TECHNION
Israel Institute
of Technology

Code is available

JUNE 18-22, 2023
CVPR
VANCOUVER, CANADA

Alias-Free Convnets:
Fractional Shift Invariance via Polynomial Activations



Hagay Michaeli Tomer Michaeli Daniel Soudry

The research of DS was funded by the European Union
(ERC, A-B-C-Deep, 101039436)

WED-PM-379

Abstract

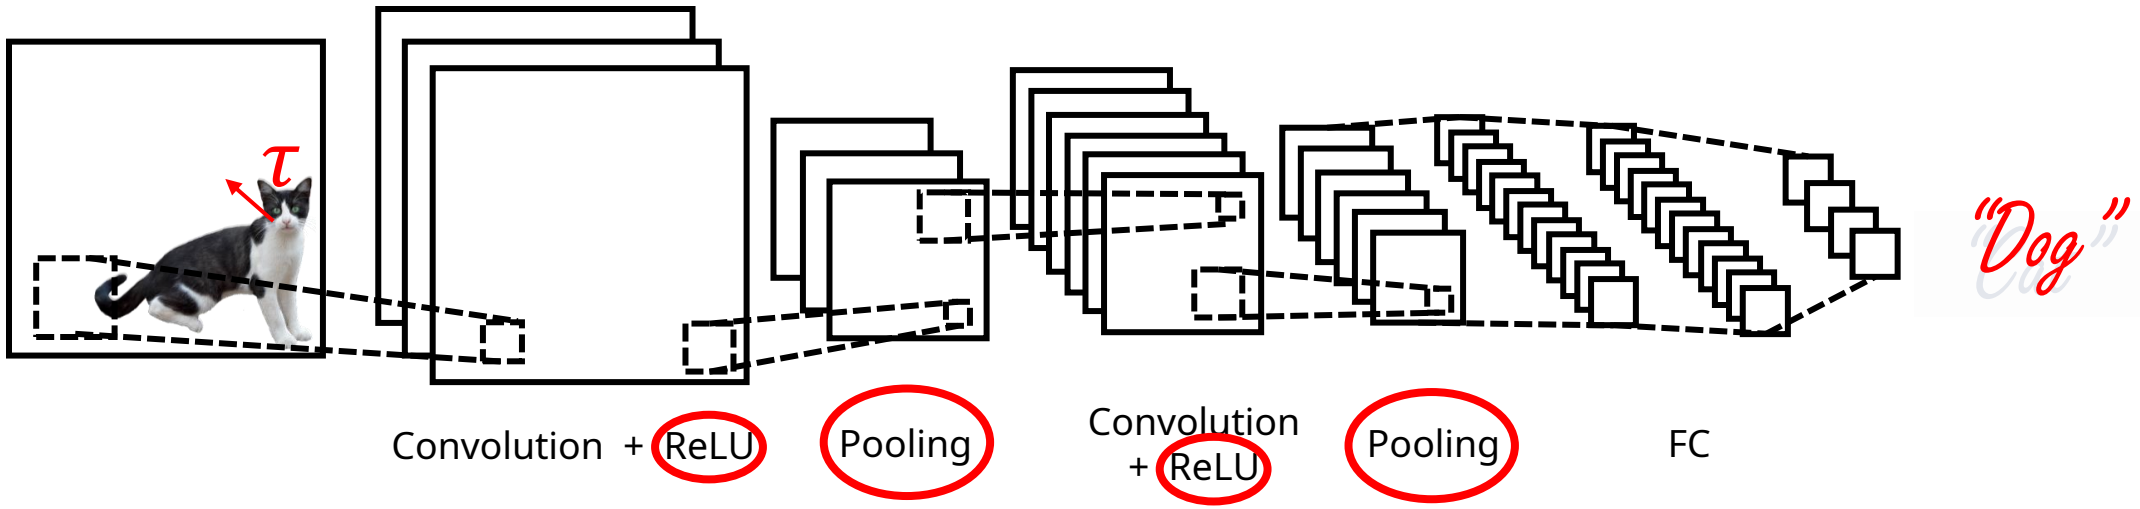
τ
"Dog"
Convolution + ReLU
Pooling
Convolution + ReLU
Pooling
FC

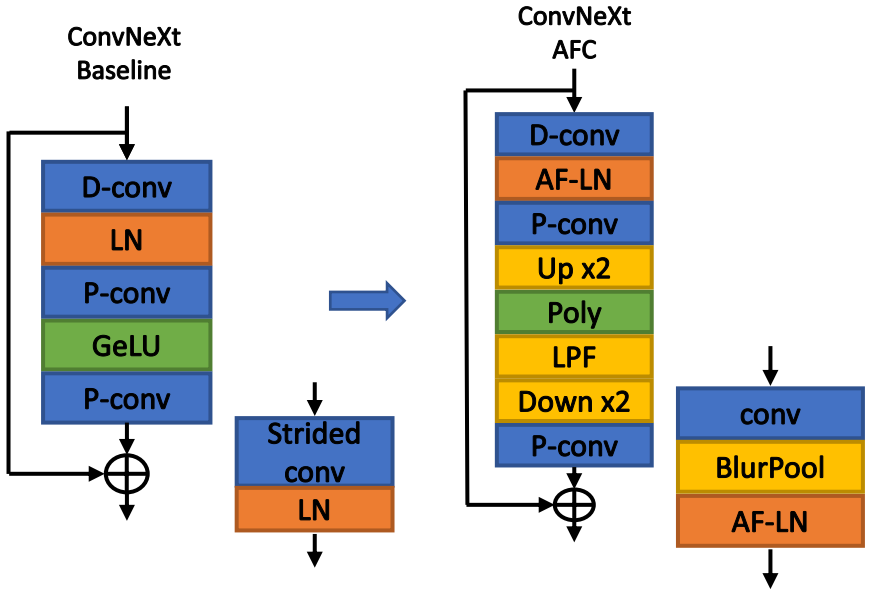
ConvNeXt Baseline
D-conv
LN
P-conv
GeLU
P-conv
Strided conv
LN
ConvNeXt AFC
D-conv
AF-LN
P-conv
Up x2
Poly
LPF
Down x2
P-conv
conv
BlurPool
AF-LN

