Rotation Invariant Learning on Contours

Plankton Workshop – IMR

Odin Hoff Gardå

PhD Student @ UiB

June 3, 2025

Rotation Invariance Visualized

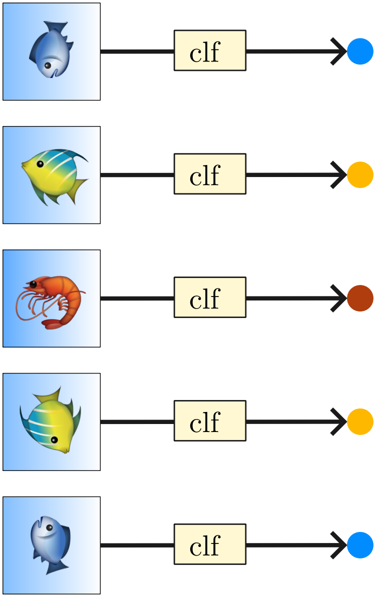

Rotation Equivariance Visualized

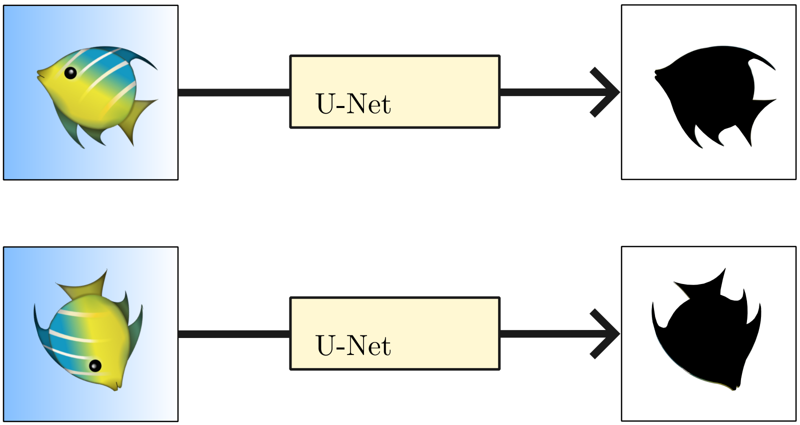

Geometric Priors

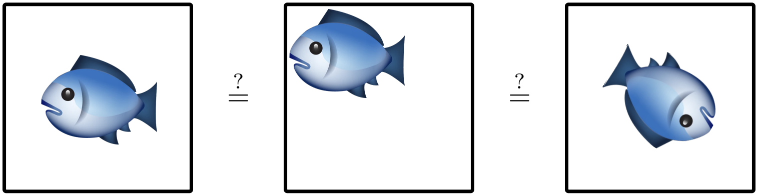

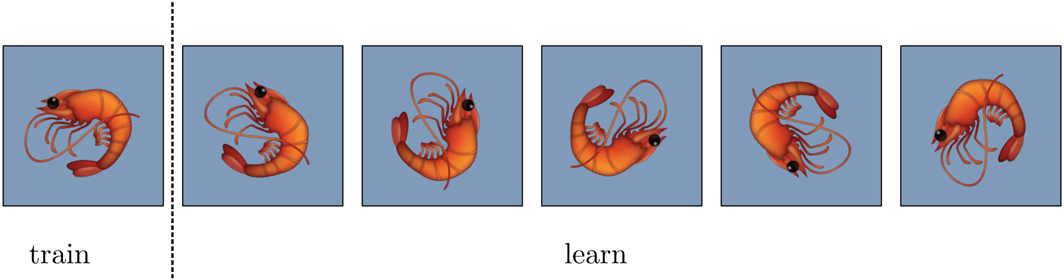

train

learn

Geometric Priors

You already use them!

Example: CNNs (Translation Equivariance)

Let f be a convolutional layer.

