

Euclidean Neural Networks...

rotation-, translation-, and
permutation-equivariant
convolutional neural networks
for 3D point clouds

**...for emulating ab initio calculations and
generating atomic geometries.**

Tess Smidt

*2018 Alvarez Postdoctoral Fellow in
Computing Sciences*




*Machine Learning for Quantum Matter
Nordita, Stockholm, Sweden
2019.08.28*



What a computational materials physicist does:


Given an atomic structure,




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

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

...where the electrons are...



...and what the electrons are doing.




Structure




Properties

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



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



What types of neural networks are best suited for these tasks?


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



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



Vectors $\Rightarrow$ *Dense NN*




Components are independent.

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




Vectors $\Rightarrow$ *Dense NN*



Components are independent.

2D images $\Rightarrow$ *Convolutional NN*




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

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




Vectors $\Rightarrow$ **Dense NN**




Components are independent.

2D images $\Rightarrow$ **Convolutional NN**



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

Text $\Rightarrow$ **Recurrent NN**




Sequential data. Next input/output depends on input/output that has come before.

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




Vectors $\Rightarrow$ **Dense NN**




Components are independent.

2D images $\Rightarrow$ **Convolutional NN**



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

Text $\Rightarrow$ **Recurrent NN**




Sequential data. Next input/output depends on input/output that has come before.

What are our data types in materials physics?
How do we build neural networks for these data types?


What assumptions do we want “built in” to our neural networks (for materials data)?

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




The properties of physical systems transform predictably under rotation.

Two point **masses** with **velocity** and **acceleration**.




Same system, with rotated coordinates.




What assumptions do we want “built in” to our neural networks (for materials data)?

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




The properties of physical systems transform predictably under rotation.

Two point **masses** with **velocity** and **acceleration**.







Same system, with rotated coordinates.




Our data types are geometry and geometric tensors.

These data types assume Euclidean symmetry (3D translations, 3D rotations, and inversion),




**Space has Euclidean symmetry, $E(3)$. Objects break that symmetry.
The broken symmetry is a subgroup of $E(3)$.**




**Space has Euclidean symmetry, $E(3)$. Objects break that symmetry.
The broken symmetry is a subgroup of $E(3)$.**




$O(3)$





Space has Euclidean symmetry, $E(3)$. Objects break that symmetry.
The broken symmetry is a subgroup of $E(3)$.




$O(3)$




$SO(2) +$
mirrors
 $(C_{\infty v})$




Space has Euclidean symmetry, $E(3)$. Objects break that symmetry.
The broken symmetry is a subgroup of $E(3)$.





$O(3)$




$SO(2) +$
mirrors
 $(C_{\infty v})$




O_h




Space has Euclidean symmetry, $E(3)$. Objects break that symmetry.
The broken symmetry is a subgroup of $E(3)$.





$O(3)$



$SO(2) +$
mirrors
($C_{\infty v}$)





O_h





$Pm-3m$
(221)


Properties of a system must be compatible with symmetry.
Which of these situations are symmetrically allowed / forbidden?




Properties of a system must be compatible with symmetry.
Which of these situations are symmetrically allowed / forbidden?



Properties of a system must be compatible with symmetry.
Which of these situations are symmetrically allowed / forbidden?



Properties of a system must be compatible with symmetry.
Which of these situations are symmetrically allowed / forbidden?




We build neural networks with Euclidean symmetry, $E(3)$ and $SE(3)$.


- What neural networks with Euclidean symmetry can do.
- How Euclidean Neural Networks work.
- Applications of Euclidean Neural Networks.

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



TRAIN




TEST



**Given a molecule and a rotated copy,
the predicted forces are the same up to rotation.
(Predicted forces are equivariant to rotation.)**




To these networks, primitive unit cells, conventional unit cells, and supercells of the same crystal will produce the same output (assuming periodic boundary conditions).




We build neural networks with Euclidean symmetry, $E(3)$ and $SE(3)$.

- What neural networks with Euclidean symmetry can do.
- How Euclidean Neural Networks work.
 - Overview
 - Input to network
 - Network operations
 - Visualizing kernels
 - Interpreting input / output
- Applications of Euclidean Neural Networks.


We use points. Images of atomic systems are sparse and imprecise.



VS.





We use continuous convolutions with atoms as convolution centers.




K. T. Schütt et al, NIPS 30 (2017).
(arXiv: 1706.08566)


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




We use points. Images of atomic systems are sparse and imprecise.



VS.





We use continuous convolutions with atoms as convolution centers.




K. T. Schütt et al, NIPS 30 (2017).
(arXiv: 1706.08566)


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




We use points. Images of atomic systems are sparse and imprecise.



VS.




We use continuous convolutions with atoms as convolution centers.




