Deep Learning for Computer Vision

Lecture 11: Building ConvNets in Keras, Pre-trained Neural Networks, Performing Network Surgery, Domain Transfer and Fine-Tuning

Peter Belhumeur

Computer Science
Columbia University

CIFAR-10

[Krizhevsky 2009]

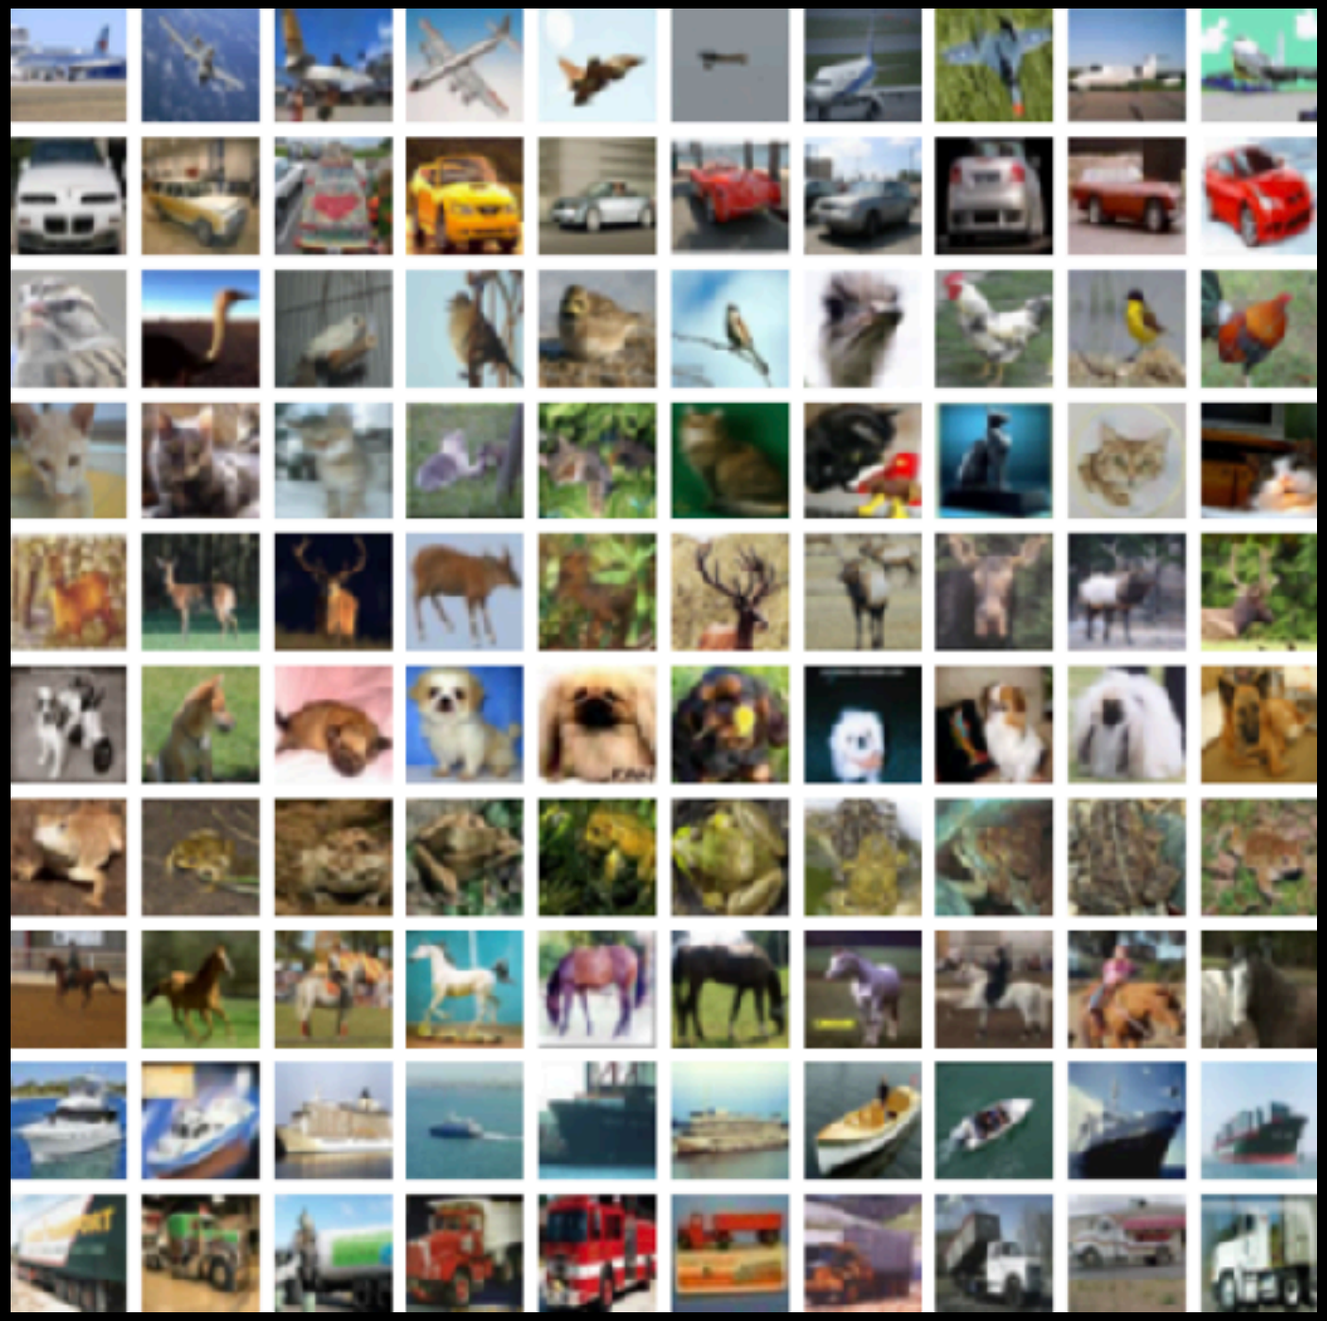

10 categories x 6,000 instances = 60,000 images (32x32x3)

Small Convolutional Net

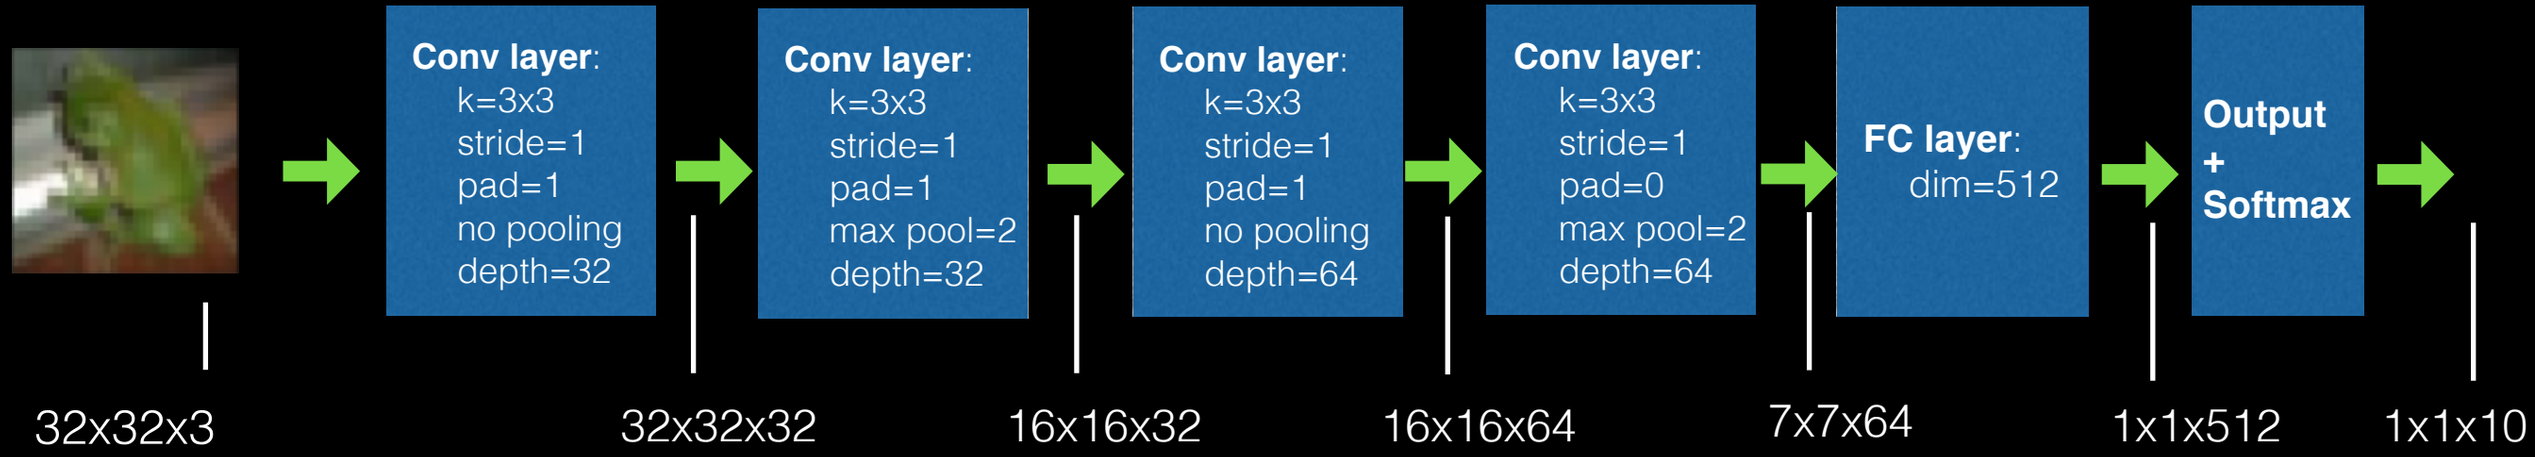

Open cifar10_keras.html in browser

Cats vs. Dogs

Data from [Kaggle.com]

vs.