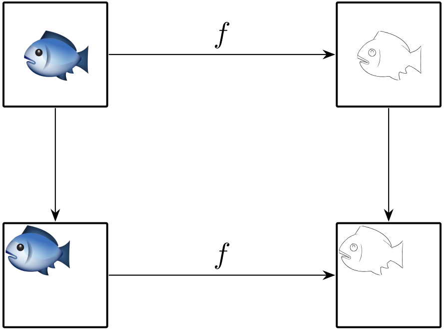

Example: CNNs (Translation Invariance)

Let F be a "fully convolutional" neural network.

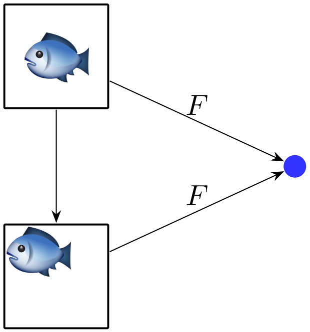

What About Rotations?

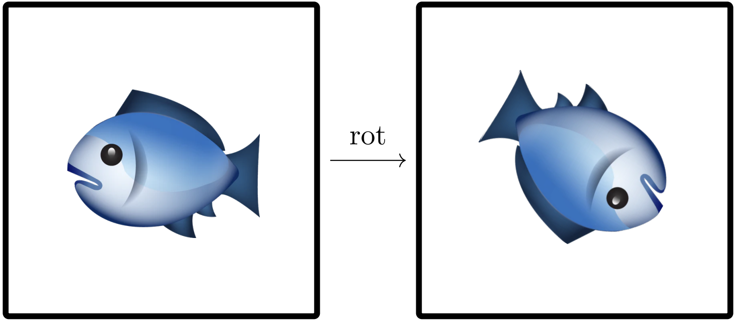

$$\text{conv}(\text{rot}(x)) \neq \text{rot}(\text{conv}(x))$$

Contours as Shape Representations

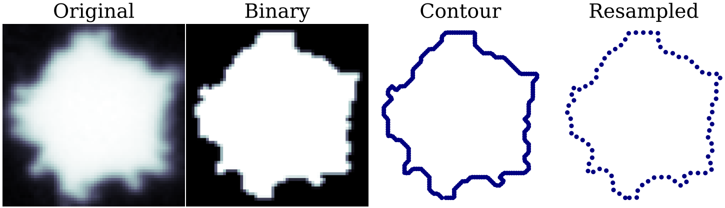

A **contour** is a sequence of points in the plane that represents the boundary of a shape.

The Complex Plane $\mathbb{C}$

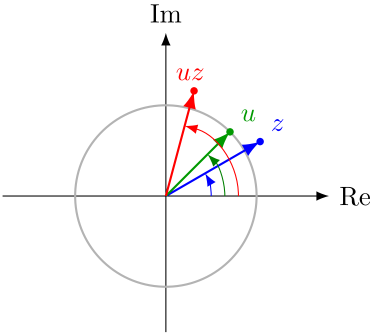

- Complex numbers = points in the plane + multiplication.
- Rotation is multiplication with unit complex numbers.

Contours as Complex Numbers

$$\mathbb{R}^2 \cong \mathbb{C}$$

- A **contour** is a sequence $(z_1, \ldots, z_n)$ of complex numbers.
- Or a function $x\colon [n] \to \mathbb{C}$, where $[n] = \{1, \ldots, n\}$.
- A **stack of** k **contours** is a function $x\colon [n] \to \mathbb{C}^k$.

Write $\mathcal{X}_n^k$ for the collection of all functions $x\colon [n] \to \mathbb{C}^k$.

Complex-valued 1D Convolution

For a filter $\phi \in \mathcal{X}_m^k$ define *circular convolution* as

$$\begin{aligned} \text{conv}_\phi \colon \mathcal{X}_n^k &\to \mathcal{X}_n^1 \\ x &\mapsto \phi \star x \end{aligned}$$

where $\phi \star x$ is 1D convolution of ϕ along x in the complex domain.

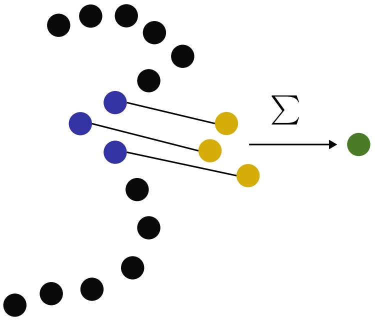

Equivariance: conv_ϕ is equivariant with respect to rotations and choice of starting point.