K. T. Schütt et al, NIPS 30 (2017).
(arXiv: 1706.08566)

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




$$g \in SE(3)$$






Translation equivariance





Rotation equivariance



Translation equivariance

Convolutional neural network ✓




Rotation equivariance?



Translation equivariance

Convolutional neural network ✓



Rotation equivariance

~~Data augmentation~~

~~Radial functions~~

Want a network that both preserves geometry and exploits symmetry.

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 translation-equivariant
(arXiv:1802.08219)

Tess Smidt*, Nathaniel Thomas*, Steven Kneip*
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



3D Steerable CNNs: Learning Rotationally
(arXiv:1807.02547)

Mario Geiger*, Maurice Weiler*, Max Welling*
Efficient framework for voxels, gated networks


**denotes equal contribution*



To be rotation-equivariant means that we can rotate our inputs
OR rotate our outputs and we get the same answer.




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



```
[[[m0]], [[m1]]],  
[[[v0x, v0y, v0z], [a0x, a0y, a0z]],  
 [[v1x, v1y, v1z], [a1x, a1y, a1z]]]
```

The input to our network is geometry and features on that geometry.
We categorize our features by how they transform under rotation.

Features have periodicity $2\pi/L$
where L is a positive integer.




	<u>Frequency</u>
Scalars	$l = 0$
Vectors	$l = 1$
3x3 Matrices	$l = 0 \oplus 1 \oplus 2$

```
{0: [[m0], [m1]],
  1: [[v0x, v0y, v0z], [a0x, a0y, a0z],
      [v1x, v1y, v1z], [a1x, a1y, a1z]]}
```

The convolutional kernels are built from functions with period $2\pi/L$

⇒ *Spherical harmonics*.



Learned Parameters

with no symmetry:

$$W(\vec{r})$$

with SO(3) symmetry:

$$R(r)Y_l^m(\hat{r})$$


Spherical harmonics

$$Y_l^m$$

angular portion of
hydrogenic wavefunctions

basis functions for $(2l + 1)$
dimensional irreducible
representations of $SO(3)$

basis functions for signals
on a sphere





Spherical harmonics of a given L transform together under rotation.


Let $\mathbf{g}$ be a 3d rotation matrix.



$\mathbf{D}$ is the Wigner-D matrix.
It has shape $[2l + 1, 2l + 1]$
and is a function of $\mathbf{g}$.



=



Features and kernels are not simply scalars.

We use tensor products with Clebsch-Gordan coefficients to combine.

Same math involved in the addition of angular momentum.

The diagram illustrates the addition of angular momentum for various combinations of two particles. Each combination is represented by a set of boxes containing Clebsch-Gordan coefficients. The combinations include $1/2 \times 1/2$, $1 \times 1/2$, $2 \times 1/2$, $3/2 \times 1/2$, 1×1 , 2×1 , $3/2 \times 1$, and $3 \times 1/2$. The boxes are arranged in a hierarchical, branching structure, with each box containing a 3x3 grid of coefficients. The coefficients are fractions and integers, representing the coupling of different angular momentum states.

Examples of tensor product: How to combine a scalar and a vector? Easy!

Angular Frequency

$$a \times \vec{b} = \vec{c}$$

1

Examples of tensor product: How to combine two vectors? Many ways.

Dot product $(a_i \quad a_j \quad a_k) \begin{pmatrix} b_i \\ b_j \\ b_k \end{pmatrix} = c$

Cross product $\vec{a} \times \vec{b} = \begin{vmatrix} \hat{i} & \hat{j} & \hat{k} \\ a_i & a_j & a_k \\ b_i & b_j & b_k \end{vmatrix} = \vec{c}$

Outer product $\begin{pmatrix} a_i \\ a_j \\ a_k \end{pmatrix} (b_i \quad b_j \quad b_k) = \begin{pmatrix} a_i b_i & a_i b_j & a_i b_k \\ a_j b_i & a_j b_j & a_j b_k \\ a_k b_i & a_k b_j & a_k b_k \end{pmatrix}$

Angular Frequency

0

1

$0 \oplus 1 \oplus 2$

Features and kernels are not simply scalars.

We use tensor products and Clebsch-Gordan coefficients to combine.




Diagram illustrating the combination of features and kernels using tensor products and Clebsch-Gordan coefficients. The diagram shows a vector space with points a (red dot) and b (green dot). Vectors $\vec{r}_{ab}$ are shown originating from a and pointing towards b . Two equations are presented, each with a handwritten label and an arrow pointing to the operation symbol:

Equation 1: $\sum_b W_{ij}(\vec{r}_{ab}) \otimes I_{bi} = O_{aj}$
Label: scalar multiply (arrow points to $\otimes$)
Label: tensor product + Clebsch-Gordan (arrow points to $\otimes$)

Equation 2: $\sum_b W_j(\vec{r}_{ab}) \otimes I_{bi} = O_{ak}$
Label: scalar multiply (arrow points to $\otimes$)
Label: tensor product + Clebsch-Gordan (arrow points to $\otimes$)

Features and kernels are not simply scalars.

