)$ with $h \notin \{h_1, h_2, h_3\}$.

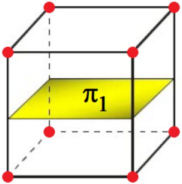

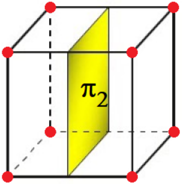

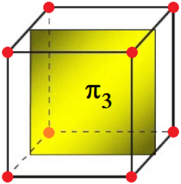

Representation Theorem for linear GENEOs

The following theorem explains the importance of the concept of a permutant measure.

Theorem (Representation Theorem for linear GENEOs)

Let us assume that the group $G \subseteq \mathrm{Bij}(X)$ *transitively acts on the finite set* X *and that* F *is a map from* $\mathbb{R}^X$ *to* $\mathbb{R}^X$. *The map* F *is a linear GENEO from* $\mathbb{R}^X$ *to* $\mathbb{R}^X$ *with respect to the identical homomorphism* $\mathrm{id}_G : g \mapsto g$ **if and only if** *a permutant measure* μ *with respect to* G *exists, such that* $F(\varphi) = \sum_{h \in \mathrm{Bij}(X)} \varphi h^{-1} \, \mu(h)$ *for every* $\varphi \in \mathbb{R}^X$, *and* $\sum_{h \in \mathrm{Bij}(X)} |\mu(h)| \leq 1$.

Building GENEOs by permutant measures

It is interesting to observe that the set $\mathrm{PM}(G)$ of permutant measures with respect to G forms a lattice. Indeed, if $\mu_1, \mu_2 \in \mathrm{PM}(G)$, then the measures μ', μ'' on $\mathrm{Bij}(X)$, defined respectively by $\mu'(h) := \min\{\mu_1(h), \mu_2(h)\}$ and $\mu''(h) := \max\{\mu_1(h), \mu_2(h)\}$, still belong to $\mathrm{PM}(G)$. Moreover, if $\mu \in \mathrm{PM}(G)$, then $|\mu| \in \mathrm{PM}(G)$.

Furthermore, $\mathrm{PM}(G)$ is closed under linear combinations. Therefore, $\mathrm{PM}(G)$ naturally carries the structure of a real vector space.

We emphasize that as the group G becomes larger, the lattice $\mathrm{PM}(G)$ becomes smaller. In other words, the theory of permutant measures becomes increasingly useful as the groups grow larger.

Building GENEOs by permutant measures

The method for building GENEOs based on permutant measures can be generalized by replacing the arithmetic mean with another symmetric function. We can indeed show that when a symmetric function and a **permutant** for the equivariance group G are available, we can easily build a (non-linear) GENEO with respect to G. First of all, let us recall the concept of *permutant*, which is equivalent to the one of *permutant measure uniformly distributed on its support*.

Definition

We say that a subset $H \subseteq \mathrm{Bij}(X)$ is a **permutant** for the group $G \subseteq \mathrm{Bij}(X)$ if either $H = \emptyset$ or $gHg^{-1} = H$ for every $g \in G$.

Note that a subset H of $\mathrm{Bij}(X)$ is a permutant for G if and only if H is a union of orbits for the conjugation action of G on $\mathrm{Bij}(X)$.

Building GENEOs by permutant measures

Let $\mathscr{S}\colon \mathbb{R}^n \to \mathbb{R}$ be a symmetric function.
Set $\mathrm{Bij}_\Phi(X) = \{g \in \mathrm{Bij}(X) : \Phi \circ g \subseteq \Phi\}$. If $H = \{h_i\}_{i=1}^n$ is a non-empty permutant for $G \subseteq \mathrm{Bij}_\Phi(X)$, then we can define an operator $\mathscr{S}_H\colon \Phi \to \mathbb{R}^X$ by setting, for any $\varphi \in \Phi$,

$$\mathscr{S}_H(\varphi) := \mathscr{S}(\varphi \circ h_1, \ldots, \varphi \circ h_n),$$

where $\mathscr{S}(\varphi \circ h_1, \ldots, \varphi \circ h_n)(x) := \mathscr{S}((\varphi \circ h_1)(x), \ldots, (\varphi \circ h_n)(x))$ for every $x \in X$.

Proposition

$\mathscr{S}_H$ is a GEO from (Φ, G) to $(\mathbb{R}^X, G)$ with respect to the identity homomorphism $\mathrm{id}_G : G \to G$. If the restriction of $\mathscr{S}$ to $\mathrm{Im}(\Phi)^n$ is non-expansive, then $\mathscr{S}_H$ is a GENEO from (Φ, G) to $(\mathbb{R}^X, G)$ with respect to id_G.

Some references about GENEOs (I)

- P. Frosini, G. Jabłoński, *Combining persistent homology and invariance groups for shape comparison*, **Discrete & Computational Geometry**, vol. 55 (2016), n. 2, pp. 373-409.
- 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, n. 9 (2019), pp. 423-433.
- G. Bocchi, S. Botteghi, M. Brasini, P. Frosini and N. Quercioli, *On the finite representation of linear group equivariant operators via permutant measures*, **Annals of Mathematics and Artificial Intelligence**, vol. 91 (2023), n. 4, pp. 465-487.