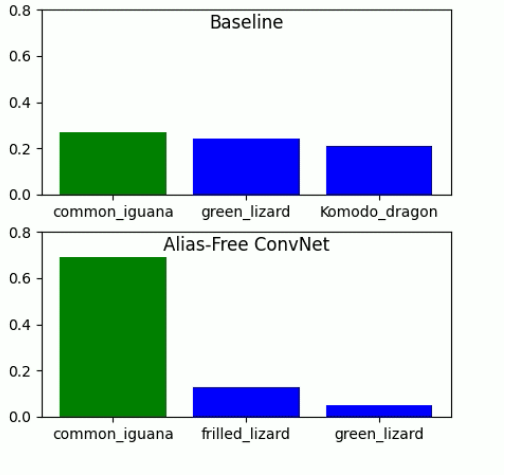
Baseline
0.8
0.6
0.4
0.2
0.0
common_iguana
green_lizard
Komodo_dragon
Alias-Free ConvNet
0.8
0.6
0.4
0.2
0.0
common_iguana
frilled_lizard
green_lizard

Convnets: built-in priors

- Label unaffected by image shift

→ Classifier should be “shift invariant”

- Features should shift with image

→ Convnet hidden layers should be “shift equivariant”

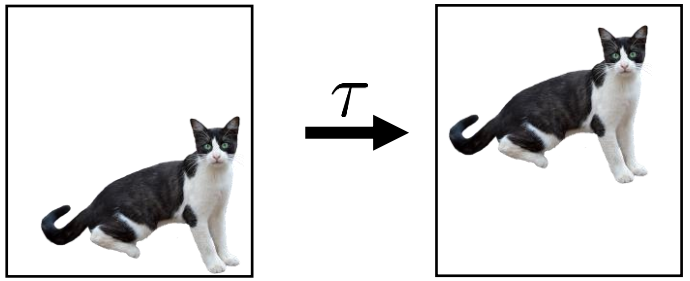

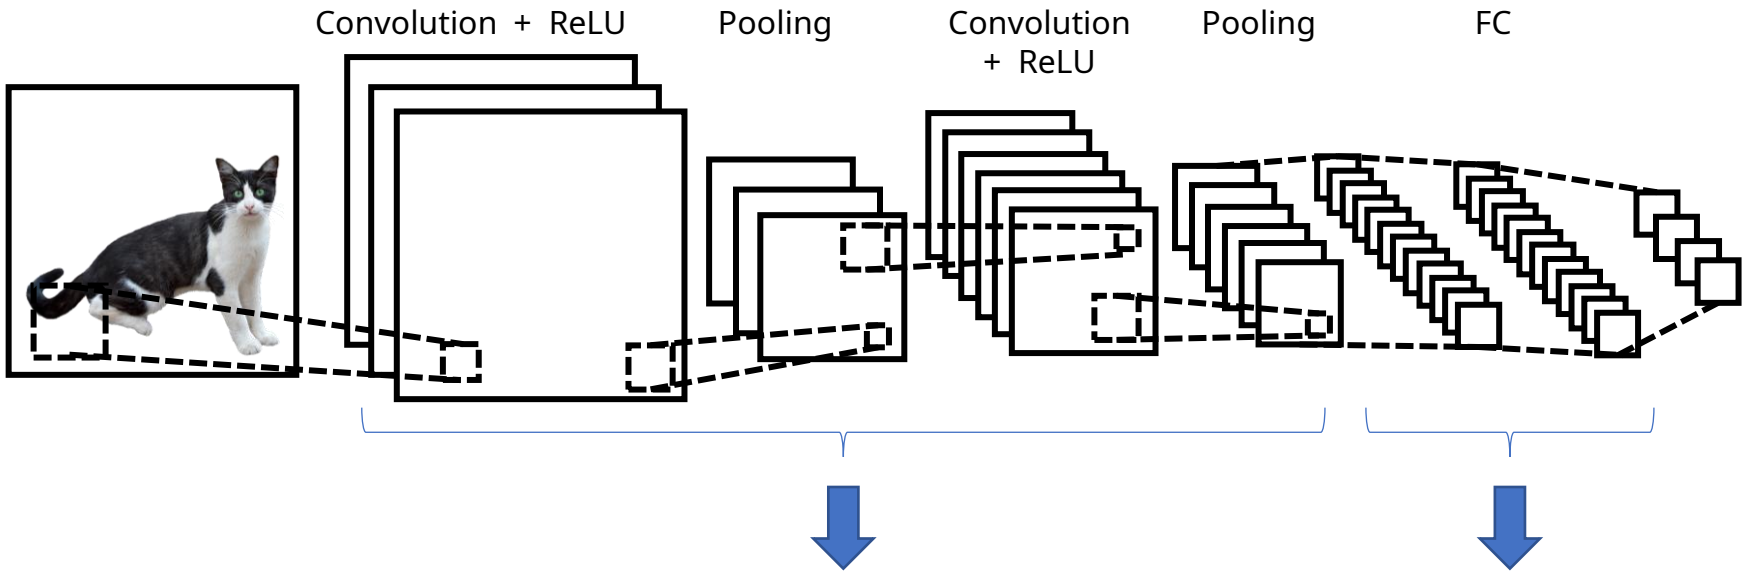

Single pixel shifts can change classification

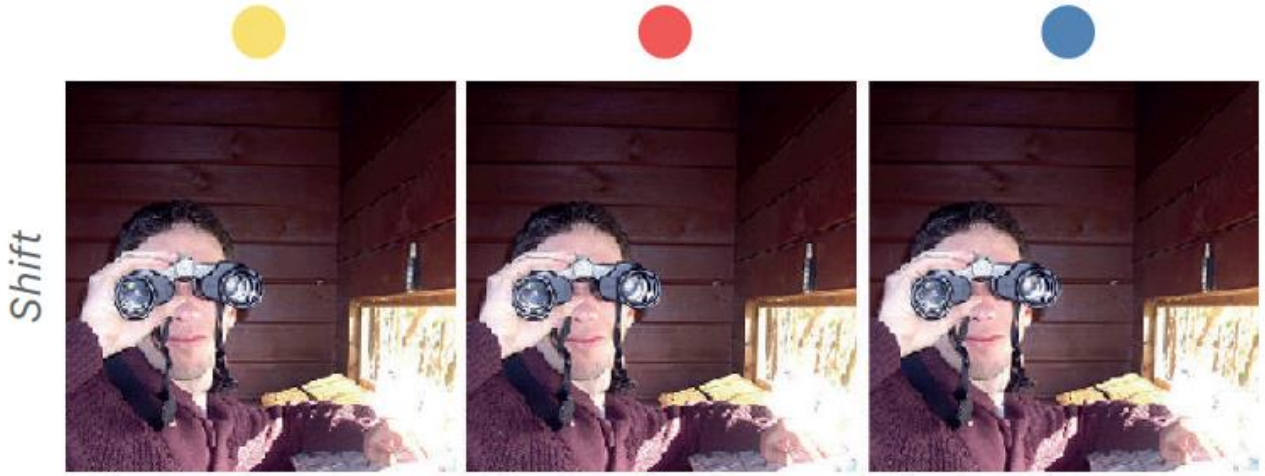

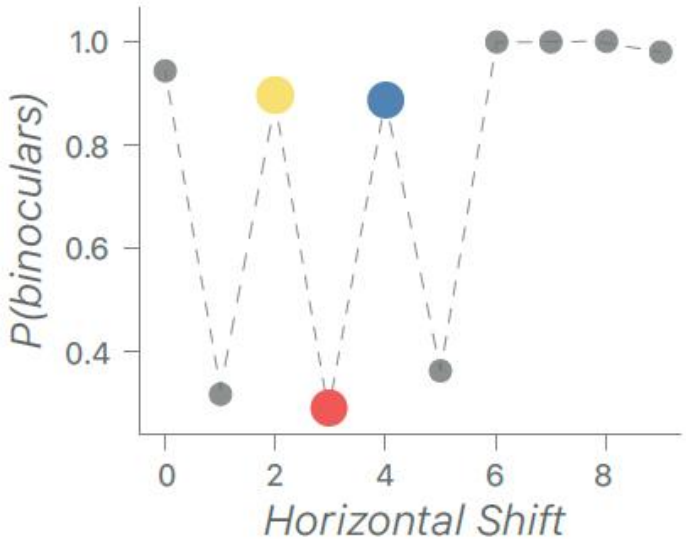

Why do deep convolutional networks generalize so poorly to small image transformations?
Azulay & Weiss, JMLR 2019

Is Feature extractor equivariant?

Building blocks:

Equivariant?

- Convolution layers — Yes
- Activation Layers (e.g., ReLU, GeLU) — Yes
- Normalization layers — Yes (for some types)

- Pooling layers — No

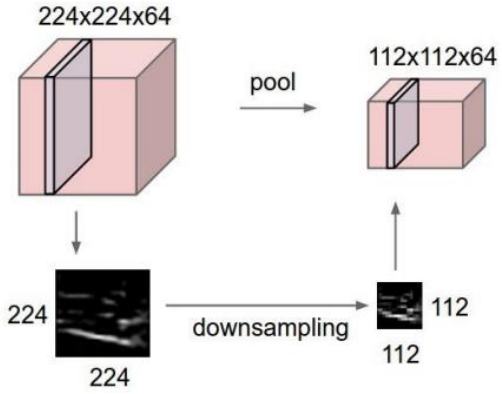

What is the problem? Aliasing!

Nyquist Theorem: Perfect reconstruction if sampling rate > 2 x bandwidth

Making Convolutional Networks Shift-Invariant Again
Zhang, ICML 2019

What is the problem? Aliasing!

