



an Equivariant Neural Network?

Lek-Heng Lim and Bradley J. Nelson

We explain equivariant neural networks, a notion underlying breakthroughs in machine learning from deep convolutional neural networks for computer vision [KSH12] to AlphaFold 2 for protein structure prediction [JEP⁺21], without assuming knowledge of equivariance or neural networks. The basic mathematical ideas are simple but are often obscured by engineering complications that come with practical realizations. We extract and focus on the mathematical aspects, and limit ourselves to a cursory treatment of the engineering issues at the end. We also include some materials with machine learning practitioners in mind.

Let \mathbb{V} and \mathbb{W} be sets, and $f: \mathbb{V} \rightarrow \mathbb{W}$ a function. If a group G acts on both \mathbb{V} and \mathbb{W} , and this action commutes with the function f :

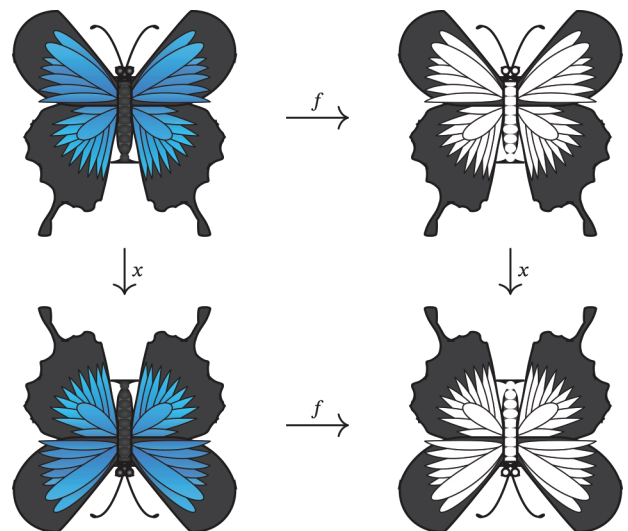
$$f(x \cdot v) = x \cdot f(v) \quad \text{for all } v \in \mathbb{V}, x \in G,$$

then we say that f is G -equivariant. The special case where G acts trivially on \mathbb{W} is called G -invariant. Linear equivariant maps are well-studied in representation theory and continuous equivariant maps are well-studied in topology. The novelty of equivariant neural networks is that they

are usually nonlinear and sometimes discontinuous, even when \mathbb{V} and \mathbb{W} are vector spaces and the actions of G are linear.

Equivariance is ubiquitous in applications where symmetries in the input space \mathbb{V} produce symmetries in the output space \mathbb{W} . We consider a simple example. An image may be regarded as a function $v: \mathbb{R}^2 \rightarrow \mathbb{R}^3$, with each pixel $p = (p_1, p_2) \in \mathbb{R}^2$ assigned some RGB color $(r, g, b) \in \mathbb{R}^3$. A simplifying assumption here is that pixels and colors can take values in a continuum. Let $\mathbb{V} = \mathbb{W}$ be the set of all images. Let the group $G = \{1, x\} \cong \mathbb{Z}/2\mathbb{Z}$ act on \mathbb{V} via top-bottom reflection, i.e., $x \cdot v$ is the image whose value at (p_1, p_2) is $v(p_1, -p_2)$. Let $\sigma: \mathbb{R}^3 \rightarrow \mathbb{R}^3$,

$$\sigma(r, g, b) = \begin{cases} (0, 0, 0) & \text{if } r = g = b = 0, \\ (255, 255, 255) & \text{otherwise.} \end{cases}$$



Lek-Heng Lim is a professor of computational and applied mathematics at the University of Chicago. His email address is lekheng@uchicago.edu.

Bradley J. Nelson is William H. Kruskal Instructor of computational and applied mathematics at the University of Chicago. His email address is bradnelson@protonmail.com.

Both authors are supported by DARPA HR00112190040; LHL is supported by NSF DMS 1854831 and ECCS 2216912; and BN is supported by NSF DMS 1547396.


Communicated by Notices Associate Editor Emilie Purvine.

For permission to reprint this article, please contact: reprint-permission@ams.org.

DOI: <https://doi.org/10.1090/noti2666>

Here $(0, 0, 0)$ and $(255, 255, 255)$ are the RGB encodings for pitch black and pure white respectively. So the map $f : \mathbb{V} \rightarrow \mathbb{V}$, $f(v) = \sigma \circ v$ transforms a color image into a black-and-white image. It does not matter whether we do a top-bottom reflection first or remove color first, the result is always the same, i.e., $f(x \cdot v) = x \cdot f(v)$ for all $v \in \mathbb{V}$. Hence the decoloring map f is G -equivariant.

Our choice of an image with left-right symmetry presents another opportunity to illustrate the notion. If we choose coordinates so that the vertical axis passes through the center of the butterfly image, then as a function $v : \mathbb{R}^2 \rightarrow \mathbb{R}^3$, it is invariant under the action of $H = \{1, s\} \cong \mathbb{Z}/2\mathbb{Z}$ on \mathbb{R}^2 via $s(p_1, p_2) = (-p_1, p_2)$, i.e., $v(s \cdot p) = v(p)$. Note that the G -equivariance of f has nothing to do with this.