Some references about GENEOs (II)

- A. Micheletti, *A new paradigm for artificial intelligence based on group equivariant non-expansive operators*, **European Mathematical Society Magazine**, 128 (2023), pp. 4–12.
- 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*, Joint Proceedings of the xAI 2024 Late-breaking Work, Demos and Doctoral Consortium co-located with the **2nd World Conference on eXplainable Artificial Intelligence** (xAI 2024), Valletta, Malta, July 17-19, 2024, (Edited by Luca Longo, Weiru Liu, Grégoire Montavon), pp. 217-224.
- G. Bocchi, M. Ferri, P. Frosini, *A novel approach to graph distinction through GENEOs and permutants*, **Scientific Reports**, 15, 6259 (2025).

Some references about GENEOs (III)

- J. J. Colombini, F. Bonchi, F. Giannini, F. Giannotti, R. Pellungrini and P. Frosini, *Mathematical foundation of interpretable equivariant surrogate models*, **3rd World Conference on Explainable Artificial Intelligence** (XAI-2025), Novel Post-hoc & Ante-hoc XAI Approaches, 09-11 July, 2025 - Istanbul, Turkey. (to appear)
- G. Bocchi, P. Frosini, A. Micheletti, A. Pedretti, G. Palermo, D. Gadioli, C. Gratteri, F. Lunghini, A. D. Biswas, P. F.W. Stouten, A. R. Beccari, A. Fava, C. Talarico, GENEOnet: *A breakthrough in protein binding pocket detection using group equivariant non-expansive operators*, **Scientific Reports**, (to appear)

End of the first part of the lecture

Thanks for your attention!

GENEOs and XAI

In this second lecture we explore how GENEOs can be applied to explainable AI.

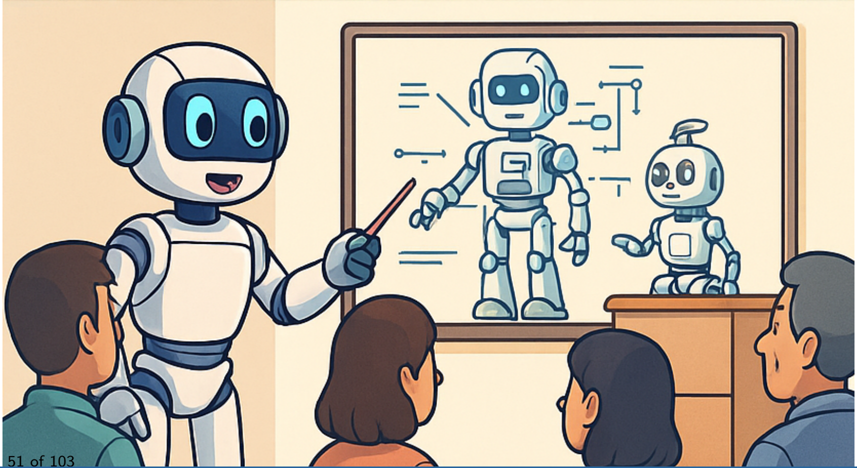

GENEOs and XAI

GENEOs as a mathematical framework to quantify distances between operators with distinct domain–codomain pairs

How close are these two operators?

Collaborators in 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)

Basic idea

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

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.

E.g., let's consider two neural networks represented as two GEOs.

Note that a GEO can take another GEO as an input.

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 have to 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 is a non-trivial challenge.

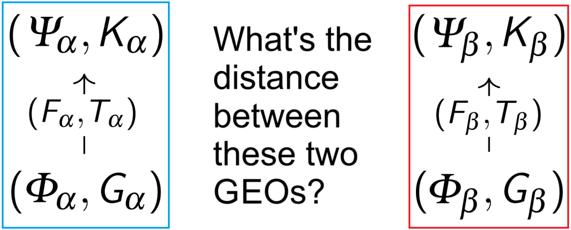

$$\begin{tikzcd}(\Psi_\alpha, K_\alpha) \\ (\Phi_\alpha, G_\alpha) \arrow[u, "{(F_\alpha, T_\alpha)}"]\end{tikzcd}$$

What's the distance between these two GEOs?

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

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.

$$\approx \quad \begin{tikzcd}
(\Psi, K) \\
(F, T) \arrow[u] \arrow[d, dash] \\
(\Phi, G)
\end{tikzcd}$$

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.

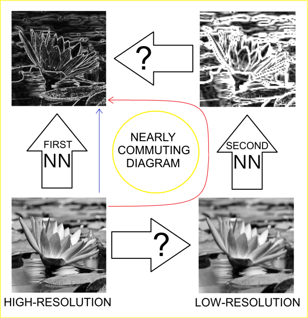

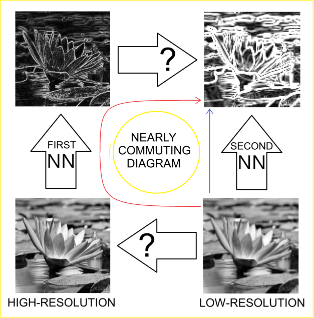

The two neural networks are considered close if there exist two pairs of horizontal GENEOs that make these diagrams "nearly commutative".

