



OPEN ACCESS

EDITED BY
Fabio Anselmi,
University of Trieste, Italy

REVIEWED BY
Maurizio Parton,
G. d'Annunzio University of Chieti and Pescara,
Italy
Filippo Maggioli,
Sapienza University of Rome, Italy

*CORRESPONDENCE
Nicola Quercioli
✉ nicola.quercioli@gmail.com

RECEIVED 04 August 2023
ACCEPTED 27 November 2023
PUBLISHED 21 December 2023

CITATION
Ferrari L, Frosini P, Quercioli N and Tombari F
(2023) A topological model for partial
equivariance in deep learning and data analysis.
Front. Artif. Intell. 6:1272619.
doi: 10.3389/frai.2023.1272619

COPYRIGHT
© 2023 Ferrari, Frosini, Quercioli and Tombari.
This is an open-access article distributed under
the terms of the [Creative Commons Attribution
License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or
reproduction in other forums is permitted,
provided the original author(s) and the
copyright owner(s) are credited and that the
original publication in this journal is cited, in
accordance with accepted academic practice.
No use, distribution or reproduction is
permitted which does not comply with these
terms.

A topological model for partial equivariance in deep learning and data analysis

Lucia Ferrari¹, Patrizio Frosini¹, Nicola Quercioli^{2*} and
Francesca Tombari³

¹Department of Mathematics, University of Bologna, Bologna, Italy, ²Department of Electrical, Electronic, and Information Engineering (DEI) and WiLab-National Laboratory for Wireless Communications, National Inter-University Consortium for Telecommunications (CNIT), University of Bologna, Bologna, Italy, ³Department of Mathematics, Royal Institute of Technology (KTH), Stockholm, Sweden

In this article, we propose a topological model to encode partial equivariance in neural networks. To this end, we introduce a class of operators, called P-GENEOs, that change data expressed by measurements, respecting the action of certain sets of transformations, in a non-expansive way. If the set of transformations acting is a group, we obtain the so-called GENEOS. We then study the spaces of measurements, whose domains are subjected to the action of certain self-maps and the space of P-GENEOs between these spaces. We define pseudo-metrics on them and show some properties of the resulting spaces. In particular, we show how such spaces have convenient approximation and convexity properties.

KEYWORDS

partial-equivariant neural network, P-GENEO, pseudo-metric space, compactness, convexity

1 Introduction

Over the past decade, several geometric techniques have been incorporated into Deep Learning (DL), giving rise to the new field of Geometric Deep Learning (GDL) (Cohen and Welling, 2016; Masci et al., 2016; Bronstein et al., 2017). This geometric approach to deep learning is exploited with a dual purpose. On one hand, geometry provides a common mathematical framework to study neural network architectures. On the other hand, a geometric bias, based on prior knowledge of the data set, can be incorporated into DL models. In this second case, GDL models take advantage of the symmetries imposed by an observer, which encode and elaborate the data. The general blueprint of many deep learning architectures is modeled by group equivariance to encode such properties. If we consider measurements on a data set and a group encoding their symmetries, i.e., transformations taking admissible measurements (for example, rotation or translation of an image), the group equivariance is the property guaranteeing that such symmetries are preserved after applying an operator (e.g., a layer in a neural network) on the observed data. In particular, let us assume that the input measurements Φ , the output measurements Ψ and, respectively, their symmetry groups G and H are given. Then the agent $F: \Phi \rightarrow \Psi$ is T -equivariant if $F(\varphi g) = F(\varphi)T(g)$, for any φ in Φ and any g in G , where T is a group homomorphism from G to H . In the theory of Group Equivariant Non-Expansive Operators (GENEOs) (Camporesi et al., 2018; Bergomi et al., 2019; Cascarano et al., 2021; Bocchi et al., 2022, 2023; Conti et al., 2022; Frosini et al., 2023; Micheletti, 2023), as in many other GDL models, the collection of all symmetries is represented by a group, but in some applications, the group axioms do not necessarily hold since real-world data rarely follow strict mathematical symmetries due to noise, incompleteness, or symmetry-breaking features. As an example, we can consider a data set that contains images of digits and the group of rotations as the group acting on it.

Rotating an image of the digit “6” by a straight angle returns an image that the user would most likely interpret as “9”. At the same time, we may want to be able to rotate the digit “6” by small angles while preserving its meaning (see Figure 1).

It is then desirable to extend the theory of GENEOS by relaxing the hypotheses on the sets of transformations. The main aim of this article is to give a generalization of the results obtained for GENEOS to a new mathematical framework, where the property of equivariance is maintained only for some transformations of the measurements, encoding a partial equivariance with respect to the action of the group of all transformations. To this end, we introduce the concept of Partial Group Equivariant Non-Expansive Operator (P-GENEO).

In this new model, there are some substantial differences with respect to the theory of GENEOS:

1. The user chooses two sets of measurements in input: the one containing the original measurements and another set that encloses the admissible variations of such measurements, defined in the same domain. For example, in the case where the function that represents the digit “6” is being observed, we define an initial space that contains this function and another space that contains certain small rotations of “6” but excludes all the others.
2. Instead of considering a group of transformations, we consider a set containing only those that do not change the meaning of our data, i.e., only those associating with each original measurement another one inside the set of its admissible variations. Therefore, by choosing the initial spaces, the user defines also which transformations of the data set, given by right composition, are admissible and which ones are not.
3. We define partial GENEOS, or P-GENEOS, as a generalization of GENEOS. P-GENEOS are operators that respect the two sets of measurements in input and the set of transformations relating them. The term “partial” refers to the fact that the set of transformations does not necessarily need to be a group.

With these assumptions in mind, we will extend the results proven in the study by Bergomi et al. (2019) and Quercioli (2021a) for GENEOS. We will define suitable pseudo-metrics on the spaces of measurements, the set of transformations, and the set of non-expansive operators. Grounding on their induced topological structures, we prove compactness and convexity of the space of P-GENEOS under the assumption that the function spaces are compact and convex. These are useful properties from a computational point of view. For example, compactness guarantees that the space can be approximated by a finite set. Moreover, convexity allows us to take the convex combination of P-GENEOS in order to generate new ones.

2 Related work

The main motivation for our study is that observed data rarely follow strict mathematical symmetries. This may be due, for example, to the presence of noise in data measurements. The idea of relaxing the hypothesis of equivariance in GDL and data analysis is not novel, as it is shown by the recent increase in the number of publications in this area (see, for example, Weiler and Cesa, 2019;

Finzi et al., 2021; Romero and Lohit, 2022; van der Ouderaa et al., 2022; Wang et al., 2022; Chachlski et al., 2023).

We identify two main ways to transform data *via* operators that are not strictly equivariant due to the lack of strict symmetries of the measurements. On one hand, one could define *approximately equivariant* operator. These are operators for which equivariance holds up to small perturbation. In this case, given two groups, G and H acting on the spaces of measurements Φ and Ψ , respectively, and a homomorphism between them, $T: G \rightarrow H$, we say that $F: \Phi \rightarrow \Psi$ is ε -equivariant if, for any $g \in G$ and $\varphi \in \Phi$, $\|F(\varphi g) - F(\varphi)T(g)\|_\infty \leq \varepsilon$. Alternatively, when defining operators transforming the measurements of certain data sets, equivariance may be substituted by *partial equivariance*. In this case, equivariance is guaranteed for a subset of the groups acting on the space of measurements, with no guarantees for this subset to be a subgroup. Among the previously cited articles about relaxing the property of equivariance in DL, the approach by Finzi et al. (2021) is closer to an approximate equivariance model. Here, the authors use a Bayesian approach to introduce an inductive bias in their network that is sensitive to approximate symmetry. The authors of Romero and Lohit (2022) utilize a partial equivariance approach, where a probability distribution is defined and associated with each group convolutional layer of the architecture, and the parameters defining it are either learnt, to achieve equivariance, or partially learnt, to achieve partial equivariance. The importance of choosing equivariance with respect to different acting groups on each layer of the CNN was actually first observed in the study by Weiler and Cesa (2019) for the group of Euclidean isometries in \mathbb{R}^2 .

The point of view of this article is closer to the latter. Our P-GENEOS are indeed operators that preserve the action of certain sets ruling the admissibility of the transformations of the measurements of our data sets. Moreover, non-expansiveness plays a crucial role in our model. This is, in fact, the feature allowing us to obtain compactness and approximability in the space of operators, distinguishing our model from the existing literature on equivariant machine learning.