A takeaway of these examples is that nonlinear and discontinuous functions may very well be equivariant. However, the best known context for discussing equivariant maps is when f is an *intertwining operator*, i.e., a linear map between vector spaces \mathbb{V} and \mathbb{W} equipped with a linear action of G . In this case, an equivalent formulation of G -equivariance takes the following form: Given linear representations of G on \mathbb{V} and \mathbb{W} , i.e., homomorphisms $\rho_1 : G \rightarrow GL(\mathbb{V})$ and $\rho_2 : G \rightarrow GL(\mathbb{W})$, a linear map $f : \mathbb{V} \rightarrow \mathbb{W}$ is said to be G -equivariant if

$$f(\rho_1(x)v) = \rho_2(x)f(v) \quad \text{for all } v \in \mathbb{V}, x \in G. \quad (1)$$

Intertwining operators preserve eigenvalues and, when G is a Lie group, the action of its Lie algebra, properties that are crucial to their use in physics [BH10].

Nevertheless the restriction to linear maps is unnecessary. The *de Rham problem* asks if $\mathbb{V} = \mathbb{W} = \mathbb{R}^n$ and $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is merely required to be a homeomorphism, then does condition (1) imply that f must be a linear map? De Rham conjectured this to be the case but it was disproved in [CS81], launching a fruitful study of *nonlinear similarity*, i.e., nonlinear homeomorphisms f with

$$f\rho_1(x)f^{-1} = \rho_2(x) \quad \text{for all } x \in G,$$

in algebraic topology and algebraic K-theory. More generally, equivariant continuous maps under continuous group actions have been thoroughly studied in equivariant topology [May96].

An equivariant neural network [CW16] is an equivariant map f constructed from alternately composing equivariant linear maps with nonlinear ones like the decoloring map above. That neural networks can be readily made equivariant is a consequence of two straightforward observations:

- (i) the composition of two G -equivariant functions $f : \mathbb{V} \rightarrow \mathbb{W}$, $g : \mathbb{U} \rightarrow \mathbb{V}$ is G -equivariant;
- (ii) the linear combination of two G -equivariant functions $f, g : \mathbb{V} \rightarrow \mathbb{W}$ is G -equivariant;

even when f, g are nonlinear. Although an equivariant neural network is nonlinear, it uses intertwining operators as building blocks, and (1) plays a key role. In some applications like the wave function on p. 622, the input \mathbb{V} or possibly some hidden layers may not be vector spaces; for simplicity we assume that they are and their G -actions are linear.

In machine learning applications, the map f is learned from data. A major advantage of requiring equivariance in a neural network f is that it allows one to greatly narrow down the search space for the parameters that define f . To demonstrate this, we begin with a simplified case that avoids group representations. A feed-forward neural network is a function $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ obtained by alternately composing affine maps $\alpha_i : \mathbb{R}^n \rightarrow \mathbb{R}^n$, $i = 1, \dots, k$, with a nonlinear function $\sigma : \mathbb{R}^n \rightarrow \mathbb{R}^n$, i.e.,

$$\mathbb{R}^n \xrightarrow{\alpha_1} \mathbb{R}^n \xrightarrow{\sigma} \mathbb{R}^n \xrightarrow{\alpha_2} \mathbb{R}^n \xrightarrow{\sigma} \dots \xrightarrow{\sigma} \mathbb{R}^n \xrightarrow{\alpha_k} \mathbb{R}^n,$$

giving $f = \alpha_k \circ \sigma \circ \alpha_{k-1} \circ \dots \circ \sigma \circ \alpha_2 \circ \sigma \circ \alpha_1$. The *depth*, also known as the number of *layers*, is k and the *width*, also known as the number of *neurons*, is n . The simplifying assumption, which will be dropped later, is that our neural network has constant width throughout all layers. The nonlinear function σ is called an *activation*, with the ReLU (rectified linear unit) function $\sigma(t) := \max(t, 0)$ for $t \in \mathbb{R}$ a standard choice. In a slight abuse of notation, the activation is extended to vector inputs $v = (v_1, \dots, v_n) \in \mathbb{R}^n$ by evaluating coordinatewise

$$\sigma(v) = (\sigma(v_1), \dots, \sigma(v_n)). \quad (2)$$

In this sense, $\sigma : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is called a *pointwise nonlinearity*. The affine function is defined by $\alpha_i(v) = A_i v + b_i$ for some $A_i \in \mathbb{R}^{n \times n}$ called the *weight* matrix and some $b_i \in \mathbb{R}^n$ called the *bias* vector. We do not include a bias b_k in the last layer.

