Convolutional Neural Networks

Milan Straka

March 25, 2019

EUROPEAN UNION
European Structural and Investment Fund
Operational Programme Research,
Development and Education

Charles University in Prague
Faculty of Mathematics and Physics
Institute of Formal and Applied Linguistics

Convergence

The training process might or might not converge. Even if it does, it might converge slowly or quickly.

We have already discussed two factors influencing it on the previous lecture:

- saturating non-linearities,
- parameter initialization strategies.

Another prominent method for dealing with slow or diverging training is *gradient clipping*.

Convergence – Gradient Clipping

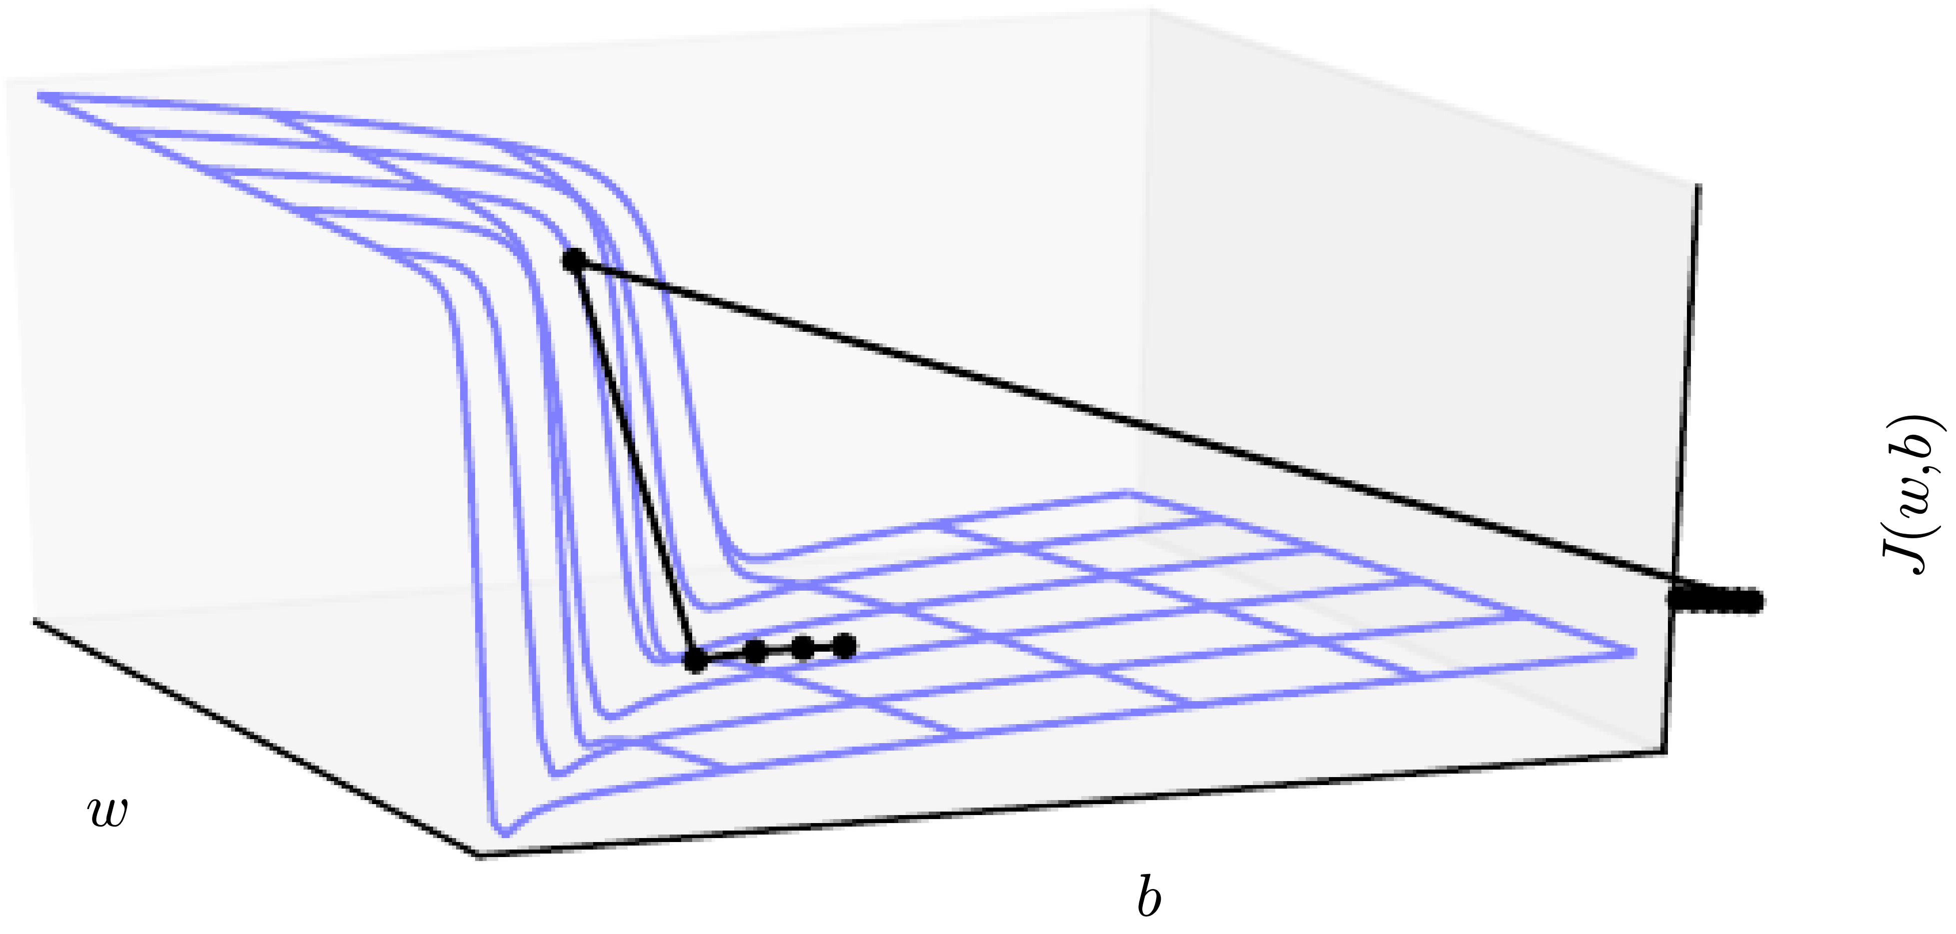

Figure 8.3, page 289 of Deep Learning Book, http://deeplearningbook.org

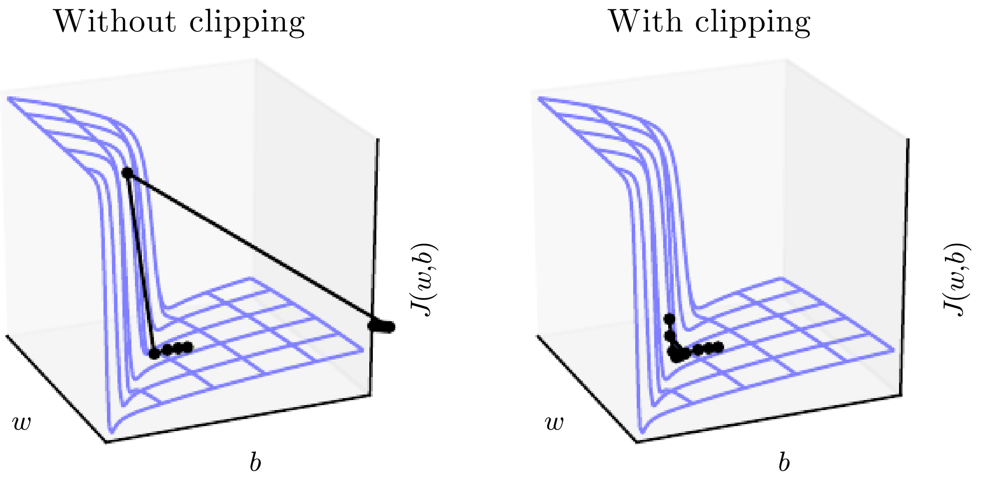

Figure 10.17, page 414 of Deep Learning Book, http://deeplearningbook.org

Using a given maximum norm, we may *clip* the gradient.

$$g \leftarrow \begin{cases} g & \text{if } ||g|| \leq c \\ c\frac{g}{||g||} & \text{if } ||g|| > c \end{cases}$$

The clipping can be per weight (`clipvalue` of `tf.keras.optimizers.Optimizer`), per variable or for the gradient as a whole (`clipnorm` of `tf.keras.optimizers.Optimizer`).

Going Deeper

Convolutional Networks

Consider data with some structure (temporal data, speech, images, …).

