Neural networks with Euclidean Symmetry for the Physical Sciences

SLIDES: https://tinyurl.com/e3nn-physics-meets-ml





Tess Smidt 2018 Alvarez Fellow in Computing Sciences

Physics ∩ *ML* 2020.11.18

Neural networks with Euclidean Symmetry for the Physical Sciences

SLIDES: https://tinyurl.com/e3nn-physics-meets-ml

All I wanted was 3D rotation equivariance and I got...

....geometric tensors, space groups, point groups, selection rules, normal modes, degeneracy, 2nd order phase transitions, and a much better understanding of physics.



Tess Smidt 2018 Alvarez Fellow in Computing Sciences

Physics ∩ *ML* 2020.11.18

The laws of physics have rotational, translational, and (unless you're a particle physicist) parity symmetry.

We want machine learning models that also obey this symmetry.

e.g. a network is our model of "physics". The input to the network is our system.



Symmetry emerges when different ways of representing something "mean" the same thing. Symmetry of representation vs. objects

Symmetry emerges when different ways of representing something "mean" the same thing. Symmetry of representation vs. objects

Euclidean symmetry, E(3):

Symmetry of 3D space The freedom to choose your coordinate system



Symmetry emerges when different ways of representing something "mean" the same thing. Symmetry of representation vs. objects

Euclidean symmetry, *E*(3):

Symmetry of 3D space The freedom to choose your coordinate system



We transform between coordinate systems with...

3D Translation



Symmetry emerges when different ways of representing something "mean" the same thing. Symmetry of representation vs. objects

Euclidean symmetry, E(3):

Symmetry of 3D space The freedom to choose your coordinate system



Symmetry of geometric objects

Looks the same under specific rotations, translations, and inversion (includes mirrors).



Symmetry emerges when different ways of representing something "mean" the same thing. Symmetry of representation vs. objects

Euclidean symmetry, E(3):

Symmetry of 3D space The freedom to choose your coordinate system



Symmetry of geometric objects

Looks the same under specific rotations, translations, and inversion (includes mirrors).



Neural networks are specially designed for different data types. Assumptions about the data type are built into how the network operates.



Arrays *⇒* Dense NN







Graph ⇔ Graph (Conv.) NN



Components are independent.

The same features can be found anywhere in an image. Locality.

Sequential data. Next input/output depends on input/output that has come before. Topological data. Nodes have features and network passes messages between nodes connected via edges.

3D physical data ⇔ Euclidean NN

Data in 3D Euclidean space. Freedom to choose coordinate system.



Neural networks are specially designed for different data types. Assumptions about the data type are built into how the network operates. Thus, symmetries are encoded by tailoring network operations.





Three ways to make models "symmetry-aware" for 3D data

e.g. How to make a model that "understands" the symmetry of atomic structures?



-0.21463	0.97837	0.33136
-0.38325	0.66317	-0.70334
-1.57552	0.03829	-1.05450
-2.34514	-0.13834	-0.29630
-1.78983	-0.36233	-2.36935
-2.72799	-0.85413	-2.64566
-0.81200	-0.13809	-3.33310
-0.98066	-0.45335	-4.36774
0.38026	0.48673	-2.98192
1.14976	0.66307	-3.74025
0.59460	0.88737	-1.66708
1.53276	1.37906	-1.39070

Coordinates are most general, but sensitive to <u>translations</u>, <u>rotations</u>, and <u>inversion</u>.

Three ways to make models "symmetry-aware" for 3D data

e.g. How to make a model that "understands" the symmetry of atomic structures?



-0.21463	0.97837	0.33136
-0.38325	0.66317	-0.70334
-1.57552	0.03829	-1.05450
-2.34514	-0.13834	-0.29630
-1.78983	-0.36233	-2.36935
-2.72799	-0.85413	-2.64566
-0.81200	-0.13809	-3.33310
-0.98066	-0.45335	-4.36774
0.38026	0.48673	-2.98192
1.14976	0.66307	-3.74025
0.59460	0.88737	-1.66708
1.53276	1.37906	-1.39070

Coordinates are most general, but sensitive to <u>translations</u>, <u>rotations</u>, and <u>inversion</u>.

Approach 1: Data Augmentation

Throw data at the problem and see what you get!

Approach 2: Invariant Inputs

Convert your data to invariant representations so the neural network can't possibly mess it up.

Approach 3: Invariant models Equivariant models

If there's no model that naturally handles coordinates, we will make one.

