

Neural Tangent Kernels for Equivariant Neural Networks

Talk at the GeUmetric Deep Learning Workshop

@Umeå

Philipp Misof

Department of Mathematical Sciences, Division of Algebra and Geometry

June 12, 2024



UNIVERSITY OF
GOTHENBURG



CHALMERS
UNIVERSITY OF TECHNOLOGY

WASP | WALLENBERG AI,
AUTONOMOUS SYSTEMS
AND SOFTWARE PROGRAM

Joint work with



Jan Gerken



Pan Kessel



arXiv:2406.06504

Outline

- 1 Motivation

Outline

- 1 Motivation
- 2 The Neural Tangent Kernel

Outline

- 1 Motivation
- 2 The Neural Tangent Kernel
- 3 Extension to Equivariant Neural Networks

Outline

- 1 Motivation
- 2 The Neural Tangent Kernel
- 3 Extension to Equivariant Neural Networks
- 4 Example: $C_4 \ltimes \mathbb{R}^2$

Outline

- 1 Motivation
- 2 The Neural Tangent Kernel
- 3 Extension to Equivariant Neural Networks
- 4 Example: $C_4 \times \mathbb{R}^2$
- 5 Experiments

Outline

- 1 Motivation
- 2 The Neural Tangent Kernel
- 3 Extension to Equivariant Neural Networks
- 4 Example: $C_4 \times \mathbb{R}^2$
- 5 Experiments
- 6 Summary

- 1 Motivation
- 2 The Neural Tangent Kernel
- 3 Extension to Equivariant Neural Networks
- 4 Example: $C_4 \times \mathbb{R}^2$
- 5 Experiments
- 6 Summary

Neural Tangent Kernel (NTK) encodes training dynamics

Neural Tangent Kernel (NTK) encodes training dynamics

- Assuming **Gradient Flow** of a NN $\mathcal{N} : \mathbb{R}^{n_{\text{in}}} \rightarrow \mathbb{R}^{n_{\text{out}}}$

$$\dot{\mathcal{N}}(x) = -\eta \sum_{i=1}^{n_{\text{train}}} \Theta_t(x, x_i) \frac{\partial \mathcal{L}}{\partial \mathcal{N}(x_i)}$$

- $\mathcal{N} \dots$ Neural Network (NN)
- $\Theta_t(x, x_i) \dots$ NTK
- $\mathcal{L} \dots$ loss function
- $\eta \dots$ learning rate


Neural Tangent Kernel (NTK) encodes training dynamics

- Assuming **Gradient Flow** of a NN $\mathcal{N} : \mathbb{R}^{n_{\text{in}}} \rightarrow \mathbb{R}^{n_{\text{out}}}$

$$\dot{\mathcal{N}}(x) = -\eta \sum_{i=1}^{n_{\text{train}}} \Theta_t(x, x_i) \frac{\partial \mathcal{L}}{\partial \mathcal{N}(x_i)}$$

- $\mathcal{N} \dots$ Neural Network (NN)
- $\mathcal{L} \dots$ loss function
- $\Theta_t(x, x_i) \dots$ NTK
- $\eta \dots$ learning rate

- Infinite width limit** $\rightarrow \Theta_t(x, x_i)$ becomes
 - deterministic
 - time-independent
- Combine with **MSE loss** \rightarrow linear ODE with **analytical solution**



What do we gain?

What do we gain?

- Analytical **study of training** independent of particular initialization
 - Influence of **hyperparameters** (*Geiger et al. 2020*)
 - Impact of **data augmentation** (*Gerken and Kessel 2024*)
 - Study of spectral **bias** of specific architectures and data distributions (*Bowman and Montufar 2022*)

What do we gain?

- Analytical **study of training** independent of particular initialization
 - Influence of **hyperparameters** (*Geiger et al. 2020*)
 - Impact of **data augmentation** (*Gerken and Kessel 2024*)
 - Study of spectral **bias** of specific architectures and data distributions (*Bowman and Montufar 2022*)
- New **kernel method** performant on small datasets (*Arora et al. 2019*)

What do we gain?