Although convenient, it is somewhat misguided to lump the bias and weight together in an affine function. Each bias b_i is intended to serve as a *threshold* for the activation σ and should be part of it, detached from the weight A_i that transforms the input. If one would like to incorporate translations, one may do so by going up one dimension, observing that $\begin{bmatrix} A & b \\ 0 & 1 \end{bmatrix} \begin{bmatrix} v \\ 1 \end{bmatrix} = \begin{bmatrix} Av+b \\ 1 \end{bmatrix}$. Hence, a better, but mathematically equivalent, description of f would be

as the composition

$$\mathbb{R}^n \xrightarrow{A_1} \mathbb{R}^n \xrightarrow{\sigma_{b_1}} \mathbb{R}^n \xrightarrow{A_2} \mathbb{R}^n \xrightarrow{\sigma_{b_2}} \dots \xrightarrow{\sigma_{b_{k-1}}} \mathbb{R}^n \xrightarrow{A_k} \mathbb{R}^n$$

where we identify $A_i \in \mathbb{R}^{n \times n}$ with the linear operator $\mathbb{R}^n \rightarrow \mathbb{R}^n$, $v \mapsto A_i v$, and for any $b \in \mathbb{R}^n$ we define $\sigma_b : \mathbb{R}^n \rightarrow \mathbb{R}^n$ by $\sigma_b(v) = \sigma(v + b) \in \mathbb{R}^n$. We will drop the composition symbol \circ to avoid clutter and write

$$f = A_k \sigma_{b_{k-1}} A_{k-1} \dots \sigma_{b_2} A_2 \sigma_{b_1} A_1$$

as if it were a product of matrices. For example, with the aforementioned ReLU as σ ,

$$\sigma_{-\theta}(t) = \max(t - \theta, 0) = \begin{cases} t - \theta & t \geq \theta, \\ 0 & t < \theta, \end{cases} \quad (3)$$

and $\theta \in \mathbb{R}$ plays the role of a threshold for activation as was intended in [Ros58, p. 392] and [MP43, p. 120].

A major computational issue with neural networks is the large number of unknown parameters, namely the $kn^2 + (k-1)n$ entries of the weights and biases, that have to be fit with data, especially for wide neural networks where n is large. To get an idea of the numbers involved in realistic situations, n may be on the order of millions of pixels for image-based tasks, whereas k is typically 10 to 50 layers deep. Computational cost aside, one may not have enough data to fit so many parameters. Thus, many successful applications of neural networks require that we identify, based on the problem at hand, an appropriate low-dimensional subset of $\mathbb{R}^{n \times n}$ from which we will find our weights A_1, \dots, A_k . For example, for a signal processing problem, we might restrict A_1, \dots, A_k to be Toeplitz matrices; the convolutional neural networks for image recognition in [KSH12], an article that launched the deep learning revolution, essentially restrict A_1, \dots, A_k to so called block-Toeplitz-Toeplitz-block or BTTB matrices. For 1D inputs with a single *channel*, i.e., inputs from \mathbb{R}^n , a general weight matrix requires n^2 parameters, whereas a Toeplitz one just needs n parameters and an m -banded Toeplitz one requires just m . For 2D inputs with c channels such as color images, i.e., inputs from \mathbb{R}^{cn^2} , a general weight matrix would have required a staggering $c^2 n^4$ parameters, whereas a BTTB one just needs $c^2 n^2$, and an $m \times m$ local BTTB one requires just $c^2 m^2$. These are all simplified versions of the convolutional layers in a convolutional neural networks, which are a quintessential example of equivariant neural networks [CW16], and in fact every equivariant neural network may be regarded as a generalized convolutional neural network in an appropriate sense [KT18].

To see how equivariance naturally restricts the range of possible A_1, \dots, A_k , let $G \subseteq \mathbb{R}^{n \times n}$ be a matrix group. Then $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is G -equivariant if

$$f(Xv) = Xf(v) \quad \text{for all } v \in \mathbb{R}^n, X \in G; \quad (4)$$

and an equivariant neural network is simply a feed-forward neural network $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ that satisfies (4). The

key to its construction is just that

$$\begin{aligned} f(Xv) &= A_k \sigma_{b_{k-1}} A_{k-1} \dots \sigma_{b_2} A_2 \sigma_{b_1} A_1 Xv \\ &= X(X^{-1} A_k X)(X^{-1} \sigma_{b_{k-1}} X)(X^{-1} A_{k-1} X) \\ &\quad \dots (X^{-1} A_2 X)(X^{-1} \sigma_{b_1} X)(X^{-1} A_1 X)v \\ &= X A'_k \sigma'_{b_{k-1}} A'_{k-1} \dots \sigma'_{b_2} A'_2 \sigma'_{b_1} A'_1 v \end{aligned}$$

and the last expression equals $Xf(v)$ if we have

$$A'_i = X^{-1} A_i X = A_i, \quad \sigma'_{b_i} = X^{-1} \sigma_{b_i} X = \sigma_{b_i} \quad (5)$$

for all $i = 1, \dots, k$, and for all $X \in G$. The condition on the right is satisfied by any pointwise nonlinearity that takes the form in (3), i.e., $b_i \in \mathbb{R}^n$ has all coordinates equal to some $\theta \in \mathbb{R}$; we will elaborate on this later. The condition on the left limits the possible weights for f to a (generally) much smaller subspace of matrices that commute with all elements of G . Finding this subspace (in fact a subalgebra) of intertwining operators,

$$\{A \in \mathbb{R}^{n \times n} : AX = XA \text{ for all } X \in G\}, \quad (6)$$

is a well-studied problem in group representation theory; a general purpose approach is to compute the null space of a matrix built from the generators of G and, if continuous, its Lie algebra [FWW21]. We caution the reader that G will generally be a very low-dimensional subset of $\mathbb{R}^{n \times n}$, as will become obvious from our example below in (8). It will be pointless to pick, say, $G = \text{SO}(n)$ as the set in (6) will then be just $\{\lambda I \in \mathbb{R}^{n \times n} : \lambda \in \mathbb{R}\}$, clearly too small to serve as meaningful weights for any neural network. Indeed, G will usually be a homomorphic image of a representation $\rho : G \rightarrow \text{GL}(n)$, i.e., the image $\rho(G)$ will play the role of G in (6). In any case, we will need to bring in group representations to address a different issue.