Three ways to make models "symmetry-aware" for 3D data

e.g. How to make a model that "understands" the symmetry of atomic structures?



Coordinates are most general, but sensitive to <u>translations</u>, <u>rotations</u>, and <u>inversion</u>.

Approach 1: Data Augmentation Approach 2: O Invariant Inputs

Throw data at the problem and see what you get!

Convert your data to invariant representations so the neural network can't possibly mess it up.

Approach 3: Invariant models Equivariant models

If there's no model that naturally handles coordinates, we will make one.

Invariance vs. Equivariance (covariance) e.g. in 3D space

Does <u>NOT</u> change *⇒* Invariant Changes <u>deterministically</u> *⇒* Equivariant



For 3D data, data augmentation is expensive, ~500 fold augmentation and you *still* don't get the guarantee of equivariance (it's only emulated).

training without rotational symmetry



training with symmetry



For a function to be <u>equivariant</u> means that we can act on our inputs with g <u>OR</u> act our outputs with g and we get the same answer (for every operation). For a function with <u>invariant</u> input (e.g. invariant models) means g is the identity (no change).



Why limit yourself to equivariant functions? You can *substantially* shrink the space of functions you need to optimize over.

This means you need less data to constrain your function.



Why <u>not</u> limit yourself to <u>invariant</u> functions? You have to <u>guarantee</u> that your input features already contain any necessary equivariant interactions (*e.g. cross-products*).



How Euclidean Neural Networks achieve equivariance to Euclidean symmetry (high level)

Euclidean Neural Networks encompass

Tensor Field Networks (arXiv:1802.08219) Clebsch-Gordon Nets (arXiv:1806.09231) 3D Steerable CNNs (arXiv:1807.02547) Cormorant (arXiv:1906.04015) SE(3)-Transformers (arXiv:2006.10503) e3nn (github.com/e3nn/e3nn) (Technically, e3nn is the only one that implements inversion)

Some relevant folks... Mario Geiger, Ben Miller, Risi Kondor, Taco Cohen, Maurice Weiler, Daniel E. Worrall, Fabian B. Fuchs, Max Welling, Nathaniel Thomas, Shubhendu Trivedi,...

Equivariant convolutional filters are based on learned radial functions and spherical harmonics...



Equivariant convolutional filters are based on learned radial functions and spherical harmonics...



Spherical harmonics of the same L transform together under rotation **g**.



Spherical harmonics transform in the same manner as the irreducible representations of SO(3).

Equivariant convolutional filters are based on learned radial functions and spherical harmonics...



Equivariant convolutional filters are based on learned radial functions and spherical harmonics...



...and geometric tensor algebra allow us to generalize scalar operations to more complex geometric tensors.



e.g. How to multiply two vectors?

- $ec{a}\cdotec{b}=c$ scalar
- $\vec{a} \times \vec{b} = \vec{c}$
 - vector

 $ec{a}\otimesec{b}=C$ $_{
m matrix}^{
m 3x3}$

Equivariant convolutional filters are based on learned radial functions and spherical harmonics...



...and geometric tensor algebra allow us to generalize scalar operations to more complex geometric tensors.



The input to our network is geometry and features on that geometry.



geometry = [[x0, y0, z0],[x1, y1, z1]]
features = [
 [m0, v0y, v0z, v0x, a0y, a0z, a0x]
 [m1, v1y, v1z, v1x, a1y, a1z, a1x]
]
....

The input to our network is geometry and features on that geometry. We categorize our features by how they transform under rotation and parity as *irreducible representations of O(3)*.



What does equivariance get you?

Given a molecule and a rotated copy, predicted forces are the same up to rotation.

(Predicted forces are equivariant to rotation.)

Additionally, networks generalize to molecules with similar motifs.



Primitive unit cells, conventional unit cells, and supercells of the same crystal produce the same output (assuming periodic boundary conditions).



E(3)NNs can express tensors of atomic orbitals and predict molecular Hamiltonians in any orientation from seeing a single example.



E(3)NNs can manipulate geometry,

which means they can be used for generative models such as autoencoders.

We can convert local *geometry* into *features* and *vice versa* via spherical harmonic projections.



E(3)NNs are extremely data efficient.

E(3)NNs are extremely data efficient.

E(3)NNs for molecular dynamics (coming soon) Water

System	TFMD, 133 data points	DeepMD, 133,500 data points
Liquid Water	29.7 (meV/Å)	40.4 (meV/Å)
Ice Ih (b)	26.7	43.3
Ice Ih (c)	16.1	26.8
Ice Ih (d)	13.3	25.4

Data from: Zhang, L. et al. E. (2018). *PRL*, *120*(14), 143001.