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.

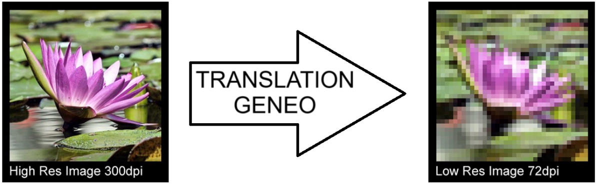

An extended pseudo-metric for *ALL* GEOs

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

$$\begin{aligned}(F_\alpha, T_\alpha) : (\Phi_\alpha, G_\alpha) \to (\Psi_\alpha, K_\alpha)\\ (F_\beta, T_\beta) : (\Phi_\beta, G_\beta) \to (\Psi_\beta, K_\beta)\end{aligned}$$

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 **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

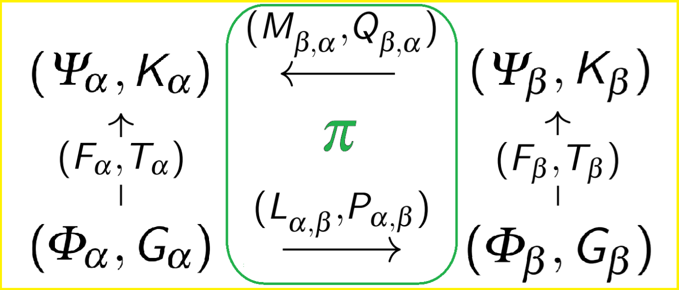

$$\begin{tikzcd}
(\Psi_\alpha, K_\alpha) & & (\Psi_\beta, K_\beta) \arrow[ll, "{(M_{\beta,\alpha}, Q_{\beta,\alpha})}"'] \\
 & \pi & \\
(\Phi_\alpha, G_\alpha) \arrow[uu, "{(F_\alpha, T_\alpha)}" description] \arrow[rr, "{(L_{\alpha,\beta}, P_{\alpha,\beta})}"] & & (\Phi_\beta, G_\beta) \arrow[uu, "{(F_\beta,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

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

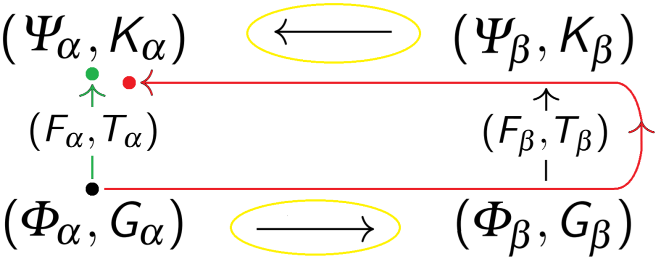

$$\begin{tikzcd}
(\Psi_\alpha, K_\alpha) & (\Psi_\beta, K_\beta) \arrow[l] \\
(\Phi_\alpha, G_\alpha) \arrow[u, mapsto, "{(F_\alpha, T_\alpha)}"] \arrow[r] & (\Phi_\beta, G_\beta) \arrow[u, mapsto, "{(F_\beta, T_\beta)}"']
\end{tikzcd}$$

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

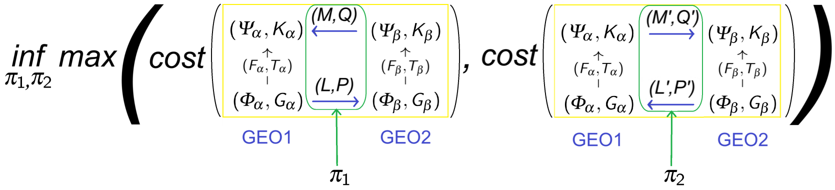

$$\inf_{\pi_1,\pi_2} \max\left(cost\left(\begin{tikzcd} (\Psi_\alpha, K_\alpha) & (\Psi_\beta, K_\beta) \arrow[l, "(M,Q)"'] \\ (\Phi_\alpha, G_\alpha) \arrow[u, "(F_\alpha, T_\alpha)"] \arrow[r, "(L,P)"] & (\Phi_\beta, G_\beta) \arrow[u, "(F_\beta, T_\beta)"'] \end{tikzcd}\right), cost\left(\begin{tikzcd} (\Psi_\alpha, K_\alpha) \arrow[r, "(M',Q')"] & (\Psi_\beta, K_\beta) \\ (\Phi_\alpha, G_\alpha) \arrow[u, "(F_\alpha, T_\alpha)"] & (\Phi_\beta, G_\beta) \arrow[u, "(F_\beta, T_\beta)"'] \arrow[l, "(L',P')"'] \end{tikzcd}\right)\right)$$

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).

The set Γ represents the elementary GEOs that we can use to build other more complex GEOs.

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 $(F, T) \otimes (F', T') \in \bar{\Gamma}$).

Complexity of GEOs

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).