- Analytical **study of training** independent of particular initialization
 - Influence of **hyperparameters** (*Geiger et al. 2020*)
 - Impact of **data augmentation** (*Gerken and Kessel 2024*)
 - Study of spectral **bias** of specific architectures and data distributions (*Bowman and Montufar 2022*)
- New **kernel method** performant on small datasets (*Arora et al. 2019*)

→ Let's extend this to **equivariant neural networks!**

- 1 Motivation
- 2 The Neural Tangent Kernel**
- 3 Extension to Equivariant Neural Networks
- 4 Example: $C_4 \times \mathbb{R}^2$
- 5 Experiments
- 6 Summary

The empirical NTK

Empirical NTK

$$\Theta_t(\mathbf{x}, \mathbf{x}_i) = \frac{\partial \mathcal{N}(\mathbf{x})}{\partial \theta} \left(\frac{\partial \mathcal{N}(\mathbf{x}_i)}{\partial \theta} \right)^\top$$

The empirical NTK

Empirical NTK

$$\Theta_t(\mathbf{x}, \mathbf{x}_i) = \frac{\partial \mathcal{N}(\mathbf{x})}{\partial \theta} \left(\frac{\partial \mathcal{N}(\mathbf{x}_i)}{\partial \theta} \right)^\top$$

- In general **complicated**

The empirical NTK

Empirical NTK

$$\Theta_t(\mathbf{x}, \mathbf{x}_i) = \frac{\partial \mathcal{N}(\mathbf{x})}{\partial \theta} \left(\frac{\partial \mathcal{N}(\mathbf{x}_i)}{\partial \theta} \right)^\top$$

- In general **complicated**
- Depends on particular **initialization** $\theta(0)$

The empirical NTK

Empirical NTK

$$\Theta_t(\mathbf{x}, \mathbf{x}_i) = \frac{\partial \mathcal{N}(\mathbf{x})}{\partial \theta} \left(\frac{\partial \mathcal{N}(\mathbf{x}_i)}{\partial \theta} \right)^\top$$

- In general **complicated**
- Depends on particular **initialization** $\theta(0)$
- **time-dependent**

NNs become Gaussian processes

- At $t = 0$, all parameters $\theta_i(0)$ are *iid*
- **law of large numbers**

$$\frac{1}{n} \sum_{i=1}^n X_i \rightarrow \mathbb{E}[X], \quad X_i \text{ iid}$$

NNs become Gaussian processes

- At $t = 0$, all parameters $\theta_i(0)$ are iid
- law of large numbers

$$\frac{1}{n} \sum_{i=1}^n X_i \rightarrow \mathbb{E}[X], \quad X_i \text{ iid}$$

Neural Network Gaussian Process

(Lee et al. 2018)

Gaussian process with zero mean and deterministic covariance

$$K(\mathbf{x}, \mathbf{x}') = \mathbb{E}[\mathcal{N}(\mathbf{x})\mathcal{N}(\mathbf{x}')^\top]$$

... NNGP kernel

The Limit NTK

The Limit NTK

Limit NTK

(Jacot, Gabriel, and Hongler 2018)

The empirical NTK converges to a deterministic kernel $\Theta_0(x, x') \xrightarrow{\text{prob}} \Theta(x, x')$

The Limit NTK

Limit NTK

(Jacot, Gabriel, and Hongler 2018)

The empirical NTK converges to a deterministic kernel $\Theta_0(x, x') \xrightarrow{\text{prob}} \Theta(x, x')$

- Only dependent on distribution of parameters θ_0

The Limit NTK

Limit NTK

(Jacot, Gabriel, and Hongler 2018)

The empirical NTK converges to a deterministic kernel $\Theta_0(\mathbf{x}, \mathbf{x}') \xrightarrow{\text{prob}} \Theta(\mathbf{x}, \mathbf{x}')$