Batzner Boris

E(3)NNs are extremely data efficient.

E(3)NNs for molecular dynamics (*coming soon*) *Water*

System	TFMD, 133 data points	DeepMD, 133,500 data points
Liquid Water	29.7 (meV/Å)	40.4 (meV/Å)
Ice Ih (b)	26.7	43.3
Ice Ih (c)	16.1	26.8
Ice Ih (d)	13.3	25.4

Data from: Zhang, L. et al. E. (2018). PRL, 120(14), 143001.

E(3)NNs for phonon density of states (arxiv:2009.05163) Training set of 1,200 crystal structures with 64 atom types. Test set includes atom types never seen. Used to find high C_v . Data from Materials Project



Simon Batzner

J_{Boris} Kozinsky

Zhantao Chen Nina Andrejevic Mingda Li

35



Features that are consequences of fully treating Euclidean symmetry...
<u>Feature 1</u>: All data (input, intermediates, output) in E(3)NNs are <u>geometric tensors</u>. Geometric tensors are the "data types" of 3D space and have <u>many</u> forms.



<u>Feature 1</u>: All data (input, intermediates, output) in E(3)NNs are <u>geometric tensors</u>. Geometric tensors are the "data types" of 3D space and have <u>many</u> forms.



Feature 2: The outputs have equal or higher symmetry than the inputs.

Curie's principle (1894): "When effects show certain asymmetry, this asymmetry must be found in the causes that gave rise to them."



<u>Feature 2</u>: The outputs have equal or higher symmetry than the inputs.

Curie's principle (1894): "When

s show certain asymmetry, this asymmetry must be found that gave rise to them."



<u>Feature 2</u>: The outputs have equal or higher symmetry than the inputs. Symmetry compiler -- can't fit a model that does symmetrically make sense



T. E. Smidt, M. Geiger, B. K. Miller. <u>https://arxiv.org/abs/2007.02005</u> (2020)

<u>Feature 2</u>: The outputs have equal or higher symmetry than the inputs.

Symmetry compiler -- can't fit a model that does symmetrically make sense



T. E. Smidt, M. Geiger, B. K. Miller. <u>https://arxiv.org/abs/2007.02005</u> (2020)

Feature 2: The outputs have equal or higher symmetry than the inputs.

Symmetry compiler -- can't fit a model that does symmetrically make sense



<u>Feature 2</u>: The outputs have equal or higher symmetry than the inputs.

Symmetry compiler -- can't fit a model that does symmetrically make sense



T. E. Smidt, M. Geiger, B. K. Miller. <u>https://arxiv.org/abs/2007.02005</u> (2020)

Feature 3: We can find data that is implied by symmetry.

Using gradients of loss wrt input we can find symmetry breaking "order parameters"



Feature 3: We can find data that is implied by symmetry.

Using gradients of loss wrt input we can find symmetry breaking "order parameters"

Octahedral tilting in perovskites (M³⁺ ⊕ R⁴⁺) ⇒

Network learns equal magnitude pseudovector order parameters on B site with proper spatial patterning.



T. E. Smidt, M. Geiger, B. K. Miller. <u>https://arxiv.org/abs/2007.02005</u> (2020)

Utilities and classes for

- building E(3) equivariant neural networks
- manipulating geometric tensors
- visualizing spherical harmonics

developers of e3nn





Ben 2019.07.08



Kostiantyn Lapchevskyi



Tess Smidt (LBNL)

Creating a basic convolution E(3) neural network

```
import torch
from e3nn import rs
from e3nn.networks import GatedConvParityNetwork
torch.set default dtype(torch.float64)
```

```
N_atom_types = 3 # For example H, C, O
Rs_in = [(N_atom_types, 0, 1)] # Input are scalars
Rs out = [(1, 1, -1)] # Predict vectors
```

```
model_kwargs = {
    'Rs_in': Rs_in, 'Rs_out': Rs_out, 'mul': 4, 'lmax': 2,
    'layers': 3, 'max_radius': r_max, 'number_of_basis': 10,
}
```

model = GatedConvParityNetwork(**model kwargs)