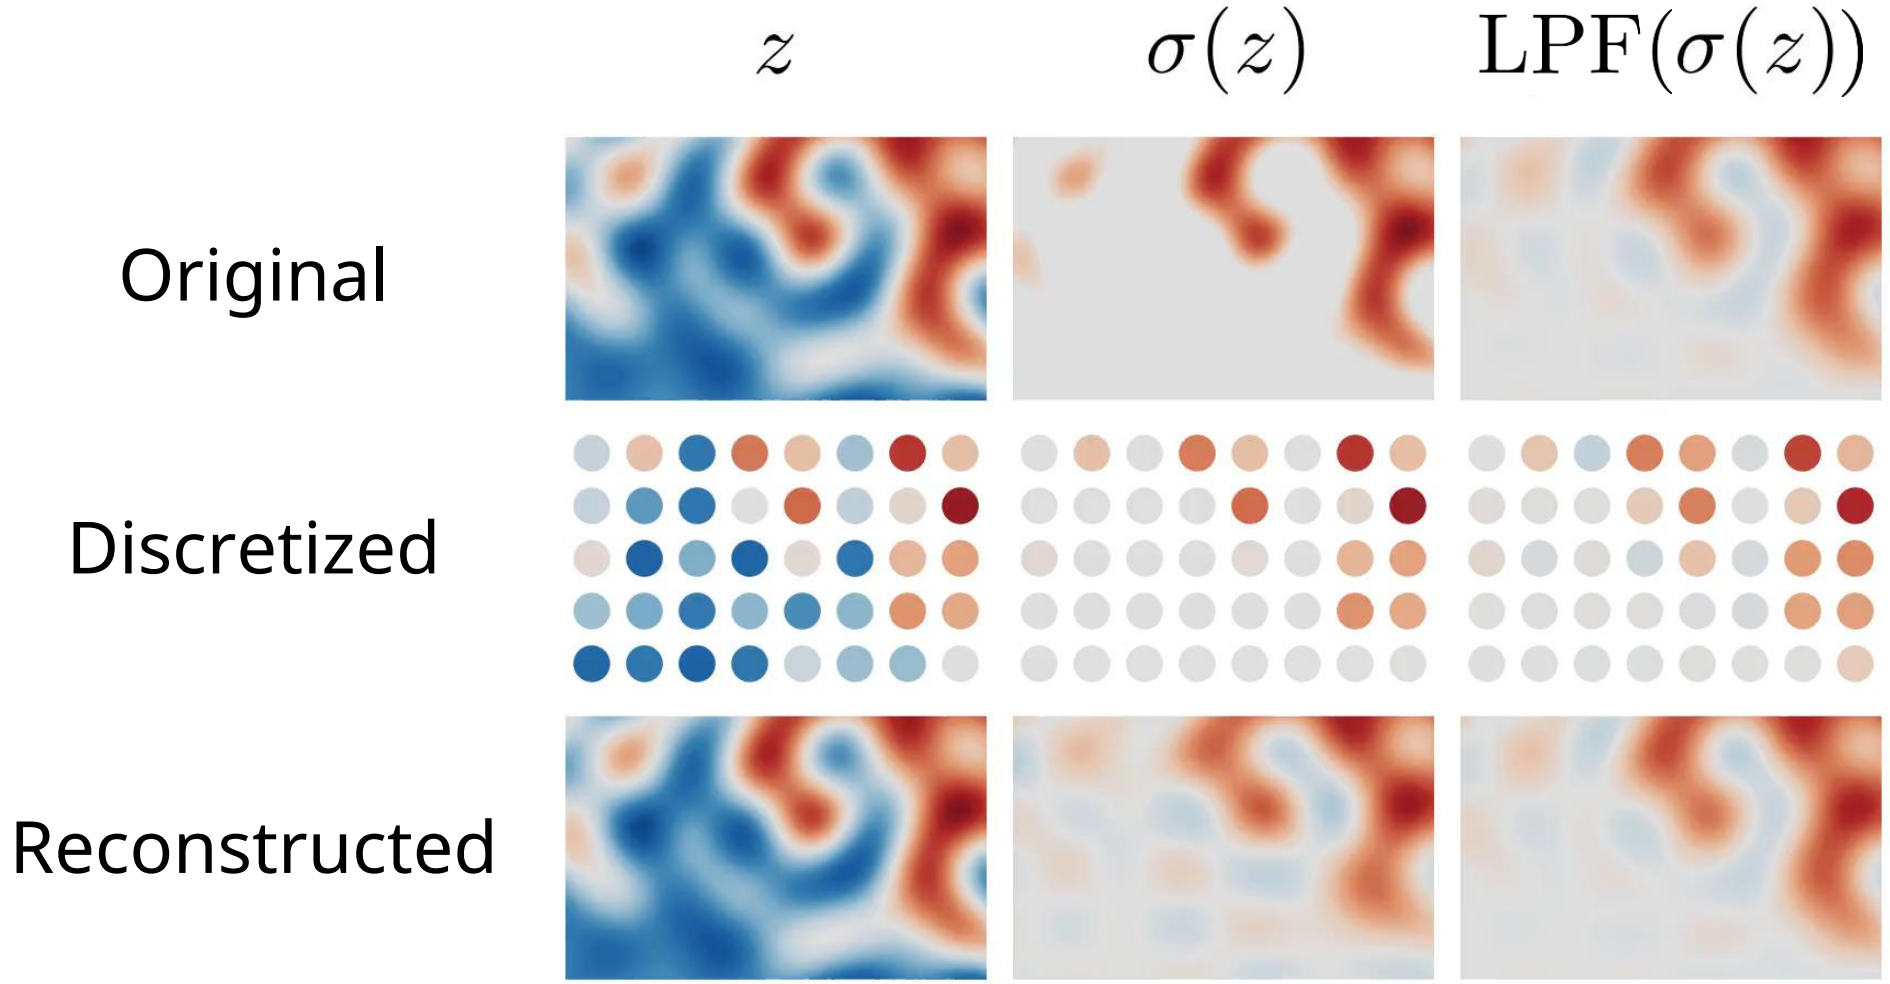

Alias-Free Generative Adversarial Networks
Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, Timo Aila, NeurIPS 2021

Alias Free Convnet (AFC)