2 categories x 11,000 instances = 22,000 images

Small Convolutional Net

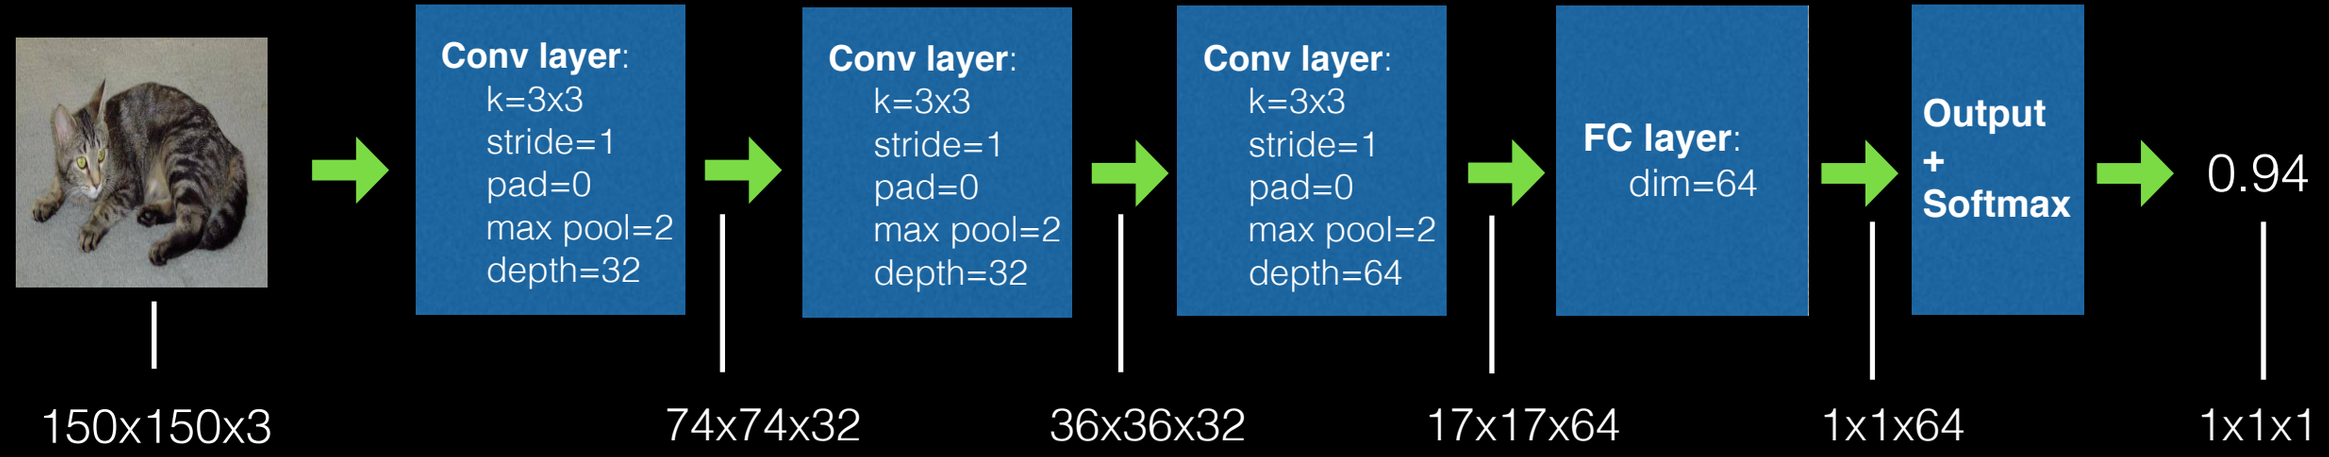

Open imagenet_vgg16_fine-tuning.html in browser

Why can’t we use one a much deeper ConvNet for this problem?

Open imagenet_vgg16_fine-tuning.html in browser

We usually don’t have enough data!

…but here is a FANTASTIC trick!

Domain Transfer

- Train a **deep** ConvNet using lots of data on a large image classification problem like ImageNet.
- Save the weights.
- Chop off the output layer (or final layers) of this ConvNet.
- Replace the output layer (or final layers) with one that fits your problem.
- Freeze the weights in the old layers and train on your data, allowing the weights in the new layer to settle.
- Unfreeze the whole network and train.

VGG16 [Simonyan and Zisserman, 2015]