In general, neural networks have different width n_i in each layer:

$$\begin{aligned} \mathbb{R}^{n_0} \xrightarrow{A_1} \mathbb{R}^{n_1} \xrightarrow{\sigma_{b_1}} \mathbb{R}^{n_1} \xrightarrow{A_2} \mathbb{R}^{n_2} \xrightarrow{\sigma_{b_2}} \dots \\ \dots \xrightarrow{\sigma_{b_{k-1}}} \mathbb{R}^{n_{k-1}} \xrightarrow{A_k} \mathbb{R}^{n_k} \end{aligned}$$

with $A_i \in \mathbb{R}^{n_{i-1} \times n_i}$, $i = 1, \dots, k$, $b_i \in \mathbb{R}^{n_i}$, $i = 1, \dots, k-1$. The simplified case treated above assumes that $n_0 = n_1 = \dots = n_k = n$. It is easy to accommodate this slight complication by introducing group representations to equip every layer with its own homomorphic copy of G . Instead of fixing G to be some subgroup of $\text{GL}(n)$, G may now be any abstract group but we introduce a homomorphism

$$\rho_i : G \rightarrow \text{GL}(n_i), \quad i = 0, 1, \dots, k,$$

in each layer, and replace the equivariant condition (5) with the more general (1), i.e.,

$$\rho_i(x)^{-1} A_i \rho_{i-1}(x) = A_i, \quad \rho_i(x)^{-1} \sigma_{b_i} \rho_i(x) = \sigma_{b_i}$$

or, equivalently,

$$A_i \rho_{i-1}(x) = \rho_i(x) A_i, \quad \sigma_{b_i} \rho_i(x) = \rho_i(x) \sigma_{b_i} \quad (7)$$

for all $x \in G$. In case (7) evokes memories of Schur’s Lemma, we would like to stress that the representations ρ_i are in general very far from being irreducible and that the map σ_{b_i} is nonlinear. Indeed the scenario described by Schur’s Lemma is undesirable for equivariant neural networks: As we pointed out earlier, we do not want to restrict our weight matrices to the form $A_i = \lambda I$ or a direct sum of these.

We summarize our discussion with a formal definition.

Definition 1. Let $A_i \in \mathbb{R}^{n_{i-1} \times n_i}$, $i = 1, \dots, k$, $b_i \in \mathbb{R}^{n_i}$, $i = 1, \dots, k - 1$, and $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ be a continuous function. Let G be a group and $\rho_i : G \rightarrow \text{GL}(n_i)$, $i = 0, \dots, k$, be its representations. The k -layer feed-forward neural network $f : \mathbb{R}^{n_0} \rightarrow \mathbb{R}^{n_k}$ given by

$$f(v) = A_k \sigma_{b_{k-1}} A_{k-1} \cdots \sigma_{b_2} A_2 \sigma_{b_1} A_1 v$$

is a G -equivariant neural network with respect to ρ_0, \dots, ρ_k if (7) holds for all $x \in G$. Here $\sigma_b : \mathbb{R}^{n_i} \rightarrow \mathbb{R}^{n_i}$, $\sigma_b(v) = \sigma(v + b)$, is a pointwise nonlinearity as in (2).

A word of caution is in order here. What we call a neural network [MP43], i.e., the alternate composition of activations with affine maps, is sometimes also called a *multilayer perceptron* [Ros58]; a standard depiction is shown in Figure 1. When it is fit with data, one would invariably feed its output into a *loss function* and that is usually not equivariant; or one might chain together multiple units of multilayer perceptrons into larger frameworks like autoencoders, generative adversarial networks, transformers, etc, that contain other nonequivariant components. In the literature, the term “neural network” sometimes refers to the entire framework collectively. In our article, it just refers to the multilayer perceptron—this is the part that is equivariant.