We use tensor products and Clebsch-Gordan coefficients to combine.

Clebsch-Gordan
↓ coeffs.

$$S_i \otimes S_j = C_{ijk} S_{ij} = S_k$$

$|i, i_m\rangle |j, j_m\rangle \langle k, k_m|$

non-zero when $-|i-j| \leq k \leq i+j$

Features and kernels are not simply scalars.

We use tensor products and Clebsch-Gordan coefficients to combine.

$\frac{1}{2}$ -Spin eigenstates

$$\begin{array}{cc} |\uparrow\rangle & \text{and} & |\downarrow\rangle \\ \frac{1}{2}, \frac{1}{2} & & \frac{1}{2}, -\frac{1}{2} \end{array}$$

2 $\frac{1}{2}$ -Spin eigenstates

$$|0,0\rangle = \frac{1}{\sqrt{2}} (|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle) \quad \text{singlet}$$

$$|1,1\rangle = |\uparrow\uparrow\rangle$$

$$|1,0\rangle = \frac{1}{\sqrt{2}} (|\uparrow\downarrow\rangle + |\downarrow\uparrow\rangle) \quad \text{triplet}$$

$$|1,-1\rangle = |\downarrow\downarrow\rangle$$

$$\frac{1}{2} \otimes \frac{1}{2} \rightarrow S_i \otimes S_j = C_{ijk} S_k \rightarrow 0 \oplus 1$$

$$\dots, | \underset{i}{\frac{1}{2}}, -\frac{1}{2} \rangle | \underset{j}{\frac{1}{2}}, \frac{1}{2} \rangle \underset{k}{\langle 0,0|}, \dots$$

Features and kernels are not simply scalars.

We use tensor products and Clebsch-Gordan coefficients to combine.


$$\mu_{ij} = \begin{bmatrix} a_{xx} & a_{xy} & a_{xz} \\ a_{yx} & a_{yy} & a_{yz} \\ a_{zx} & a_{zy} & a_{zz} \end{bmatrix}$$

$$\mu_{ij} = C_{ijk} S_k$$

$$\dots, |1,0\rangle |1,0\rangle \langle 2,0|, \dots$$

$z \quad z \quad z^2 - x^2 - y^2$

$$1 \otimes 1 = 0 \oplus 1 \oplus 2$$



$$\begin{aligned} a_{x^2} &= \frac{1}{\sqrt{3}} Y_{0,0} - \frac{1}{\sqrt{30}} Y_{2,0} + \frac{1}{\sqrt{10}} Y_{2,2} \\ a_{y^2} &= \frac{1}{\sqrt{3}} Y_{0,0} - \frac{1}{\sqrt{30}} Y_{2,0} - \frac{1}{\sqrt{10}} Y_{2,2} \\ a_{z^2} &= \frac{1}{\sqrt{3}} Y_{0,0} + \frac{2}{\sqrt{30}} Y_{2,0} \end{aligned}$$

$$\begin{aligned} a_{xy/yz} &= \frac{1}{\sqrt{10}} Y_{2,-2} \pm \frac{1}{\sqrt{6}} Y_{1,0} \\ a_{xz/zx} &= \frac{1}{\sqrt{10}} Y_{2,1} \pm \frac{1}{\sqrt{6}} Y_{1,-1} \\ a_{yz/zy} &= \frac{1}{\sqrt{10}} Y_{2,-1} \pm \frac{1}{\sqrt{6}} Y_{1,1} \end{aligned}$$

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.

$$R(r) Y_l^m(\hat{r}) C_{lij} = K_{ij}$$



$$R_0(r) \overset{\text{L}=0}{\text{[spherical]}} + R_1(r) \overset{\text{L}=1}{\text{[dipole]}} + R_2(r) \overset{\text{L}=2}{\text{[quadrupole]}} = K_{ij}$$

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.

$$R(r)$$


$L=0$

$$R_0(r)$$



$$K_{ij}$$

$L=2$






$$= K_{ij}$$

We can interpret our outputs as numerical features or geometry.




We can interpret our outputs as numerical features or geometry.

Example of $L < 6$ coefficients interpreted as geometry for randomly initialized network applied to a tetrahedron with a center.



We can interpret our outputs as numerical features or geometry.

Example of $L < 6$ coefficients interpreted as geometry for randomly initialized network applied to a tetrahedron with a center.



We can generate point sets from using peaks of spherical harmonic signals!
Sets == permutation invariant. Difficult for neural networks.



We build neural networks with Euclidean symmetry, $E(3)$ and $SE(3)$.