Other forms of composition of GEOs can be added to the model.

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 explained **at a level ε** by the action of another agent of **complexity less than k** represented by a GEO $(F_\beta, T_\beta) \in \bar{\Gamma}$ when $d_E((F_\alpha, T_\alpha), (F_\beta, T_\beta)) \leq \varepsilon$.

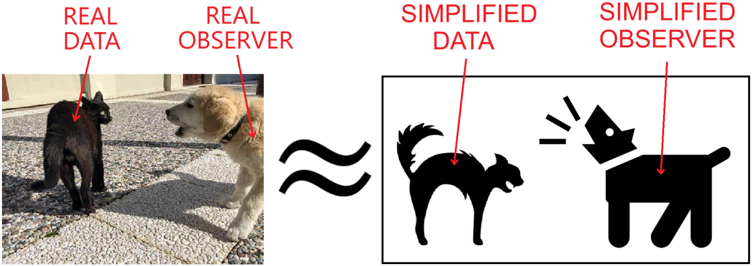

Explainable signal reconstruction

Reconstruction of sparse urban wireless signals via GENEOs

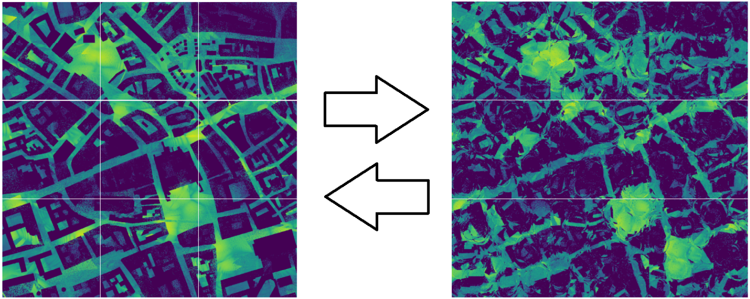

Collaborators in this research

- Lorenzo Mario Amorosa (University of Bologna & WiLab-CNIT)
- Francesco Conti (Université Côte d'Azur)
- Yiqun Ge (Huawei Technologies Canada)
- Tayebeh Lotfi Mahyari (Huawei Technologies Canada)
- Nicola Quercioli (CINECA)
- Flavio Zabini (University of Bologna & WiLab-CNIT)

Explainable signal reconstruction

We propose a GENEO-based approach for reconstructing radio signals in urban wireless networks from extremely sparse samples.

We describe the signal to be reconstructed as a function $\varphi : \mathbb{R}^2 \to [0,1]$. The function φ represents the intensity of the signal over a 2D grid, and is often referred to as ground truth or GT. We also consider the function $\psi : \mathbb{R}^2 \to [0,1]$ taking each point p to the reliability $\psi(p)$ of the value $\varphi(p)$ measured at the point $p = (x,y) \in \mathbb{R}^2$.

We aim at reconstructing φ by starting from a very poor sampling of φ (tipically 2% or 3% of the GT), by using GENEOs.

Explainable signal reconstruction

The main idea is to build GENEOs able to look for meaningful patterns in the signal.

In our model, a pattern is a pair of functions $P = (h, \chi_h)$, where:

- $h \colon \mathbb{R}^2 \to [0,1]$ represents a shape we are looking for in the signals;
- $\chi_h \colon \mathbb{R}^2 \to \{0,1\}$ establishes where the values expressed by h are reliable.

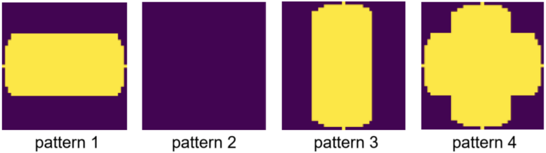

Simple illustrative examples of patterns.

Explainable signal reconstruction

We consider a pair $S = (f, \psi_f)$ of functions from $\mathbb{R}^2$ to $[0,1]$, representing the available sampling and its reliability.
Then, for each possible pattern $P = (h, \chi_h)$, we can define:

$$\begin{aligned}\mathscr{S}_{S,P}(x,y) = \int_{\mathbb{R}^2} &|f(x+\xi, y+\eta) - h(\xi,\eta)| \\ &\cdot \psi_f(x+\xi, y+\eta)\chi_h(\xi,\eta)\ d\mu(\xi,\eta)\ ,\end{aligned}$$

which tells us how much S differs from P in a neighborhood of the point (x,y); and

$$\mathscr{A}_{S,P}(x,y) = \int_{\mathbb{R}^2} \psi_f(x+\xi, y+\eta)\chi_h(\xi,\eta)\ d\mu(\xi,\eta)\ ,$$

which quantifies the reliability of our data in the region where we can perform the comparison between S and P.