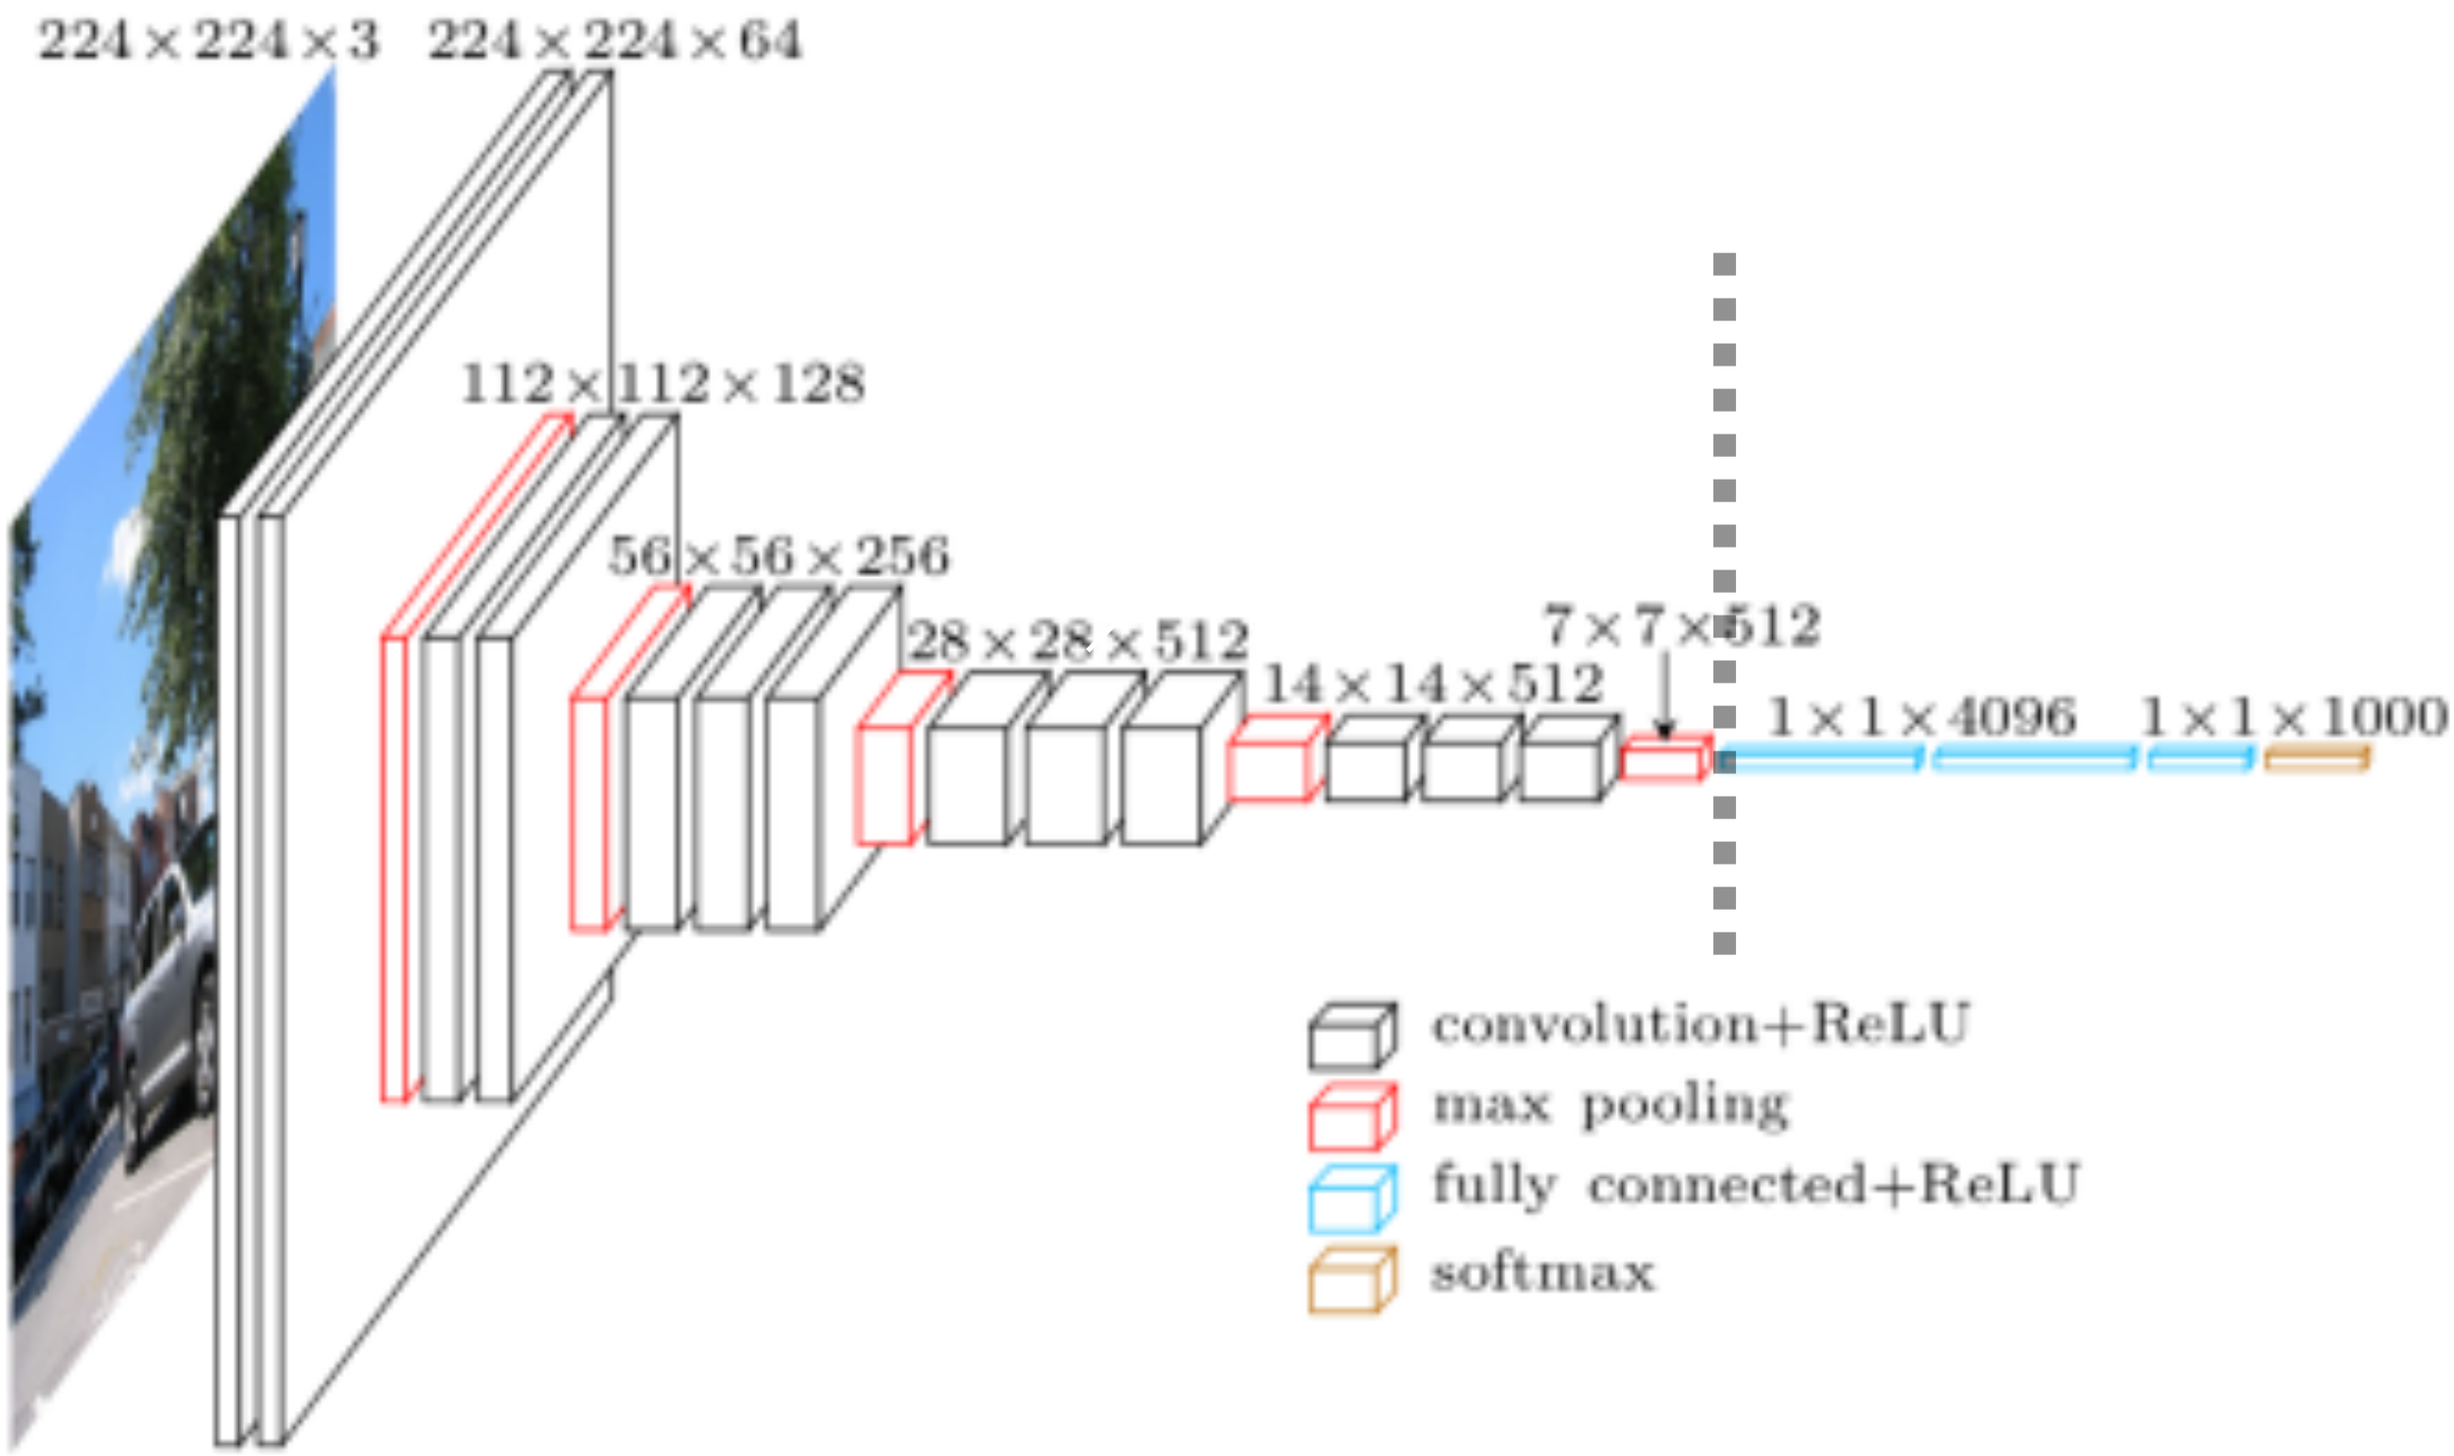

ImageNet

[Deng et al. 2009]