- Only dependent on distribution of parameters θ_0
- Can be computed recursively

$$\begin{aligned}K^{(\ell+1)} &= f(K^{(\ell)}) \\ \Theta^{(\ell+1)} &= g(\Theta^{(\ell)}, K^{(\ell+1)}, \dot{K}^{(\ell+1)})\end{aligned}$$

where

$$\dot{K}^{(\ell+1)}(\mathbf{x}, \mathbf{x}') = \mathbb{E}[\sigma'(\mathcal{N}^{(\ell)}(\mathbf{x}))\sigma'(\mathcal{N}^{(\ell)}(\mathbf{x}'))],$$

$\sigma \dots$ non-linearity

The limit NTK is frozen

= time-independent

The limit NTK is frozen

= time-independent

if non-linearity σ is Lipschitz, \mathcal{C}^2 and has bounded second derivative

Frozen NTK

(Jacot, Gabriel, and Hongler 2018)

$$\Theta_t(\mathbf{x}, \mathbf{x}') \rightarrow \Theta(\mathbf{x}, \mathbf{x}')$$

uniformly in t when taking the limit layer by layer

The limit NTK is frozen

= time-independent

if non-linearity σ is Lipschitz, \mathcal{C}^2 and has bounded second derivative

Frozen NTK

(Jacot, Gabriel, and Hongler 2018)

$$\Theta_t(\mathbf{x}, \mathbf{x}') \rightarrow \Theta(\mathbf{x}, \mathbf{x}')$$

uniformly in t when taking the limit layer by layer

- Greatly simplifies ODE!

The limit NTK is frozen

= time-independent

if non-linearity σ is Lipschitz, \mathcal{C}^2 and has bounded second derivative

Frozen NTK

(Jacot, Gabriel, and Hongler 2018)

$$\Theta_t(\mathbf{x}, \mathbf{x}') \rightarrow \Theta(\mathbf{x}, \mathbf{x}')$$

uniformly in t when taking the limit layer by layer

- Greatly simplifies ODE!
- MSE loss \rightarrow linear ODE with analytical solution

NTK as a kernel method

NTK as a kernel method

Kernel $\Phi(\mathbf{x}, \mathbf{x}') : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ similarity measure

Kernel regression:

$$f(\mathbf{x}) = \alpha(\mathcal{Y})\Phi(\mathbf{x}, \mathcal{X})$$

with $(\mathcal{X}, \mathcal{Y}) \dots$ whole training set, $\alpha(\mathcal{X}, \mathcal{Y}) \dots$ coefficients of the regression

NTK as a kernel method

Kernel $\Phi(\mathbf{x}, \mathbf{x}') : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ similarity measure

Kernel regression:

$$f(\mathbf{x}) = \alpha(\mathcal{Y})\Phi(\mathbf{x}, \mathcal{X})$$

with $(\mathcal{X}, \mathcal{Y}) \dots$ whole training set, $\alpha(\mathcal{X}, \mathcal{Y}) \dots$ coefficients of the regression

Take **infinite width infinite time limit** and **MSE** loss \rightarrow Mean of Gaussian process is

$$\mu(\mathbf{x}) = \Theta(\mathbf{x}, \mathcal{X})\Theta(\mathcal{X}, \mathcal{X})^{-1}\mathcal{Y}$$

\rightarrow NTK yields kernel method

- 1 Motivation
- 2 The Neural Tangent Kernel
- 3 Extension to Equivariant Neural Networks**
- 4 Example: $C_4 \ltimes \mathbb{R}^2$
- 5 Experiments
- 6 Summary

Ingredients for Equivariant Neural Networks

Equivariance under group G

$$\mathcal{N}(\rho_{\text{reg}}(\mathbf{g})f) = \rho_{\text{reg}}(\mathbf{g})\mathcal{N}(f), \quad \mathbf{g} \in G$$

→ **group convolutional neural networks (GCNN)**

Ingredients for Equivariant Neural Networks

Equivariance under group G

$$\mathcal{N}(\rho_{\text{reg}}(g)f) = \rho_{\text{reg}}(g)\mathcal{N}(f), \quad g \in G$$