Explainable signal reconstruction

In the definition of $\mathscr{S}_{S,P}(x,y)$ and $\mathscr{A}_{S,P}(x,y)$, we consider the standard Lebesgue measure $\bar{\mu}$ on $\mathbb{R}^2$, and set $\mu(A) = \bar{\mu}(A)/\bar{\mu}(D_h)$ for each measurable $A \subseteq \mathbb{R}^2$, where the support D_h of χ_h is assumed to have finite positive measure. Hence, we can define the following GENEO F, which maps the function f to the function $\hat{c}_{S,P}$ given by

$$\hat{c}_{S,P}(x,y) = \mathscr{A}_{S,P}(x,y) - \mathscr{S}_{S,P}(x,y)\ .$$

The value $\hat{c}_{S,P}(x,y)$ quantifies the similarity between the pattern described by h and the structure of f in the neighborhood of (x,y).

Explainable signal reconstruction

Proposition

$0 \leq \hat{c}_{S,P}(x,y) \leq \mathscr{A}_{S,P}(x,y) \leq 1$ for every $(x,y) \in \mathbb{R}^2$.

Proposition

Assume that two signals $S_1 = (f_1, \psi_{f_1})$, $S_2 = (f_2, \psi_{f_2})$ and a pattern $P = (h, \chi_h)$ are given. If $\psi_{f_1} \equiv \psi_{f_2}$, it holds that $\mathscr{A} := \mathscr{A}_{S_1,P} = \mathscr{A}_{S_2,P}$ and

$$\|\hat{c}_{S_1,P} - \hat{c}_{S_2,P}\|_\infty \leq \|\mathscr{A}\|_\infty \|f_1 - f_2\|_\infty.$$

Proposition

The map F taking $S = (f, \psi_f)$ to the function $\hat{c}_{S,P}$ is a GENEO with respect to the group of planar translations.

Explainable signal reconstruction

Remark

Depending on the geometry of D_h and the choice of patterns, we can allow F to be equivariant with respect to a larger group, i.e., the group of isometries of the plane.

In the following, we will assume that our samplings and our patterns are reliable everywhere, i.e., $\psi \equiv \psi_{f_1} \equiv \psi_{f_2} \equiv 1$ for every f_1, f_2.

Given a sampled signal $\hat{S} = (\hat{\varphi}, \hat{\psi})$, we have to consider a *library* of N patterns to reconstruct φ, that is

$$\{P_1 = (h_1, \chi_{h_1}), \dots, P_N = (h_N, \chi_{h_N})\} \ .$$

Explainable signal reconstruction

Therefore, we can compute $\{\hat{c}_{\hat{S},P_i}\}_{i=1}^{N}$.

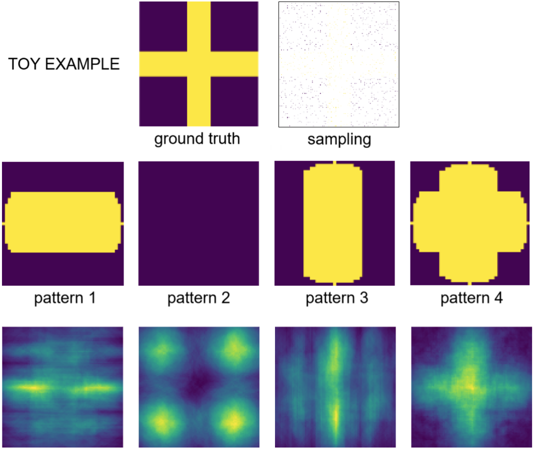

Heatmaps of the corresponding functions $\hat{c}_{S,P}(x,y)$

Explainable signal reconstruction

Now, for any index i we define the similarity coefficient of the pattern P_i at the point p, when P_i is centered at another point $q \in \mathbb{R}^2$, as

$$\mathrm{sim}(p, P_i, q) = \hat{c}_{\hat{S}, P_i}(q) \chi_{h_i}(p - q) \ .$$

In plain words, the value $\mathrm{sim}(p, P_i, q)$ quantifies how plausible it is that at the point p the dominant pattern is P_i, centered at the point q. The term $\chi_{h_i}(p - q)$ is needed because we are interested only in the points of P_i at which the pattern is reliable (i.e., $\chi_{h_i}(p - q) \neq 0$).

Explainable signal reconstruction

The last step in our procedure consists of using the functions $\mathrm{sim}(\cdot, P_i, q)$, varying P_i and q, to produce the reconstructed signal $\varphi_{\mathrm{rec}} \colon \mathbb{R}^2 \to [0,1]$. To do this, considering

$$\left(\,\bar{q}(p), \bar{\imath}(p)\,\right) = \underset{\substack{i \in \{1,\dots,N\} \\ q \in \mathbb{R}^2}}{\mathrm{argmax}}\ \mathrm{sim}\big(p, P_i, q\big)\ ,$$

we reconstruct the point p using the $\bar{\imath}(p)$-th pattern with a shift equal to $\bar{q}(p)$ as

$$\varphi_{\mathrm{rec}}(p) := h_{\bar{\imath}(p)}\Big(p - \bar{q}(p)\Big)\ .$$

Explainable signal reconstruction

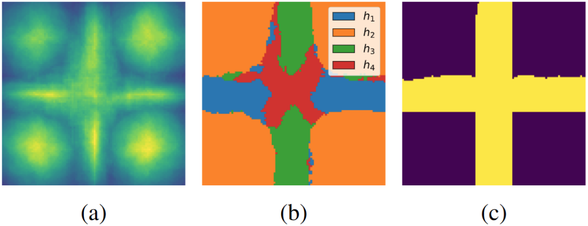

(a) (b) (c)