20,000+ categories x ~1000 instances = 14,000,000+ images

Domain Transfer + Fine-Tuning

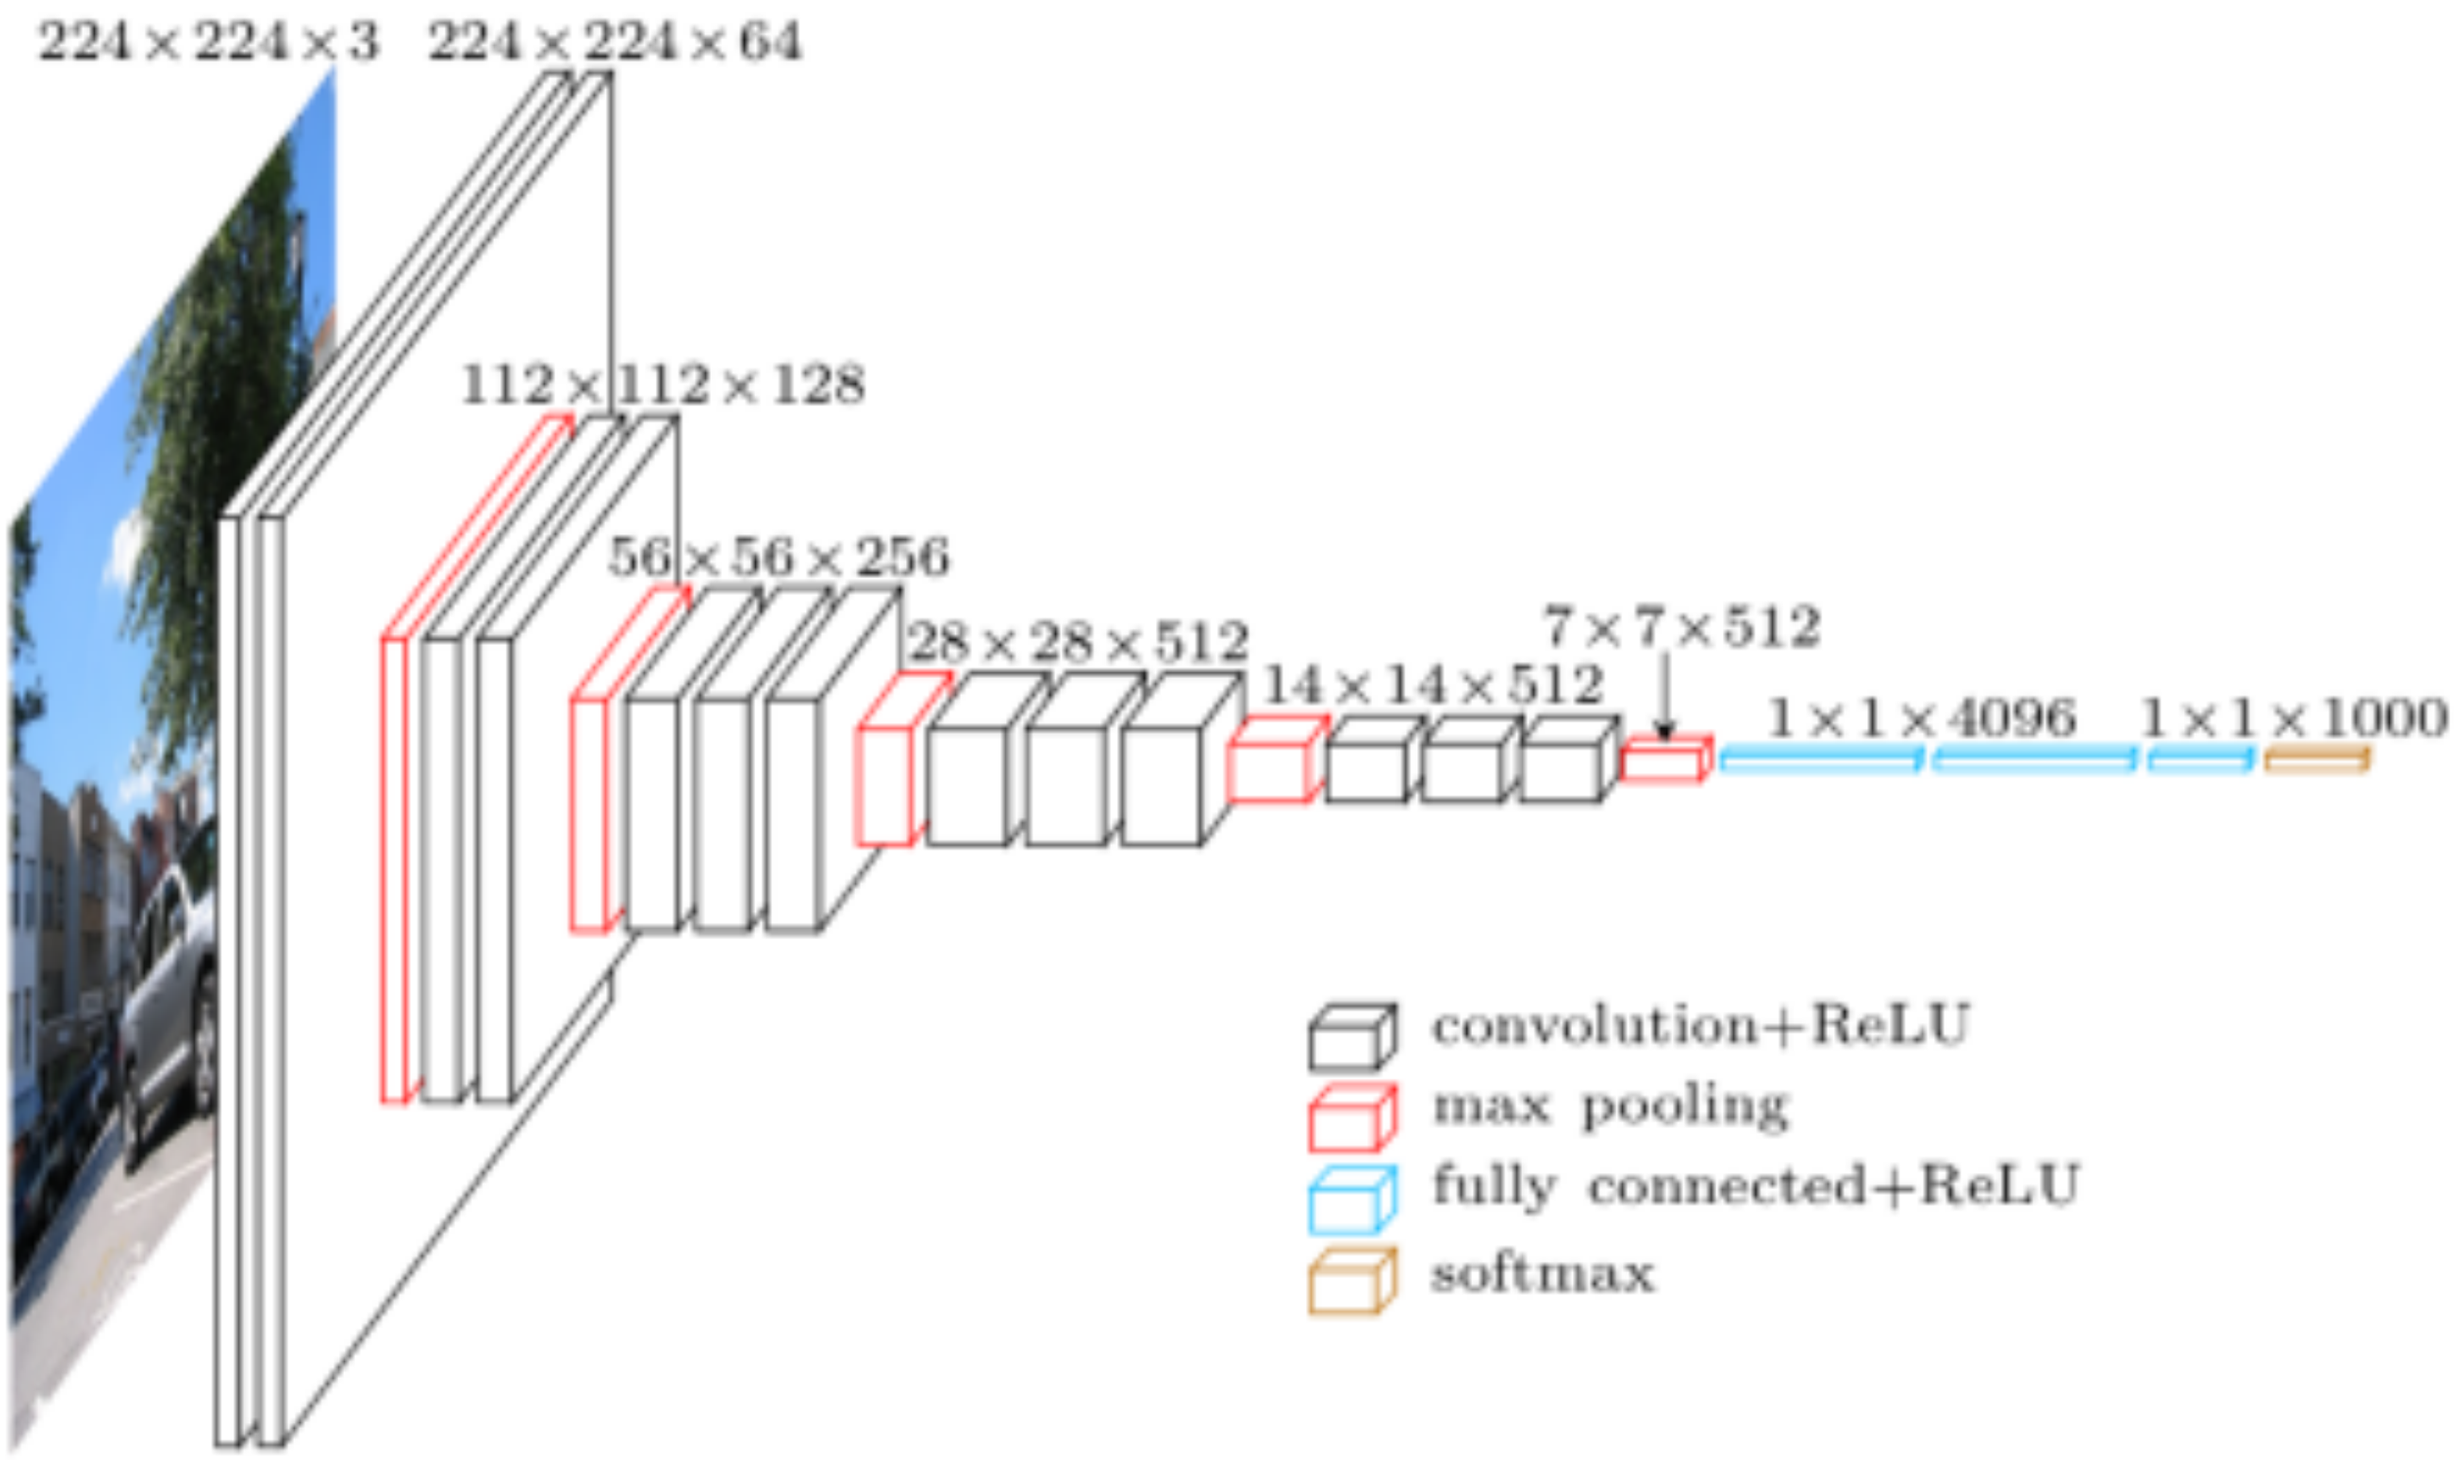

Network Surgery

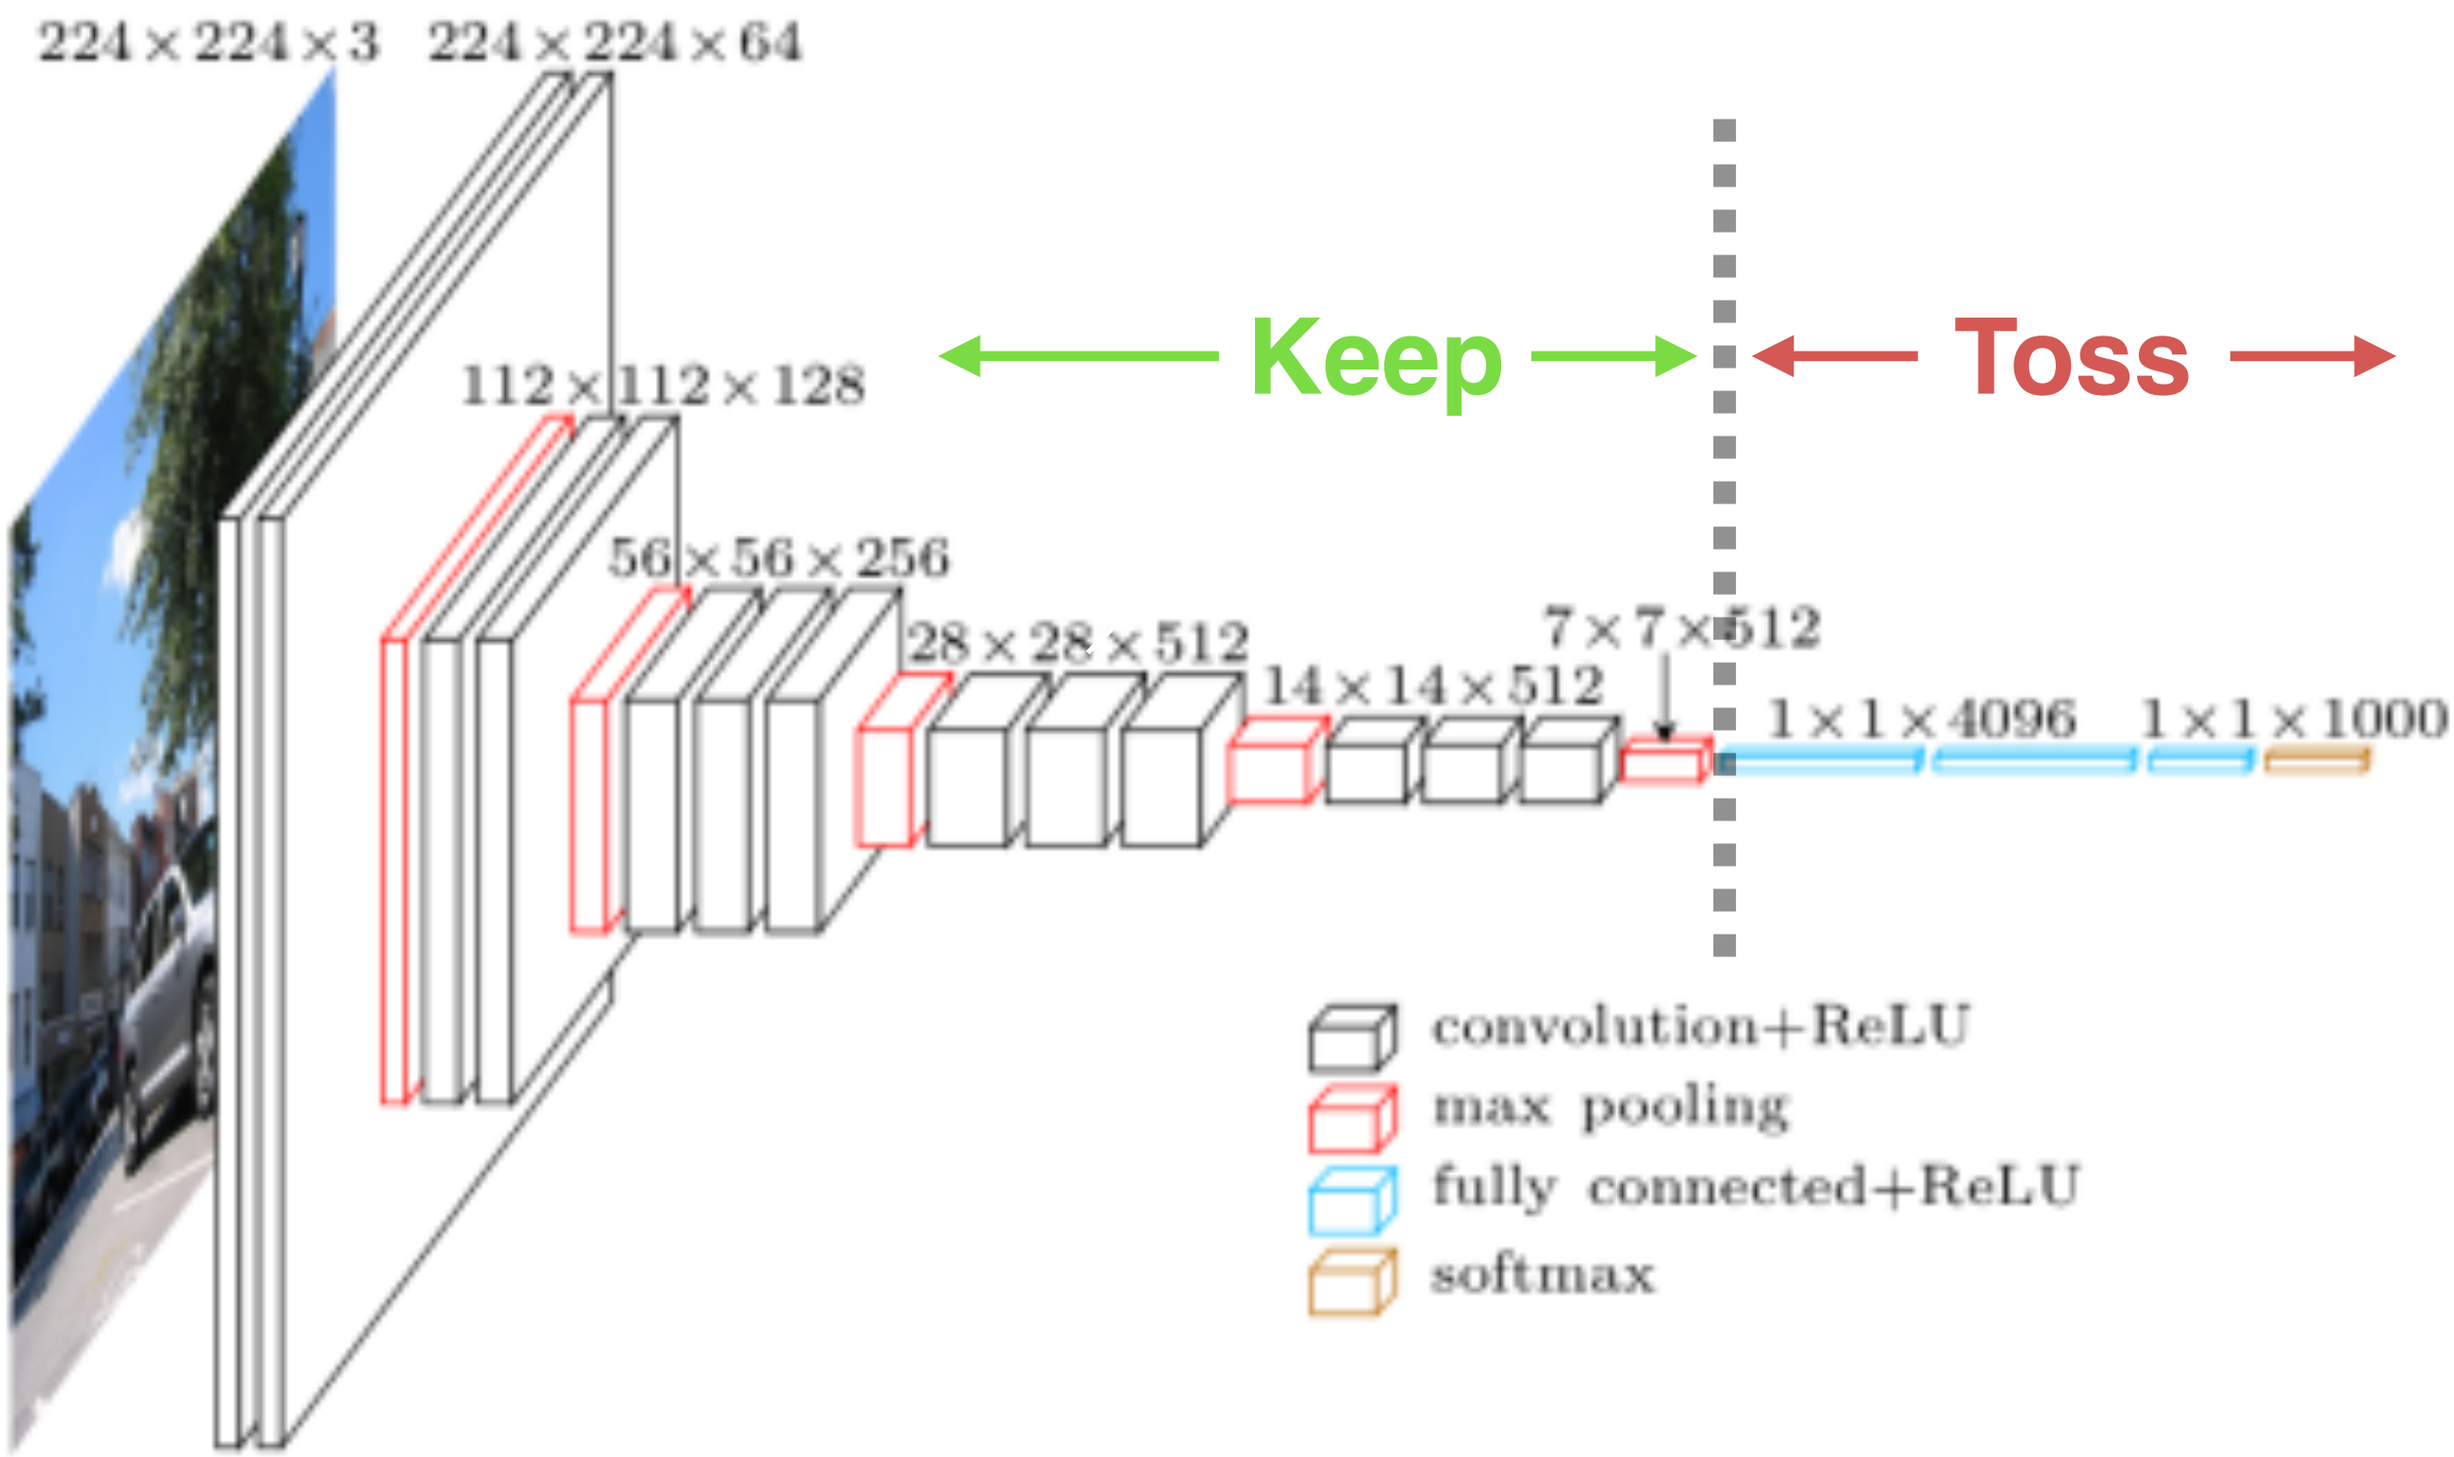