We will use an insightful toy example as illustration. Let $\mathbb{V} = (\mathbb{R}^3)^m = \mathbb{R}^{3m}$ be the set of possible positions of m unit-weight masses, $\mathbb{W} = \mathbb{R}^3$, and $f : \mathbb{V} \rightarrow \mathbb{W}$ compute the center of mass

$$f(y_1, \dots, y_m) = \frac{1}{m} \sum_{i=1}^m y_i \tag{8}$$

with $y_1, \dots, y_m \in \mathbb{R}^3$. We use the same system of coordinates for each copy of \mathbb{R}^3 in \mathbb{V} and \mathbb{W} . If we work in a different coordinate system, the position of the center of mass remains unchanged but its coordinates will change accordingly. For simplicity, we consider a linear change of coordinates, represented by the action of a matrix $X \in \text{GL}(3)$ on each point in \mathbb{R}^3 . By linearity,

$$f(X(y_1, \dots, y_m)) = \frac{1}{m} \sum_{i=1}^m X y_i = X f(y_1, \dots, y_m),$$

so f is $\text{GL}(3)$ -equivariant. Since each mass has the same unit weight, f is also invariant under permutations of the input points. Let $\pi \in S_m$, which acts on \mathbb{V} via

$\pi(y_1, \dots, y_m) = (y_{\pi(1)}, \dots, y_{\pi(m)})$ and acts trivially on \mathbb{W} via $\pi(y) = y$. As the sum in (8) is permutation invariant,

$$f(\pi(y_1, \dots, y_m)) = \pi f(y_1, \dots, y_m),$$

so f is S_m -invariant. Combining our two group actions, we see that f is $(\text{GL}(3) \times S_m)$ -equivariant. Note that the group here is $G = \text{GL}(3) \times S_m$, which has much lower dimension than $\text{GL}(\mathbb{V}) = \text{GL}(3m)$ for large m . This is typical in equivariant neural networks.

In this simple example, we not only know f but have an explicit expression for it. In general, there are many functions that we know should be equivariant or invariant to certain group actions, but for which we do not know any simple closed-form expression; and this is where it helps to assume that f is given by some neural network whose parameters could be determined by fitting it with data, or, if it is used as an ansatz, by plugging into some differential or integral equations. A simple data-fitting example is provided by semantic segmentation in images, which seeks to classify pixels as belonging to one of several types of objects. If we rotate or mirror an image, we expect that pixel labels should follow the pixels. A more realistic version of the center of mass example would be a molecule represented by positions of its atoms, which comes up in chemical property or drug response predictions. Here we want equivariance with respect to coordinate transformations, but we wish to preserve pairwise distances between atoms and chirality, so the natural group to use is $\text{SO}(3)$ [KLT18] or the special Euclidean group $\text{SE}(3)$ [WGW⁺18, FFWW20]. The much-publicized protein structure prediction engine of DeepMind’s AlphaFold 2 relies on an $\text{SE}(3)$ -equivariant neural network and an $\text{SE}(3)$ -invariant attention module [JEP⁺21]. In [TEW⁺21], $\text{SE}(3)$ -equivariant convolution is used to improve accuracy assessments of RNA structure models.

Another straightforward example comes from computational quantum chemistry, where one seeks a solution to a Schrödinger equation: if we write $v_i \in \mathbb{R}^3 \times \{-1/2, 1/2\}$, then the wave function of m identical spin-1/2 fermions is antisymmetric, i.e.,

$$f(v_{\pi(1)}, v_{\pi(2)}, \dots, v_{\pi(m)}) = (-1)^{\text{sgn}(\pi)} f(v_1, v_2, \dots, v_m)$$

for all $\pi \in S_m$. In other words, the increasingly popular antisymmetric neural networks [HSN20] are S_m -equivariant neural networks. Even without going into the details, the reader could well imagine that restricting to neural networks that are antisymmetric is a savings from having to consider all possible neural networks. More esoteric examples in particle physics call for Lorentz groups of various stripes like $\text{O}(1, 3)$, $\text{O}^+(1, 3)$, $\text{SO}(1, 3)$, or $\text{SO}^+(1, 3)$, which are used in Lorentz-equivariant neural networks to identify top quarks in data from high-energy physics experiments [BAO⁺20].

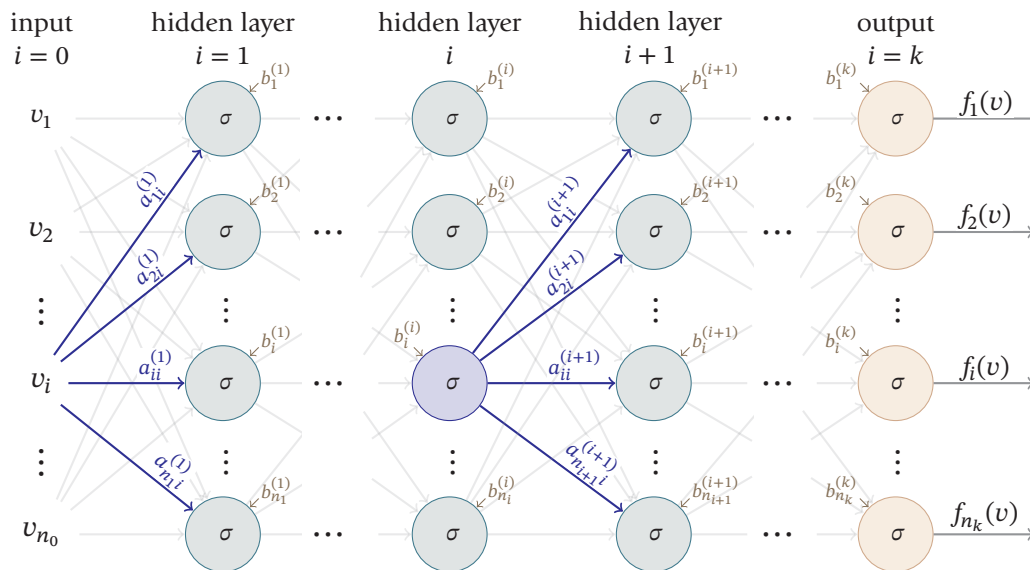


Figure 1. A k -layer σ -activated feed-forward neural network, also known as a multilayer perceptron.