(a) Heatmap of the maximum confidence $\max_i \hat{c}_{S,P_i}(p)$ at each point p, indicating how well the best-matching pattern $P_{\bar{\imath}}$ locally agrees with the original image.
(b) Map of indices $\bar{\imath}(p) = \operatorname{argmax}_i \hat{c}_{S,P_i}(x,y)$, showing which pattern $P_{\bar{\imath}}$ achieves the highest confidence at each point.
(c) Reconstructed output $\varphi_{\text{rec}}(p)$ produced by the GENEO-based pattern matching.

Explainable signal reconstruction

In our numerical experiments, we focus on reconstructing the signal over a two-dimensional area. We generated signal measurements with the Sionna RT ray-tracing simulator, using its built-in outdoor urban scenarios for Munich and Paris. Each scenario is discretized into a $L \times L$ grid of $1\,\mathrm{m}^2$ pixels $\{p_j\}_{j=1}^{L^2}$.

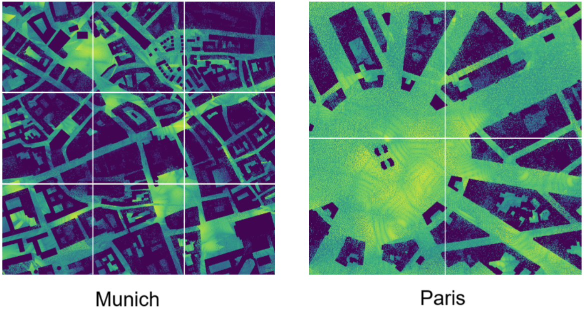
Munich

Paris

Explainable signal reconstruction

Each $L \times L$ signal map φ is partitioned using circular tiling, covering the full 270×270 m^2 area and keeping minimal overlap between circles. This process yields 67 patterns per area. In this paper, each χ_i is modelled as a fixed, rotation-invariant mask with a support of diameter of 45 m:

$$\chi(x, y) = \begin{cases} 1 & \text{if } x^2 + y^2 \leq 22^2, \\ 0 & \text{otherwise.} \end{cases}$$

Each pattern is therefore a pair $P_i = \left(h_i, \chi\right)$, with $i = 1, \ldots, 67$. To enrich the library, we generate 24 rotated variants of every pattern by using 15° increments, such that pattern $h_{\frac{i\,\theta}{15}+i}$ is equal to pattern h_i rotated by θ degrees. Thus each area contributes $67 \times 24 = 1608$ patterns.

Explainable signal reconstruction

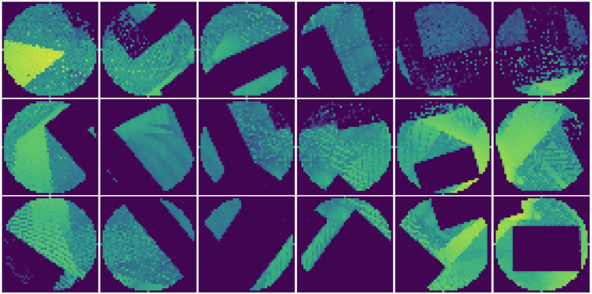

Examples of patterns extracted from Munich signal images.

Explainable signal reconstruction

We sample only $M\%$ of the signal values, selected uniformly at random, and mark those as "known." We define

$$\hat{\psi} : \{p_j\}_{j=1}^{L^2} \to \{0,1\} \ ,$$
$$\sum_j \hat{\psi}(p_j) = \frac{M}{100} L^2 \ ,$$

where $\hat{\psi}(p_j) = 1$ if the measurement $\hat{\varphi}(p_j)$ is retained, and 0 otherwise. We consider $M \in \{1,2,3\}$.

Explainable signal reconstruction

To simulate noisy or corrupted measurements, we replace a further $Q\%$ of those retained pixels with uniform random noise in $[0,1]$. Concretely, out of the $\frac{M}{100}L^2$ locations with $\hat{\psi}=1$, we choose $Q\%$ uniformly at random and for each such pixel $\bar{p}_j$ set

$$\hat{\varphi}(\bar{p}_j) \leftarrow u \sim \mathscr{U}(0,1) ,$$

with $Q \in \{15,30\}$.
This process yields incomplete and noisy observations of the true normalized signal map, reflecting realistic measurement limitations.

Explainable signal reconstruction

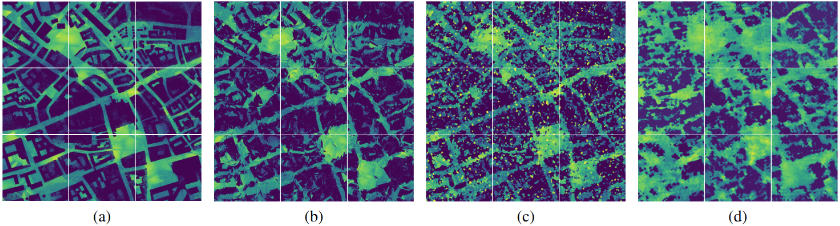

Comparison of reconstruction methods on the Munich scenario, with only $M = 3\%$ of signals known, whose $Q = 15\%$ of them is featuring errors. (a) ground truth, (b) GENEO, (c) 1-KNN, and (d) U-Net.

