Learning Steerable Function for Efficient Image Resampling

Jiacheng Li*   Chang Chen*   Wei Huang   Zhiqiang Lang   Fenglong Song   Youliang Yan   Zhiwei Xiong✉

*Equal contribution   ✉Correspondence author

https://lerf.pages.dev

JUNE 18-22, 2023
CVPR
VANCOUVER, CANADA

Background: Interpolation

Typically, resampling through interpolation can be implemented as:
① **Obtain relative offsets** via projecting target coordinates back to the source coordinate space.
② **Predict resampling weights** with fixed resampling function.
③ **Aggregate source pixels** with weighted summation to obtain pixel value at the target coordinate.

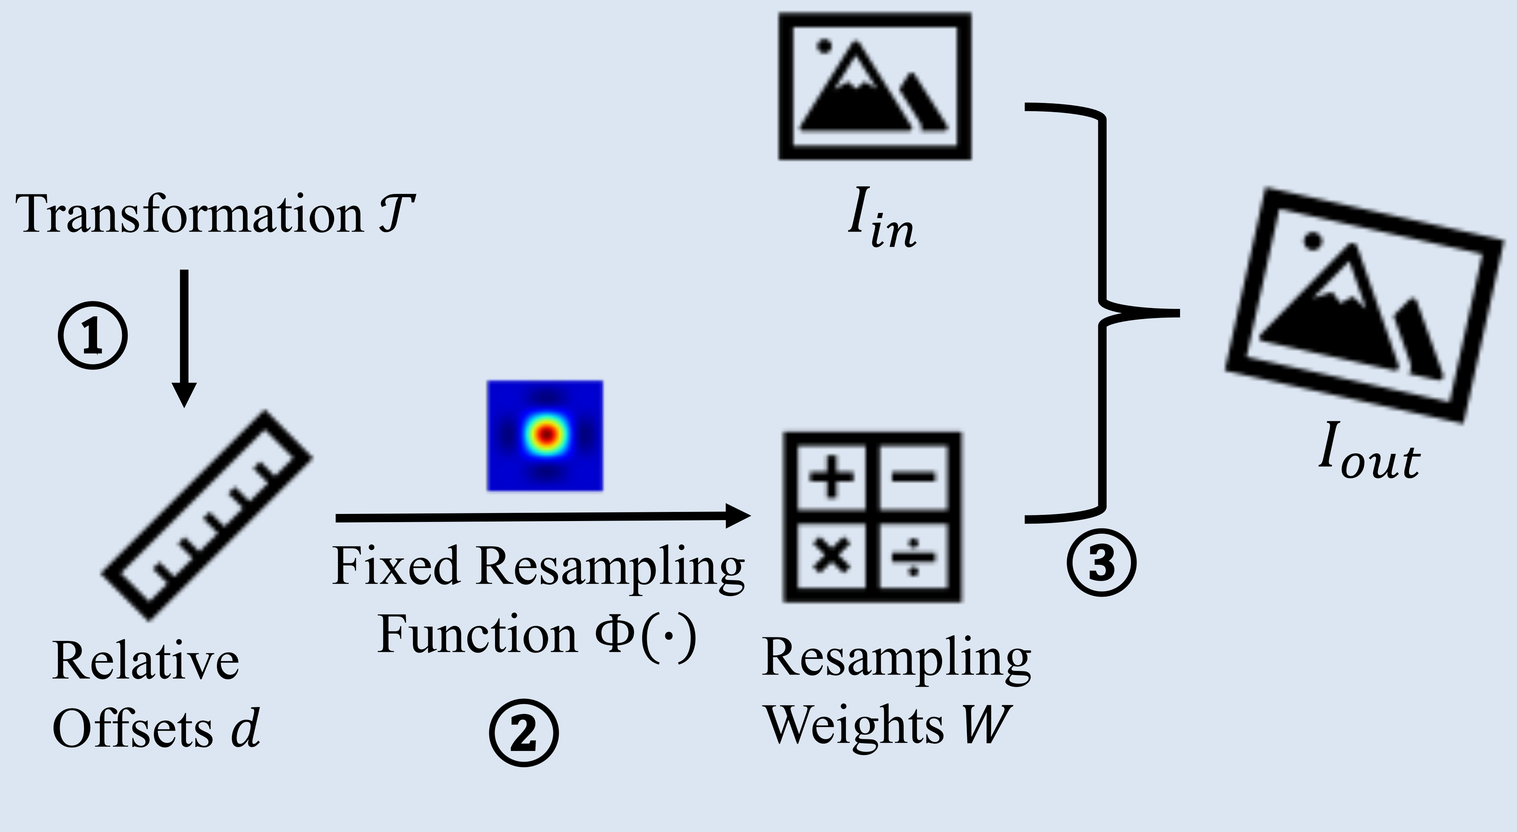

Learning Resampling Function

In LeRF, we assign **spatially-varying** steerable resampling functions to image pixels, whose orientations are parameterized with several hyper-parameters, and we train a neural network to learn these hyper-parameters for each pixel in an end-to-end manner.