Network Surgery

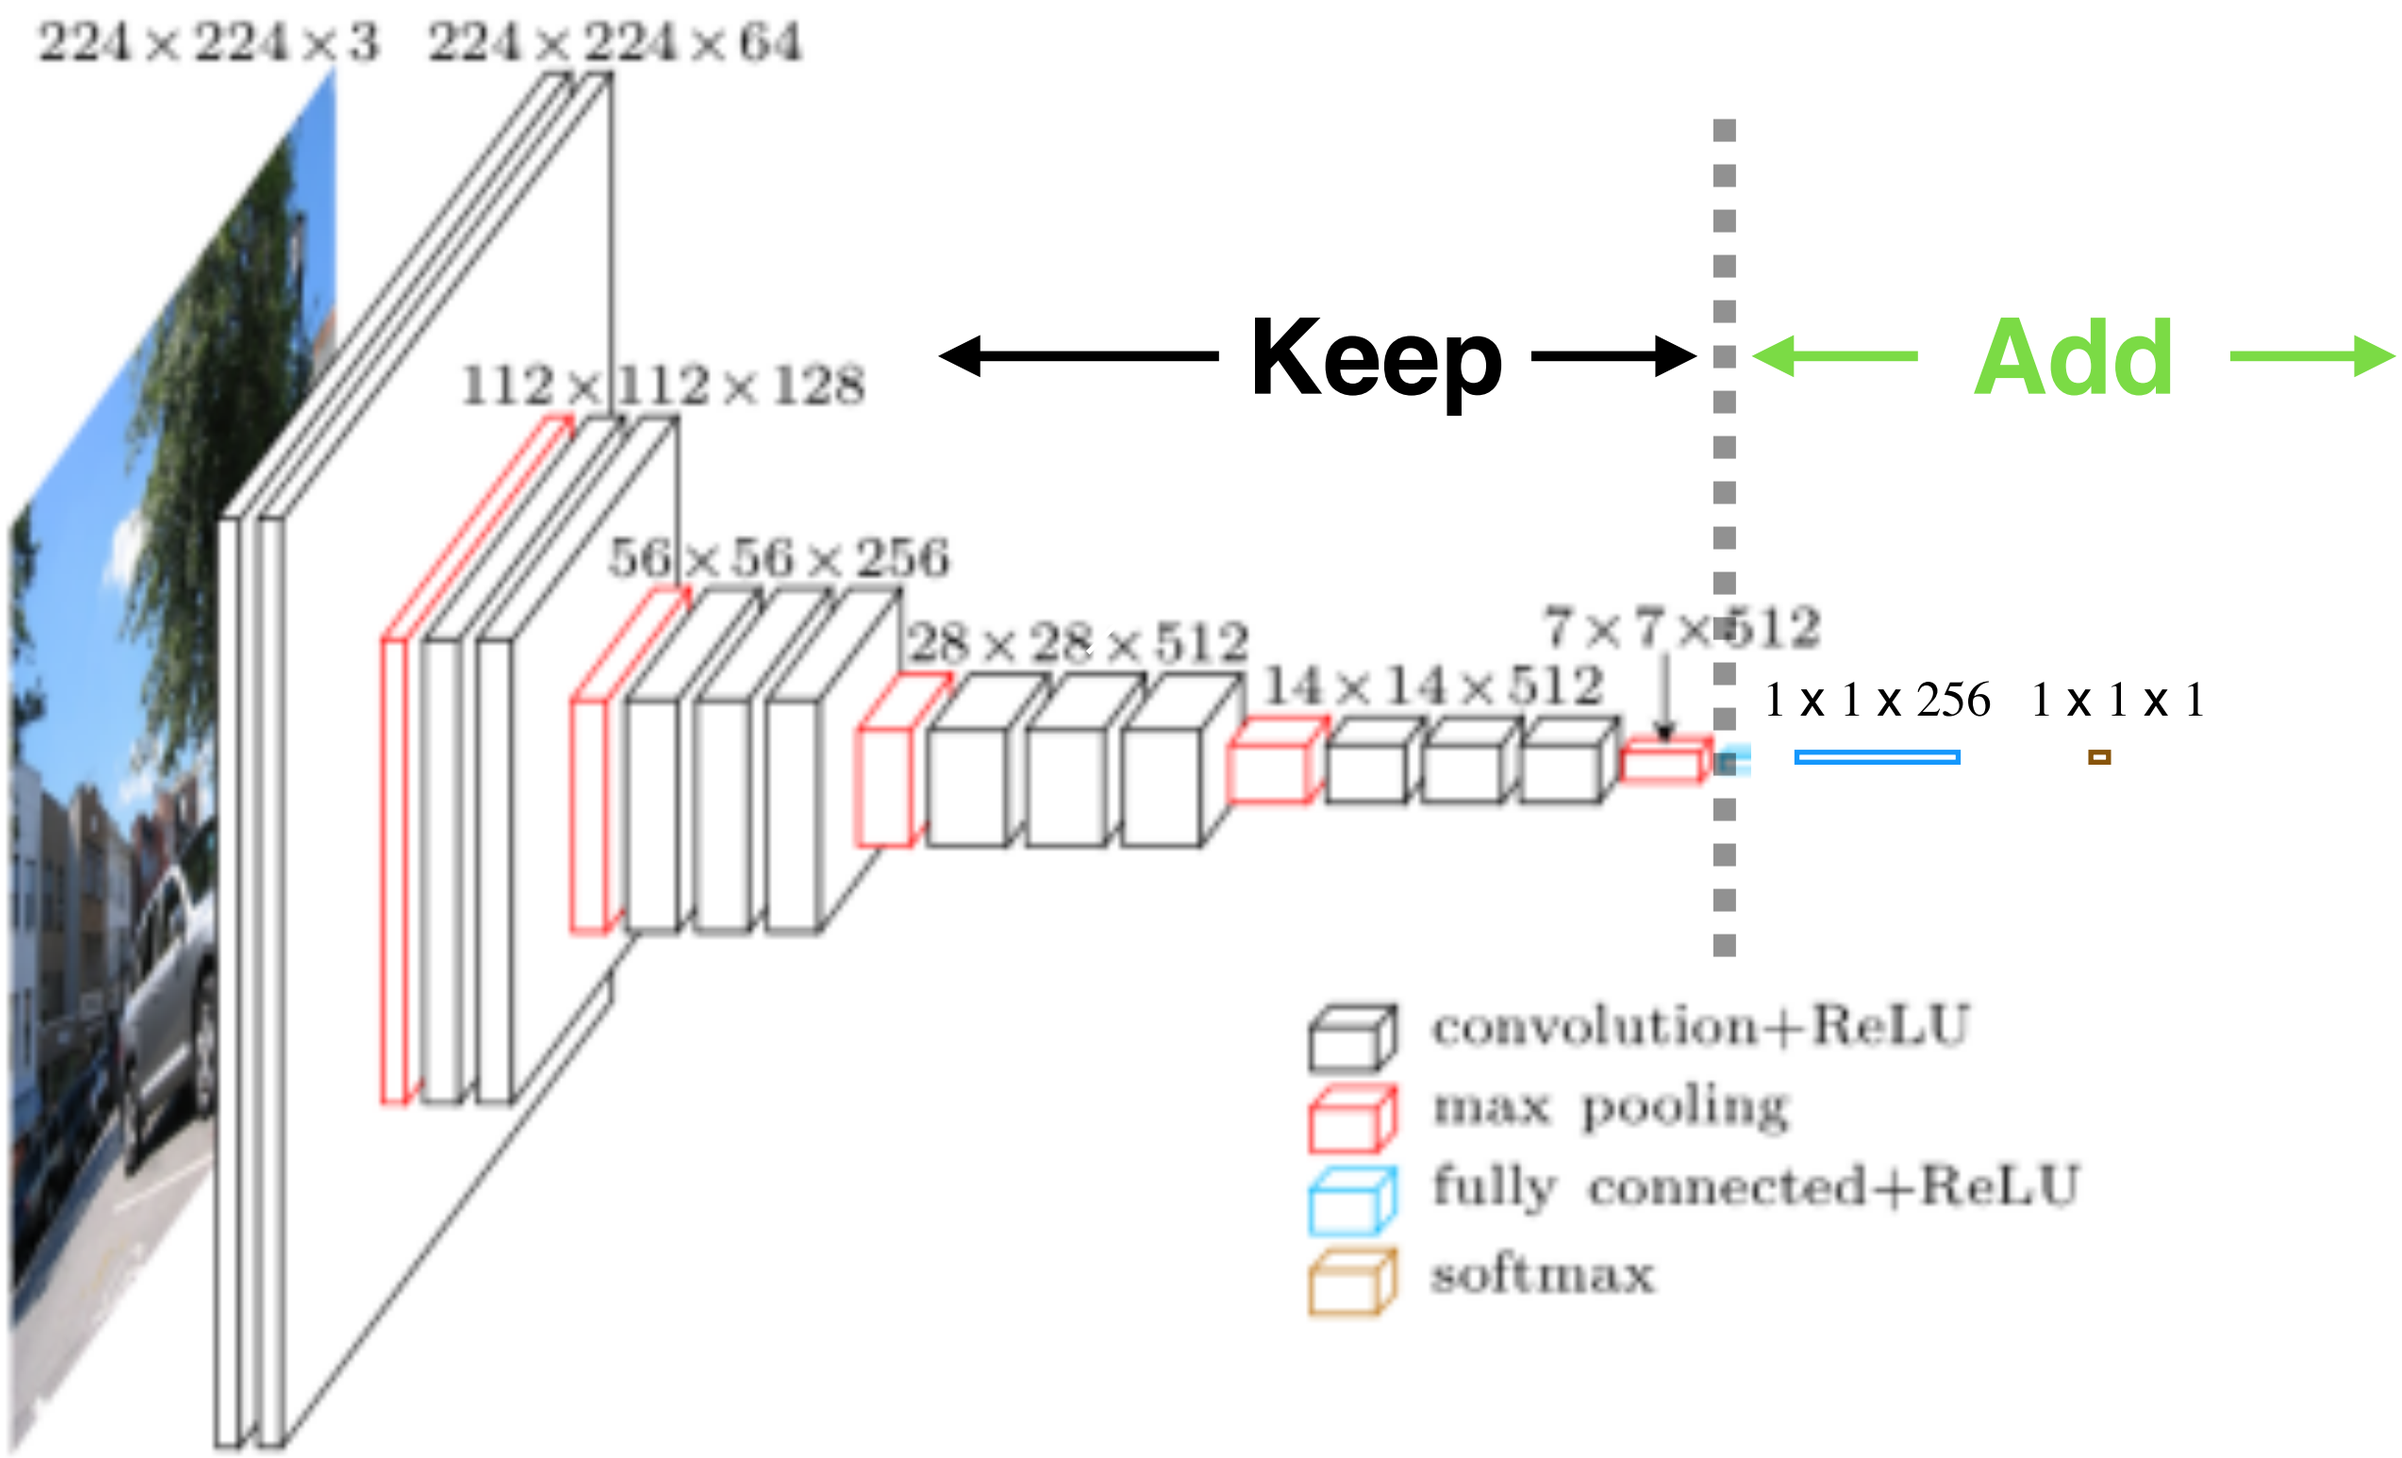

First, train