3 Mathematical setting

3.1 Data sets and operations

Consider a set X and the normed vector space $(\mathbb{R}_b^X, \|\cdot\|_\infty)$, where \mathbb{R}_b^X is the space of all bounded real-valued functions on X and $\|\cdot\|_\infty$ is the usual uniform norm, i.e., for any $f \in \mathbb{R}_b^X$, $\|f\|_\infty := \sup_{x \in X} |f(x)|$. On the set X , the space of transformations is given by elements of $\text{Aut}(X)$, i.e., the group of bijections from X to itself. Then, we can consider the right group action \mathcal{R} defined as follows (we represent composition as a juxtaposition of functions):

$$\mathcal{R}: \mathbb{R}_b^X \times \text{Aut}(X) \rightarrow \mathbb{R}_b^X, \quad (\varphi, s) \mapsto \varphi s.$$

Remark 3.1. For every $s \in \text{Aut}(X)$, the map $\mathcal{R}_s: \mathbb{R}_b^X \rightarrow \mathbb{R}_b^X$, with $\mathcal{R}_s(\varphi) := \varphi s$ preserves the distances. In fact, for any $\varphi_1, \varphi_2 \in \mathbb{R}_b^X$,




FIGURE 1 Example of a symmetry breaking feature. Applying a rotation g of $\pi/4$, the digit “6” preserves its meaning (**left**). The rotation g^4 of π is, instead, not admissible, since it transforms the digit “6” into the digit “9” (**right**).

by bijectivity of s , we have that

$$\begin{aligned} \|\mathcal{R}_s(\varphi_1) - \mathcal{R}_s(\varphi_2)\|_\infty &= \sup_{x \in X} |\varphi_1 s(x) - \varphi_2 s(x)| \\ &= \sup_{y \in X} |\varphi_1(y) - \varphi_2(y)| \\ &= \|\varphi_1 - \varphi_2\|_\infty. \end{aligned}$$

In our model, our data sets are represented as two sets Φ and Φ' of bounded real-valued measurements on X . In particular, X represents the space where the measurements can be made, Φ is the space of permissible measurements, and Φ' is a space which Φ can be transformed into, without changing the interpretation of its measurements after a transformation is applied. In other words, we want to be able to apply some admissible transformations on the space X so that the resulting changes in the measurements in Φ are contained in the space Φ' . Thus, in our model, we consider operations on X in the following way:

Definition 3.2. A (Φ, Φ') -operation is an element s of $\text{Aut}(X)$ such that, for any measurement $\varphi \in \Phi$, the composition φs belongs to Φ' . The set of all (Φ, Φ') operations is denoted by $\text{Aut}_{\Phi, \Phi'}(X)$.

Remark 3.3. We can observe that the identity function id_X is an element of $\text{Aut}_{\Phi, \Phi'}(X)$ if $\Phi \subseteq \Phi'$.

For any $s \in \text{Aut}_{\Phi, \Phi'}(X)$, the restriction to $\Phi \times \text{Aut}_{\Phi, \Phi'}(X)$ of the map \mathcal{R}_s takes values in Φ' since $\mathcal{R}_s(\varphi) := \varphi s \in \Phi'$ for any $\varphi \in \Phi$. We can consider the restriction of the map \mathcal{R} (for simplicity, we will continue to use the same symbol to denote this restriction):

$$\mathcal{R}: \Phi \times \text{Aut}_{\Phi, \Phi'}(X) \rightarrow \Phi', \quad (\varphi, s) \mapsto \varphi s$$

where $\mathcal{R}(\varphi, s) = \mathcal{R}_s(\varphi)$, for every $s \in \text{Aut}_{\Phi, \Phi'}(X)$ and every $\varphi \in \Phi$.

Definition 3.4. Let X be a set. A **perception triple** is a triple (Φ, Φ', S) with $\Phi, \Phi' \subseteq \mathbb{R}_b^X$ and $S \subseteq \text{Aut}_{\Phi, \Phi'}(X)$. The set X is called the **domain** of the perception triple and is denoted by $\text{dom}(\Phi, \Phi', S)$.

Example 3.5. Given $X = \mathbb{R}^2$, consider two rectangles R and R' in X . Assume $\Phi := \{\varphi: X \rightarrow [0, 1]: \text{supp}(\varphi) \subseteq R\}$ and $\Phi' := \{\varphi': X \rightarrow [0, 1]: \text{supp}(\varphi') \subseteq R'\}$. We recall that, if we consider a function $f: X \rightarrow \mathbb{R}$, the *support* of f is the set of points in the domain, where the function does not vanish, i.e., $\text{supp}(f) = \{x \in X \mid f(x) \neq 0\}$. Consider S as the set of translations that bring R into R' . The triple (Φ, Φ', S) is a perception triple. If Φ represents a set of gray level images, S determines which translations can be applied to our pictures.

3.2 Pseudo-metrics on data sets

In our model, considering a generic set X , data are represented by a space $\Omega \subseteq \mathbb{R}_b^X$ of bounded real-valued functions. We endow the real line \mathbb{R} with the usual Euclidean metric and the space X with an extended pseudo-metric induced by Ω :

$$D_X^\Omega(x_1, x_2) = \sup_{\omega \in \Omega} |\omega(x_1) - \omega(x_2)|$$

for every $x_1, x_2 \in X$. The choice of this pseudo-metric over X means that two points can only be distinguished if they assume different values for some measurements. For example, if Φ contains only a constant function and X contains at least two points, the distance between any two points of X is always null.

The pseudo-metric space $X_\Omega := (X, D_X^\Omega)$ can be considered as a topological space with the basis

$$\mathcal{B}_\Omega = \{B_\Omega(x_0, r)\}_{x_0 \in X, r \in \mathbb{R}^+} = \{\{x \in X : D_X^\Omega(x, x_0) < r\}\}_{x_0 \in X, r \in \mathbb{R}^+},$$

and the induced topology is denoted by τ_Ω . The reason for considering a topological space X , rather than just a set, follows from the need of formalizing the assumption that data are stable under small perturbations.

Remark 3.6. In our case, there are two collections of functions Φ and Φ' in \mathbb{R}_b^X representing our data, both of which induce a topology on X . Hence, in the model, we consider two pseudo-metric spaces X_Φ and $X_{\Phi'}$ with the same underlying set X . If $\Phi \subseteq \Phi' \subseteq \mathbb{R}_b^X$, the topologies τ_Φ and $\tau_{\Phi'}$ are comparable and, in particular, $\tau_{\Phi'}$ is finer than τ_Φ .

Now, given a set $\Omega \subseteq \mathbb{R}_b^X$, we will prove a result about the compactness of the pseudo-metric space X_Ω . Before proceeding, let us recall the following lemma (e.g., see [Gaal, 1964](#)):

Lemma 3.7. Let (P, d) be a pseudo-metric space. The following conditions are equivalent:

1. P is totally bounded;
2. Every sequence in P admits a Cauchy subsequence.

Theorem 3.8. If Ω is totally bounded, X_Ω is totally bounded.

Proof: By Lemma 3.7, it will suffice to prove that every sequence in X admits a Cauchy subsequence with respect to the pseudo-metric D_X^Ω . A sequence $(x_i)_{i \in \mathbb{N}}$ in X_Ω is considered and a real

number $\varepsilon > 0$ is taken. Since Ω is totally bounded, we can find a finite subset $\Omega_\varepsilon = \{\omega_1, \dots, \omega_n\}$ such that for every $\omega \in \Omega$, there exists $\omega_r \in \Omega$ for which $\|\omega - \omega_r\|_\infty < \varepsilon$. We can consider now the real sequence $(\omega_1(x_i))_{i \in \mathbb{N}}$, which is bounded since $\omega_1 \in \mathbb{R}_b^X$. From Bolzano-Weierstrass Theorem, it follows that we can extract a convergent subsequence $(\omega_1(x_{i_h}))_{h \in \mathbb{N}}$. Again, we can extract from $(\omega_2(x_{i_h}))_{h \in \mathbb{N}}$ another convergent subsequence $(\omega_2(x_{i_{h_t}}))_{t \in \mathbb{N}}$. Repeating the process, we are able to extract a subsequence of $(x_i)_{i \in \mathbb{N}}$, that for simplicity of notation we can indicate as $(x_{i_j})_{j \in \mathbb{N}}$, such that $(\omega_k(x_{i_j}))_{j \in \mathbb{N}}$ is a convergent subsequence in \mathbb{R} , and hence a Cauchy sequence in \mathbb{R} , for every $k \in \{1, \dots, n\}$. By construction, Ω_ε is finite, then we can find an index \bar{j} such that for any $k \in \{1, \dots, n\}$

$$|\omega_k(x_{i_\ell}) - \omega_k(x_{i_m})| \leq \varepsilon, \quad \text{for every } \ell, m \geq \bar{j}.$$

Furthermore, we have that, for any $\omega \in \Omega$, any $\omega_k \in \Omega_\varepsilon$, and any $\ell, m \in \mathbb{N}$

$$\begin{aligned} |\omega(x_{i_\ell}) - \omega(x_{i_m})| &\leq |\omega(x_{i_\ell}) - \omega_k(x_{i_\ell})| + \\ &\quad |\omega_k(x_{i_\ell}) - \omega_k(x_{i_m})| + |\omega_k(x_{i_m}) - \omega(x_{i_m})| \\ &\leq \|\omega - \omega_k\|_\infty + |\omega_k(x_{i_\ell}) - \omega_k(x_{i_m})| + \\ &\quad \|\omega_k - \omega\|_\infty. \end{aligned}$$

We observe that the choice of \bar{j} depends only on ε and Ω_ε not on k . Then, choosing a $\omega_k \in \Omega_\varepsilon$ such that $\|\omega_k - \omega\|_\infty < \varepsilon$, we get $\|\omega(x_{i_\ell}) - \omega(x_{i_m})\|_\infty < 3\varepsilon$ for every $\omega \in \Omega$ and every $\ell, m \geq \bar{j}$. Then,

$$D_X^\Omega(x_{i_\ell}, x_{i_m}) = \sup_{\omega \in \Omega} |\omega(x_{i_\ell}) - \omega(x_{i_m})| < 3\varepsilon \quad \text{for every } \ell, m \geq \bar{j}.$$

Then $(x_{i_j})_{j \in \mathbb{N}}$ is a Cauchy sequence in X_Ω . For Lemma 3.7 the statement holds.