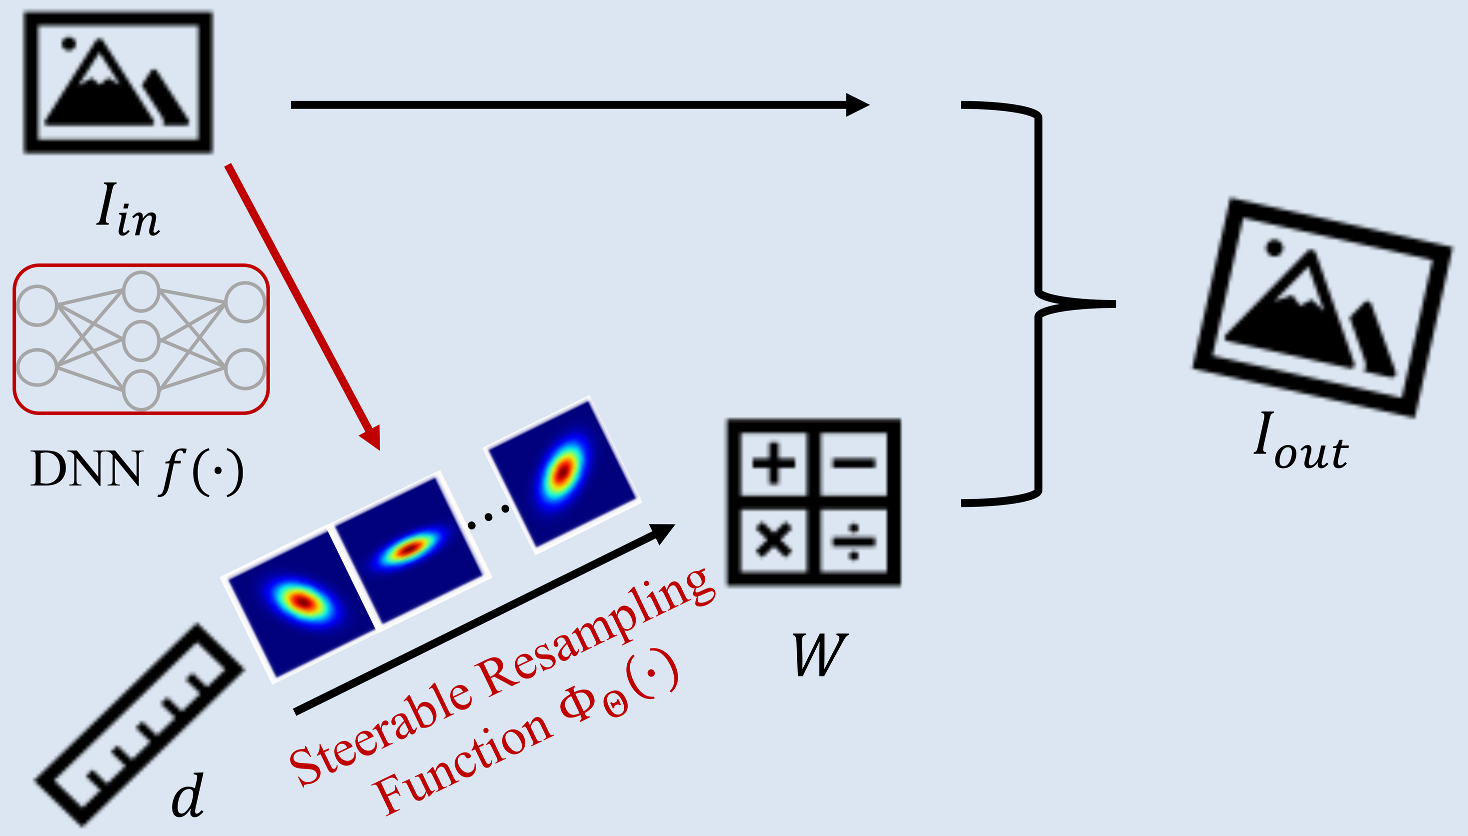

Evaluation

At a similar running time, LeRF outperforms interpolation significantly.

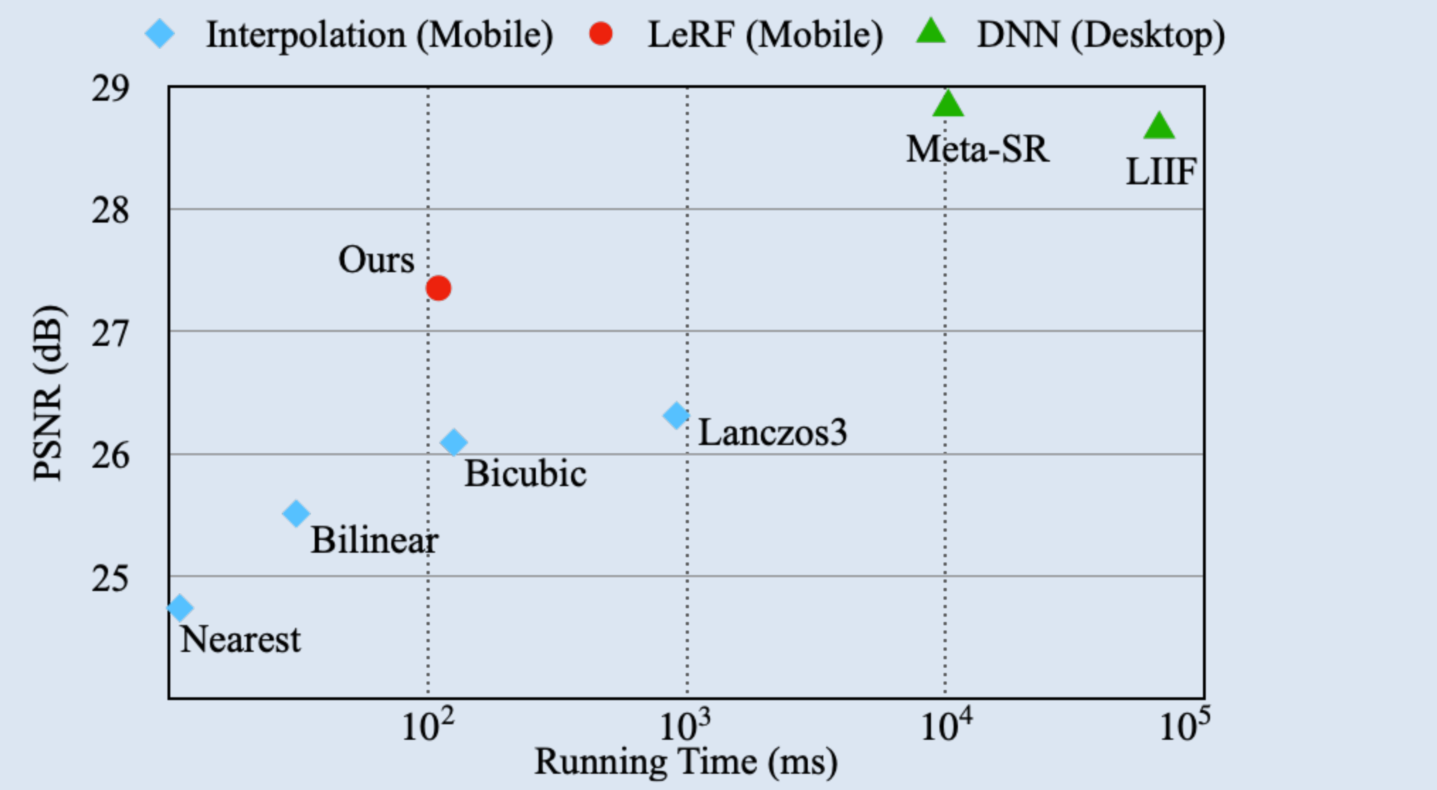

Visualized Learned Resampling Functions

As shown below, the shapes of resampling functions are well adapted to corners, flat surfaces, and edges with various orientations.