Activation Functions

A function $f\colon \mathbb{C} \to \mathbb{C}$ is rotation equivariant if and only if $f(z) = g(|z|)z$ for some $g\colon [0, \infty) \to \mathbb{C}$.

Examples:

- Siglog: $(|z| + 1)^{-1}z$.
- Amplitude-phase-type: $\tanh(|z|)z|z|^{-1}$.
- ModReLU: $\mathrm{ReLU}(|z| + b)z|z|^{-1}$ with learnable $b \in \mathbb{R}$.

Local Spatial Pooling (Coarsening)

$$\mathcal{X}_n^k \to \mathcal{X}_m^k \quad (m \leq n)$$

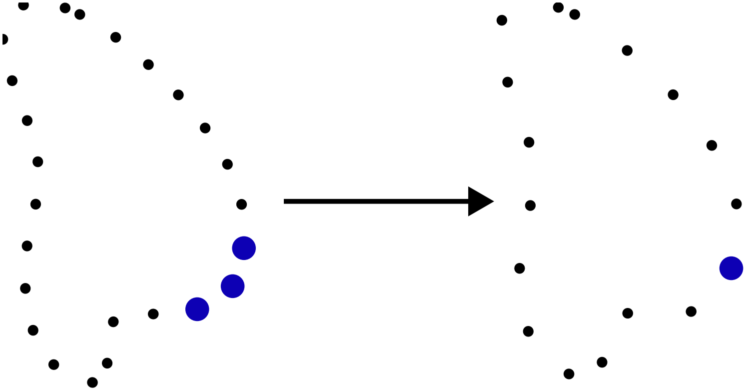

Reduce the number of points.

Global Pooling (Invariant Layer)

Rotation invariant function $\mathcal{X}_n^k \to \mathbb{R}^k$.

For example, any function depending only on radii such as mean or max, or a learnable function.

Plankton Classification Dataset

5 classes from Mesozooplankton→Animalia→Arthropoda
→Crustacea→Copepoda→Calanoida.

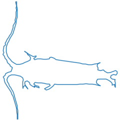

Class	#
Temoridae/Temora/Temora spp	257
Metridinidae/Metridia_late	271
Clausocalanidae/Microcalanus_late	421
Centropagidae/Centropages_late	345
Calanidae/Calanus_late	644
Total	1938

Split 50/50 stratified. Image size 64×64, contour length 128.

Plankton Results

Model	Accuracy (test)	# Params
CNN	0.642 ± 0.022	≈114k
CNN+Aug	0.896 ± 0.004	≈114k
Ours	0.879 ± 0.008	≈100k
Ours+RH	0.905 ± 0.012	≈101k

Table: Classification metrics on dataset. Average over 10 runs.

- **CNN** Standard CNN on gray-scale images. With and without rotation augmentation.
- **Ours** Complex-valued 1d convolutions and global pooling.
- **Ours+RH** With a simple rotation-invariant texture descriptor (radial histogram).

Fashion MNIST Contours

Contours based on the Fashion MNIST dataset.

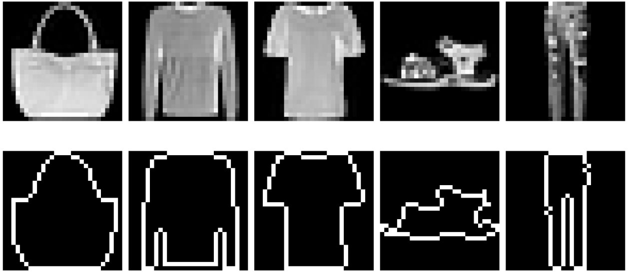

Model	Accuracy	# Params
CNN+Aug	0.860 ± 0.001	≈294k
Ours	0.878 ± 0.001	≈77k

Table: Performance on Fashion MNIST. Average over 30 runs.

The End

Questions?