Corollary 3.9. *If Ω is totally bounded and X_Ω is complete, X_Ω is compact.*

Proof: From Theorem 3.8, we have that X_Ω is totally bounded, and since by hypothesis it is also complete, it is compact.

Now, we will prove that the choice of the pseudo-metric D_X^Ω on X makes the functions in Ω non-expansive.

Definition 3.10. Two pseudo-metric spaces (P, d_P) and (Q, d_Q) are considered. A **non-expansive** function from (P, d_P) to (Q, d_Q) is a function $f: P \rightarrow Q$ such that $d_Q(f(p_1), f(p_2)) \leq d_P(p_1, p_2)$ for any $p_1, p_2 \in P$.

We denote as $NE(P, Q)$ the space of all non-expansive functions from (P, d_P) to (Q, d_Q) .

Proposition 3.11. $\Omega \subseteq NE(X_\Omega, \mathbb{R})$.

Proof: For any $x_1, x_2 \in X$, we have that

$$|\omega(x_1) - \omega(x_2)| \leq \sup_{\omega \in \Omega} |\omega(x_1) - \omega(x_2)| = D_X^\Omega(x_1, x_2).$$

Then, the topology on X induced by D_X^Ω naturally makes the measurements in Ω continuous. In particular, since the previous results hold for a generic $\Omega \subseteq \mathbb{R}_b^X$, they are also true for Φ and Φ' in our model.

Remark 3.12. Assuming that (Φ, Φ', S) is a perception triple. A function $\varphi' \in \Phi'$ may not be continuous from X_Φ to \mathbb{R} and a function $\varphi \in \Phi$ may not be continuous from $X_{\Phi'}$ to \mathbb{R} . In other words, the topology on X induced by the pseudo-metric of one of the function spaces does not make the functions in the other continuous.

Example 3.13. Assuming $X = \mathbb{R}$, for every $a, b \in \mathbb{R}$ the functions $\varphi_a: X \rightarrow \mathbb{R}$ and $\varphi'_b: X \rightarrow \mathbb{R}$ are defined by setting

$$\varphi_a(x) = \begin{cases} 0 & \text{if } x \geq a \\ 1 & \text{otherwise} \end{cases}, \quad \varphi'_b(x) = \begin{cases} 0 & \text{if } x \leq b \\ 1 & \text{otherwise} \end{cases}.$$

Suppose $\Phi := \{\varphi_a : a \geq 0\}$ and $\Phi' := \{\varphi'_b : b \leq 0\}$. Consider the symmetry with respect to the y-axis, i.e., the map $s(x) = -x$. Surely, $s \in \text{Aut}_{\Phi, \Phi'}(X)$. We can observe that the function $\varphi_1 \in \Phi$ is not continuous from $X_{\Phi'}$ to \mathbb{R} ; indeed $D_X^{\Phi'}(0, 2) = 0$, but $|\varphi_1(0) - \varphi_1(2)| = 1$.

However, if $\Phi \subseteq \Phi'$, we have that the functions in Φ are also continuous on $X_{\Phi'}$, indeed:

Corollary 3.14. *If $\Phi \subseteq \Phi'$, then $\Phi \subseteq NE(X_{\Phi'}, \mathbb{R})$.*

Proof: By Proposition 3.11, the statement trivially holds since $\Phi \subseteq \Phi' \subseteq NE(X_{\Phi'}, \mathbb{R})$.

3.3 Pseudo-metrics on the space of operations

Proposition 3.15. *Every element of $\text{Aut}_{\Phi, \Phi'}(X)$ is non-expansive from $X_{\Phi'}$ to X_Φ .*

Proof: Considering a bijection $s \in \text{Aut}_{\Phi, \Phi'}(X)$ we have that

$$\begin{aligned} D_X^\Phi(s(x_1), s(x_2)) &= \sup_{\varphi \in \Phi} |\varphi(s(x_1)) - \varphi(s(x_2))| \\ &= \sup_{\varphi \in \Phi_s} |\varphi(x_1) - \varphi(x_2)| \\ &\leq \sup_{\varphi' \in \Phi'} |\varphi'(x_1) - \varphi'(x_2)| = D_X^{\Phi'}(x_1, x_2) \end{aligned}$$

for every $x_1, x_2 \in X$, where $\Phi_s = \{\varphi s, \varphi \in \Phi\}$. Then, $s \in NE(X_{\Phi'}, X_\Phi)$ and the statement is proved.

Now, we are ready to put more structure on $\text{Aut}_{\Phi, \Phi'}(X)$. Considering a set $\Omega \subseteq \mathbb{R}_b^X$ of bounded real-valued functions, we can endow the set $\text{Aut}(X)$ with a pseudo-metric inherited from Ω

$$D_{\text{Aut}}^\Omega(s_1, s_2) := \sup_{\omega \in \Omega} \|\omega s_1 - \omega s_2\|_\infty$$

for any s_1, s_2 in $\text{Aut}(X)$.

Remark 3.16. Analogously to what happens in Remark 3.6 for X , the sets Φ and Φ' can endow $\text{Aut}(X)$ with two possibly different pseudo-metrics D_{Aut}^Φ and $D_{\text{Aut}}^{\Phi'}$. In particular, we can consider $\text{Aut}_{\Phi, \Phi'}(X)$ as a pseudo-metric subspace of $\text{Aut}(X)$ with the induced pseudo-metrics.

Remark 3.17. We observe that, for any s_1, s_2 in $\text{Aut}(X)$,

$$\begin{aligned} D_{\text{Aut}}^\Omega(s_1, s_2) &:= \sup_{\omega \in \Omega} \|\omega s_1 - \omega s_2\|_\infty \\ &= \sup_{x \in X} \sup_{\omega \in \Omega} |\omega(s_1(x)) - \omega(s_2(x))| \\ &= \sup_{x \in X} D_X^\Omega(s_1(x), s_2(x)). \end{aligned} \tag{3.3.1}$$

In other words, the pseudo-metric D_{Aut}^Ω , which is based on the action of the elements of $\text{Aut}(X)$ on the set Ω , is exactly the usual uniform pseudo-metric on X_Ω .

3.4 The space of operations

Since we are only interested in transformations of functions in Φ , it would be natural to just endow $\text{Aut}_{\Phi, \Phi'}(X)$ with the pseudo-metric D_{Aut}^Φ . However, it is sometimes necessary to consider the pseudo-metric $D_{\text{Aut}}^{\Phi'}$ in order to guarantee the continuity of the composition of elements in $\text{Aut}_{\Phi, \Phi'}(X)$, whenever it is admissible. Considering two elements s, t in $\text{Aut}_{\Phi, \Phi'}(X)$ such that st is still an element of $\text{Aut}_{\Phi, \Phi'}(X)$, i.e., for every function $\varphi \in \Phi$, we have that $\varphi st \in \Phi'$. Then, for any $\varphi \in \Phi$ we have that

$$\varphi' := \varphi s \in \Phi s \subseteq \Phi', \quad \varphi' t \in \Phi'.$$

Therefore, t is also an element of $\text{Aut}_{\Phi, \Phi'}(X)$. By definition, Φs is contained in Φ' for every $s \in \text{Aut}_{\Phi, \Phi'}(X)$, and this justifies the choice of considering in $\text{Aut}_{\Phi, \Phi'}(X)$ also the pseudo-metric $D_{\text{Aut}}^{\Phi'}$. We have shown, in particular, that if s, t are elements of $\text{Aut}_{\Phi, \Phi'}(X)$ such that st is still an element of $\text{Aut}_{\Phi, \Phi'}(X)$, t is an element of $\text{Aut}_{\Phi, \Phi'}(X)$, which is an implication of the following proposition:

Proposition 3.18. *Let $s, t \in \text{Aut}_{\Phi, \Phi'}(X)$. Then, $st \in \text{Aut}_{\Phi, \Phi'}(X)$ if $t \in \text{Aut}_{\Phi, \Phi'}(X)$.*

Proof: If the composition st belongs to $\text{Aut}_{\Phi, \Phi'}(X)$, we have already proven that $t \in \text{Aut}_{\Phi, \Phi'}(X)$. On the other hand, if $t \in \text{Aut}_{\Phi, \Phi'}(X)$, we have that $\bar{\varphi}t \in \Phi'$ for every $\bar{\varphi} \in \Phi s$. Since $\varphi(st) = (\varphi s)t$, it follows that $\varphi(st) \in \Phi'$ for every $\varphi \in \Phi$. Therefore, $st \in \text{Aut}_{\Phi, \Phi'}(X)$ and the statement is proved.