→ **group convolutional neural networks (GCNN)**

- Group convolutional layers
- Lifting layer
- elementwise non-linearity σ
- group pooling layer (to achieve **invariance**)

→ Need recursive relations for the NTK for all of them

Group Convolution (GCNN)

- Data given in form of feature maps $f : X \rightarrow \mathbb{R}^{n_{\text{in}}}$
- For grey-scale images: $f : \mathbb{Z}^2 \rightarrow \mathbb{R}$

Group Convolution (GCNN)

- Data given in form of feature maps $f : X \rightarrow \mathbb{R}^{n_{\text{in}}}$
- For grey-scale images: $f : \mathbb{Z}^2 \rightarrow \mathbb{R}$
- group convolutional **lifting layer**

$$[\mathcal{N}^{(1)}(f)](g) = [\kappa * f](g) = \frac{1}{\sqrt{n_{\text{in}} \text{vol}(\mathcal{S}_{\kappa})}} \int_X dx \kappa(\rho(g^{-1})x) f(x)$$

$\kappa : X \rightarrow \mathbb{R}^{n_{\text{in}} \times n_1} \dots$ filter with support \mathcal{S}_{κ}

Group Convolution (GCNN)

- Data given in form of feature maps $f : X \rightarrow \mathbb{R}^{n_{\text{in}}}$
- For grey-scale images: $f : \mathbb{Z}^2 \rightarrow \mathbb{R}$
- group convolutional **lifting layer**

$$[\mathcal{N}^{(1)}(f)](g) = [\kappa * f](g) = \frac{1}{\sqrt{n_{\text{in}} \text{vol}(\mathcal{S}_{\kappa})}} \int_X dx \kappa(\rho(g^{-1})x) f(x)$$

$\kappa : X \rightarrow \mathbb{R}^{n_{\text{in}} \times n_1} \dots$ filter with support \mathcal{S}_{κ}

- consecutive **group convolutional layer**

$$[\mathcal{N}^{(\ell+1)}(f)](g) = [\kappa * \mathcal{N}^{(\ell)}(f)](g) = \frac{1}{\sqrt{n_{\ell} \text{vol}(\mathcal{S}_{\kappa})}} \int_G dh \kappa(g^{-1}h) [\mathcal{N}^{(\ell)}(f)](h)$$

$\kappa : G \rightarrow \mathbb{R}^{n_{\ell} \times n_{\ell+1}} \dots$ filter

Group Convolution (GCNN)

- Data given in form of feature maps $f : X \rightarrow \mathbb{R}^{n_{\text{in}}}$
- For grey-scale images: $f : \mathbb{Z}^2 \rightarrow \mathbb{R}$
- group convolutional **lifting layer**

$$[\mathcal{N}^{(1)}(f)](g) = [\kappa * f](g) = \frac{1}{\sqrt{n_{\text{in}} \text{vol}(\mathcal{S}_{\kappa})}} \int_X dx \kappa(\rho(g^{-1})x) f(x)$$

$\kappa : X \rightarrow \mathbb{R}^{n_{\text{in}} \times n_1}$... filter with support \mathcal{S}_{κ}

- consecutive **group convolutional layer**

$$[\mathcal{N}^{(\ell+1)}(f)](g) = [\kappa * \mathcal{N}^{(\ell)}(f)](g) = \frac{1}{\sqrt{n_{\ell} \text{vol}(\mathcal{S}_{\kappa})}} \int_G dh \kappa(g^{-1}h) [\mathcal{N}^{(\ell)}(f)](h)$$

$\kappa : G \rightarrow \mathbb{R}^{n_{\ell} \times n_{\ell+1}}$... filter

- **infinite width limit:** channels $n_1, \dots, n_{L-1} \rightarrow \infty$

Equivariant NTK

GCNN layers

At each layer ℓ

$$K_{g,g'}^{(\ell)}(f,f') = \mathbb{E} \left[[\mathcal{N}^{(\ell)}(f)](g) [\mathcal{N}^{(\ell)}(f')](g')^\top \right]$$