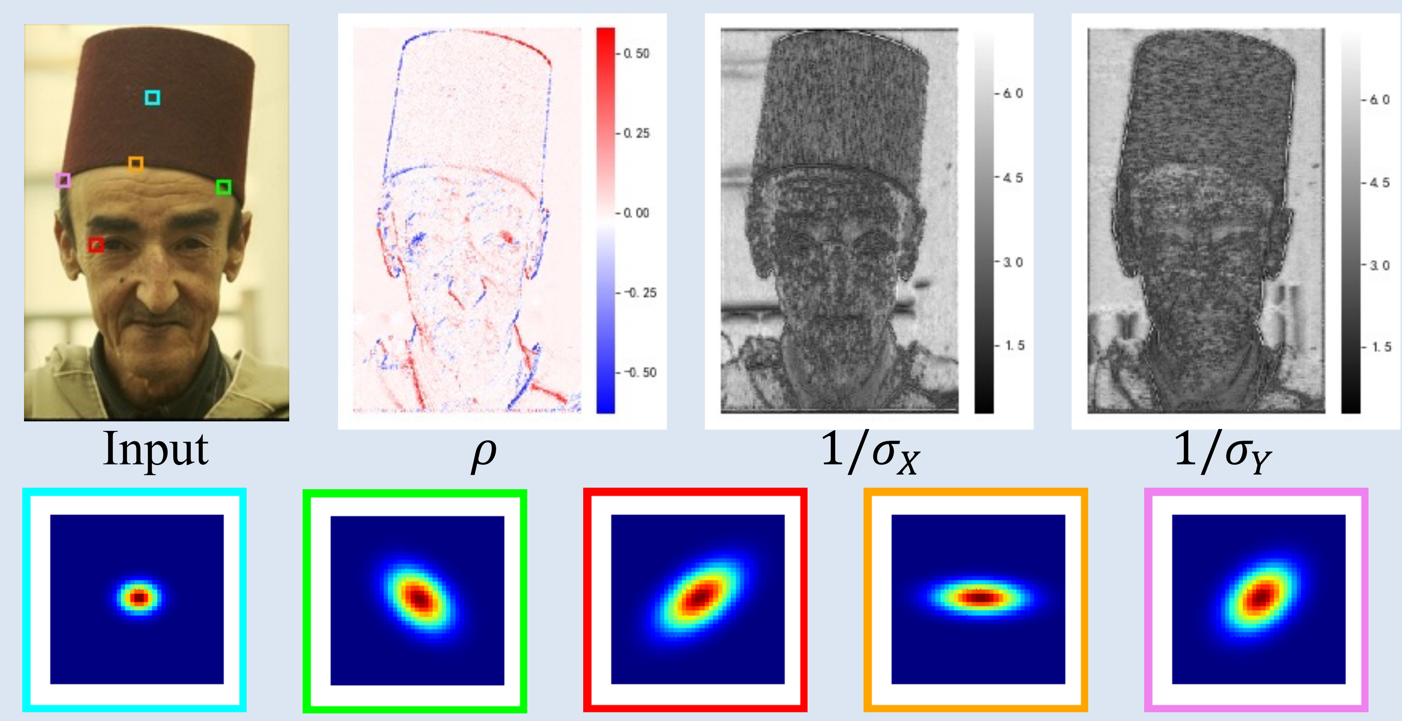

Resampling Function: Fixed to Learned

The assumption on **local continuity** of interpolation allows for continuous resampling under arbitrary transformations, yet leads to blurry results due to ignoring different local structures. In contrast, we introduce steerable resampling functions and learn them in a **data-driven** way.