Remark 3.19. Let $t \in \text{Aut}_{\Phi, \Phi'}(X)$. We can observe that if $s \in \text{Aut}_\Phi(X)$, $\Phi s \subseteq \Phi$ and $st \in \text{Aut}_{\Phi, \Phi'}(X)$.

Lemma 3.20. *Consider $r, s, t \in \text{Aut}(X)$. For any $\Omega \subseteq \mathbb{R}_b^X$, it holds that*

$$D_{\text{Aut}}^\Omega(rt, st) = D_{\text{Aut}}^\Omega(r, s).$$

Proof: Since \mathcal{R}_t preserves the distances, we have that:

$$\begin{aligned} D_{\text{Aut}}^\Omega(rt, st) &:= \sup_{\omega \in \Omega} \|\omega rt - \omega st\|_\infty \\ &= \sup_{\omega \in \Omega} \|\omega r - \omega s\|_\infty \\ &= D_{\text{Aut}}^\Omega(r, s). \end{aligned}$$

Lemma 3.21. *Consider $r, s \in \text{Aut}(X)$ and $t \in \text{Aut}_{\Phi, \Phi'}(X)$. It holds that*

$$D_{\text{Aut}}^\Phi(tr, ts) \leq D_{\text{Aut}}^{\Phi'}(r, s).$$

Proof: Since $\Phi t \subseteq \Phi'$, we have that:

$$\begin{aligned} D_{\text{Aut}}^\Phi(tr, ts) &= \sup_{\varphi \in \Phi} \|\varphi tr - \varphi ts\|_\infty \\ &= \sup_{\varphi' \in \Phi t} \|\varphi' r - \varphi' s\|_\infty \\ &\leq \sup_{\varphi' \in \Phi'} \|\varphi' r - \varphi' s\|_\infty \\ &= D_{\text{Aut}}^{\Phi'}(r, s). \end{aligned}$$

Let Π be the set of all pairs (s, t) such that $s, t, st \in \text{Aut}_{\Phi, \Phi'}(X)$. We endow Π with the pseudo-metric

$$D_\Pi((s_1, t_1), (s_2, t_2)) := D_{\text{Aut}}^\Phi(s_1, s_2) + D_{\text{Aut}}^{\Phi'}(t_1, t_2)$$

and the corresponding topology.

Proposition 3.22. *The function $\circ: \Pi \rightarrow (\text{Aut}_{\Phi, \Phi'}(X), D_{\text{Aut}}^\Phi)$ that maps (s, t) to st is non-expansive and hence continuous.*

Proof: Consider two elements $(s_1, t_1), (s_2, t_2)$ of Π . By Lemma 3.20 and Lemma 3.21,

$$\begin{aligned} D_{\text{Aut}}^\Phi(s_1 t_1, s_2 t_2) &\leq D_{\text{Aut}}^\Phi(s_1 t_1, s_2 t_1) + D_{\text{Aut}}^\Phi(s_2 t_1, s_2 t_2) \\ &\leq D_{\text{Aut}}^\Phi(s_1, s_2) + D_{\text{Aut}}^{\Phi'}(t_1, t_2) \\ &= D_\Pi((s_1, t_1), (s_2, t_2)). \end{aligned}$$

Therefore, the statement is proved.

Let Υ be the set of all s with $s, s^{-1} \in \text{Aut}_{\Phi, \Phi'}(X)$.

Proposition 3.23. *The function $(\cdot)^{-1}: (\Upsilon, D_{\text{Aut}}^{\Phi'}) \rightarrow (\text{Aut}_{\Phi, \Phi'}(X), D_{\text{Aut}}^\Phi)$, that maps s to s^{-1} , is non-expansive, and hence continuous.*

Proof: Consider two bijections $s_1, s_2 \in \Upsilon$. Because of Lemma 3.20 and Lemma 3.21, we obtain that

$$\begin{aligned} D_{\text{Aut}}^\Phi(s_1^{-1}, s_2^{-1}) &= D_{\text{Aut}}^\Phi(s_1^{-1} s_2, s_2^{-1} s_2) \\ &= D_{\text{Aut}}^\Phi(s_1^{-1} s_2, \text{id}_X) \\ &= D_{\text{Aut}}^\Phi(s_1^{-1} s_2, s_1^{-1} s_1) \\ &\leq D_{\text{Aut}}^{\Phi'}(s_2, s_1) = D_{\text{Aut}}^{\Phi'}(s_1, s_2). \end{aligned}$$

We have previously defined the map

$$\mathcal{R}: \Phi \times \text{Aut}_{\Phi, \Phi'}(X) \rightarrow \Phi', \quad (\varphi, s) \mapsto \varphi s$$

where $\mathcal{R}(\Phi, s) = \mathcal{R}_s(\Phi)$, for every $s \in \text{Aut}_{\Phi, \Phi'}(X)$.

Proposition 3.24. *The function \mathcal{R} is continuous by choosing the pseudo-metric D_{Aut}^Φ on $\text{Aut}_{\Phi, \Phi'}(X)$.*

Proof: We have that

$$\begin{aligned} \|\mathcal{R}(\varphi, t) - \mathcal{R}(\bar{\varphi}, s)\|_\infty &= \|\varphi t - \bar{\varphi} s\|_\infty \\ &\leq \|\varphi t - \varphi s\|_\infty + \|\varphi s - \bar{\varphi} s\|_\infty \\ &= \|\varphi t - \varphi s\|_\infty + \|\varphi - \bar{\varphi}\|_\infty \\ &\leq D_{\text{Aut}}^\Phi(t, s) + \|\varphi - \bar{\varphi}\|_\infty \end{aligned}$$

for any $\varphi, \bar{\varphi} \in \Phi$ and any $t, s \in \text{Aut}_{\Phi, \Phi'}(X)$. This proves that \mathcal{R} is continuous.

Now, we can give a result about the compactness of $(\text{Aut}_{\Phi, \Phi'}(X), D_{\text{Aut}}^\Phi)$ under suitable assumptions.

Proposition 3.25. *If Φ and Φ' are totally bounded, $(\text{Aut}_{\Phi, \Phi'}(X), D_{\text{Aut}}^\Phi)$ is totally bounded.*

Proof: Consider a sequence $(s_i)_{i \in \mathbb{N}}$ in $\text{Aut}_{\Phi, \Phi'}(X)$ and a real number $\varepsilon > 0$. Since Φ is totally bounded, we can find a finite subset $\Phi_\varepsilon = \{\varphi_1, \dots, \varphi_n\}$ such that for every $\varphi \in \Phi$, there exists $\varphi_r \in \Phi_\varepsilon$ for which $\|\varphi - \varphi_r\|_\infty < \varepsilon$. Now, consider the sequence $(\varphi_1 s_i)_{i \in \mathbb{N}}$ in Φ' . Since also Φ' is totally bounded, from Lemma 3.7, it follows that we can extract a Cauchy subsequence $(\varphi_1 s_{i_h})_{h \in \mathbb{N}}$. Again, we can extract another Cauchy subsequence $(\varphi_2 s_{i_{h_t}})_{t \in \mathbb{N}}$. Repeating the process for every $k \in \{1, \dots, n\}$, we are able to extract a subsequence of $(s_i)_{i \in \mathbb{N}}$, that for simplicity of notation we can indicate as $(s_j)_{j \in \mathbb{N}}$, such that $(\varphi_k s_{i_j})_{j \in \mathbb{N}}$ is a Cauchy sequence in Φ' for every $k \in \{1, \dots, n\}$.

