


Revisiting convolutions



Definition


Definition The convolution of two functions f and g is denoted by $*$ as the integral of the product of the two functions after one is reversed and shifted

$$(f * g)(t) \stackrel{\text{def}}{=} \int_{-\infty}^{\infty} f(\tau)g(t - \tau) d\tau = \int_{-\infty}^{\infty} f(t - \tau)g(\tau) d\tau$$

- For images $a_{rc} = \mathbf{x} * \mathbf{w} = \sum_{i=-a}^a \sum_{j=-b}^b x_{r-i, c-j} \cdot w_{ij}$
- To generalize to graphs, we must understand convolutions to their core
- Let's check some properties of what makes convolution a convolution

Convolutions matrices


- Fully connected \Leftrightarrow Full matrix multiplication




$$y_j = w_{j1}x_1 + \dots + w_{jn}x_n$$

$$= \sum_i w_{ji}x_i$$

Weight matrix w




- Convolution \Leftrightarrow Block diagonal matrix multiplication
 - Sharing weights after shifting them



$$y_j = w_{j,i-1}x_{i-1} + w_{j,i}x_i + w_{j,i+1}x_{i+1}$$

Convolutional weight matrix w




Convolutions as circulant matrices

- Convolutional matrices $C(w)$ are circulant
- Multidiagonal matrices
 - Each column (row) as above but shifted once to the right (below)

$$C(w) = \begin{bmatrix} \boxed{c_1} & c_3 & c_2 \\ c_2 & c_1 & c_3 \\ c_3 & c_2 & c_1 \end{bmatrix}$$

w



<https://towardsdatascience.com/deriving-convolution-from-first-principles-4ff124888028>

Circulant matrices commute


- In general

$$A \cdot B \neq B \cdot A$$

- For circulant matrices

$$C(w) \cdot C(u) = C(u) \cdot C(w)$$

- \Rightarrow Convolutions commute



The shift operator

- For $w = [0, 1, \dots, 0] \Rightarrow C(w)$ the right-shift operator
 - Similar to convolution: 'shift once' to the right
 - Transpose for left-shift
- The shift operator is an orthogonal matrix
- The shift operator is also circulant

The diagram illustrates the shift operator S and its transpose S^T using 4x4 matrices and vectors.

Top row: $y = Sx$ and $y = S^T x$. The vector x is $[0, 1, 0, 0]^T$ (yellow, blue, green, yellow). The vector y is $[0, 0, 1, 0]^T$ (blue, green, yellow, blue). The matrix S is a right-shift operator, and S^T is its transpose, a left-shift operator.

Bottom row: $S S^T = S^T S = I$, where I is the 4x4 identity matrix.

Circulant matrices \Leftrightarrow Translation equivariance

- Circulant matrices enable translation equivariance to convolutions
 - Change the location of the input
 - The results will be the same (but shifted)





Illustration of shift equivariance as the interchangeability of shift and blur operations.

All circulant matrices have the same eigenvectors!

<https://github.com/mitmath/1806/blob/master/lectures/Circulant-Matrices.ipynb>

- The eigenvalues of the circulant matrices are different
- But the eigenvectors always the same!
 - The eigenvectors of the “translation transformation/operator”

All circulant matrices have the same eigenvectors!

<https://github.com/mitmath/1806/blob/master/lectures/Circulant-Matrices.ipynb>

- The eigenvalues of the circulant matrices are different
- But the eigenvectors always the same!
 - The eigenvectors of the “translation transformation/operator”

```
In [26]: A = circulant([-1, 2, 1, 0, 0])
print('Circulant matrix')
print(A)

eigvals, eigvecs = np.linalg.eig(A)
print('\nEigenvalues')
print(eigvals)

print('\nEigenvectors')
print(eigvecs)
```

Circulant matrix
[[-1 0 0 1 2]
[2 -1 0 0 1]
[1 2 -1 0 0]
[0 1 2 -1 0]
[0 0 1 2 -1]]

Eigenvalues
[2. +0.j -1.19+2.49j -1.19-2.49j -2.31+0.22j -2.31-0.22j]

Eigenvectors
[[0.45+0.j 0.14-0.43j 0.14+0.43j -0.36+0.26j -0.36-0.26j]
[0.45+0.j -0.36-0.26j -0.36+0.26j 0.45+0.j 0.45-0.j]
[0.45+0.j -0.36+0.26j -0.36-0.26j -0.36-0.26j -0.36+0.26j]
[0.45+0.j 0.14+0.43j 0.14-0.43j 0.14+0.43j 0.14-0.43j]
[0.45+0.j 0.45+0.j 0.45-0.j 0.14-0.43j 0.14+0.43j]]

```
In [40]: v = np.random.rand(3)
z = np.zeros(2)
A = circulant(np.append(v, z))
print('Circulant matrix')
print(A)

eigvals, eigvecs = np.linalg.eig(A)
print('\nEigenvalues')
print(eigvals)

print('\nEigenvectors')
print(eigvecs)
```

Circulant matrix
[[0.87 0. 0. 0.61 0.13]
[0.13 0.87 0. 0. 0.61]
[0.61 0.13 0.87 0. 0.]
[0. 0.61 0.13 0.87 0.]
[0. 0. 0.61 0.13 0.87]]

Eigenvalues
[1.61+0.j 0.42+0.48j 0.42-0.48j 0.95+0.5j 0.95-0.5j]

Eigenvectors
[[0.45+0.j 0.14-0.43j 0.14+0.43j -0.36-0.26j -0.36+0.26j]
[0.45+0.j -0.36-0.26j -0.36+0.26j 0.45+0.j 0.45-0.j]
[0.45+0.j -0.36+0.26j -0.36-0.26j -0.36+0.26j -0.36-0.26j]
[0.45+0.j 0.14+0.43j 0.14-0.43j 0.14-0.43j 0.14+0.43j]
[0.45+0.j 0.45+0.j 0.45-0.j 0.14+0.43j 0.14-0.43j]]

Circulant eigenvectors \Leftrightarrow Shift eigenvectors

- All circulant matrices have the same eigenvectors (or better eigenspace)
 - The shift operator is a circulant matrix
- The circulant eigenvectors are the eigenvectors of shift
 - No wonder they are the same: shift is always the same
- Any convolution with any filter \mathbf{w} involves the same eigenvectors!


