ELASTIC: IMPROVING CNNS WITH DYNAMIC SCALING POLICIES

https://github.com/allenai/elastic

Huiyu Wang[1] Aniruddha Kembhavi[2] Ali Farhadi[2,3,4] Alan Yuille[1] Mohammad Rastegari[2,4]
[1]Johns Hopkins University [2]Allen Institute for Artificial Intelligence [3]University of Washington [4]Xnor.ai

MOTIVATION

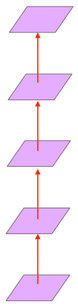

Fixed

- Previous CNNs use fixed scaling policies.
- We learn **dynamic scaling policies** from data.

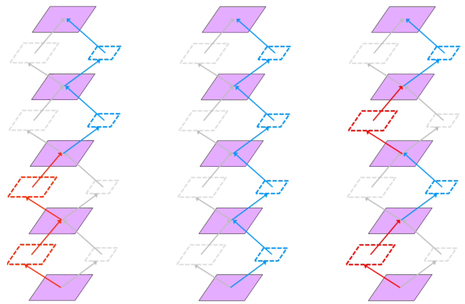

Dynamic

ANALYSIS

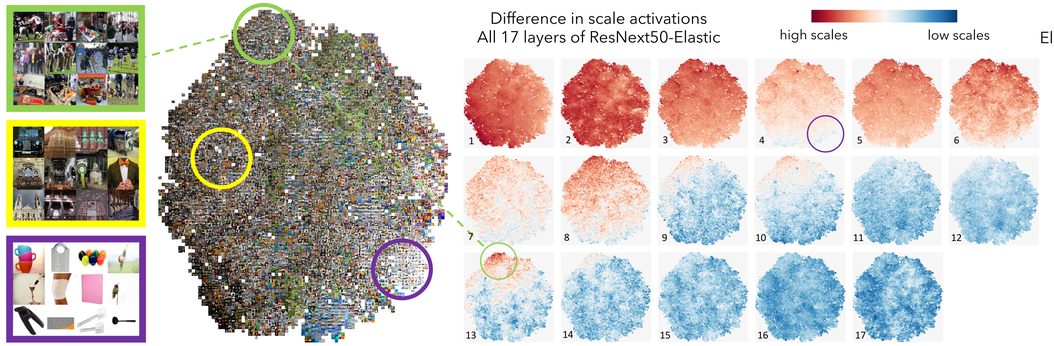

PYRAMIDS VS. ELASTIC

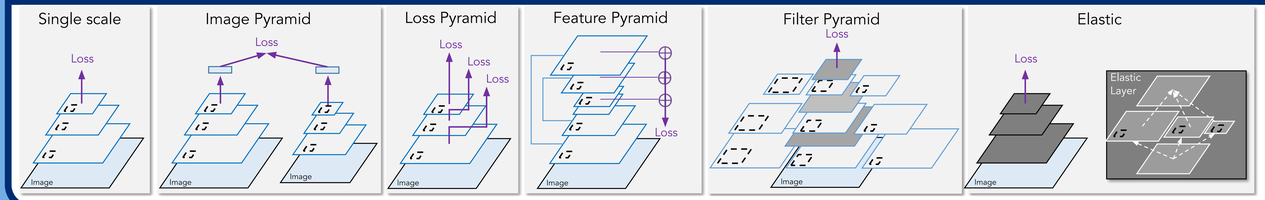

IMPLEMENTATION

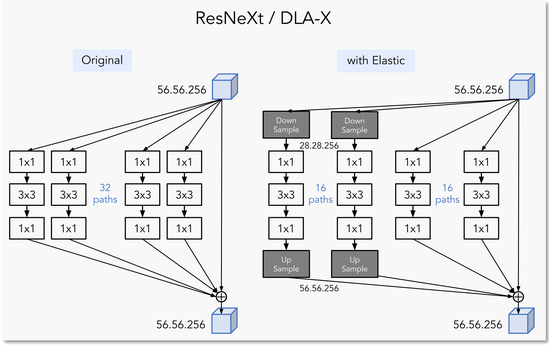

DYNAMIC SCALING POLICIES

ANALYSIS

- **Which category requires more detail?**

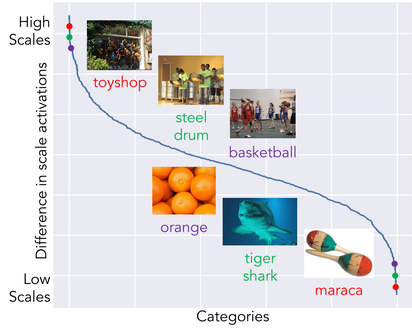

- **Does ELASTIC generalize to out-of-distribution scales?**

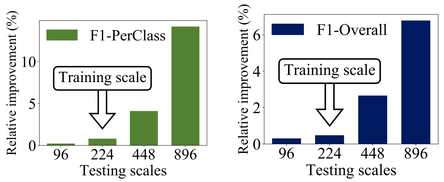

RESULTS

- **ImageNet classification**

Model	# Params	FLOPs	Top-1	Top-5
DLA-X60*	17.6M	3.6B	21.92	6.03
DLA-X60+Elastic	17.6M	3.2B	**21.25**	**5.71**
DLA-X102	26.8M	6.0B	21.5	-
DLA-X102+Elastic	25.0M	6.0B	**20.71**	**5.38**
ResNeXt50*	25.0M	4.2B	22.23	6.25
ResNeXt50+Elastic	25.2M	4.2B	**21.56**	**5.83**
SE-ResNeXt50*	27.6M	4.2B	21.87	5.93
SE-ResNeXt50+Elastic	27.8M	4.2B	**21.38**	**5.86**
ResNeXt101*	44.2M	8.0B	21.18	5.83
ResNeXt101+Elastic	44.3M	7.9B	**20.83**	**5.41**

- **MS COCO multi-label classification**

Model	F1-PerClass	F1-Overall
DLA-X60*	70.79	75.41
DLA-X60+Elastic	**71.35**	**75.77**
ResNeXt50*	70.12	74.52
ResNeXt50+Elastic	**71.08**	**75.37**
ResNeXt101*	70.95	75.21
ResNeXt101+Elastic	**71.83**	**75.93**

- **PASCAL VOC semantic segmentation**

Model	Original	Elastic
DLA-X60*	69.96	**73.59**
ResNeXt50*	75.29	**77.70**
ResNeXt101*	77.47	**78.51**