Since Φ_ε is finite, we can find an index \bar{j} such that for any $k \in \{1, \dots, n\}$

$$\|\varphi_k s_{i_\ell} - \varphi_k s_{i_m}\|_\infty \leq \varepsilon, \quad \text{for every } \ell, m \geq \bar{j}. \tag{3.4.1}$$

Furthermore, we have that for any $\varphi \in \Phi$, any $\varphi_k \in \Phi_\varepsilon$, and any $\ell, m \in \mathbb{N}$

$$\begin{aligned} \|\varphi s_{i_\ell} - \varphi s_{i_m}\|_\infty &\leq \|\varphi s_{i_\ell} - \varphi_k s_{i_\ell}\|_\infty + \|\varphi_k s_{i_\ell} - \varphi_k s_{i_m}\|_\infty + \|\varphi_k s_{i_m} - \varphi s_{i_m}\|_\infty \\ &= \|\varphi - \varphi_k\|_\infty + \|\varphi_k s_{i_\ell} - \varphi_k s_{i_m}\|_\infty + \|\varphi_k - \varphi\|_\infty. \end{aligned}$$

We observe that the choice of \bar{j} in (3.4.1) depends only on ε and Φ_ε not on φ . Then, choosing a $\varphi_k \in \Phi_\varepsilon$ such that $\|\varphi_k - \varphi\|_\infty < \varepsilon$, we get $\|\varphi s_{i_\ell} - \varphi s_{i_m}\|_\infty < 3\varepsilon$ for every $\varphi \in \Phi$ and every $\ell, m \geq \bar{j}$. Hence, for every $\ell, m \in \mathbb{N}$

$$D_{\text{Aut}}^\Phi(s_{i_\ell}, s_{i_m}) = \sup_{\varphi \in \Phi} \|\varphi s_{i_\ell} - \varphi s_{i_m}\|_\infty < 3\varepsilon$$

Therefore, $(s_j)_{j \in \mathbb{N}}$ is a Cauchy sequence in $\text{Aut}_{\Phi, \Phi'}(X)$. For Lemma 3.7, the statement holds.

Corollary 3.26. *Assume that $S \subseteq \text{Aut}_{\Phi, \Phi'}(X)$. If Φ and Φ' are totally bounded and (S, D_{Aut}^Φ) is complete, it is also compact.*

Proof: From Proposition 3.25, we have that S is totally bounded, and since by hypothesis it is also complete, the statement holds.

4 The space of P-GENEOs

In this section, we introduce the concept of Partial Group Equivariant Non-Expansive Operator (P-GENEO). P-GENEOs allow us to transform data sets, preserving symmetries and distances and maintaining the acceptability conditions of the transformations. We will also describe some topological results about the structure of the space of P-GENEOs and some techniques used for defining new P-GENEOs in order to populate the space of P-GENEOs.

Definition 4.1. Let X, Y be sets and $(\Phi, \Phi', S), (\Psi, \Psi', Q)$ be perception triples with domains X and Y , respectively. Consider a triple of functions (F, F', T) with the following properties:

- $F: \Phi \rightarrow \Psi, F': \Phi' \rightarrow \Psi', T: S \rightarrow Q;$

- For any $s, t \in S$ such that $st \in S$ it holds that $T(st) = T(s)T(t);$
- For any $s \in S$ such that $s^{-1} \in S$ it holds that $T(s^{-1}) = T(s)^{-1};$
- (F, F', T) is equivariant, i.e., $F'(\varphi s) = F(\varphi)T(s)$ for every $\varphi \in \Phi, s \in S.$

The triple (F, F', T) is called a **perception map** or a **Partial Group Equivariant Operator (P-GEO)** from (Φ, Φ', S) to (Ψ, Ψ', Q) .

In Remark 3.3, we observed that $\text{id}_X \in \text{Aut}_{\Phi, \Phi'}(X)$ if $\Phi \subseteq \Phi'$. Then, we can consider a perception triple (Φ, Φ', S) with $\Phi \subseteq \Phi'$ and $\text{id}_X \in S \subseteq \text{Aut}_{\Phi, \Phi'}(X)$. Now, we will show how a P-GEO from this perception triple behaves.

Lemma 4.2. *Consider two perception triples (Φ, Φ', S) and (Ψ, Ψ', Q) with domains X and Y , respectively, and with $\text{id}_X \in S \subseteq \text{Aut}_{\Phi, \Phi'}(X)$. Let (F, F', T) be a P-GEO from (Φ, Φ', S) to (Ψ, Ψ', Q) . Then, $\Psi \subseteq \Psi'$ and $\text{id}_Y \in Q \subseteq \text{Aut}_{\Psi, \Psi'}(Y)$.*

Proof: Since (F, F', T) is a P-GEO, by definition, we have that, for any $s, t \in S$ such that $st \in S, T(st) = T(s)T(t)$. Since $\text{id}_X \in S$, then

$$T(\text{id}_X) = T(\text{id}_X \text{id}_X) = T(\text{id}_X)T(\text{id}_X)$$

and hence $T(\text{id}_X) = \text{id}_Y \in Q \subseteq \text{Aut}_{\Psi, \Psi'}(Y)$. Moreover, for Remark 3.3, we have that $\Psi \subseteq \Psi'$.

Proposition 4.3. *Consider two perception triples (Φ, Φ', S) and (Ψ, Ψ', Q) with domains X and Y , respectively, and with $\text{id}_X \in S \subseteq \text{Aut}_{\Phi, \Phi'}(X)$. Let (F, F', T) be a P-GEO from (Φ, Φ', S) to (Ψ, Ψ', Q) . Then $F'|_\Phi = F$.*

Proof: Since (F, F', T) is a P-GEO, it is equivariant, and by Lemma 4.2, we have that

$$F'(\varphi) = F'(\varphi \text{id}_X) = F(\varphi)T(\text{id}_X) = F(\varphi)\text{id}_Y = F(\varphi)$$

for every $\varphi \in \Phi$.

Definition 4.4. Assume that (Φ, Φ', S) and (Ψ, Ψ', Q) are perception triples. If (F, F', T) is a perception map from (Φ, Φ', S) to (Ψ, Ψ', Q) and F, F' are non-expansive, i.e.,

$$\begin{aligned} \|F(\varphi_1) - F(\varphi_2)\|_\infty &\leq \|\varphi_1 - \varphi_2\|_\infty, \\ \|F'(\varphi'_1) - F'(\varphi'_2)\|_\infty &\leq \|\varphi'_1 - \varphi'_2\|_\infty \end{aligned}$$

for every $\varphi_1, \varphi_2 \in \Phi, \varphi'_1, \varphi'_2 \in \Phi', (F, F', T)$ is called a **Partial Group Equivariant Non-Expansive Operator (P-GENEO)**.

In other words, a P-GENEO is a triple (F, F', T) such that F, F' are non-expansive and the following diagram commutes for every $s \in S$:

$$\begin{array}{ccc} \Phi & \xrightarrow{\mathcal{R}_s} & \Phi' \\ F \downarrow & & \downarrow F' \\ \Psi & \xrightarrow{\mathcal{R}_{T(s)}} & \Psi' \end{array}$$




FIGURE 2
Squares used in Example 4.6.

Remark 4.5. We can observe that a GENEIO (see Bergomi et al., 2019) can be represented as a special case of P-GENEIO, considering two perception triples $(\Phi, \Phi', S), (\Psi, \Psi', Q)$ such that $\Phi = \Phi', \Psi = \Psi'$, and the subsets containing the invariant transformations S and Q are groups (and then the map $T: S \rightarrow Q$ is a homomorphism). In this setting, a P-GENEIO (F, F', T) is a triple where the operators F, F' are equal to each other (because of Proposition 4.3), and the map T is a homomorphism. Hence, instead of the triple, we can simply write the pair (F, T) that is a GENEIO.

Considering two perception triples, we typically want to study the space of all P-GENEIOs between them with the map T fixed. Therefore, when the map T is fixed and specified, we will simply consider pairs of operators (F, F') instead of triples (F, F', T) , and we say that (F, F') is a P-GENEIO associated with or with respect to the map T . Moreover, in this case, we indicate the property of equivariance of the triple (F, F', T) writing that the pair (F, F') is T -equivariant.

Example 4.6. Let $X = \mathbb{R}^2$. Take a real number $l > 0$. In X , consider the square $Q_1 := [0, l] \times [0, l]$ and its translation s_a of a vector $a = (a_1, a_2) \in \mathbb{R}^2, Q_1' := [a_1, l + a_1] \times [a_2, l + a_2]$. Analogously, let us consider a real number $0 < \epsilon < l$ and two squares inside Q_1 and $Q_1', Q_2 := [\epsilon, l - \epsilon] \times [\epsilon, l - \epsilon]$ and $Q_2' := [a_1 + \epsilon, l + a_1 - \epsilon] \times [a_2 + \epsilon, l + a_2 - \epsilon]$, as shown in Figure 2.