Convert between Cartesian tensors (with symmetric indices) and Irrep tensors and calculate degrees of freedom (e.g. elasticity tensor)

```
import torch
from e3nn import rs
from e3nn.tensor import CartesianTensor
torch.set default dtype(torch.float64)
```

```
rank4 = torch.zeros(3, 3, 3, 3) # Placeholder
Rs, Q = CartesianTensor(rank4, 'ijkl=jikl=klij').to_irrep_transformation()
print("Representations: ", Rs)
print("Degrees of freedom: ", rs.dim(Rs))
```

```
>> Representations: [(2, 0, 1), (2, 2, 1), (1, 4, 1)]
>> Degrees of freedom: 21
```

Plot 3x3 matrix as linear combination of spherical harmonics.

```
# Symmetric Matrix
M = torch.randn(3,3)
M = M + M.transpose(0, 1)
 Plot matrix
px.imshow(M)
                                                                     1.5
matrix = CartesianTensor(M, formula='ij=ji').to irrep tensor()
                                                                     1
r, f = SphericalTensor.from irrep tensor(matrix).plot()
                                                                     0.5
                                                                     0
                                                                     -0.5
# Plot SH signal
                                                                      -1
                                                                     -1.5
                                                                      1.5
surface plot = lambda r, f: go.Surface(
                                                                       0
    x=r[..., 0], y=r[..., 1], z=r[..., 2],
                                                                          surfacecolor=f, showscale=False)
                                                                        1.5
go.Figure([surface plot(r, f)])
                                                                                     50
```

collaborators of e3nn





Thomas Hardin



Frank Noé



Josh Rackers



Claire West



Simon Batzner

Boris Kozinsky









Eugene Kwan



Mingda L

A Quick Recap!

3D Euclidean symmetry: rotations, translation, inversion Different coordinate systems ⇔ same physical system

Euclidean Neural Networks are equivariant to E(3) Convolutional filters ⇔ learned radial functions and spherical harmonics

Geometric tensor algebra Equivariant nonlinearities (did not discuss)

Equivariance can have unintended features.

1) Symmetry specific data types

2) Output symmetry equal to inputs

- Implement group equivariance and get all subgroups for FREE!
- Symmetry compilers

3) Grad loss wrt input can break symmetry

Feel free to reach out ifTess Smidtyou have any questions!tsmidt@lbl.gov

A Quick Recap!

3D Euclidean symmetry: rotations, translation, inversion Different coordinate systems

Euclidean Neural Networks are equivariant to E(3)

Convolutional filters

Iearned radial functions and spherical harmonics

Geometric tensor algebra Equivariant nonlinearities (did not discuss)

Equivariance can have unintended features.

- 1) Symmetry specific data types
- 2) Output symmetry equal to inputs
 - Implement group equivariance and get all subgroups for FREE!
 - Symmetry compilers
- 3) Grad loss wrt input can break symmetry

Feel free to reach out if you have any questions!

Tess Smidt tsmidt@lbl.gov

Resources on Euclidean neural networks: <u>e3nn.org</u>

e3nn Code (PyTorch): http://github.com/e3nn/e3nn "quick" tutorial: https://tinyurl.com/e3nn-quick-tutorial-202011 e3nn_tutorial: http://blondegeek.github.io/e3nn_tutorial/ Papers:

> Tensor Field Networks (arXiv:1802.08219) Clebsch-Gordon Nets (arXiv:1806.09231) 3D Steerable CNNs (arXiv:1807.02547) Cormorant (arXiv:1906.04015) SE(3)-Transformers (arXiv:2006.10503) tfnns on proteins (arXiv:2006.09275) e3nn on QM9 (arXiv:2008.08461) e3nn for symm breaking (arXiv:2008.08461) E(3) and equivariance in ML (chemrxiv.12935198.v1) e3nn for phonon DOS (arxiv:2009.05163)

My past talks (look for video / slide links): https://blondegeek.github.io/talks

Calling in backup (slides)!



Applications so far...

- Finding order parameters of 2nd order structural phase transitions
- Molecular dynamics (Harvard)
- Molecule and crystal property prediction (FU Berlin)
- Inverting invariant representations of atomic geometries (Sandia)
- Autoencoding Geometry
- Predicting molecular <u>Hamiltonians</u> (TU Berlin)
- Long range interactions (FU Berlin, TU Berlin)
- Electron density prediction for large molecules (Sandia)
- Predicting chemical shifts for <u>NMR</u> (*Merck, MIT*)
- Conditional protein design (UW)
- Inverse design of <u>optical properties</u> of nanoparticle assemblies (*LBL and UW*)
- Phonon properties of crystal structures (MIT)
- Anharmonic elastic properties of crystal structures (UTEP)

55

Spherical harmonics of a given L transform together under rotation.