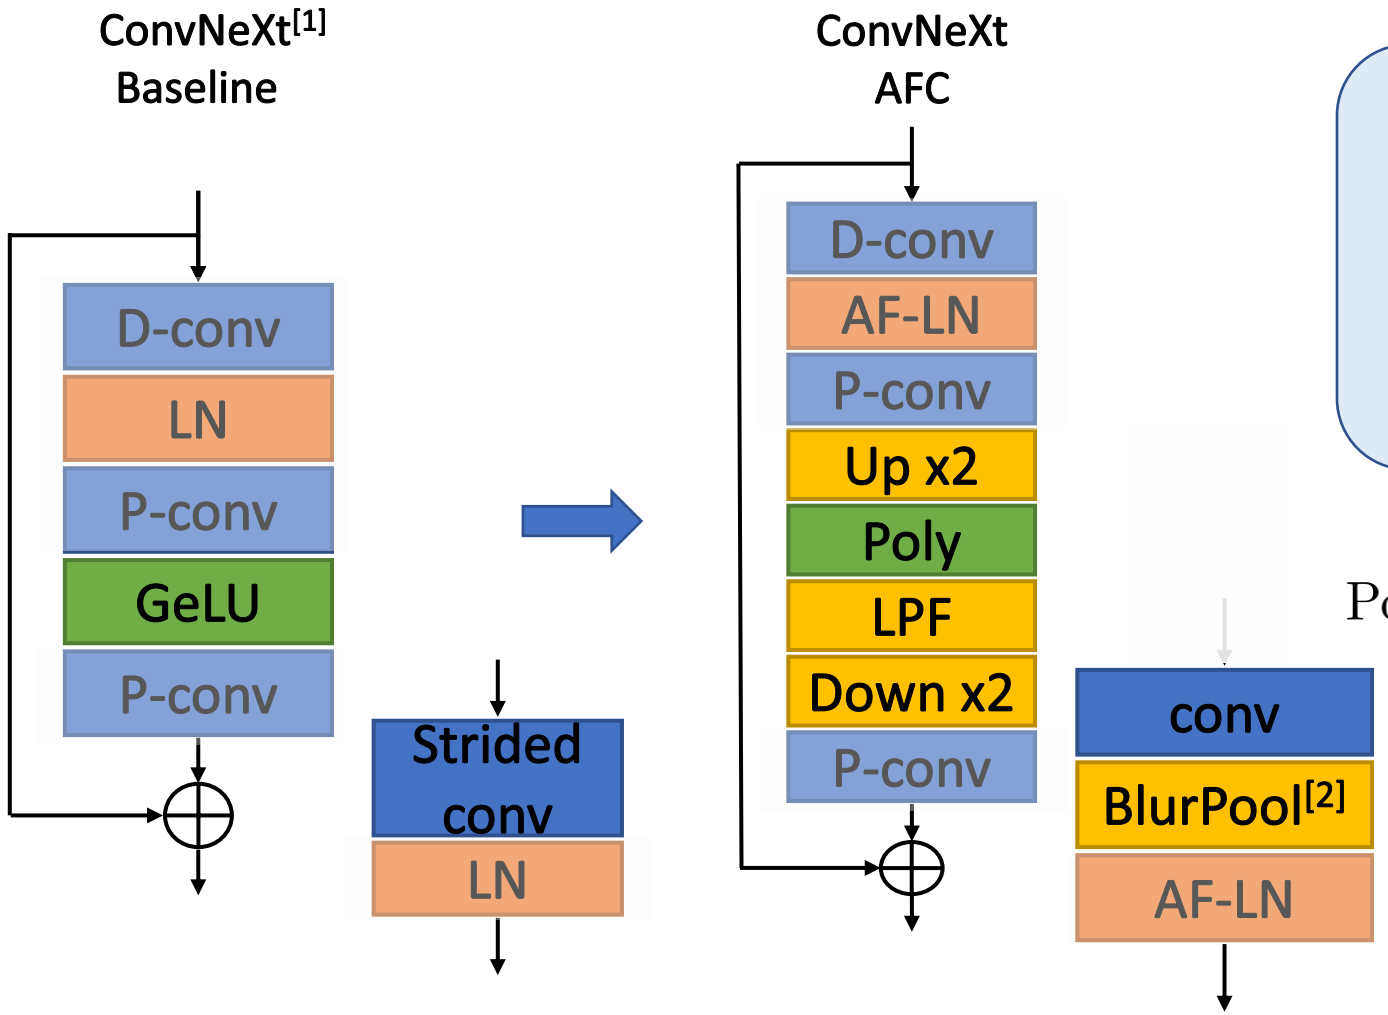

Theorem: AFC have perfect

1) feature equivariance

2) classifier invariance

$$\mathrm{Poly}(x) = a_0 + a_1 x + a_2 x^2$$

[1] **A ConvNet for the 2020s**
Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, Saining Xie, CVPR 2022