Consider the following function spaces in \mathbb{R}_b^X :

$$\begin{aligned} \Phi &:= \{\varphi: X \rightarrow \mathbb{R} \mid \text{supp}(\varphi) \subseteq Q_1\} \\ \Phi' &:= \{\varphi': X \rightarrow \mathbb{R} \mid \text{supp}(\varphi') \subseteq Q_1'\} \\ \Psi &:= \{\psi: X \rightarrow \mathbb{R} \mid \text{supp}(\psi) \subseteq Q_2\} \\ \Psi' &:= \{\psi': X \rightarrow \mathbb{R} \mid \text{supp}(\psi') \subseteq Q_2'\}. \end{aligned}$$

Let $S := \{s_a^{-1}\}$, where s_a is the translation by the vector $a = (a_1, a_2)$. The triples (Φ, Φ', S) and (Ψ, Ψ', S) are perception

triples. This example could model the translation of two nested gray-scale images. We want to build now an operator between these images in order to obtain a transformation that commutes with the selected translation. We can consider the triple of functions (F, F', T) defined as follows: $F: \Phi \rightarrow \Psi$ is the operator that maintains the output of functions in Φ at points of Q_2 and set them to zero outside it; analogously, $F': \Phi' \rightarrow \Psi'$ is the operator that maintains the output of functions in Φ' at points of Q_2' and set them to zero outside it and $T = \text{id}_S$. Therefore, the triple (F, F', T) is a P-GENEIO from (Φ, Φ', S) to (Ψ, Ψ', S) . It turns out that the maps are non-expansive, and the equivariance holds

$$F'(\varphi s_a^{-1}) = F(\varphi)T(s_a^{-1}) = F(\varphi)s_a^{-1}$$

for any $\varphi \in \Phi$. From the point of view of application, we are considering two square images and their translations, and we apply an operator that “cuts” the images, taking into account only the part of the image that interests the observer. This example justifies the definition of P-GENEIO as a triple of operators (F, F', T) , without requiring F and F' to be equal in the possibly non-empty intersection of their domains. In fact, if φ is a function contained in $\Phi \cap \Phi'$, its image via F and F' may be different.

4.1 Methods to construct P-GENEIOs

Starting from a finite number of P-GENEIOs, we will illustrate some methods to construct new P-GENEIOs. First of all, the composition of two P-GENEIOs is still a P-GENEIO.

Proposition 4.7. Given two composable P-GENEIOs, $(F_1, F_1', T_1): (\Phi, \Phi', S) \rightarrow (\Psi, \Psi', Q)$ and $(F_2, F_2', T_2): (\Psi, \Psi', Q) \rightarrow (\Omega, \Omega', K)$, their composition defined as

$$(F, F', T) := (F_2 \circ F_1, F_2' \circ F_1', T_2 \circ T_1): (\Phi, \Phi', S) \rightarrow (\Omega, \Omega', K)$$

is a P-GENEIO.

Proof: First, one could easily check that the map $T = T_2 \circ T_1$ respects the second and the third property of Definition 4.1. Therefore, it remains to verify that $F(\Phi) \subseteq \Omega, F'(\Phi') \subseteq \Omega'$ and the properties of equivariance and non-expansiveness are maintained.

1. Since $F_1(\Phi) \subseteq \Psi$ and $F_2(\Psi) \subseteq \Omega$, we have that $F(\Phi) = (F_2 \circ F_1)(\Phi) = F_2(F_1(\Phi)) \subseteq F_2(\Psi) \subseteq \Omega$. Analogously, $F'(\Phi') \subseteq \Omega'$.
2. Since (F_1, F_1', T_1) and (F_2, F_2', T_2) are equivariant, (F, F', T) is equivariant. Indeed, for every $\varphi \in \Phi$, we have that

$$\begin{aligned} F'(\varphi s) &= (F_2' \circ F_1')(\varphi s) = F_2'(F_1'(\varphi s)) \\ &= F_2'(F_1(\varphi)T_1(s)) = F_2(F_1(\varphi))T_2(T_1(s)) \\ &= (F_2 \circ F_1)(\varphi)(T_2 \circ T_1)(s) = F(\varphi)T(s). \end{aligned}$$

3. Since F_1 and F_2 are non-expansive, F is non-expansive; indeed for every $\varphi_1, \varphi_2 \in \Phi$, we have that

$$\begin{aligned} \|F(\varphi_1) - F(\varphi_2)\|_\infty &= \|(F_2 \circ F_1)(\varphi_1) - (F_2 \circ F_1)(\varphi_2)\|_\infty \\ &= \|F_2(F_1(\varphi_1)) - F_2(F_1(\varphi_2))\|_\infty \\ &\leq \|F_1(\varphi_1) - F_1(\varphi_2)\|_\infty \\ &\leq \|\varphi_1 - \varphi_2\|_\infty. \end{aligned}$$

Analogously, F' is non-expansive.

Given a finite number of P-GENEOs with respect to the same map T , we illustrate a general method to construct a new operator as a combination of them. Given two sets X and Y , consider a finite set $\{H_1, \dots, H_n\}$ of functions from $\Omega \subseteq \mathbb{R}_b^X$ to \mathbb{R}_b^Y and a map $\mathcal{L}: \mathbb{R}^n \rightarrow \mathbb{R}$, where \mathbb{R}^n is endowed with the norm $\|(x_1, \dots, x_n)\|_\infty := \max_{1 \leq i \leq n} |x_i|$. We define $\mathcal{L}^*(H_1, \dots, H_n): \Omega \rightarrow \mathbb{R}_b^Y$ as

$$\mathcal{L}^*(H_1, \dots, H_n)(\omega) := [\mathcal{L}(H_1(\omega), \dots, H_n(\omega))],$$

for any $\omega \in \Omega$, where $[\mathcal{L}(H_1(\omega), \dots, H_n(\omega))]: Y \rightarrow \mathbb{R}$ is defined by setting

$$[\mathcal{L}(H_1(\omega), \dots, H_n(\omega))](y) := \mathcal{L}(H_1(\omega)(y), \dots, H_n(\omega)(y))$$

for any $y \in Y$. Now, consider two perception triples (Φ, Φ', S) and (Ψ, Ψ', Q) with domains X and Y , respectively, and a finite set of P-GENEOs $(F_1, F'_1), \dots, (F_n, F'_n)$ between them associated with the map $T: S \rightarrow Q$. We can consider the functions $\mathcal{L}^*(F_1, \dots, F_n): \Phi \rightarrow \mathbb{R}_b^Y$ and $\mathcal{L}^*(F'_1, \dots, F'_n): \Phi' \rightarrow \mathbb{R}_b^Y$, defined as before and state the following result.

Proposition 4.8. Assume that $\mathcal{L}: \mathbb{R}^n \rightarrow \mathbb{R}$ is non-expansive. If $\mathcal{L}^*(F_1, \dots, F_n)(\Phi) \subseteq \Psi$ and $\mathcal{L}^*(F'_1, \dots, F'_n)(\Phi') \subseteq \Psi'$, $(\mathcal{L}^*(F_1, \dots, F_n), \mathcal{L}^*(F'_1, \dots, F'_n))$ is a P-GENEO from (Φ, Φ', S) to (Ψ, Ψ', Q) with respect to T .

Proof: By hypothesis, $\mathcal{L}^*(F_1, \dots, F_n)(\Phi) \subseteq \Psi$ and $\mathcal{L}^*(F'_1, \dots, F'_n)(\Phi') \subseteq \Psi'$, so we just need to verify the properties of equivariance and non-expansiveness.

1. Since $(F_1, F'_1), \dots, (F_n, F'_n)$ are T -equivariant, for any $\varphi \in \Phi$ and any $s \in S$, we have that:

$$\begin{aligned} \mathcal{L}^*(F'_1, \dots, F'_n)(\varphi s) &= [\mathcal{L}(F'_1(\varphi s), \dots, F'_n(\varphi s))] \\ &= [\mathcal{L}(F_1(\varphi)T(s), \dots, F_n(\varphi)T(s))] \\ &= [\mathcal{L}(F_1(\varphi), \dots, F_n(\varphi))]T(s) \\ &= \mathcal{L}^*(F_1, \dots, F_n)(\varphi)T(s). \end{aligned}$$

Therefore, $(\mathcal{L}^*(F_1, \dots, F_n), \mathcal{L}^*(F'_1, \dots, F'_n))$ is T -equivariant.

2. Since F_1, \dots, F_n and \mathcal{L} are non-expansive, for any $\varphi_1, \varphi_2 \in \Phi$, we have that:

$$\begin{aligned} &\|\mathcal{L}^*(F_1, \dots, F_n)(\varphi_1) - \mathcal{L}^*(F_1, \dots, F_n)(\varphi_2)\|_\infty \\ &= \max_{y \in Y} |[\mathcal{L}(F_1(\varphi_1), \dots, F_n(\varphi_1))](y) \\ &\quad - [\mathcal{L}(F_1(\varphi_2), \dots, F_n(\varphi_2))](y)| \\ &= \max_{y \in Y} |\mathcal{L}(F_1(\varphi_1)(y), \dots, F_n(\varphi_1)(y)) \\ &\quad - \mathcal{L}(F_1(\varphi_2)(y), \dots, F_n(\varphi_2)(y))| \\ &\leq \max_{y \in Y} \|(F_1(\varphi_1)(y) - F_1(\varphi_2)(y), \dots, F_n(\varphi_1)(y) - F_n(\varphi_2)(y))\|_\infty \\ &= \max_{y \in Y} \max_{1 \leq i \leq n} |F_i(\varphi_1)(y) - F_i(\varphi_2)(y)| \\ &= \max_{1 \leq i \leq n} \|F_i(\varphi_1) - F_i(\varphi_2)\|_\infty \\ &\leq \|\varphi_1 - \varphi_2\|_\infty. \end{aligned}$$

Hence, $\mathcal{L}^*(F_1, \dots, F_n)$ is non-expansive. Analogously, since F'_1, \dots, F'_n and \mathcal{L} are non-expansive, $\mathcal{L}^*(F'_1, \dots, F'_n)$ is non-expansive.