the top layers

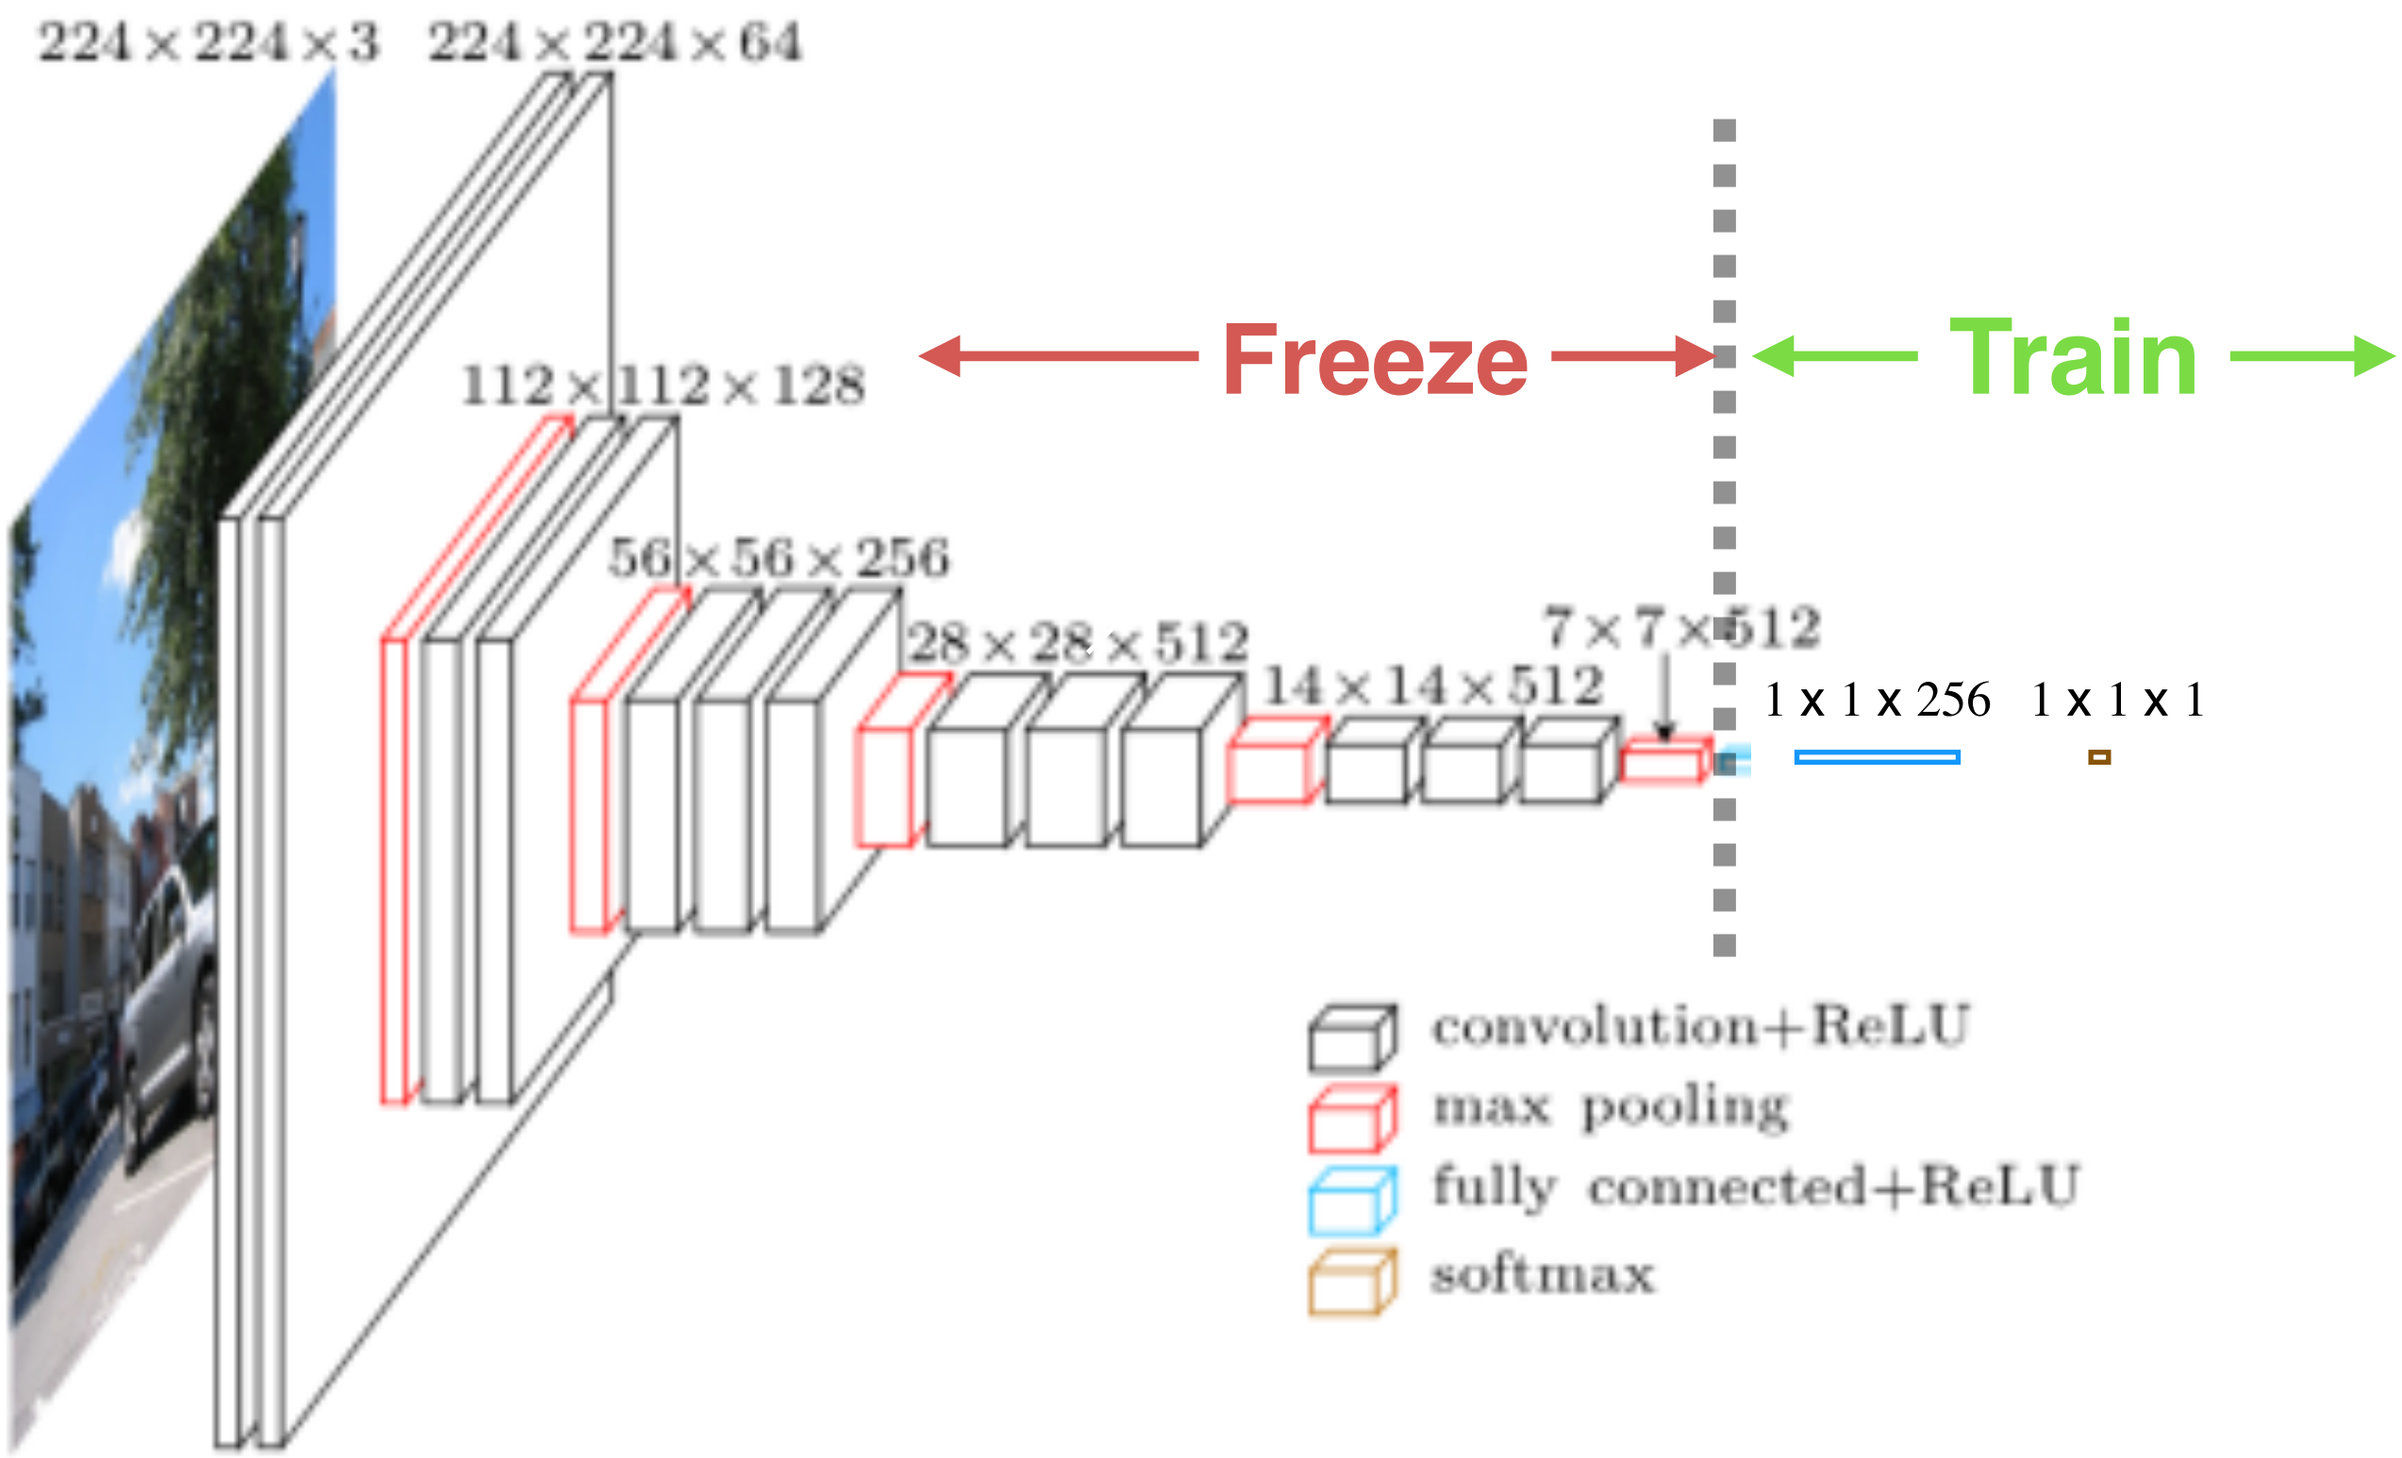

Then, train the whole network