[2] **Making Convolutional Networks Shift-Invariant Again**
Zhang, ICML 2019

Why this works? A look in the frequency domain

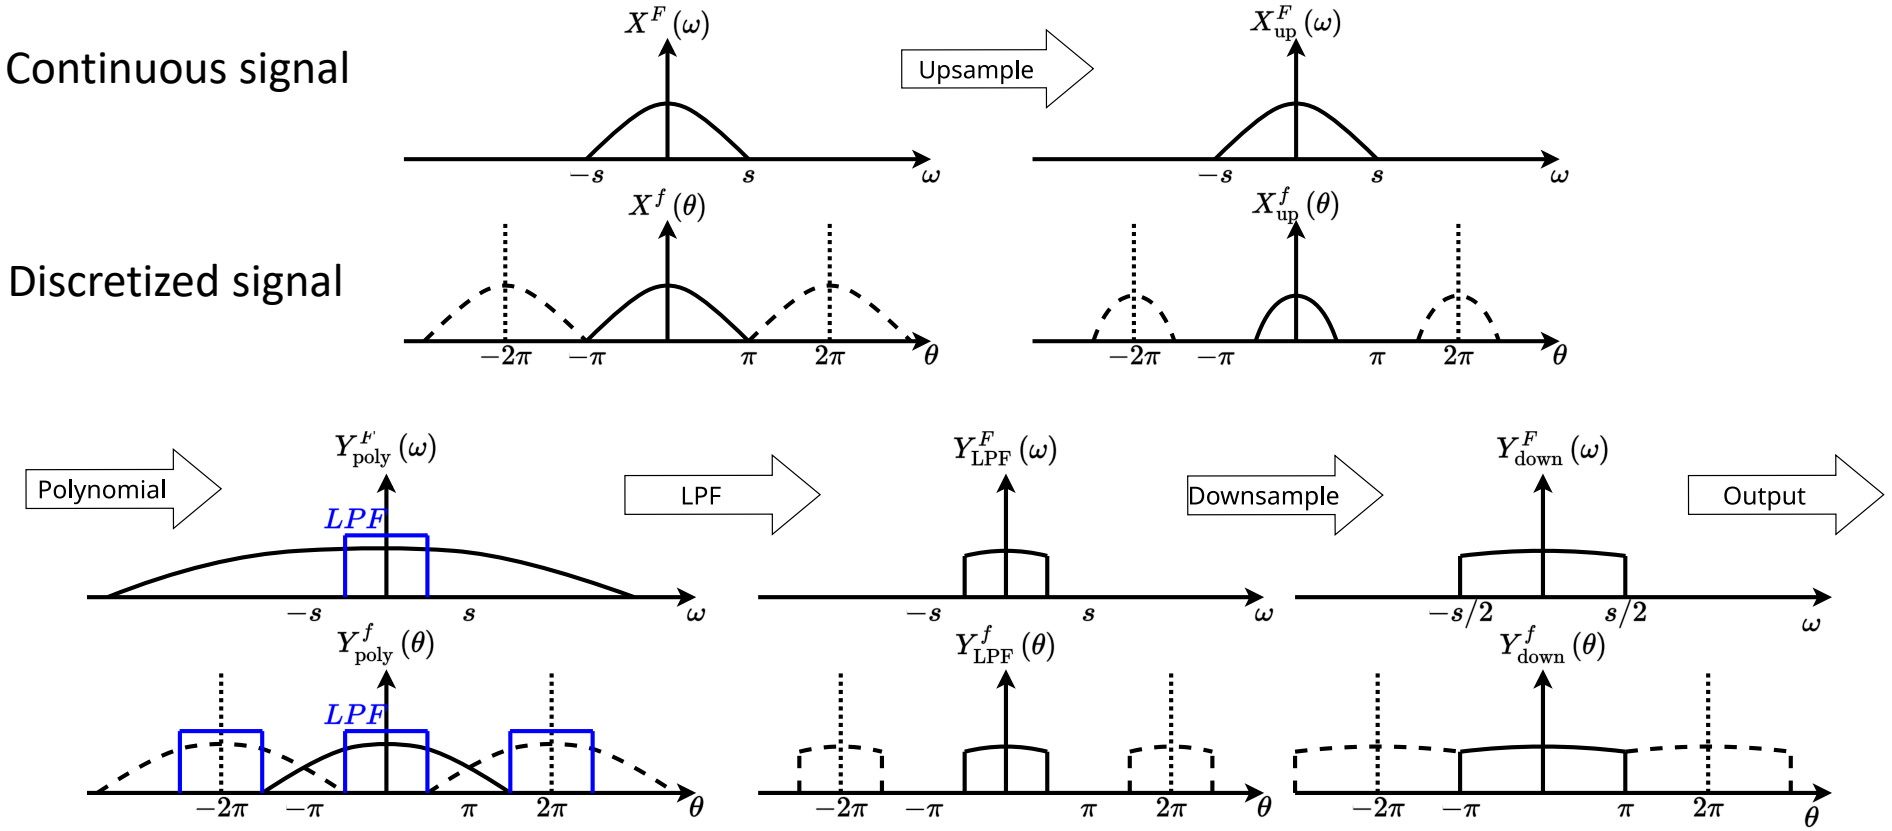

Alias-Free Generative Adversarial Networks
Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, Timo Aila, NeurIPS 2021

What's the problem without polynomials?

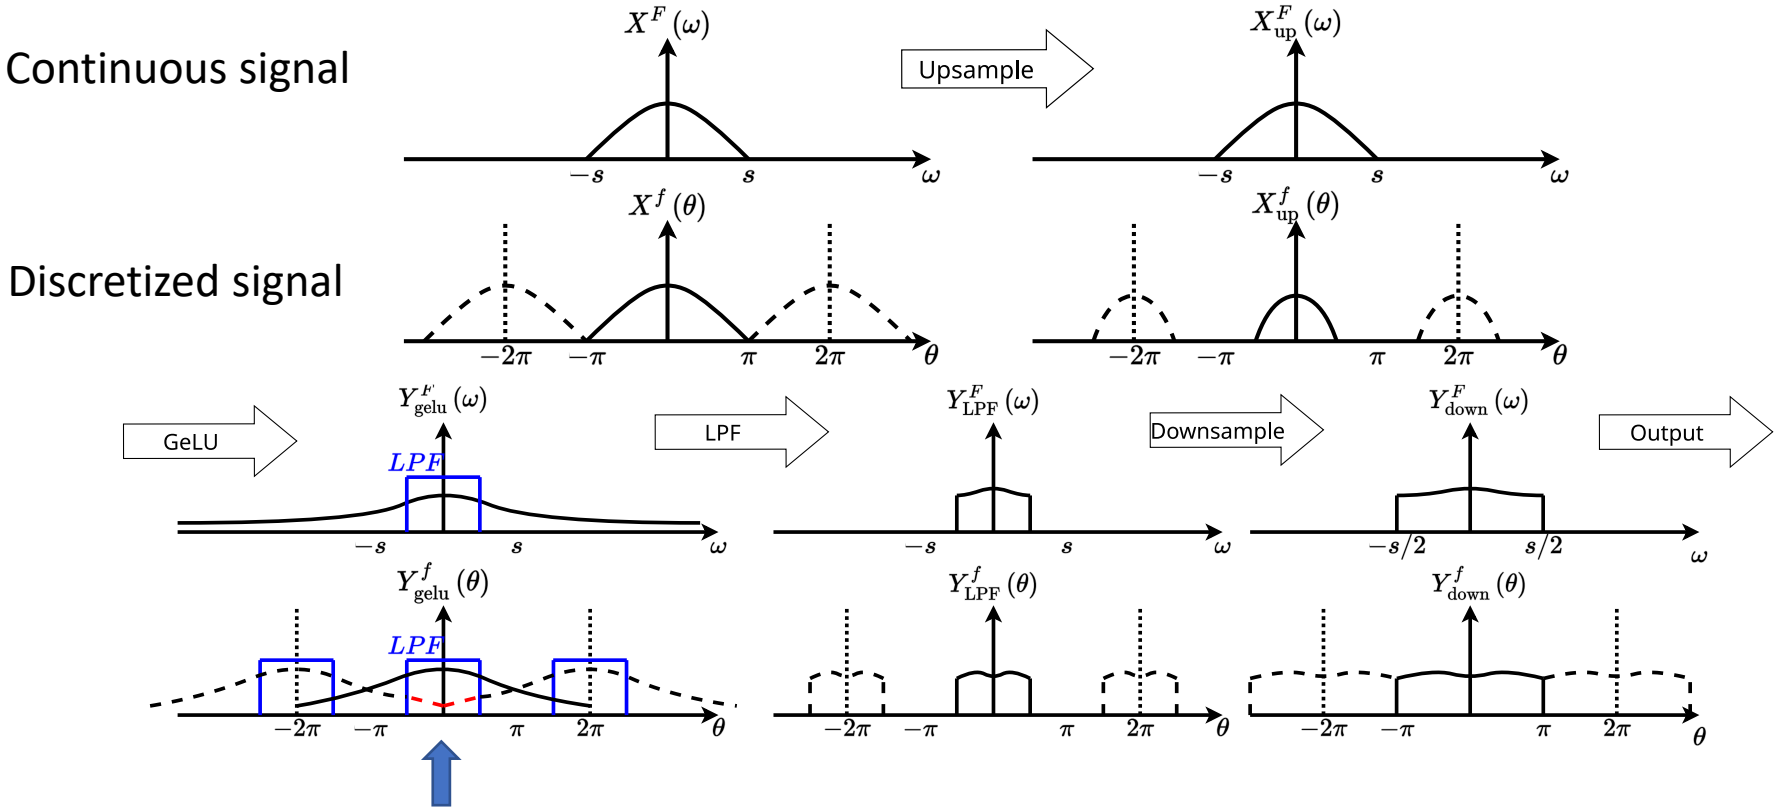

Still some aliasing….