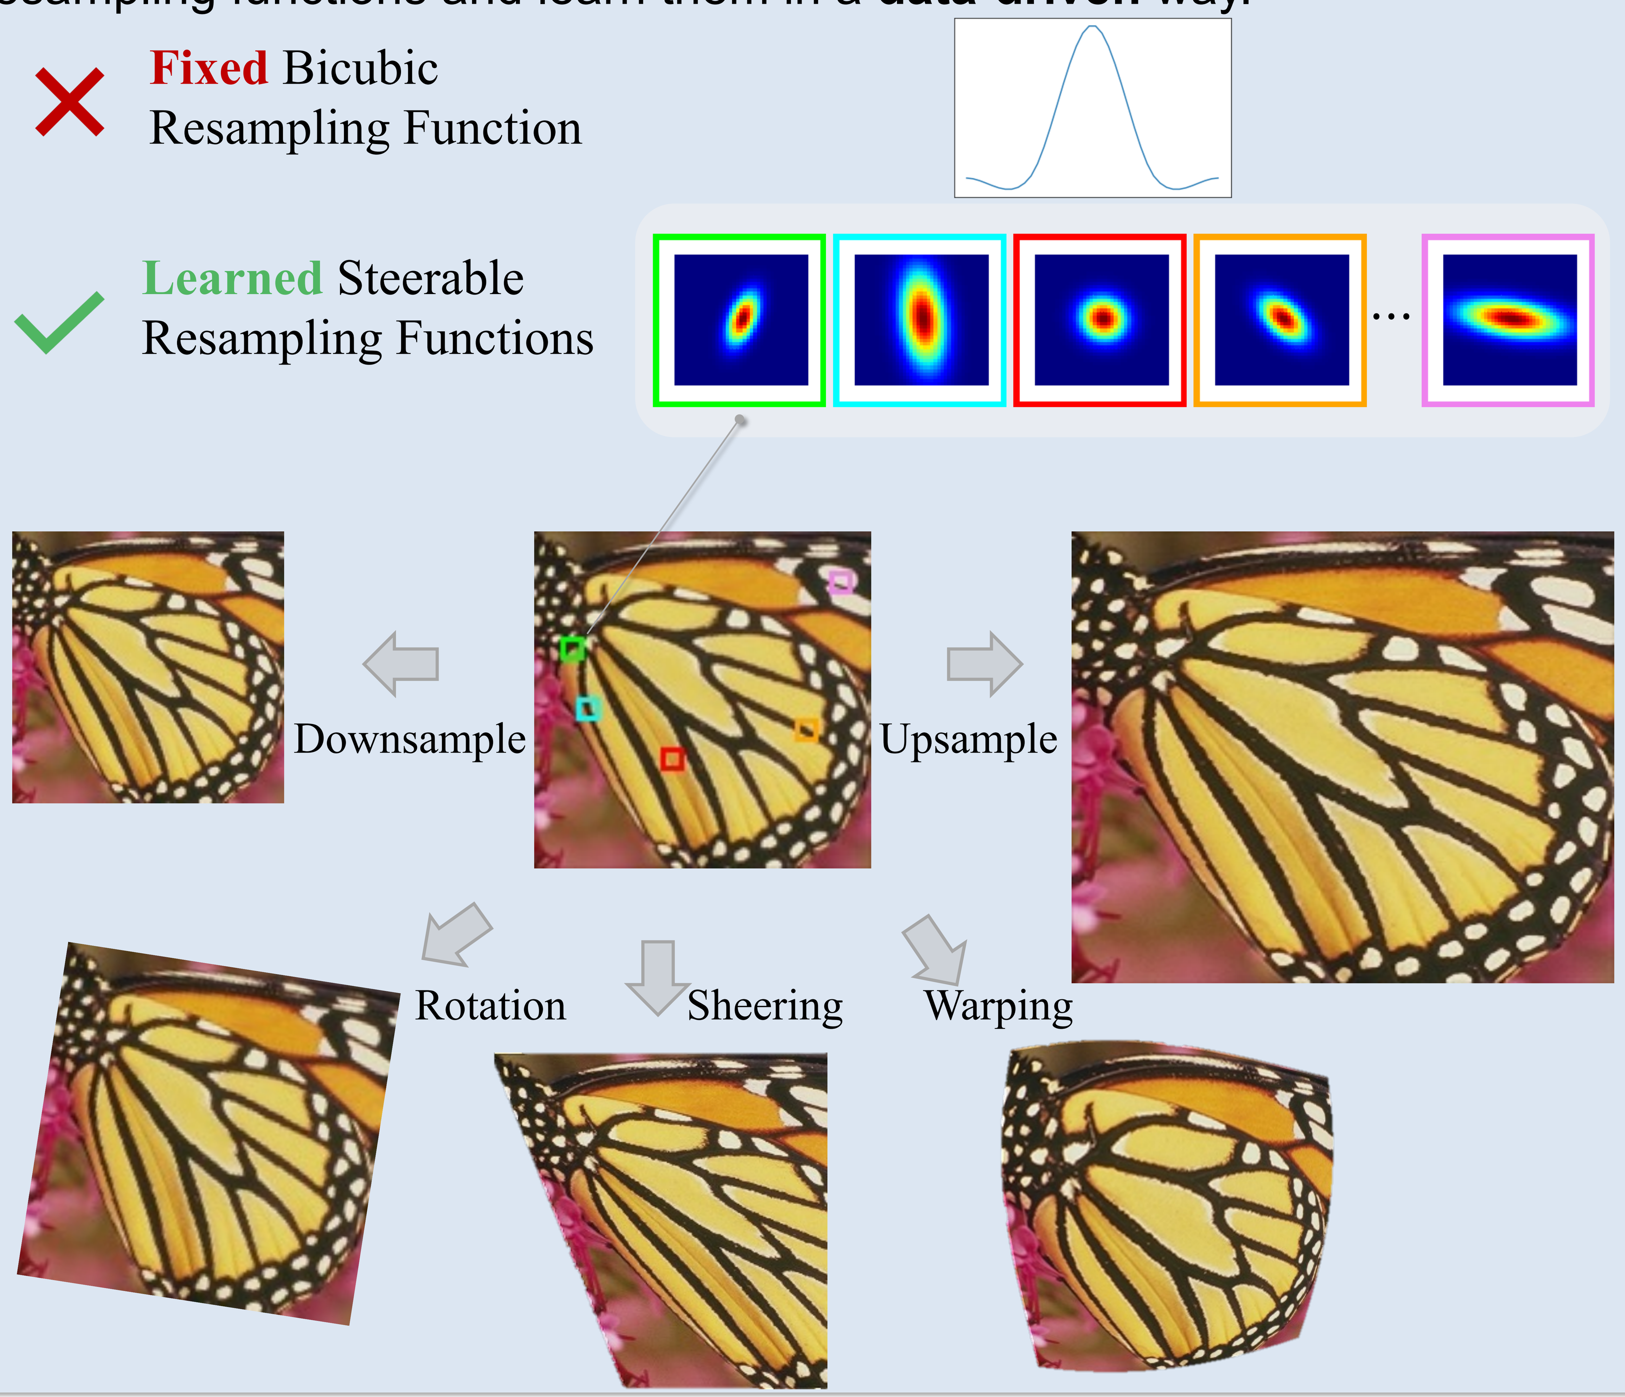

Steerable Function

We utilize the anisotropic gaussian, which can be parameterized with 3 hyper-parameters: ρ, σ_X, σ_Y.

$$\boldsymbol{\Phi}_{\boldsymbol{\Sigma}}(\boldsymbol{x}-\boldsymbol{\mu}) = \frac{1}{2\pi|\boldsymbol{\Sigma}|^{\frac{1}{2}}} exp\{(\boldsymbol{x}-\boldsymbol{\mu})^T\boldsymbol{\Sigma}^{-1}(\boldsymbol{x}-\boldsymbol{\mu})\}, \boldsymbol{\Sigma} = \begin{pmatrix} \sigma_X^2 & \rho\sigma_X\sigma_Y \\ \rho\sigma_X\sigma_Y & \sigma_Y^2 \end{pmatrix}$$