Explainable signal reconstruction

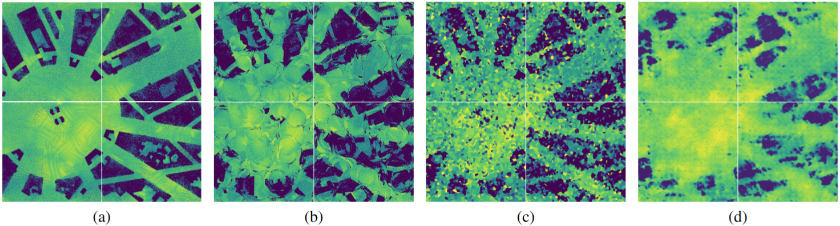

Comparison of reconstruction methods on the Paris scenario, with only $M = 3\%$ of signals known, whose $Q = 15\%$ of them is featuring errors. (a) ground truth, (b) GENEO, (c) 1-KNN, and (d) U-Net.

Explainable signal reconstruction

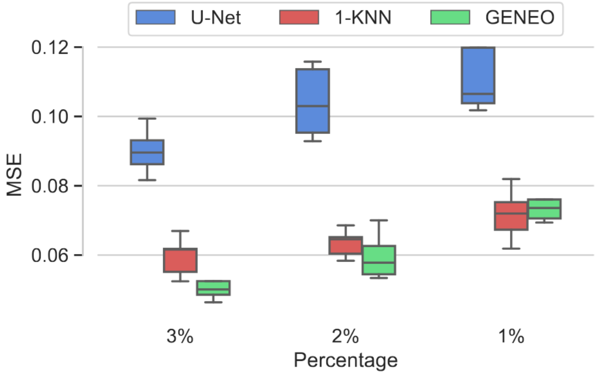

MSE achieved by GENEO, U-Net, and 1-KNN for normalized signal reconstruction in Munich scenario, where $M \in \{1, 2, 3\}$ and $Q = 15$.

Explainable signal reconstruction

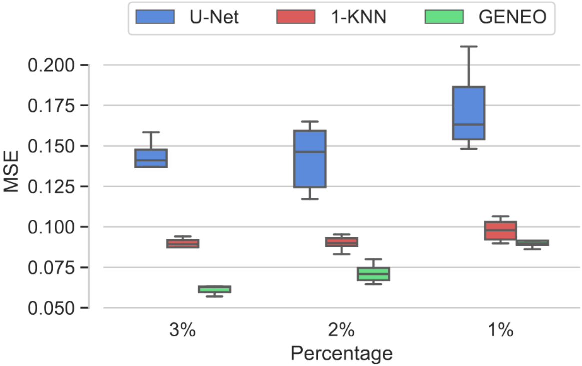

MSE achieved by GENEO, U-Net, and 1-KNN for normalized signal reconstruction in Munich scenario, where $M \in \{1, 2, 3\}$ and $Q = 30$.

Explainable signal reconstruction

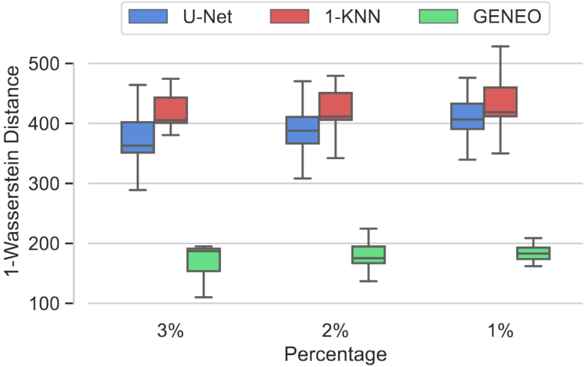

1-Wasserstein achieved by GENEO, U-Net, and 1-KNN for normalized signal reconstruction in Munich scenario, where $M \in \{1, 2, 3\}$ and $Q = 15$.

Explainable signal reconstruction

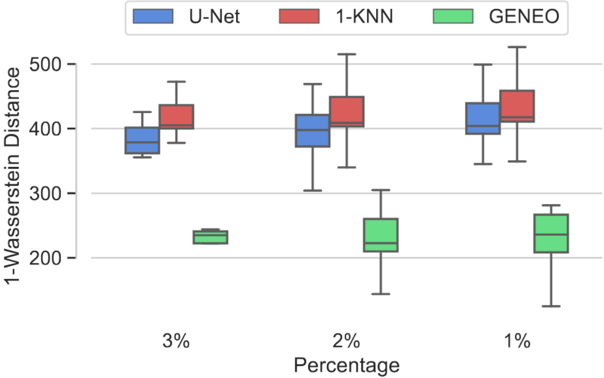

1-Wasserstein achieved by GENEO, U-Net, and 1-KNN for normalized signal reconstruction in Munich scenario, where $M \in \{1,2,3\}$ and $Q = 30$.