Therefore, $(\mathcal{L}^*(F_1, \dots, F_n), \mathcal{L}^*(F'_1, \dots, F'_n))$ is a P-GENEO from (Φ, Φ', S) to (Ψ, Ψ', Q) with respect to T .

Remark 4.9. The above result describes a general method to build new P-GENEOs, starting from a finite number of known P-GENEOs via non-expansive maps. Some examples of such non-expansive maps are the maximum function, the power mean, and the convex combination (for further details, see Frosini and Quercioli, 2017; Quercioli, 2021a,b).

4.2 Compactness and convexity of the space of P-GENEOs

Given two perception triples, under some assumptions on the data sets, it is possible to show two useful features in applications: compactness and convexity. These two properties guarantee, on the one hand, that the space of P-GENEOs can be approximated by a finite subset of them, and, on the other hand, that a convex combination of P-GENEOs is again a P-GENEO.

First, we define a metric on the space of P-GENEOs. Let X, Y be sets and consider two sets $\Omega \subseteq \mathbb{R}_b^X, \Delta \subseteq \mathbb{R}_b^Y$, we can define the distance

$$D_{\text{NE}}^\Omega(F_1, F_2) := \sup_{\omega \in \Omega} \|F_1(\omega) - F_2(\omega)\|_\infty$$

for every $F_1, F_2 \in \text{NE}(\Omega, \Delta)$.

The metric $D_{\text{P-GENEO}}$ on the space $\mathcal{F}_T^{\text{all}}$ of all the P-GENEOs between the perception triples (Φ, Φ', S) and (Ψ, Ψ', Q) associated with the map T is defined as

$$\begin{aligned} D_{\text{P-GENEO}}((F_1, F'_1), (F_2, F'_2)) &:= \max\{D_{\text{NE}}^\Phi(F_1, F_2), D_{\text{NE}}^{\Phi'}(F'_1, F'_2)\} \\ &= \max\{\sup_{\varphi \in \Phi} \|F_1(\varphi) - F_2(\varphi)\|_\infty, \sup_{\varphi' \in \Phi'} \|F'_1(\varphi') - F'_2(\varphi')\|_\infty\} \end{aligned}$$

for every $(F_1, F'_1), (F_2, F'_2) \in \mathcal{F}_T^{\text{all}}$.

4.2.1 Compactness

Before proceeding, we need to prove that the following result holds:

Lemma 4.10. If $(P, d_P), (Q, d_Q)$ are compact metric spaces, $\text{NE}(P, Q)$ is compact.

Proof: Theorem 5 in the study by Li et al. (2012) implies that $\text{NE}(P, Q)$ is relatively compact, since it is a equicontinuous space of maps. Hence, it will suffice to show that $\text{NE}(P, Q)$ is closed. Considering a sequence $(F_i)_{i \in \mathbb{N}}$ in $\text{NE}(P, Q)$ such that $\lim_{i \rightarrow \infty} F_i = F$, we have that

$$d_Q(F(p_1), F(p_2)) = \lim_{i \rightarrow \infty} d_Q(F_i(p_1), F_i(p_2)) \leq d_P(p_1, p_2)$$

for every $p_1, p_2 \in P$. Therefore, $F \in \text{NE}(P, Q)$. It follows that $\text{NE}(P, Q)$ is closed.

Consider two perception triples (Φ, Φ', S) and (Ψ, Ψ', Q) , with domains X and Y , respectively, and the space \mathcal{F}_T^{all} of P-GENEOs between them associated with the map $T: S \rightarrow Q$. The following result holds:

Theorem 4.11. *If Φ, Φ', Ψ and Ψ' are compact, \mathcal{F}_T^{all} is compact with respect to the metric $D_{P-GENEO}$.*

Proof: By definition, $\mathcal{F}_T^{all} \subseteq \mathbf{NE}(\Phi, \Psi) \times \mathbf{NE}(\Phi', \Psi')$. Since Φ, Φ', Ψ and Ψ' are compact, for Lemma 4.10, the spaces $\mathbf{NE}(\Phi, \Psi)$ and $\mathbf{NE}(\Phi', \Psi')$ are also compact, and then, by Tychonoff's Theorem, the product $\mathbf{NE}(\Phi, \Psi) \times \mathbf{NE}(\Phi', \Psi')$ is also compact, with respect to the product topology. Hence, to prove our statement, it suffices to show that \mathcal{F}_T^{all} is closed. Let us consider a sequence $((F_i, F'_i))_{i \in \mathbb{N}}$ of P-GENEOs, converging to a pair $(F, F') \in \mathbf{NE}(\Phi, \Psi) \times \mathbf{NE}(\Phi', \Psi')$. Since (F_i, F'_i) is T -equivariant for every $i \in \mathbb{N}$ and the action of Q on Ψ is continuous (see Proposition 3.24), (F, F') belongs to \mathcal{F}_T^{all} . Indeed, we have that

$$F'(\varphi s) = \lim_{i \rightarrow \infty} F'_i(\varphi s) = \lim_{i \rightarrow \infty} F_i(\varphi)T(s) = F(\varphi)T(s)$$

for every $s \in S$ and every $\varphi \in \Phi$. Hence, \mathcal{F}_T^{all} is a closed subset of a compact set, and then, it is also compact.

4.2.2 Convexity

Assume that Ψ, Ψ' are convex. Let $(F_1, F'_1), \dots, (F_n, F'_n) \in \mathcal{F}_T^{all}$ and consider an n -tuple $(a_1, \dots, a_n) \in \mathbb{R}^n$ with $a_i \geq 0$ for every $i \in \{1, \dots, n\}$ and $\sum_{i=1}^n a_i = 1$. We can define two operators $F_\Sigma: \Phi \rightarrow \Psi$ and $F'_\Sigma: \Phi' \rightarrow \Psi'$ as

$$F_\Sigma(\varphi) := \sum_{i=1}^n a_i F_i(\varphi), \text{ and } F'_\Sigma(\varphi') := \sum_{i=1}^n a_i F'_i(\varphi')$$

for every $\varphi \in \Phi, \varphi' \in \Phi'$. We notice that the convexity of Ψ and Ψ' guarantees that F_Σ and F'_Σ are well defined.

Proposition 4.12. *(F_Σ, F'_Σ) belongs to \mathcal{F}_T^{all} .*

Proof: By hypothesis, for every $i \in \{1, \dots, n\}$, (F_i, F'_i) is a perception map, and then:

$$\begin{aligned} F'_\Sigma(\varphi s) &= \sum_{i=1}^n a_i F'_i(\varphi s) = \sum_{i=1}^n a_i (F_i(\varphi)T(s)) \\ &= \left(\sum_{i=1}^n a_i F_i(\varphi) \right) T(s) \\ &= F_\Sigma(\varphi)T(s) \end{aligned}$$

for every $\varphi \in \Phi$ and every $s \in S$. Furthermore, since for every $i \in \{1, \dots, n\}$, $F_i(\Phi) \subseteq \Psi$ and Ψ are convex, also $F_\Sigma(\Phi) \subseteq \Psi$. Analogously, the convexity of Ψ' implies that $F'_\Sigma(\Phi') \subseteq \Psi'$.

Therefore (F_Σ, F'_Σ) is a P-GEO. It remains to show the non-expansiveness of F_Σ and F'_Σ . Since F_i is non-expansive for any i , for every $\varphi_1, \varphi_2 \in \Phi$, we have that

$$\begin{aligned} \|F_\Sigma(\varphi_1) - F_\Sigma(\varphi_2)\|_\infty &= \left\| \sum_{i=1}^n a_i F_i(\varphi_1) - \sum_{i=1}^n a_i F_i(\varphi_2) \right\|_\infty \\ &= \left\| \sum_{i=1}^n a_i (F_i(\varphi_1) - F_i(\varphi_2)) \right\|_\infty \\ &\leq \sum_{i=1}^n |a_i| \|F_i(\varphi_1) - F_i(\varphi_2)\|_\infty \\ &\leq \sum_{i=1}^n |a_i| \|\varphi_1 - \varphi_2\|_\infty = \|\varphi_1 - \varphi_2\|_\infty. \end{aligned}$$

Analogously, since every F'_i is non-expansive, for every $\varphi'_1, \varphi'_2 \in \Phi'$, we have that

$$\|F'_\Sigma(\varphi'_1) - F'_\Sigma(\varphi'_2)\|_\infty \leq \sum_{i=1}^n |a_i| \|\varphi'_1 - \varphi'_2\|_\infty = \|\varphi'_1 - \varphi'_2\|_\infty.$$

Therefore, we have proven that (F_Σ, F'_Σ) is a P-GEO with F_Σ and F'_Σ non-expansive. Hence it is a P-GENEO.