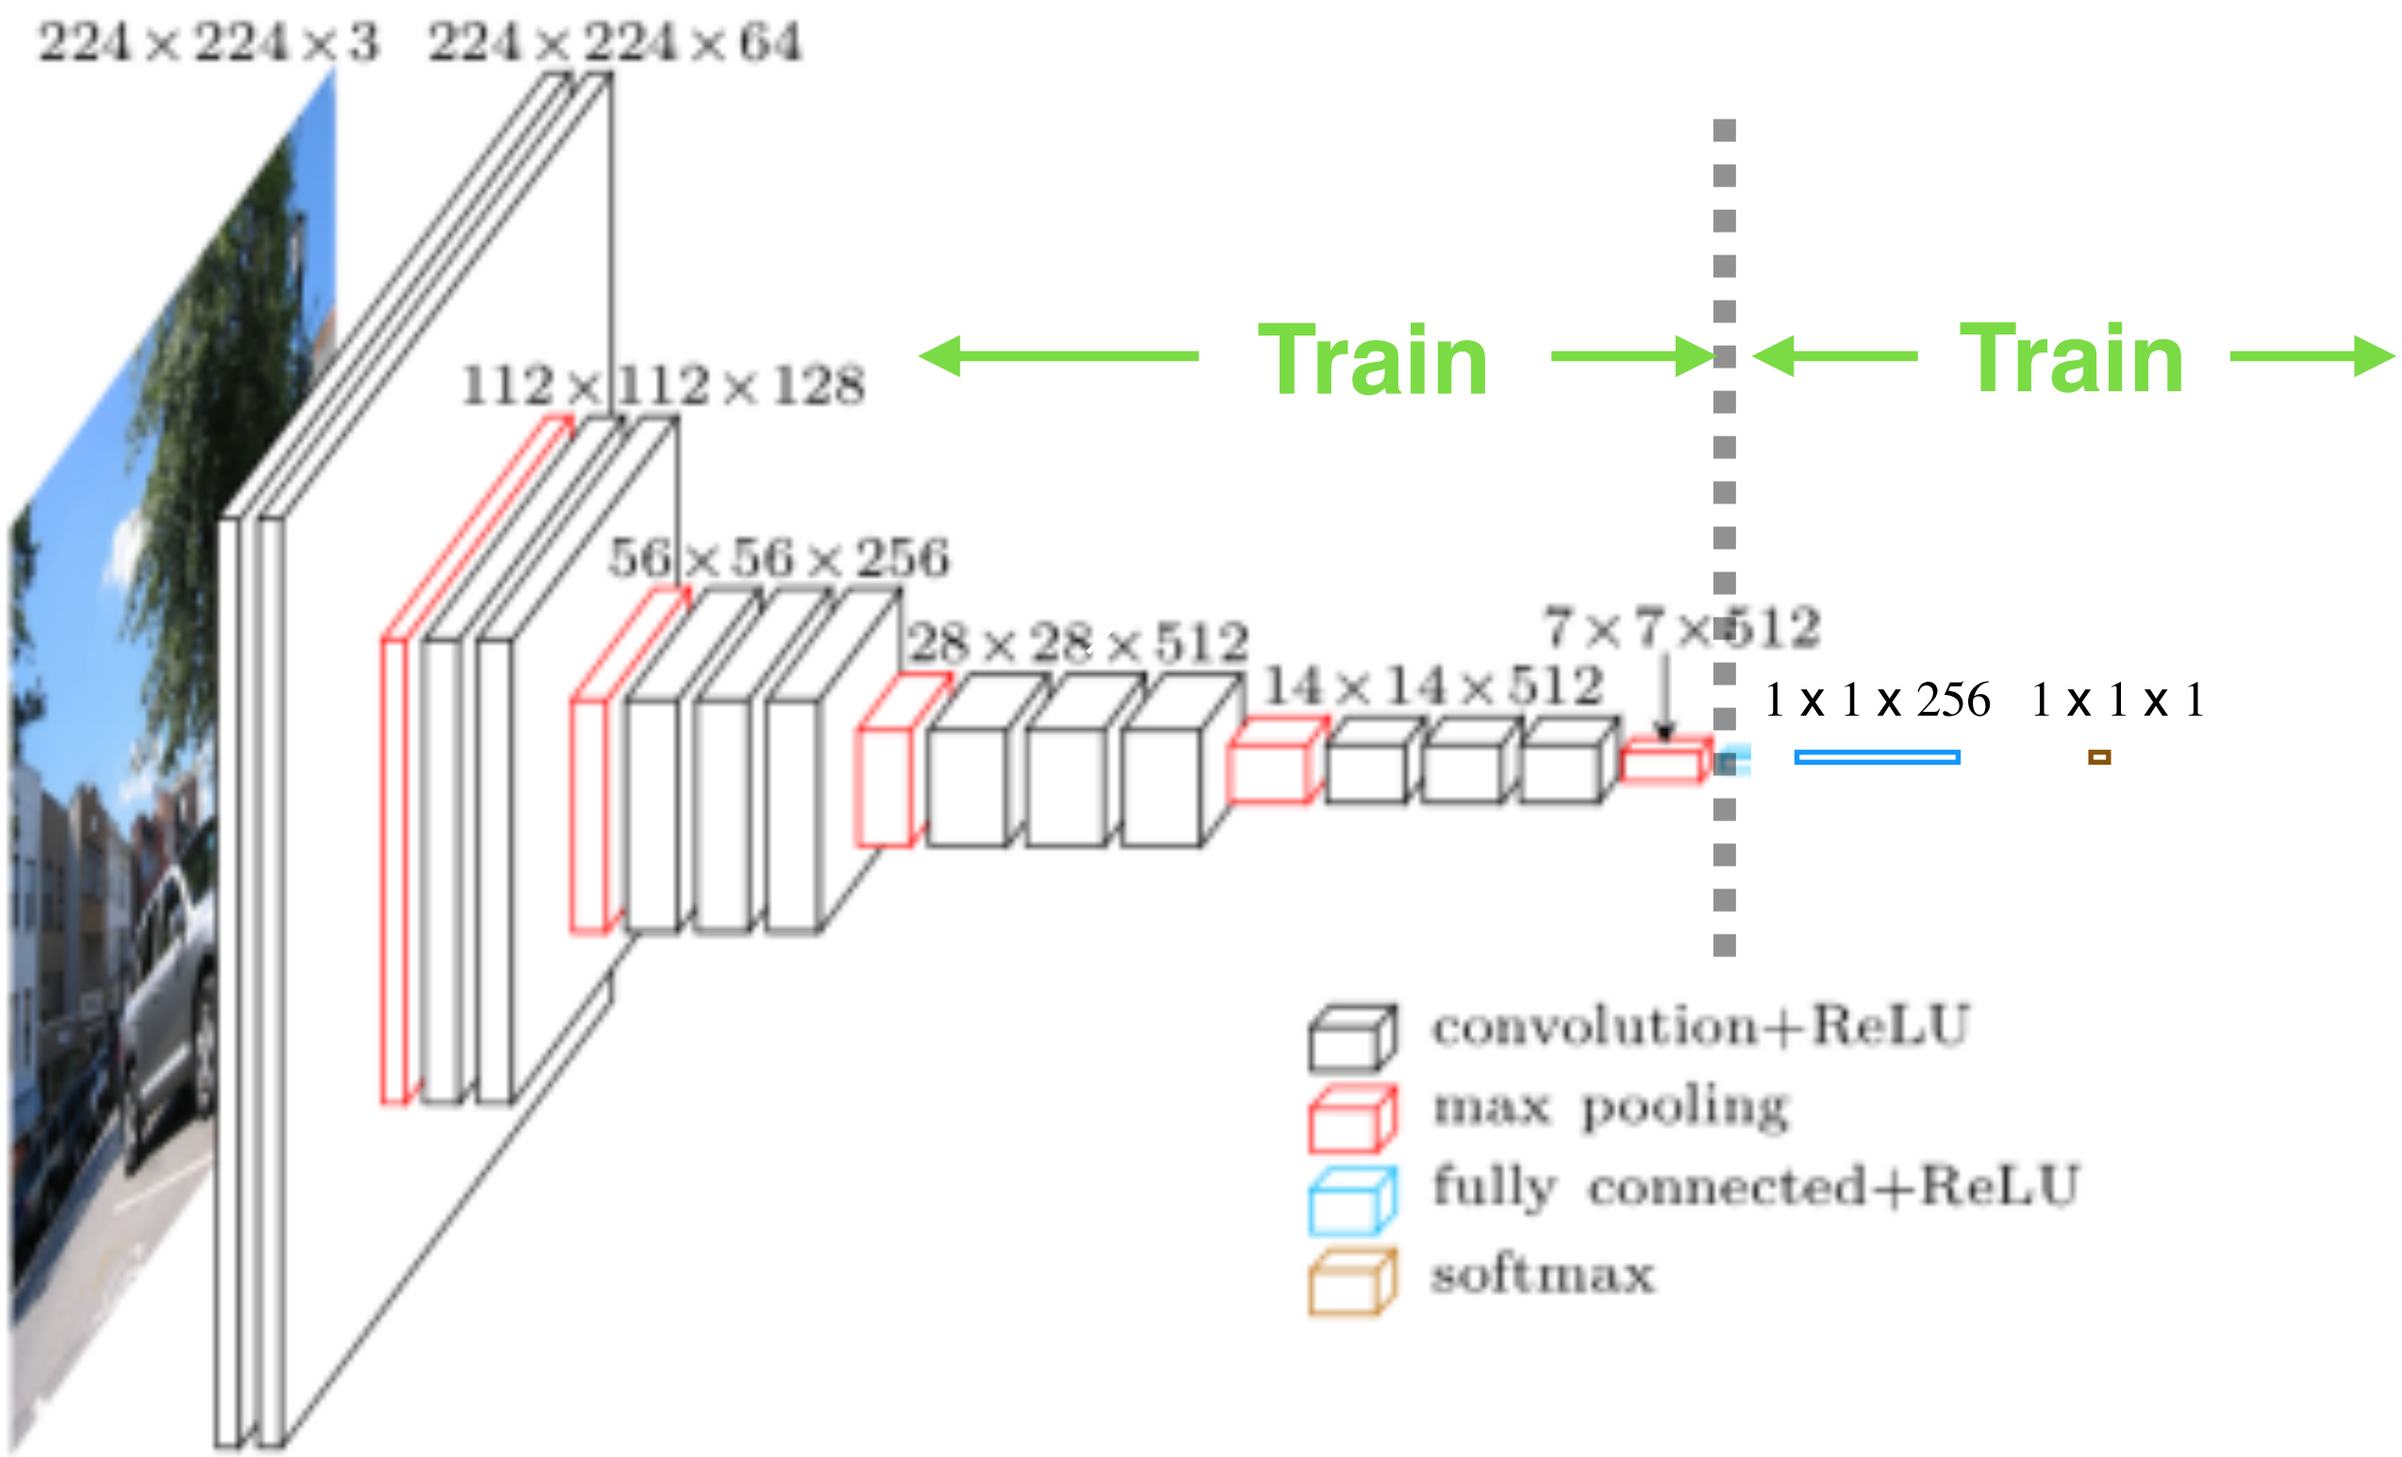

Open imagenet_vgg16_fine-tuning.html in browser