Equivariant NTK

GCNN layers

At each layer ℓ

$$K_{g,g'}^{(\ell)}(f,f') = \mathbb{E} \left[[\mathcal{N}^{(\ell)}(f)](g) [\mathcal{N}^{(\ell)}(f')](g')^\top \right]$$

Recursion for GCNN layer

$$K_{g,g'}^{(\ell+1)}(f,f') = \frac{1}{\text{vol}(\mathcal{S}_\kappa)} \int_{\mathcal{S}_\kappa} dh K_{gh,g'h}^{(\ell)}(f,f')$$
$$\Theta_{g,g'}^{(\ell+1)}(f,f') = K_{g,g'}^{(\ell+1)}(f,f') + \frac{1}{\text{vol}(\mathcal{S}_\kappa)} \int_{\mathcal{S}_\kappa} dh \Theta_{gh,g'h}^{(\ell)}(f,f')$$

Equivariant NTK

GCNN layers

Equivariant NTK

GCNN layers

Recursion for GCNN lifting layer

$$K_{g,g'}^{(\ell+1)}(f,f') = \frac{1}{\text{vol}(\mathcal{S}_\kappa)} \int_{\mathcal{S}_\kappa} dx K_{\rho(g)x, \rho(g')x'}^{(\ell)}(f,f')$$

$$\Theta_{g,g'}^{(\ell+1)}(f,f') = K_{g,g'}^{(\ell+1)}(f,f') + \frac{1}{\text{vol}(\mathcal{S}_\kappa)} \int_{\mathcal{S}_\kappa} dx \Theta_{\rho(g)x, \rho(g')x'}^{(\ell)}(f,f')$$

Equivariant NTK

GCNN layers

Recursion for GCNN lifting layer

$$K_{g,g'}^{(\ell+1)}(f,f') = \frac{1}{\text{vol}(\mathcal{S}_\kappa)} \int_{\mathcal{S}_\kappa} dx K_{\rho(g)x, \rho(g')x'}^{(\ell)}(f,f')$$

$$\Theta_{g,g'}^{(\ell+1)}(f,f') = K_{g,g'}^{(\ell+1)}(f,f') + \frac{1}{\text{vol}(\mathcal{S}_\kappa)} \int_{\mathcal{S}_\kappa} dx \Theta_{\rho(g)x, \rho(g')x'}^{(\ell)}(f,f')$$

If first layer \rightarrow

$$K_{x,x'}^{(0)}(f,f') = f(x)f'(x'), \quad \Theta_{x,x'}^{(0)}(f,f') = 0$$

Equivariant NTK

Non-linearity

Elementwise non-linearity σ : $\mathcal{N}^{(\ell+1)}(\mathbf{g}) = \sigma(\mathcal{N}^{(\ell)})(\mathbf{g})$

Equivariant NTK

Non-linearity

Elementwise non-linearity $\sigma: \mathcal{N}^{(\ell+1)}(\mathbf{g}) = \sigma(\mathcal{N}^{(\ell)})(\mathbf{g})$

Recursion for non-linearity

$$\Lambda^{(\ell)}(f, f') = \begin{pmatrix} K^{(\ell)}(f, f) & K^{(\ell)}(f, f') \\ K^{(\ell)}(f', f) & K^{(\ell)}(f', f') \end{pmatrix}$$

$$K^{(\ell+1)}(f, f') = \mathbb{E}_{(u, v) \sim \mathcal{N}(0, \Lambda^{(\ell)}(f, f'))} [\sigma(u)\sigma(v)]$$

$$\dot{K}^{(\ell+1)}(f, f') = \mathbb{E}_{(u, v) \sim \mathcal{N}(0, \Lambda^{(\ell)}(f, f'))} [\sigma'(u)\sigma'(v)]$$

$$\Theta^{(\ell+1)}(f, f') = \dot{K}^{(\ell+1)}(f, f')\Theta^{(\ell)}(f, f')$$

Equivariant NTK

Non-linearity

Elementwise non-linearity $\sigma: \mathcal{N}^{(\ell+1)}(\mathbf{g}) = \sigma(\mathcal{N}^{(\ell)})(\mathbf{g})$