We now discuss the equivariant condition for pointwise nonlinearities $X^{-1}\sigma_b X = \sigma_b$. It is instructive to look at a simple numerical example. Suppose we apply a pointwise nonlinearity σ and a permutation matrix X given by

$$\sigma(t) = \begin{cases} +1 & t \geq 3.0, \\ -1 & t < 3.0, \end{cases} \quad X = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix},$$

to a vector $v = (2.1, 3.4, 0.2) \in \mathbb{R}^3$. We see that $\sigma(Xv) = X\sigma(v)$:

$$\begin{bmatrix} 2.1 \\ 3.4 \\ 0.2 \end{bmatrix} \xrightarrow{X} \begin{bmatrix} 3.4 \\ 0.2 \\ 2.1 \end{bmatrix} \xrightarrow{\sigma} \begin{bmatrix} +1 \\ -1 \\ -1 \end{bmatrix} \xrightarrow{X^{-1}} \begin{bmatrix} -1 \\ +1 \\ -1 \end{bmatrix}$$

which clearly holds more generally, i.e., $X^{-1}\sigma X = \sigma$ for any permutation matrix X and any pointwise nonlinearity σ . The bottom line is that the permutation matrix $X = \rho(\pi)$ comes from a representation $\rho : S_n \rightarrow \text{GL}(n)$; and since $\pi \in S_n$ acts on the indices of v and $\sigma : \mathbb{R}^n \rightarrow \mathbb{R}^n$ acts on the values of v , the two actions are always independent of each other. More generally, it is easy to see that if we include a bias term $b \in \mathbb{R}^n$, then $\sigma_b : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is S_n -equivariant as long as b has all coordinates equal [CW16]. This does not necessarily hold for more general b : Take the example above and set the bias to be $b = (-1, 0, 0)$.

$$\begin{bmatrix} 2.1 \\ 3.4 \\ 0.2 \end{bmatrix} \xrightarrow{X} \begin{bmatrix} 3.4 \\ 0.2 \\ 2.1 \end{bmatrix} \xrightarrow{\sigma_b} \begin{bmatrix} -1 \\ -1 \\ -1 \end{bmatrix} \xrightarrow{X^{-1}} \begin{bmatrix} -1 \\ -1 \\ -1 \end{bmatrix}$$

but $\sigma_b(v) = \sigma(v + b) = (-1, +1, -1)$. So $X^{-1}\sigma_b X \neq \sigma_b$. Going beyond pointwise nonlinearity is a nontrivial issue and is crucial when the neural network requires more than just S_n -equivariance. We will say a few words about this below.

The mathematical ideas that we have described are all fairly straightforward. Indeed the technical challenges in equivariant neural networks are mostly about getting these mathematical ideas to work in real-life situations, what we have swept under the “engineering complications” rug. We will discuss a few of these but as engineering complications go, they invariably depend on the problem at hand and every case is different.

The butterfly image example presented at the beginning already concealed several difficulties. While we have assumed that images are functions $v : \mathbb{R}^2 \rightarrow \mathbb{R}^3$, in real life they are sampled on a grid, i.e., pixels are discrete, and a more realistic model would be $v : \mathbb{Z}^2 \rightarrow \mathbb{R}^3$. Instead of a straightforward $\text{SO}(2)$ -equivariance as one might expect for imaging problems, one instead finds discussions of equivariance [CW16] with respect to wallpaper groups like

$$G_1 = \left\{ \begin{bmatrix} 1 & 0 & m_1 \\ 0 & 1 & m_2 \\ 0 & 0 & 1 \end{bmatrix} \in \mathbb{R}^{3 \times 3} : m_1, m_2 \in \mathbb{Z} \right\}$$

for translation in \mathbb{Z}^2 ; or

$$G_2 = \left\{ \begin{bmatrix} \cos(k\pi/2) & -\sin(k\pi/2) & m_1 \\ \sin(k\pi/2) & \cos(k\pi/2) & m_2 \\ 0 & 0 & 1 \end{bmatrix} \in \mathbb{R}^{3 \times 3} : \right. \\ \left. k = 0, 1, 2, 3; m_1, m_2 \in \mathbb{Z} \right\}$$

that augments G_1 with right-angle rotations; or

$$G_3 = \left\{ \begin{bmatrix} (-1)^j \cos(k\pi/2) & (-1)^{j+1} \sin(k\pi/2) & m_1 \\ \sin(k\pi/2) & \cos(k\pi/2) & m_2 \\ 0 & 0 & 1 \end{bmatrix} \in \mathbb{R}^{3 \times 3} : \right. \\ \left. k = 0, 1, 2, 3; \right. \\ \left. j = 0, 1; m_1, m_2 \in \mathbb{Z} \right\}$$

that further augments G_2 with reflections. The reason for these choices is that they have to send pixels to pixels.