Then, the following result holds:


Corollary 4.13. *If Ψ, Ψ' are convex, the set \mathcal{F}_T^{all} is convex.*

Proof: It is sufficient to apply Proposition 4.12 for $n = 2$ by setting $a_1 = t, a_2 = 1 - t$ for $0 \leq t \leq 1$.

5 P-GENEOs in applications

The importance of equivariance with respect to a group is becoming clear and widespread in many machine learning applications used for drug design, traffic forecasting, object recognition, and detection (see, e.g., Bronstein et al., 2021; Gerken et al., 2023). In some situations, however, requiring equivariance with respect to a whole group could even become an obstacle in the correct learning process of an equivariant neural network. In the following, we describe a possible application to optical character recognition (OCR), in which partial equivariance might be better suited than equivariance. Consider a planar transformation that deforms characters. One may notice that if such transformation is performed too many times, the letter may lose or change its meaning, as shown in Figure 3. Another example is given by a reparameterization of the domain of a sound message. While a limited contraction or dilation of the domain can preserve the meaning attributed to the sound, an iterated application of the same transformation can radically change the perceived message.

Furthermore, experiments performed in the study by Weiler and Cesa (2019) have shown that tuning the level of equivariance in each layer of a neural network may increase the performance of the model. This tuning is, however, performed manually. The successive step, conducted in the study by Romero and Lohit (2022), is to learn the level of equivariance of each layer directly from data, possibly restricting to certain subsets whenever



the full equivariance prevents a good classification performance. The authors of [Romero and Lohit \(2022\)](#) test their result on MNIST. In applications of this type, the use of P-GENEOs could allow partial equivariance to be framed within a precise mathematical model.

6 Conclusion

In this article, we proposed a generalization of some known results in the theory of GENEOS to a new mathematical framework, where the collection of all symmetries is represented by a subset of a group of transformations. We introduced P-GENEOs and showed that they are a generalization of GENEOS. We defined pseudo-metrics on the space of measurements and on the space of P-GENEOs and studied their induced topological structures. Under the assumption that the function spaces are compact and convex, we showed compactness and convexity of the space of P-GENEOs. In particular, compactness guarantees that any operator can be approximated by a finite number of operators belonging to the same space, while convexity allows us to build new P-GENEOs by taking convex combinations of P-GENEOs. Compactness and convexity together ensure that every strictly convex loss function on the space of P-GENEOs admits a unique global minimum. Given a collection of P-GENEOs, we presented a general method to construct new P-GENEOs as combinations of the initial ones.

References

- Bergomi, M. G., Frosini, P., Giorgi, D., and Quercioli, N. (2019). Towards a topological-geometrical theory of group equivariant non-expansive operators for data analysis and machine learning. *Nat. Mach. Intellig.* 1, 423–433. doi: 10.1038/s42256-019-0087-3
- Bocchi, G., Botteghi, S., Brasini, M., Frosini, P., and Quercioli, N. (2023). On the finite representation of linear group equivariant operators via permutants measures. *Ann. Mathem. Artif. Intellig.* 91, 465–487. doi: 10.1007/s10472-022-09830-1
- Bocchi, G., Frosini, P., Micheletti, A., Pedretti, A., Gratteri, C., Lunghini, F., et al. (2022). GENEOnet: a new machine learning paradigm based on Group Equivariant Non-Expansive Operators. An application to protein pocket detection. *arXiv*.
- Bronstein, M. M., Bruna, J., LeCun, Y., Szlam, A., and Vandergheynst, P. (2017). Geometric deep learning: going beyond euclidean data. *IEEE Signal Process. Mag.* 34, 18–42. doi: 10.1109/MSP.2017.2693418
- Bronstein, M. M., Bruna, J., Cohen, T., and Velivcković, P. (2021). Geometric deep learning: grids, groups, graphs, geodesics, and gauges. *arXiv*.
- Camporesi, F., Frosini, P., and Quercioli, N. (2018). “On a new method to build group equivariant operators by means of permutants,” in *Machine Learning and Knowledge Extraction: Second IFIP TC 5, TC 8/WG 8.4, 8.9, TC 12/WG 12.9 International Cross-Domain Conference, CD-MAKE 2018*. Hamburg, Germany: Springer, 265–272.
- Cascarano, P., Frosini, P., Quercioli, N., and Saki, A. (2021). On the geometric and Riemannian structure of the spaces of group equivariant non-expansive operators. *arXiv*.
- Chachlowski, W., De Gregorio, A., Quercioli, N., and Tombari, F. (2023). Symmetries of data sets and functoriality of persistent homology. *Theory Appl. Categor.* 39, 667–686.

Data availability statement

The original contributions presented in the study are included in the article/supplementary material, further inquiries can be directed to the corresponding author.

Author contributions

LF: Writing – original draft. PF: Writing – original draft, Writing – review & editing. NQ: Writing – original draft, Writing – review & editing. FT: Writing – original draft, Writing – review & editing.

Funding

The author(s) declare financial support was received for the research, authorship, and/or publication of this article. This research has been partially supported by INdAM-GNSAGA. FT was supported by the Wallenberg AI, Autonomous System and Software Program (WASP) funded by Knut and Alice Wallenberg Foundation. PF and NQ carried out this work in the framework of the CNIT National Laboratory WiLab and the WiLab-Huawei Joint Innovation Center.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher’s note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

- Cohen, T., and Welling, M. (2016). "Group equivariant convolutional networks," in *International Conference on Machine Learning*. PMLR, 2990–2999. Available online at: jmlr.org
- Conti, F., Frosini, P., and Quercioli, N. (2022). On the construction of group equivariant non-expansive operators via permutants and symmetric functions. *Front. Artif. Intellig.* 5, 16. doi: 10.3389/frai.2022.786091
- Finzi, M., Benton, G., and Wilson, A. G. (2021). "Residual pathway priors for soft equivariance constraints," in *Advances in Neural Information Processing Systems* (Cambridge, MA: The MIT Press), 30037–30049.
- Frosini, P., Gridelli, I., and Pascucci, A. (2023). A probabilistic result on impulsive noise reduction in topological data analysis through group equivariant non-expansive operators. *Entropy* 25, 1150. doi: 10.3390/e25081150
- Frosini, P., and Quercioli, N. (2017). "Some remarks on the algebraic properties of group invariant operators in persistent homology," in *Machine Learning and Knowledge Extraction*, eds. A. Holzinger, P. Kieseberg, A. M. Tjoa, and E. Weippl. Cham: Springer International Publishing, 14–24.
- Gaal, S. (1964). *Point Set Topology*. Amsterdam: Elsevier Science.
- Gerken, J., Aronsson, J., Carlsson, O., Linander, H., Ohlsson, F., Petersson, C., et al. (2023). Geometric deep learning and equivariant neural networks. *Artif. Intell. Rev.* 56, 1–58. doi: 10.1007/s10462-023-10502-7
- Li, R., Zhong, S., and Swartz, C. (2012). An improvement of the Arzel-Ascoli theorem. *Topol. Appl.* 159, 2058–2061. doi: 10.1016/j.topol.2012.01.014
- Masci, J., Rodolà, E., Boscaini, D., Bronstein, M. M., and Li, H. (2016). "Geometric deep learning," in *SIGGRAPH ASIA 2016 Courses*. New York, NY: Association for Computing Machinery, 1–50.
- Micheletti, A. (2023). A new paradigm for artificial intelligence based on group equivariant non-expansive operators. *Eur. Math. Soc. Mag.* 128, 4–12. doi: 10.4171/mag/133
- Quercioli, N. (2021a). *On the Topological Theory of Group Equivariant Non-Expansive Operators* (PhD thesis). Bologna: Alma Mater Studiorum - Università di Bologna. Available online at: <http://amsdottorato.unibo.it/9770/>. (accessed July 12, 2023).
- Quercioli, N. (2021b). "Some new methods to build group equivariant non-expansive operators in TDA," in *Topological Dynamics and Topological Data Analysis*, 229–238.
- Romero, D. W., and Lohit, S. (2022). Learning partial equivariances from data. *Adv. Neural Inform. Proc. Syst.* 35, 36466–36478.
- van der Ouderaa, T., Romero, D. W., and van der Wilk, M. (2022). Relaxing equivariance constraints with non-stationary continuous filters. *Adv. Neural Inform. Proc. Syst.* 35, 33818–33830.
- Wang, R., Walters, R., and Yu, R. (2022). "Approximately equivariant networks for imperfectly symmetric dynamics," in *Proceedings of the 39th International Conference on Machine Learning*. PMLR, 23078–23091. Available online at: jmlr.org
- Weiler, M., and Cesa, G. (2019). "General E(2)-equivariant steerable CNNs," in *Advances in Neural Information Processing Systems*, eds. H. Wallach, H. Larochelle, A. Beygelzimer, F. d Alché-Buc, E. Fox, and R. Garnett. New York: Curran Associates, Inc, 32.