Predict ab initio forces for molecular dynamics

Preliminary results originally presented at APS March Meeting 2019. Paper in progress.

Testing on liquid water, Euclidean neural networks (*Tensor-Field Molecular Dynamics*) require less data to train than traditional networks to get state of the art results.

	MAE [meV/A]	RMSE [meV/A]
TFMD, 100	27.9	38.20
TFMD, 1000	11.29	14.82
Deep-MD, 133,500	not reported	40.0





Data set from: [1] Zhang, L. et al. E. (2018). *PRL*, *120*(14), 143001.

Euclidean neural networks can manipulate geometry,

which means they can be used for generative models such as autoencoders.

Euclidean neural networks can manipulate geometry,

which means they can be used for generative models such as autoencoders.

To encode/decode, we have to be able to convert *geometry* into *features* and *vice versa*. We do this via spherical harmonic projections.



Equivariant neural networks can learn to invert invariant representations.



We can also build an autoencoder for geometry: e.g. Autoencoder on 3D Tetris



Centers deleted

We can also build an autoencoder for geometry: e.g. Autoencoder on 3D Tetris



Euclidean Neural Networks are similar to convolutional neural networks...



We encode the symmetries of 3D Euclidean space (3D translation- and 3D rotation-equivariance).

 $g \in SE(3)$



Euclidean Neural Networks are similar to convolutional neural networks...

We use points. Images of atomic systems are sparse and imprecise. We use continuous convolutions with atoms as convolution centers. VS. VS. VS.

We encode the symmetries of 3D Euclidean space (3D translation- and 3D rotation-equivariance).

 $g \in SE(3)$

Euclidean Neural Networks are similar to convolutional neural networks...

We use points. Images of atomic systems are sparse and imprecise. We use continuous convolutions with atoms as convolution centers. Other atoms Convolution center

We encode the symmetries of 3D Euclidean space (3D translation- and 3D rotation-equivariance).







Translation equivariance Convolutional neural network ✓





Rotation equivariance

Data augmentation Radial functions (invariant) Want a network that both preserves geometry and exploits symmetry.

(3)

Invariant featurizations can be very expressive if well-crafted Many *invariant* featurizations use *equivariant* operations e.g. a (simplified) SOAP kernel for ethane molecule C₂H₆

(1)

(2)

L = 0 L = 1 L = 2 L = 3 L = 4 L = 5L = 6

- (1) Project neighbors of given atom onto spherical harmonics (<u>equivariant</u> quantity).
- Interact signals from different atoms via tensor dot product (<u>equivariant</u> operation) to produce scalars (<u>invariant</u> quantity).
- (3) Give scalars to model.

For a function to be <u>equivariant</u> means that we can act on our inputs with g <u>OR</u> act our outputs with g and we get the same answer (for every operation). For a function to be <u>invariant</u> means g is the identity (no change).



Why limit yourself to equivariant functions? You can *substantially* shrink the space of functions you need to optimize over.

This means you need less data to constrain your function.



Why not limit yourself to invariant functions? You have to <u>guarantee</u> that your input features already contain any necessary equivariant interactions (*e.g. cross-products*).



Neural networks are specially designed for different data types. Assumptions about the data type are built into how the network operates.



Arrays *⇒* Dense NN







Graph ⇔ Graph (Conv.) NN



Components are independent.

The same features can be found anywhere in an image. Locality.

Sequential data. Next input/output depends on input/output that has come before. Topological data. Nodes have features and network passes messages between nodes connected via edges.

3D physical data ⇔ Euclidean NN

Data in 3D Euclidean space. Freedom to choose coordinate system.



Neural networks are specially designed for different data types. Assumptions about the data type are built into how the network operates. Symmetries emerge from these assumptions.




Neural networks can't mess up invariant representations.

You can use *ANY* neural network with an invariant representation. Invariant representations can be used for other machine learning algorithms (e.g. kernel methods).



If you can craft a good representation -- great!

But deep learning's specialty is feature learning.

So, maybe use a different machine learning approach (e.g. kernel methods).

Analogous to... the laws of (non-relativistic) physics have Euclidean symmetry, even if systems do not.

The **network** is our model of "**physics**". The **input** to the network is our **system**.

A Euclidean symmetry preserving network produces *outputs* that preserve the subset of symmetries induced by the *input*.