Explainable signal reconstruction

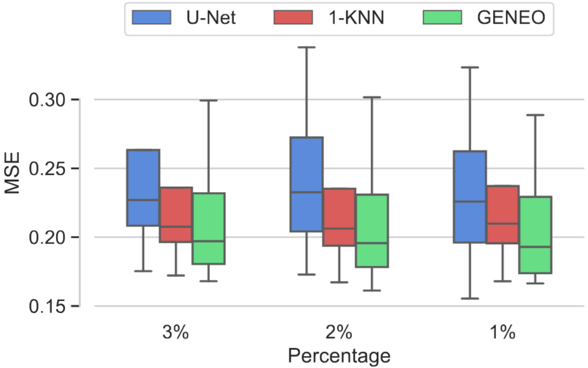

MSE achieved by GENEO, U-Net, and 1-KNN for normalized signal reconstruction in Paris scenario, where $M \in \{1, 2, 3\}$ and $Q = 15$.

Explainable signal reconstruction

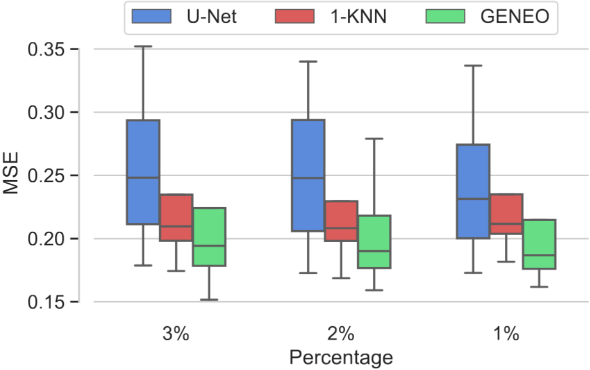

MSE achieved by GENEO, U-Net, and 1-KNN for normalized signal reconstruction in Paris scenario, where $M \in \{1, 2, 3\}$ and $Q = 30$.

Explainable signal reconstruction

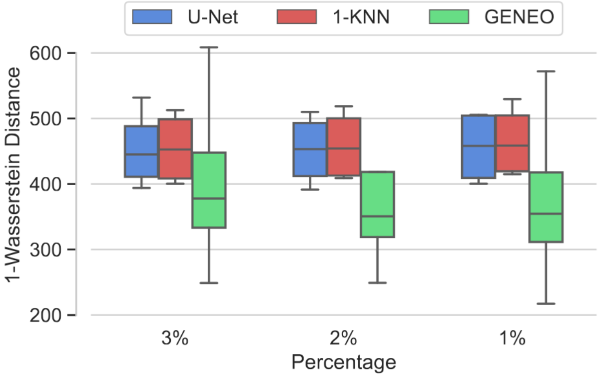

1-Wasserstein achieved by GENEO, U-Net, and 1-KNN for normalized signal reconstruction in Paris scenario, where $M \in \{1, 2, 3\}$ and $Q = 15$.

Explainable signal reconstruction

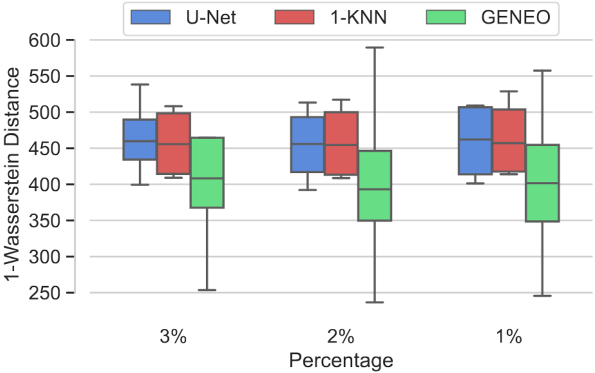

1-Wasserstein achieved by GENEO, U-Net, and 1-KNN for normalized signal reconstruction in Paris scenario, where $M \in \{1, 2, 3\}$ and $Q = 30$.

Explainable signal reconstruction

MSE Performance: For $Q = 15$, GENEO (green) achieves the lowest reconstruction error in terms of MSE across all sampling ratios M. Under heavier corruption (Q=30), GENEO's advantage persists: it consistently outperforms 1-KNN and U-Net, with the largest margin at M=1.

1-Wasserstein Performance: For $Q = 15$, GENEO again yields substantially lower topological error, improving over 1-KNN and U-Net across all M. For Q=30, GENEO maintains its lead in topological fidelity, while both baselines exhibit worse performance in terms of signal reconstruction.

TAKE-AWAY MESSAGE

To sum up, GENEOs are novel mathematical tools designed to approximate equivariant neural networks using a compositional approach. GENEOs are generally interpretable, making them potentially beneficial for explainable artificial intelligence (XAI).


Thanks for your attention!
LAKE COMO SCHOOL OF ADVANCED STUDIES
REASON IS CHOICE