Recursion for non-linearity

$$\Lambda_{g,g'}^{(\ell)}(f,f') = \begin{pmatrix} K_{g,g}^{(\ell)}(f,f) & K_{g,g'}^{(\ell)}(f,f') \\ K_{g',g}^{(\ell)}(f',f) & K_{g',g'}^{(\ell)}(f',f') \end{pmatrix}$$

$$K_{g,g'}^{(\ell+1)}(f,f') = \mathbb{E}_{(u,v) \sim \mathcal{N}(0, \Lambda_{g,g'}^{(\ell)}(f,f'))} [\sigma(u)\sigma(v)]$$

$$\dot{K}_{g,g'}^{(\ell+1)}(f,f') = \mathbb{E}_{(u,v) \sim \mathcal{N}(0, \Lambda_{g,g'}^{(\ell)}(f,f'))} [\sigma'(u)\sigma'(v)]$$

$$\Theta_{g,g'}^{(\ell+1)}(f,f') = \dot{K}_{g,g'}^{(\ell+1)}(f,f') \Theta_{g,g'}^{(\ell)}(f,f')$$

Equivariant NTK

Group pooling layer

Making equivariant network **invariant** \rightarrow **group pooling**

$$\mathcal{N}^{(\ell+1)}(\mathbf{f}) = \frac{1}{\sqrt{\text{vol}(G)}} \int_G dg [\mathcal{N}^{(\ell)}(\mathbf{f})](g)$$

Equivariant NTK

Group pooling layer

Making equivariant network **invariant** \rightarrow **group pooling**

$$\mathcal{N}^{(\ell+1)}(\mathbf{f}) = \frac{1}{\sqrt{\text{vol}(G)}} \int_G dg [\mathcal{N}^{(\ell)}(\mathbf{f})](g)$$

Recursion for group pooling layer

$$K^{(\ell+1)}(\mathbf{f}, \mathbf{f}') = \frac{1}{\text{vol}(G)} \int_G dg dg' K_{g,g'}^{(\ell)}(\mathbf{f}, \mathbf{f}')$$
$$\Theta^{(\ell+1)}(\mathbf{f}, \mathbf{f}') = \frac{1}{\text{vol}(G)} \int_G dg dg' \Theta_{g,g'}^{(\ell)}(\mathbf{f}, \mathbf{f}')$$

Symmetries of the NTK

$$K_{g,g'}^{(\ell)}(\rho_{\text{reg}}(h)f, \rho_{\text{reg}}(h')f') = K_{h^{-1}g, h'^{-1}g'}^{(\ell)}(f, f')$$
$$\Theta_{g,g'}^{(\ell)}(\rho_{\text{reg}}(h)f, \rho_{\text{reg}}(h')f') = \Theta_{h^{-1}g, h'^{-1}g'}^{(\ell)}(f, f')$$

- 1 Motivation
- 2 The Neural Tangent Kernel
- 3 Extension to Equivariant Neural Networks
- 4 Example: $C_4 \times \mathbb{R}^2$**
- 5 Experiments
- 6 Summary

$$C_n \times \mathbb{R}^2$$

Discrete roto-translations in the plane

- $g = tr$, where $t \in \mathbb{R}^2$ and $r \in C_n$

$C_n \times \mathbb{R}^2$

Discrete roto-translations in the plane

- $g = tr$, where $t \in \mathbb{R}^2$ and $r \in C_n$
- Keep t as spatial indices, for r extend to nn_ℓ channel indices in layer $\ell \rightarrow$

$$K_{g=tr, g'=t'r'}(f, f') = [K_{rr}(f, f')](t, t'), \quad \Theta_{g=tr, g'=t'r'}(f, f') = [\Theta_{rr'}(f, f')](t, t')$$

$C_n \times \mathbb{R}^2$

Discrete roto-translations in the plane

- $g = tr$, where $t \in \mathbb{R}^2$ and $r \in C_n$
- Keep t as spatial indices, for r extend to nn_ℓ channel indices in layer $\ell \rightarrow$

$$K_{g=tr, g'=t'r'}(f, f') = [K_{rr}(f, f')](t, t'), \quad \Theta_{g=tr, g'=t'r'}(f, f') = [\Theta_{rr'}(f, f')](t, t')$$