Equivariance can have unintuitive consequences.

Partition graph with permutation equivariant function into two sets using ordered labels.

Predict node labels [0, 1] vs. [1, 0]



Equivariance can have unintuitive consequences.

Partition graph with permutation equivariant function into two sets using ordered labels. You can't *due to degeneracy*.



There's nothing to distinguish one partition to be "first" vs. "second".

Convolutions: Local vs. Global Symmetry

Convolutions capture local symmetry. Interaction of features in later layers yields global symmetry. e.g. Coordination environments in crystals

Atomic systems form geometric motifs that can appear at multiple locations and orientations.

Space group: Symmetry of unit cell (Global symmetry)



Symmetry emerges when different ways of representing something "mean" the same thing.

Representation can have symmetry, operations can preserve symmetry, and objects can have symmetry.

Translation symmetry in 2D:

Features "mean" the same thing in any location.



Symmetry emerges when different ways of representing something "mean" the same thing. *Representation can have symmetry, operations can preserve symmetry, and objects can have symmetry.*

Translation symmetry in 2D:

Features "mean" the same thing in any location.



Symmetry of 2D objects Boundaries "break" global translation symmetry.



Periodic boundary conditions preserve discrete translation symmetry.



Symmetry emerges when different ways of representing something "mean" the same thing.

Representation can have symmetry, operations can preserve symmetry, and objects can have symmetry.

Permutation symmetry, S_N:

Symmetry of sets The freedom to list things in any order



Symmetry emerges when different ways of representing something "mean" the same thing. *Representation can have symmetry, operations can preserve symmetry, and objects can have symmetry.*

Permutation symmetry, S_N :

Symmetry of sets The freedom to list things in any order Symmetry of elements of a graph Graph automorphism, specific nodes are indistinguishable (same global connectivity)



vector in vector space

 $x \in X$ inputs $y\in Y$ outputs $w \in W$ weights function (neural network)...

...which is equivalent to writing. f(x, w) = y $f: X, W \to Y$

vector in vector space

 $x \in X$ inputs $y \in Y$ outputs $w \in W$ weights

element of group $q \in G$

function (neural network)... f(x, w) = y

representation of g acting on vector space

 $D_{\boldsymbol{x}}(\boldsymbol{g}) \\ D_{\boldsymbol{y}}(\boldsymbol{g})$

tor space $D_{oldsymbol{x}}(g):X o X$ $D_{oldsymbol{y}}(g):Y o Y$

...which is equivalent to writing.

 $f: X, W \to Y$

vector in vector space $x \in X$ inputs $y \in Y$ outputs $w \in W$ weights element of group $q \in G$

...which is equivalent to writing. function (neural network)... $f: X, W \to Y$ f(x, w) = yrepresentation of g acting on vector space

 $D_{\boldsymbol{x}}(\boldsymbol{q})$ $D_{\boldsymbol{u}}(\boldsymbol{g})$ $D_x(g): X \to X$ $D_u(g): Y \to Y$

equivariant to x if

 $f(D_{\boldsymbol{x}}(\boldsymbol{g})\boldsymbol{x}, D_{\boldsymbol{w}}(\boldsymbol{g})\boldsymbol{w})$







M. Zaheer et al, Deep Sets, NeurIPS 2017

Theorem 2 A function f(X) operating on a set X having elements from a countable universe, is a valid set function, i.e., **invariant** to the permutation of instances in X, iff it can be decomposed in the form $\rho\left(\sum_{x \in X} \phi(x)\right)$, for suitable transformations ϕ and ρ .

The extension to case when \mathfrak{X} is uncountable, like $\mathfrak{X} = \mathbb{R}$, we could only prove that $f(X) = \rho\left(\sum_{x \in X} \phi(x)\right)$ holds for sets of fixed size. The proofs and difficulties in handling the uncountable case, are discussed in Appendix A. However, we still conjecture that exact equality holds in general.

Next, we analyze the **equivariant** case when $\mathfrak{X} = \mathcal{Y} = \mathbb{R}$ and \mathbf{f} is restricted to be a neural network layer. The standard neural network layer is represented as $\mathbf{f}_{\Theta}(\mathbf{x}) = \boldsymbol{\sigma}(\Theta \mathbf{x})$ where $\Theta \in \mathbb{R}^{M \times M}$ is the weight vector and $\sigma : \mathbb{R} \to \mathbb{R}$ is a nonlinearity such as sigmoid function. The following lemma states the necessary and sufficient conditions for permutation-equivariance in this type of function.