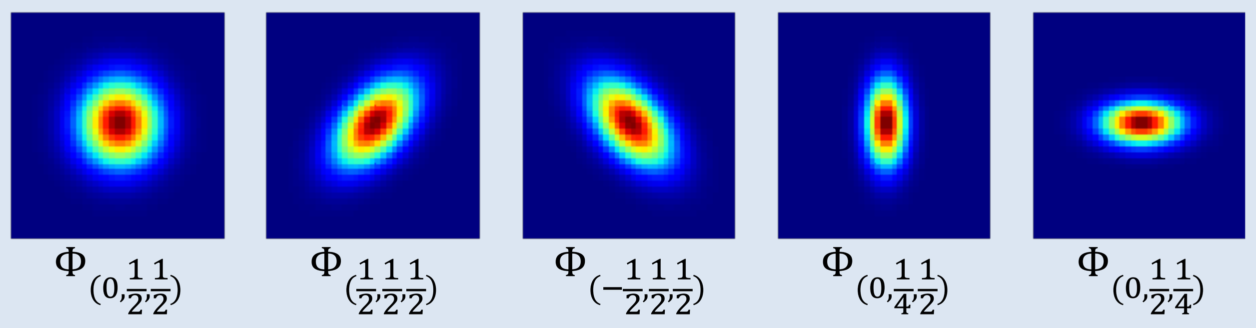

LUT Acceleration

Furthermore, we present an efficient implementation, where the inference of the learned neural network is accelerated with **look-up tables** (LUTs). We further design a directional ensemble strategy and edge-sensitive indexing patterns to better capture local structures in images.

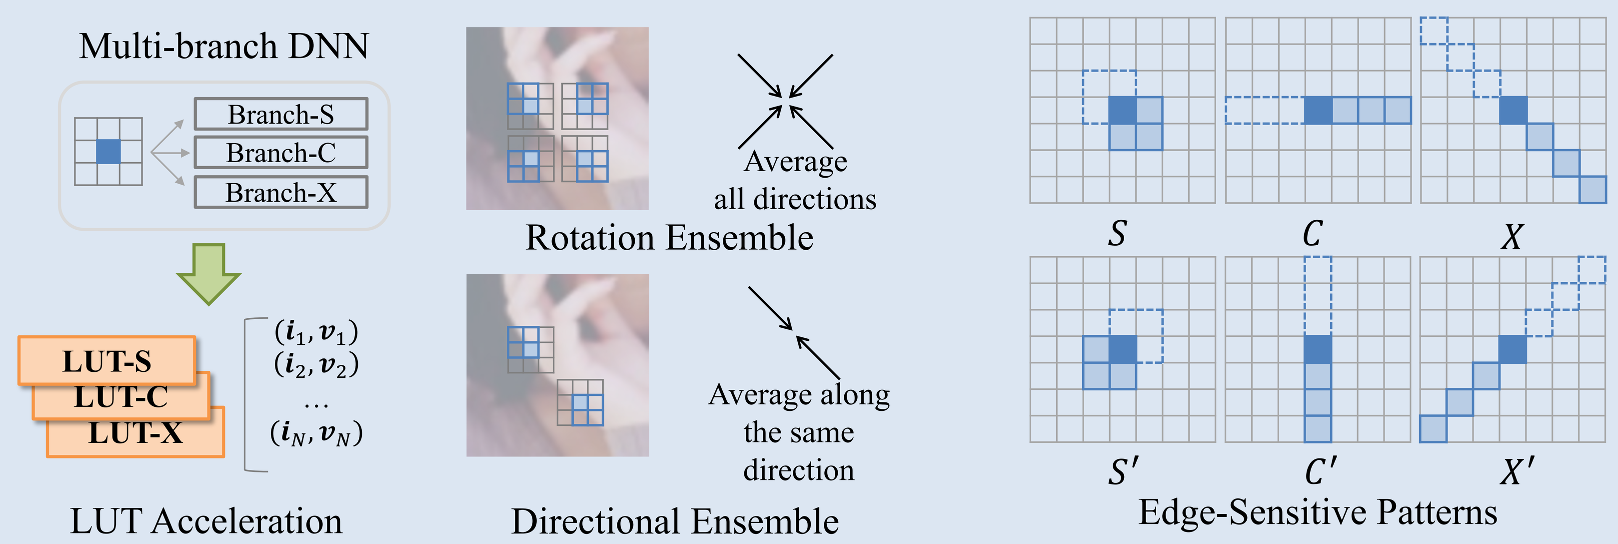

Learn more about LUT acceleration at https://mulut.pages.dev.

Generalization

The learned resampling functions generalize well to unseen deformations, e.g., downsampling, rotation, sheering, and even arbitrary warping (according to a barrel-shaped distortion or optical flow).