- For CNN, neural-tangents library provides the following highly optimized operator

$$[\mathcal{A}_{S_\kappa}(K)](t, t') = \frac{1}{S_\kappa} \int_{S_\kappa} d\tilde{t} K(t + \tilde{t}, t' + \tilde{t})$$

$C_n \times \mathbb{R}^2$

Discrete roto-translations in the plane

- $g = tr$, where $t \in \mathbb{R}^2$ and $r \in C_n$
- Keep t as spatial indices, for r extend to nn_ℓ channel indices in layer $\ell \rightarrow$

$$K_{g=tr, g'=t'r'}(f, f') = [K_{rr}(f, f')](t, t'), \quad \Theta_{g=tr, g'=t'r'}(f, f') = [\Theta_{rr'}(f, f')](t, t')$$

- For CNN, neural-tangents library provides the following highly optimized operator

$$[\mathcal{A}_{S_\kappa}(K)](t, t') = \frac{1}{S_\kappa} \int_{S_\kappa} d\tilde{t} K(t + \tilde{t}, t' + \tilde{t})$$

- Can this be exploited for $C_n \times \mathbb{R}^2$ as well?

$$C_n \times \mathbb{R}^2$$

Discrete roto-translations in the plane

→ **Yes!**

$C_n \times \mathbb{R}^2$ Discrete roto-translations in the plane

→ **Yes!**

$$[K_{rr'}^{(\ell+1)}(f, f')](t, t') = \sum_{\tilde{r} \in C_n} [\mathcal{A}_{\rho(r)\delta_\kappa}(\tilde{K}_{r\tilde{r}, r'\tilde{r}}^{(\ell)}(f, f'))](t, \rho(rr'^{-1})t')$$

$$[\Theta_{rr'}^{(\ell+1)}(f, f')](t, t') = [K_{rr'}^{(\ell+1)}(f, f')](t, t') + \sum_{\tilde{r} \in C_n} [\mathcal{A}_{\rho(r)\delta_\kappa}(\tilde{\Theta}_{r\tilde{r}, r'\tilde{r}}^{(\ell)}(f, f'))](t, \rho(rr'^{-1})t'),$$

$C_n \times \mathbb{R}^2$

Discrete roto-translations in the plane

→ **Yes!**

$$[K_{rr'}^{(\ell+1)}(f, f')](t, t') = \sum_{\tilde{r} \in C_n} [\mathcal{A}_{\rho(r)\mathcal{S}_\kappa}(\tilde{K}_{r\tilde{r}, r'\tilde{r}}^{(\ell)}(f, f'))](t, \rho(rr'^{-1})t')$$

$$[\Theta_{rr'}^{(\ell+1)}(f, f')](t, t') = [K_{rr'}^{(\ell+1)}(f, f')](t, t') + \sum_{\tilde{r} \in C_n} [\mathcal{A}_{\rho(r)\mathcal{S}_\kappa}(\tilde{\Theta}_{r\tilde{r}, r'\tilde{r}}^{(\ell)}(f, f'))](t, \rho(rr'^{-1})t'),$$

where

$$[\tilde{K}_{r r'}^{(\ell)}(f, f')](t, t') = [K_{r r'}^{(\ell)}(f, f')](t, \rho(r' r^{-1})t')$$

$$\tilde{\Theta}_{r r'}^{(\ell)}(f, f') = [\Theta_{r r'}^{(\ell)}(f, f')](t, \rho(r' r^{-1})t')$$



- 1 Motivation
- 2 The Neural Tangent Kernel
- 3 Extension to Equivariant Neural Networks
- 4 Example: $C_4 \times \mathbb{R}^2$
- 5 Experiments**
- 6 Summary