Alias-Free Generative Adversarial Networks
Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, Timo Aila, NeurIPS 2021

Results (1): Perfect Equivariance

- Representations change after ½ pixel shift:

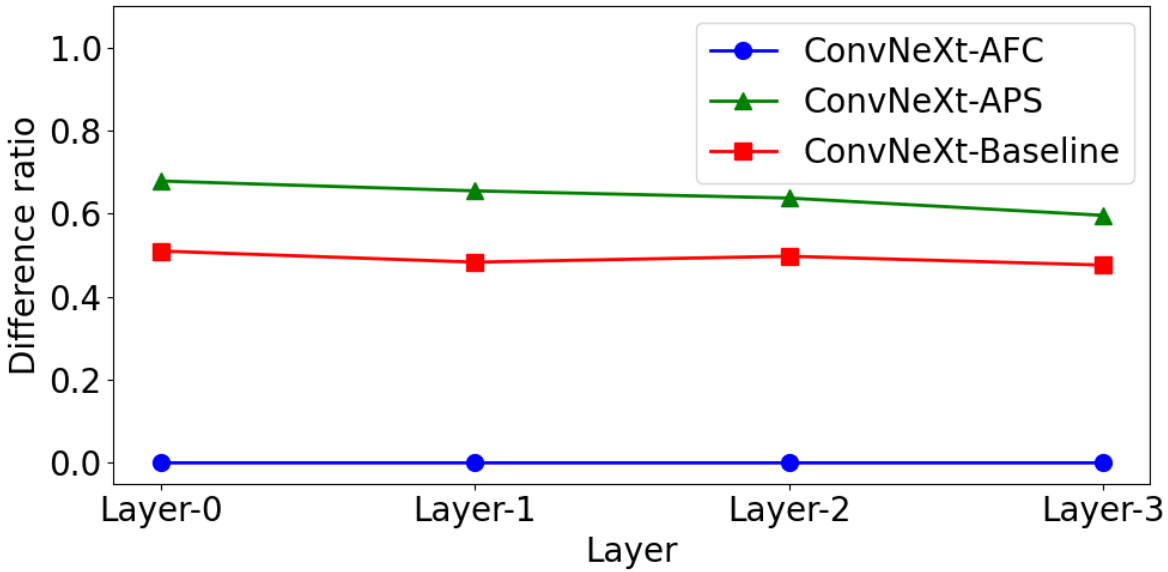

$$\left(\text{diff} \triangleq \frac{1}{CHW}\sum_{c,i,j}\frac{\left|y^0_{c,i,j}-y^1_{c,i,j}\right|}{\max\left(\left|y^0_{c,i,j}\right|,\left|y^1_{c,i,j}\right|\right)+\varepsilon}\right)$$

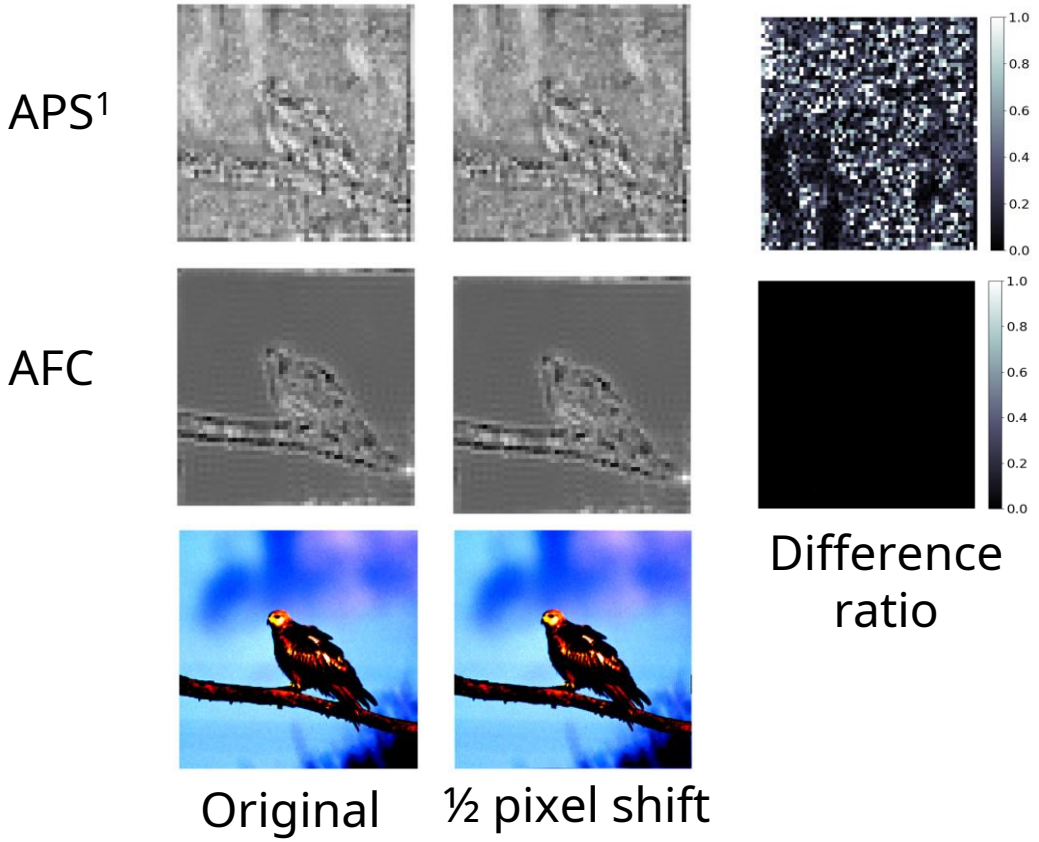

[1] **Truly shift-equivariant convolutional neural networks with adaptive polyphase upsampling**
Anadi Chaman, Ivan Dokmanić

