


R-CNN: Region-based CNN

A framework for object detection with ConvNets



Input image

Extract region proposals

Compute CNN features

Classify regions

Any proposal method
(boxes or masks)

Any architecture,
not just AlexNet
(including GoogLeNet)

SoftMax, SVM

R-CNN-ZF baseline on ILSVR14

	val2	test
ILSVRC13 + AN	29.7%	
ILSVRC14 + AN	31.2%	
ILSVRC14 + ZF	32.1%	
ILSVRC14 + ZF + BBOX	33.4%	
ILSVRC14 + ZF + BBOX		34.5%

Take away from *other* entries in ILSVRC14:
huge gains come from improving the CNN module