Eigenvectors of circulant matrices

- The k -th eigenvector of $n \times n$ circulant

$$e^{(k)} = \begin{bmatrix} \omega_n^{0 \cdot k} \\ \omega_n^{1 \cdot k} \\ \omega_n^{2 \cdot k} \\ \vdots \\ \omega_n^{(n-1) \cdot k} \end{bmatrix}, \text{ where } \omega_n = \exp\left(\frac{2\pi \cdot i}{n}\right)$$

- Collecting all eigenvectors

$$\Phi = [e^{(0)} \quad e^{(1)} \quad e^{(2)} \quad \dots \quad e^{(n-1)}]$$



Circulant matrix eigenvectors \Rightarrow Discrete Fourier Transform

- This looks a lot like Discrete Fourier Transform
 - The computer friendly Fourier Transform

$$\begin{aligned} X_k &= \sum_{n=0}^{N-1} x_n \cdot e^{-\frac{i2\pi}{N}kn} \\ &= \sum_{n=0}^{N-1} x_n \cdot \left[\cos\left(\frac{2\pi}{N}kn\right) - i \cdot \sin\left(\frac{2\pi}{N}kn\right) \right], \end{aligned} \quad (\text{Eq.1})$$

- Convolution \Leftrightarrow Discrete Fourier Transform

Matrix diagonalization \rightarrow

$$\begin{aligned} \mathbf{x} * \mathbf{w} &= \mathcal{C}(\mathbf{w}) \cdot \mathbf{x} \\ &= (\Phi \cdot \Lambda(\mathbf{w}) \cdot \Phi^*) \cdot \mathbf{x} \\ &= \Phi \cdot \underbrace{(\Lambda(\mathbf{w}) \cdot \underbrace{(\Phi^* \cdot \mathbf{x})}_{\text{Discrete Fourier Transform (with DFT matrix)}}))}_{\text{Inner product with eigenvalues of weight matrix}} \leftarrow \text{Convolution theorem} \\ &\quad \underbrace{\hspace{10em}}_{\text{Inverse Discrete Fourier Transform (with Inverse DFT matrix)}} \end{aligned}$$

Convolution theorem


- The Fourier of a convolution equal to dot product of individual Fourier

$$\mathcal{F}\{f * g\} = \mathcal{F}\{f\} \odot \mathcal{F}\{g\} \Rightarrow f * g = \mathcal{F}^{-1}\{\mathcal{F}\{f\} \odot \mathcal{F}\{g\}\}$$

- Convolution in “time/space” domain is equivalent to matrix multiplication in “frequency/spectral” domain
 - Frequency defined by Fourier bases $\exp(-\frac{i2\pi}{N} \cdot kn)$
- Discrete case $\mathcal{F}\{f\}$ becomes a matrix multiplication with shift DFT matrix

$$\mathbf{w} * \mathbf{x} = \underbrace{\Phi^{-1} \left(\underbrace{\Lambda(\mathbf{w})}_{\mathcal{F}\{f\}} \cdot \underbrace{(\Phi \cdot \mathbf{x})}_{\mathcal{F}\{g\}} \right)}_{\mathcal{F}^{-1}}$$


Convolution theorem: $\mathbf{x} * \mathbf{w} = \Phi \cdot (\Lambda(\mathbf{w}) \cdot (\Phi^* \cdot \mathbf{x}))$



<https://towardsdatascience.com/deriving-convolution-from-first-principles-4ff124888028>

Convolution theorem: $\mathbf{x} * \mathbf{w} = \Phi \cdot (\Lambda(\mathbf{w}) \cdot (\Phi^* \cdot \mathbf{x}))$

- If we can compute (inverse) Fourier transforms and their inverse fast, then we are game
- Fast Fourier Transform (FFT): A faster version of DFT
 - $O(n \log n)$ vs $O(n^2)$
 - Replace sliding window convolutions with very fast matrix multiplications
- Convolution as diagonalization of convolutional circulant matrix



<https://towardsdatascience.com/deriving-convolution-from-first-principles-4ff124888028>

So what?


- A more core understanding of what actually convolutions do
- Can we generalize to other equivariances beyond translation?

Group Equivariant Deep Learning


Group convolutional neural networks¹ (G-CNNs) improve over classical CNNs by:

- Allowing weight sharing beyond just translations
- Making geometric data augmentations obsolete
- Data efficiency (one example at a some pose is enough)
- Deal with context (relative poses, like capsule nets)


Symmetries in audio⁵ (translation, scale/pitch)




Symmetries in computer vision^{3,4} (translation, scale, rotation, perspective)



Symmetries in medical image analysis^{2,3} (translation, rotation, scale)



Molecular and Physical systems⁶ (translation, rotation, reflection)



So what?

- A more core understanding of what actually convolutions do
- Can we generalize to other equivariances beyond translation?
- Can we generalize to other structures, like graphs?