There is also the important issue of *aliasing* [Zha19]. When pixels are discrete, rotation will involve interpolation, and pointwise nonlinearities introduce higher order harmonics that produce aliasing and break equivariance [FW21]. This can happen even with discrete translations like those in the groups G_1, G_2, G_3 above for standard convolutional neural networks. Dealing with aliasing and choosing equivariant activations that do not compromise expressive power [MFSL19] are important problems that cannot be underemphasized; and dealing with these issues constitute a mainstay of the research and development in equivariance neural networks.

In reality the pixels of an image are not just discrete but also finite in number. So instead of $v : \mathbb{Z}^2 \rightarrow \mathbb{R}^3$, a p -pixel image is more accurately a function $v : \{x_1, \dots, x_p\} \rightarrow \mathbb{R}^3$ on some discrete finite subset of points $x_1, \dots, x_p \in \mathbb{R}^2$. Since these p points are fixed, we may conveniently regard $v \in \mathbb{R}^p \oplus \mathbb{R}^p \oplus \mathbb{R}^p$ with each copy of \mathbb{R}^p representing one of three color channels. In such cases the output of each layer should not be treated simply as a vector space \mathbb{R}^{n_i} but a direct sum $\mathbb{R}^{n_i} = \mathbb{R}^{p_1} \oplus \dots \oplus \mathbb{R}^{p_m}$, with p_1, \dots, p_m depending on i and $n_i = p_1 + \dots + p_m$. The weight matrix $A_i : \mathbb{R}^{n_{i-1}} \rightarrow \mathbb{R}^{n_i}$ would then have a corresponding block structure and the representation $\rho_i : G \rightarrow GL(n_i)$ takes the form $\rho_i = \bigoplus_{j=1}^m \rho_{ij}$ with $\rho_{ij} : G \rightarrow GL(p_j)$.

For molecular structure prediction problems, in [FWFW20, TEW⁺21], the input is a collection of points y_1, \dots, y_m augmented with various information in addition to location coordinates, giving the input layer \mathbb{R}^{n_0} a direct sum structure $\mathbb{R}^{p_1} \oplus \dots \oplus \mathbb{R}^{p_m}$ that propagates through later layers. Just to give a flavor of what is involved, in [TEW⁺21], the inputs y_1, \dots, y_m are atom positions in a model of an RNA molecule, with an encoding of atom type, and the output is an estimate of the root mean square error of the model's structure; in one example in [FWFW20], the inputs y_1, \dots, y_m encode position, velocity, and charge of m particles, and the output is an estimate of the location and velocity of each particle after some amount of time. In both examples, the function f that maps inputs to outputs has no known expression but is known to be equivariant with respect to $SE(3)$, i.e., translations and rotations of the coordinate system. The weight matrix $A : \mathbb{R}^{q_1} \oplus \dots \oplus \mathbb{R}^{q_m} \rightarrow \mathbb{R}^{p_1} \oplus \dots \oplus \mathbb{R}^{p_m}$ has a block structure $A = [A_{ij}]_{i,j=1}^m$, $A_{ij} \in \mathbb{R}^{p_i \times q_j}$, and is equivariant if each block $A_{ij} : \mathbb{R}^{q_j} \rightarrow \mathbb{R}^{p_i}$ is equivariant. Equivariance constrains each block A_{ij} to depend entirely on the relative input locations $y_i - y_j$, and the permitted matrices can be expressed in terms of radial kernels, spherical harmonics, and Clebsch–Gordan coefficients [WGW⁺18].

The engineering aspects of equivariant neural networks are many and varied. While we have selectively discussed a few that are more common and mathematical in nature, we have also ignored many that are specific to the application at hand and often messy. We avoided most jargon used in the original literature as it tends to be mathematically imprecise or application specific. Nevertheless, we stress that equivariant neural networks are ultimately used in an engineering context and a large part of their success has to do with overcoming real engineering challenges.

ACKNOWLEDGMENT. We would like to thank the two anonymous reviewers and the handling editor for their many useful suggestions and comments. We would also like to thank Yuehaw Khoo, Risi Kondor, Zehua Lai, and Jiangying Zhou for helpful discussions. The simplified constant width case was adapted from [Lim21, Example 2.16] where it was used to illustrate tensor transformation rules. We thank StackExchange user 'Ulysses' for his TikZ butterfly figure in <https://tex.stackexchange.com/a/495243>.