Lemma 3 The function $\mathbf{f}_{\Theta} : \mathbb{R}^M \to \mathbb{R}^M$ defined above is permutation **equivariant** iff all the offdiagonal elements of Θ are tied together and all the diagonal elements are equal as well. That is,

 $\Theta = \lambda \mathbf{I} + \gamma \ (\mathbf{1}\mathbf{1}^{\mathsf{T}}) \qquad \lambda, \gamma \in \mathbb{R} \quad \mathbf{1} = [1, \dots, 1]^{\mathsf{T}} \in \mathbb{R}^{M} \qquad \mathbf{I} \in \mathbb{R}^{M \times M} \text{ is the identity matrix}$

This result can be easily extended to higher dimensions, *i.e.*, $\mathfrak{X} = \mathbb{R}^d$ when λ, γ can be matrices.

Convolutional neural networks can "cheat" by being sensitive to "boundaries".

(e.g. Predict geodesics on projected maps with and without periodic boundary conditions)



Boundaries break symmetry.

User: Stebe https://en.wikipedia.org/wiki/Gall-Peters_projection Pixels cannot be distinguished due to translation equivariance.

Nodes can be distinguished due to differing topology by latitude (e.g. poles)!

94

In the physical sciences... What our our data types? 3D geometry and geometric tensors...

...which transform predictably under 3D rotation, translation, and inversion.

These data types <u>assume</u> Euclidean symmetry. *⇒* Thus, we need neural networks that preserve Euclidean symmetry.

Geometric tensors take many forms. They are a general data type beyond materials.



Our unit test: Trained on 3D Tetris shapes in one orientation, these network can perfectly identify these shapes in any orientation.



Several groups converged on similar ideas around the same time.

Tensor field networks: Rotation- and translation-equivariant neural networks for 3D point clouds (arXiv:1802.08219) Tess Smidt*, Nathaniel Thomas*, Steven Kearnes, Lusann Yang, Li Li, Kai Kohlhoff, Patrick Riley Points, nonlinearity on norm of tensors

Clebsch-Gordan Nets: a Fully Fourier Space Spherical Convolutional Neural Network (arXiv:1806.09231)

Risi Kondor, Zhen Lin, Shubhendu Trivedi

Only use tensor product as nonlinearity, no radial function

 3D Steerable CNNs: Learning Rotationally Equivariant Features in Volumetric Data (arXiv:1807.02547)
Mario Geiger*, Maurice Weiler*, Max Welling, Wouter Boomsma, Taco Cohen Efficient framework for voxels, gated nonlinearity

*denotes equal contribution

Several groups converged on similar ideas around the same time.

Tensor field networks: Rotation- and translati (arXiv:1802.08219) Tess Smidt*, Nathaniel Thomas*, Steven Kearne Points, nonlinearity on norm of tensors

Clebsch-Gordan Nets: a Fully Fourier Space (*arXiv:1806.09231*) Risi Kondor, Zhen Lin, Shubhendu Trivedi

Only use tensor product as nonlinearity, no

3D Steerable CNNs: Learning Rotationally Eq (arXiv:1807.02547) <u>Mario Geiger*</u>, Maurice Weiler*, Max Welling, W Efficient framework for voxels, gated nonlii

*denotes equal contribution

e3nn / e3nn rked from mariogeiger/se3cnn		n - 0 ★ Star 2 ÿ Fork 26
⇔ Code ⊕ Issues 2 n Pull rec ⇔ Settings	Tensor field networ = Euclidean neura	rks + 3D steerable CNN al networks (e3nn)
uclidean Neural Networks		Edit
lanage topics		
⊕ 793 commits	🕆 0 packages 🛛 🗞 2 releases	🛚 8 contributors 🛛 🕸 MIT
Branch: point - New pull reque	Create new file Upload files	Find file Clone or download -
This branch is 20 commits ahead,	1 commit behind mariogeiger:point.	n Pull request 🗈 Compare
🆍 mariogeiger Merge pull reques	t #4 from bkmi/point	Latest commit 917dcb9 yesterday
🖿 e3nn	Merge pull request #4 from bkmi/point	yesterday
examples	refactor into e3nn	6 days ago
src/real_spherical_harmonics	rsh: extended to handle float32	2 months ago
e tests	GatedBlock(Op, Rs_out,)	last month
.gitignore	change directories structure	6 months ago
	Create LICENSE	6 months ago
README.md	refactor e3nn	10 days ago
setup.py	rename e3nn	last month

The group E(3) is the group of 3 dimensional rotations, translations and mirror. This library aims to create E(3) equivariant convolutional neural networks.