Kernel convergence

- Comparing limit kernels with mean of finite-width kernels
- $G = C_4 \times \mathbb{R}^2$ and random input data
- Relative error averaged over all components of the Gram matrix
- Implemented in the JAX-based `neural-tangents` library¹

Kernel convergence

- Comparing limit kernels with mean of finite-width kernels
- $G = C_4 \times \mathbb{R}^2$ and random input data
- Relative error averaged over all components of the Gram matrix
- Implemented in the JAX-based `neural-tangents` library¹



Histological image classification

The data set

NCT-CRC-HE-100K (*Kather, Halama, and Marx 2018*)

- 100 000 histological images of human colorectal cancer (CRC) and normal tissue
- 9 classes, 2 are cancerous
- 224×224 pixels in color

Histological image classification


The data set

NCT-CRC-HE-100K (*Kather, Halama, and Marx 2018*)

- 100 000 histological images of human colorectal cancer (CRC) and normal tissue
- 9 classes, 2 are cancerous
- 224×224 pixels in color

Here:

- downsampled to 32×32 pixels
- Only up to 1000 training samples used



Histological image classification

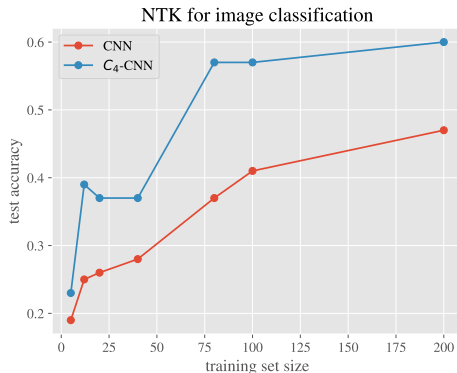
Performance for small training sets

- NTK kernel method
- Corresponds to training with MSE loss
- Comparison
 - CNN
 - $\mathcal{C}_4 \times \mathbb{R}^2$

Histological image classification

Performance for small training sets

- NTK kernel method
- Corresponds to training with MSE loss
- Comparison
 - CNN
 - $\mathcal{C}_4 \times \mathbb{R}^2$



- 1 Motivation
- 2 The Neural Tangent Kernel
- 3 Extension to Equivariant Neural Networks
- 4 Example: $C_4 \times \mathbb{R}^2$
- 5 Experiments
- 6 Summary

Summary

Summary

- Derived **all necessary recursive relations** for the NTK of GCNNs

Summary

- Derived **all necessary recursive relations** for the NTK of GCNNs
- Enabled **analytical study** of training dynamics for equivariant neural networks in the infinite width limit

Summary

- Derived **all necessary recursive relations** for the NTK of GCNNs
- Enabled **analytical study** of training dynamics for equivariant neural networks in the infinite width limit

Summary

- Derived **all necessary recursive relations** for the NTK of GCNNs
- Enabled **analytical study** of training dynamics for equivariant neural networks in the infinite width limit

What's next?

- GCNNs over compact groups utilizing the Fourier domain

Summary

- Derived **all necessary recursive relations** for the NTK of GCNNs
- Enabled **analytical study** of training dynamics for equivariant neural networks in the infinite width limit

What's next?

- GCNNs over compact groups utilizing the Fourier domain
- equivariant graph NNs

Summary

- Derived **all necessary recursive relations** for the NTK of GCNNs
- Enabled **analytical study** of training dynamics for equivariant neural networks in the infinite width limit

What's next?

- GCNNs over compact groups utilizing the Fourier domain
- equivariant graph NNs
- Trainability and generalization regimes of equivariant NNs (*Xiao, Pennington, and Schoenholz 2020*)

Summary

- Derived **all necessary recursive relations** for the NTK of GCNNs
- Enabled **analytical study** of training dynamics for equivariant neural networks in the infinite width limit

What's next?

- GCNNs over compact groups utilizing the Fourier domain
- equivariant graph NNs
- Trainability and generalization regimes of equivariant NNs (*Xiao, Pennington, and Schoenholz 2020*)
- Relations to Quantum Field Theory (*Banta et al. 2023*)

Thanks for your attention!



arXiv:2406.06504