References

- [BAO⁺20] Alexander Bogatskiy, Brandon Anderson, Jan Olfmann, Marwah Roussi, David Miller, and Risi Kondor, *Lorentz group equivariant neural network for particle physics*, Proceedings of the 37th International Conference on Machine Learning (ICML), 2020, pp. 992–1002.
- [BH10] John Baez and John Huerta, *The algebra of grand unified theories*, Bull. Amer. Math. Soc. (N.S.) 47 (2010), no. 3, 483–552.
- [CS81] Sylvain E. Cappell and Julius L. Shaneson, *Nonlinear similarity*, Ann. of Math. 113 (1981), no. 2, 315–355.
- [CW16] Taco Cohen and Max Welling, *Group equivariant convolutional networks*, Proceedings of the 33rd International Conference on Machine Learning (ICML), 2016, pp. 2990–2999.
- [FW21] Daniel Franzen and Michael Wand, *General nonlinearities in SO(2)-equivariant CNNs*, Advances in Neural Information Processing Systems (NeurIPS), 2021, pp. 1970–1981.
- [FWFW20] Fabian Fuchs, Daniel E. Worrall, Volker Fischer, and Max Welling, *SE(3)-transformers: 3D roto-translation equivariant attention networks*, Advances in Neural Information Processing Systems (NeurIPS), 2020, pp. 1970–1981.
- [FWW21] Marc Finzi, Max Welling, and Andrew Gordon Wilson, *A practical method for constructing equivariant multi-layer perceptrons for arbitrary matrix groups*, Proceedings of the 38th International Conference on Machine Learning (ICML), 2021, pp. 3318–3328.
- [HSN20] Jan Hermann, Zeno Schätzle, and Frank Noé, *Deep-neural-network solution of the electronic Schrödinger equation*, Nat. Chem. 12 (2020), 891–897.

- [JEP⁺21] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Židek, Anna Potapenko, Alex Bridgland, Clemens Meyer, Simon A. A. Kohl, Andrew J. Ballard, Andrew Cowie, Bernardino Romera-Paredes, Stanislav Nikolov, Rishub Jain, Jonas Adler, Trevor Back, Stig Petersen, David Reiman, Ellen Clancy, Michal Zielinski, Martin Steinegger, Michalina Pacholska, Tamas Berghammer, Sebastian Bodenstern, David Silver, Oriol Vinyals, Andrew W. Senior, Koray Kavukcuoglu, Pushmeet Kohli, and Demis Hassabis, *Highly accurate protein structure prediction with AlphaFold*, *Nature* **596** (2021), 583–589.
- [KLT18] Risi Kondor, Zhen Lin, and Shubhendu Trivedi, *Clebsch–Gordan nets: a fully Fourier space spherical convolutional neural network*, *Advances in Neural Information Processing Systems (NeurIPS)*, 2018, pp. 10117–10126.
- [KSH12] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton, *ImageNet classification with deep convolutional neural networks*, *Advances in Neural Information Processing Systems (NIPS)*, 2012, pp. 1097–1105.
- [KT18] Risi Kondor and Shubhendu Trivedi, *On the generalization of equivariance and convolution in neural networks to the action of compact groups*, *Proceedings of the 35th International Conference on Machine Learning (ICML)*, 2018, pp. 2747–2755.
- [Lim21] Lek-Heng Lim, *Tensors in computations*, *Acta Numer.* **30** (2021), 555–764.
- [May96] J. P. May, *Equivariant homotopy and cohomology theory*, CBMS Regional Conference Series in Mathematics, vol. 91, Published for the Conference Board of the Mathematical Sciences, Washington, DC; by the American Mathematical Society, Providence, RI, 1996. With contributions by M. Cole, G. Comezana, S. Costenoble, A. D. Elmendorf, J. P. C. Greenlees, L. G. Lewis, Jr., R. J. Piacenza, G. Triantafillou, and S. Waner, DOI 10.1090/cbms/091. MR1413302
- [MFSL19] Haggai Maron, Ethan Fetaya, Nimrod Segol, and Yaron Lipman, *On the universality of invariant networks*, *Proceedings of the 36th International Conference on Machine Learning (ICML)*, 2019, pp. 4363–4371.
- [MP43] Warren S. McCulloch and Walter Pitts, *A logical calculus of the ideas immanent in nervous activity*, *Bull. Math. Biophys.* **5** (1943), 115–133, DOI 10.1007/bf02478259. MR10388
- [Ros58] Frank Rosenblatt, *The perceptron: A probabilistic model for information storage and organization in the brain*, *Psychol. Rev.* **65** (1958), 386–408.
- [TEW⁺21] Raphael J. L. Townshend, Stephan Eismann, Andrew M. Watkins, Ramya Rangan, Maria Karelina, Rhiju Das, and Ron O. Dror, *Geometric deep learning of RNA structure*, *Science* **373** (2021), no. 6558, 1047–1051.
- [WGW⁺18] Maurice Weiler, Mario Geiger, Max Welling, Wouter Boomsma, and Taco S. Cohen, *3D Steerable CNNs: Learning rotationally equivariant features in volumetric data*, *Advances in Neural Information Processing Systems (NeurIPS)*, 2018, pp. 10381–10392.
- [Zha19] Richard Zhang, *Making convolutional networks shift-invariant again*, *Proceedings of the 36th International Conference on Machine Learning (ICML)*, 2019, pp. 7324–7334.



Lek-Heng Lim



Bradley J. Nelson

Credits

Butterfly images are courtesy of Ulysses via StackExchange. Figure 1 is courtesy of Lek-Heng Lim and Bradley J. Nelson. Photo of Lek-Heng Lim is courtesy of Sou-Cheng Choi. Photo of Bradley J. Nelson is courtesy of Xiaotong Suo.