Spherical harmonics of a given L transform together under rotation.



How to encode (Pooling layer). Recursively convert geometry to features.



How to decode (Unpooling layer). Recursively convert features to geometry.



We want to convert geometric information (3D coordinates of atomic positions) into features on a trivial geometry (a single point) and back again.



We want to convert geometric information (3D coordinates of atomic positions) into features on a trivial geometry (a single point) and back again.



We want to convert geometric information (3D coordinates of atomic positions) into features on a trivial geometry (a single point) and back again.



Atomic structures are <u>hierarchical</u> and can be constructed from <u>recurring geometric</u> <u>motifs</u>.

- + Encode geometry
- + Encode hierarchy
- + Decode geometry
- + Decode hierarchy

(Need to do this in a recursive manner)

To autoencode, we have to be able to convert geometry into features and vice versa. We do this via spherical harmonic projections.





What a computational materials physicist does:

Given an atomic structure,



...where the electrons are...



...use quantum theory and supercomputers to determine...

 $\hat{H} \left| \psi \right\rangle = E \left| \psi \right\rangle$

...and what the electrons are doing.



Structure Properties We want to use deep learning to speed up calculations, hypothesize new structures, perform inverse design, and organize these relations.




We want to use deep learning to speed up calculations, hypothesize new structures, perform inverse design, and organize these relations.



Given a single example of a degenerate solution, it knows what other solutions are possible by symmetry.

(Useful for ensuring you're not biasing your sampling.)

To be rotation-equivariant means that we can rotate our inputs <u>OR</u> rotate our outputs and we get the same answer *(for every operation)*.

in
$$\rightarrow$$
 Rot \rightarrow Layer \rightarrow out
= in \rightarrow Layer \rightarrow Rot \rightarrow out

For L=1 \Rightarrow L=1, the filters will be a learned, radially-dependent linear combinations of the L = 0, 1, and 2 spherical harmonics.



Random filters for L=1 ⇒ L=1... (3 in L=1 channels by 3 out L=1 channels)

... as a function of increasing r. Time showing filter for varying r, where $0 \le r \le r_{max}$



Radial distance is magnitude as a function of angle



Predictions for O_h symmetry

4

0.2

0

_0.2

.0.4

.0.4

Y

_0.2



Prediction of network trained with symmetry breaking input and given symmetry breaking input along z.





deep learning \subset machine learning \subset artificial intelligence

model | deep learning | data | cost function | way to update parameters | conv. nets ¹¹⁵

model ("neural network"): Function with learnable parameters.



model | deep learning | data | cost function | way to update parameters | conv. nets ¹¹⁶

model ("neural network"): Function with learnable parameters.

Ex: "Fully-connected" network



model ("neural network"): Function with learnable parameters.

Ex: "Fully-connected" network

118

Neural networks with multiple layers can learn more complicated functions.

model | deep learning | data | cost function | way to update parameters | conv. nets



model | deep learning | data | cost function | way to update parameters | conv. nets

119

deep learning: Add more layers.



model | deep learning | data | cost function | way to update parameters | conv. nets ¹²⁰

data:

Want lots of it. Model has many parameters. Don't want to easily overfit.



model | deep learning | data | cost function | way to update parameters | conv. nets

121

cost function:

A metric to assess how well the model is performing. The cost function is evaluated on the output of the model. Also called the **loss** or **error**.

way to update parameters:

Construct a model that is differentiable

Easiest to do with differentiable programming frameworks: e.g. Torch, TensorFlow, JAX, ... Take derivatives of the cost function (loss or error) wrt to learnable parameters. This is called backpropogation (aka the chain rule).

$$\Delta W_{ij} = -\eta \frac{\partial \operatorname{error}(f(W, x), y)}{\partial W_{ij}}$$



convolutional neural networks:

Used for images. In each layer, scan over image with learned filters.



http://deeplearning.stanford.edu/wiki/index.php/Feature_extraction_using_convolution

model | deep learning | data | cost function | way to update parameters | conv. nets

convolutional neural networks:

Used for images. In each layer, scan over image with learned filters.



http://cs.nyu.edu/~fergus/tutorials/deep_learning_cvpr12/

model | deep learning | data | cost function | way to update parameters | conv. nets

125