Results (2): Perfect Invariance

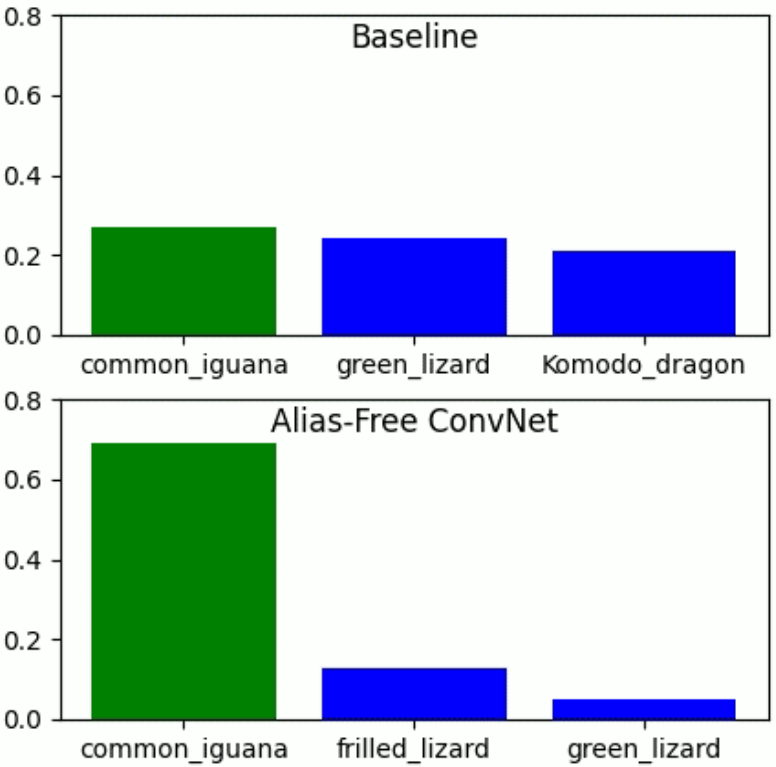

Results (3): Certified adversarial accuracy

Model (ImageNet)	Test Accuracy	Integer grid	Half-pixel grid	Fractional grid
ConvNeXt-Baseline	**82.12**	76.63	73.65	77.82
ConvNeXt-APS	82.11	**82.11**	79.68	76.31
ConvNeXT-AFC (ours)	81.04	81.04	**81.04**	**81.04**

$$T_{\text{integer}} = \{(i,\ j)\ |1 \leq i, j \leq 31\}$$

$$T_{\text{half}} = \left\{\left(\frac{i}{2},\ \frac{j}{2}\right)\ |1 \leq i, j \leq 63\right\}$$

$$T_{\text{frac}} = \left\{\left(\frac{m_1}{n_1}, \frac{m_2}{n_2}\right)\ |1 \leq m_{1,2} \leq n_{1,2} \leq 12\right\}$$

Results (4): Improved robustness to “camera translations”

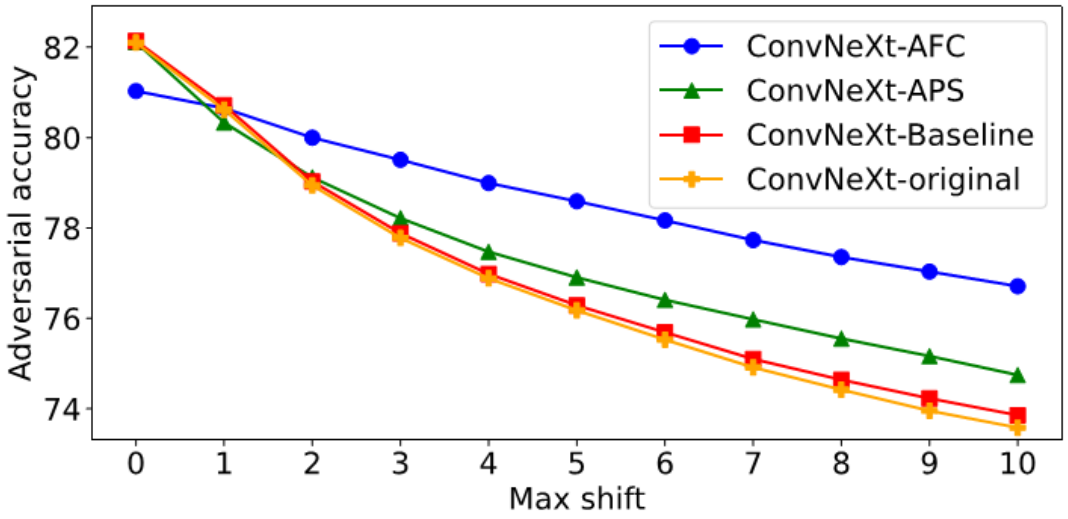

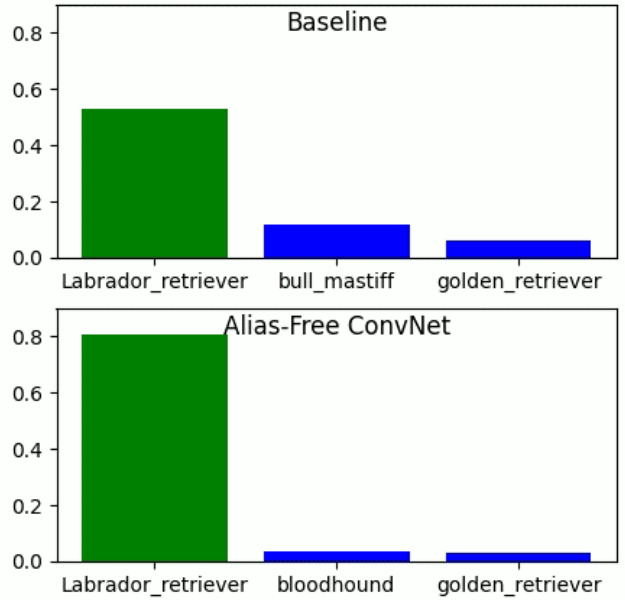

Summary: “Alias free convnets”

- Perfect invariance and equivariance to fractional shifts
- SOTA accuracy with adversarial fractional shifts
- First polynomial network with near-SOTA on ImageNet

Thank you!