- What neural networks with Euclidean symmetry can do.
- How Euclidean Neural Networks work.
- **Applications of Euclidean Neural Networks.**

Applications: Predicting ab initio forces for molecular dynamics

Simon Batzner (MIT/Harvard) and Boris Kozinsky (Harvard)

Presented at APS March Meeting 2019



*Direct prediction
of forces rather
than gradient of
scalar energy.*

Applications: Predicting ab initio forces for molecular dynamics

Simon Batzner (MIT/Harvard) and Boris Kozinsky (Harvard)

Presented at APS March Meeting 2019

Force predictions for MD17 dataset

comparison to SchNet values in [1], Force MAE in kcal/(mol Angstrom)

Training Set: 1,000 examples	Molecule	Direct Prediction	SchNet
	Benzene	0.19	0.31
	Toluene	0.39	0.57
	Malonaldehyde	1.00	0.66
	Salicylic acid	0.78	0.85
	Aspirin	1.09	1.35
	Ethanol	0.73	0.39
	Uracil	0.60	0.56
	Naphthalene	0.30	0.58



Training Set: 50,000 examples	Molecule	Direct Prediction	SchNet
	Benzene	0.16	0.17
	Toluene	0.08	0.09
	Malonaldehyde	0.08	0.08
	Salicylic acid	0.16	0.19
	Aspirin	0.24	0.33
	Ethanol	0.06	0.05
	Uracil	0.06	0.11
	Naphthalene	0.06	0.11



*Direct prediction
of forces rather
than gradient of
scalar energy.*



Applications: Predicting molecular Hamiltonians with atom-centered basis sets

K. T. Schütt, M. Gastegger,
A. Tkatchenko, K.-R. Müller,
R. J. Maurer.
arXiv:1906.10033 (2019)



Applications: Predicting molecular Hamiltonians with atom-centered basis sets

K. T. Schütt, M. Gastegger,
A. Tkatchenko, K.-R. Müller,
R. J. Maurer.
arXiv:1906.10033 (2019)



Problem! Hamiltonian matrix looks very different depending on coordinate system -- traditionally requires augmenting data.





Mario 2019.07.02

Applications: Predicting molecular Hamiltonians with atom-centered basis sets

With Euclidean neural networks -- we only need to learn on one example and the output is guaranteed to be equivariant!

$$RHR^T$$


$$R\vec{v}$$




Mario 2019.07.02

Applications: Reconstructing atomic positions from coarse-grained geometries for molecular dynamics


W. Wang, R. Gómez-Bombarelli. arXiv:1812.02706 (revised 2019)



original molecule



coarse grained molecule




decoded molecule




Ben 2019.07.08

Applications: Reconstructing atomic positions from coarse-grained geometries for molecular dynamics


W. Wang, R. Gómez-Bombarelli. arXiv:1812.02706 (revised 2019)



original molecule




coarse grained molecule





decoded molecule

Test problem



Centers of a tetrahedral chain.




Tetrahedral chain.



Ben 2019.07.08


Applications: Reconstructing atomic positions from coarse-grained geometries for molecular dynamics



Peaks of spherical harmonic signal represent point locations.


Model preserves degeneracy of symmetric configurations

We can break symmetry with latent variable sampling.




Ben 2019.07.08

Applications: Creating an autoencoder for discrete geometry



Applications: Creating an autoencoder for discrete geometry





Atomic structures are hierarchical and can be constructed from geometric motifs.

- + Encode geometry ✓
- + Encode hierarchy ?
- + Decode geometry ?
- + Decode hierarchy ?



(Need to do this in a recursive manner)

How to encode: Recursively convert geometry to a vector



**Edges are shown for visualization. May not be included.*

How to decode: Recursively convert a vector to geometry



**Edges are shown for visualization. May not be included.*

atomic architects summer 2019



Tess Smidt



Hashim Piracha



Mario Geiger



Ben Miller

tensor field networks

Stanford



Nate
Thomas

Google Accelerated Science Team



Patrick
Riley



Steve
Kearnes



Lusann
Yang



Li
Li



Kai
Kohlhoff

- ***Euclidean neural networks*** operate on points/voxels and have symmetries of $E(3)$.
- Inputs to the network lower this symmetry to a subgroup of $E(3)$.
- Symmetry of outputs are constrained to the symmetry of the inputs.
- The inputs and outputs of our network are geometry and geometric tensors.
- Convolutional filters are built from spherical harmonics with a learned radial function.

Applications: Molecular dynamics, predicting Hamiltonians, coarse-graining, autoencoders...

We expect these networks to be generally useful for physics, chemistry, and geometry.

Reach out to me if you are interested and/or have any questions!

se3cnn Code (PyTorch):

<https://github.com/mariogeiger/se3cnn>

Tensor Field Networks (arXiv:1802.08219)

3D Steerable CNNs (arXiv:1807.02547)

Tess Smidt
tsmidt@lbl.gov

Calling in backup (slides)!