Unlike densely connected layers, we might want:

- local interactions;
- parameter sharing (equal response everywhere);
- shift invariance.

Convolution Operation

For a functions x and w, *convolution* $x * w$ is defined as

$$(x * w)(t) = \int x(a)w(t-a)\,\mathrm{d}a.$$

For vectors, we have

$$(\boldsymbol{x} * \boldsymbol{w})_t = \sum\nolimits_i x_i w_{t-i}.$$

Convolution operation can be generalized to two dimensions by

$$(\boldsymbol{I} * \boldsymbol{K})_{i,j} = \sum\nolimits_{m,n} \boldsymbol{I}_{m,n}\boldsymbol{K}_{i-m,j-n}.$$

Closely related is *cross-corellation*, where K is flipped:

$$S_{i,j} = \sum\nolimits_{m,n} \boldsymbol{I}_{i+m,j+n}\boldsymbol{K}_{m,n}.$$

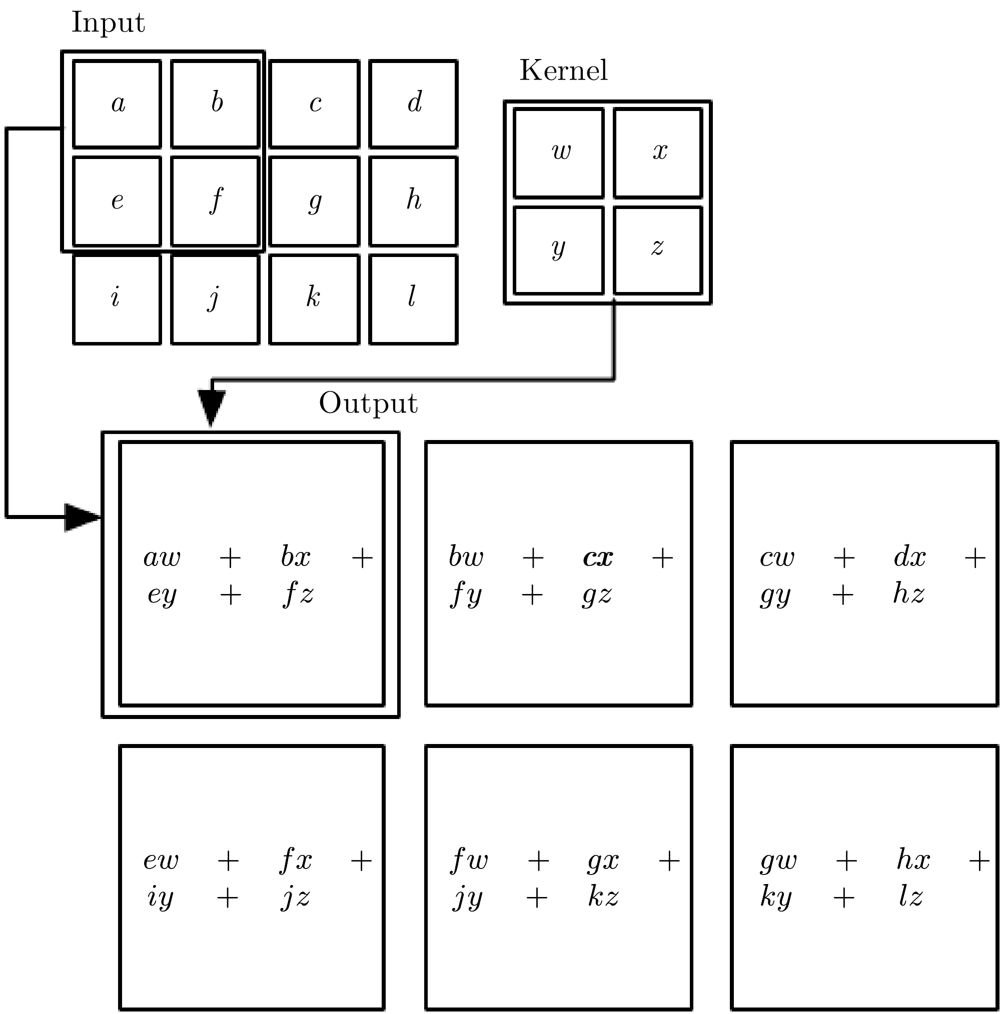

Figure 9.1, page 334 of Deep Learning Book, http://deeplearningbook.org

Convolutional Networks

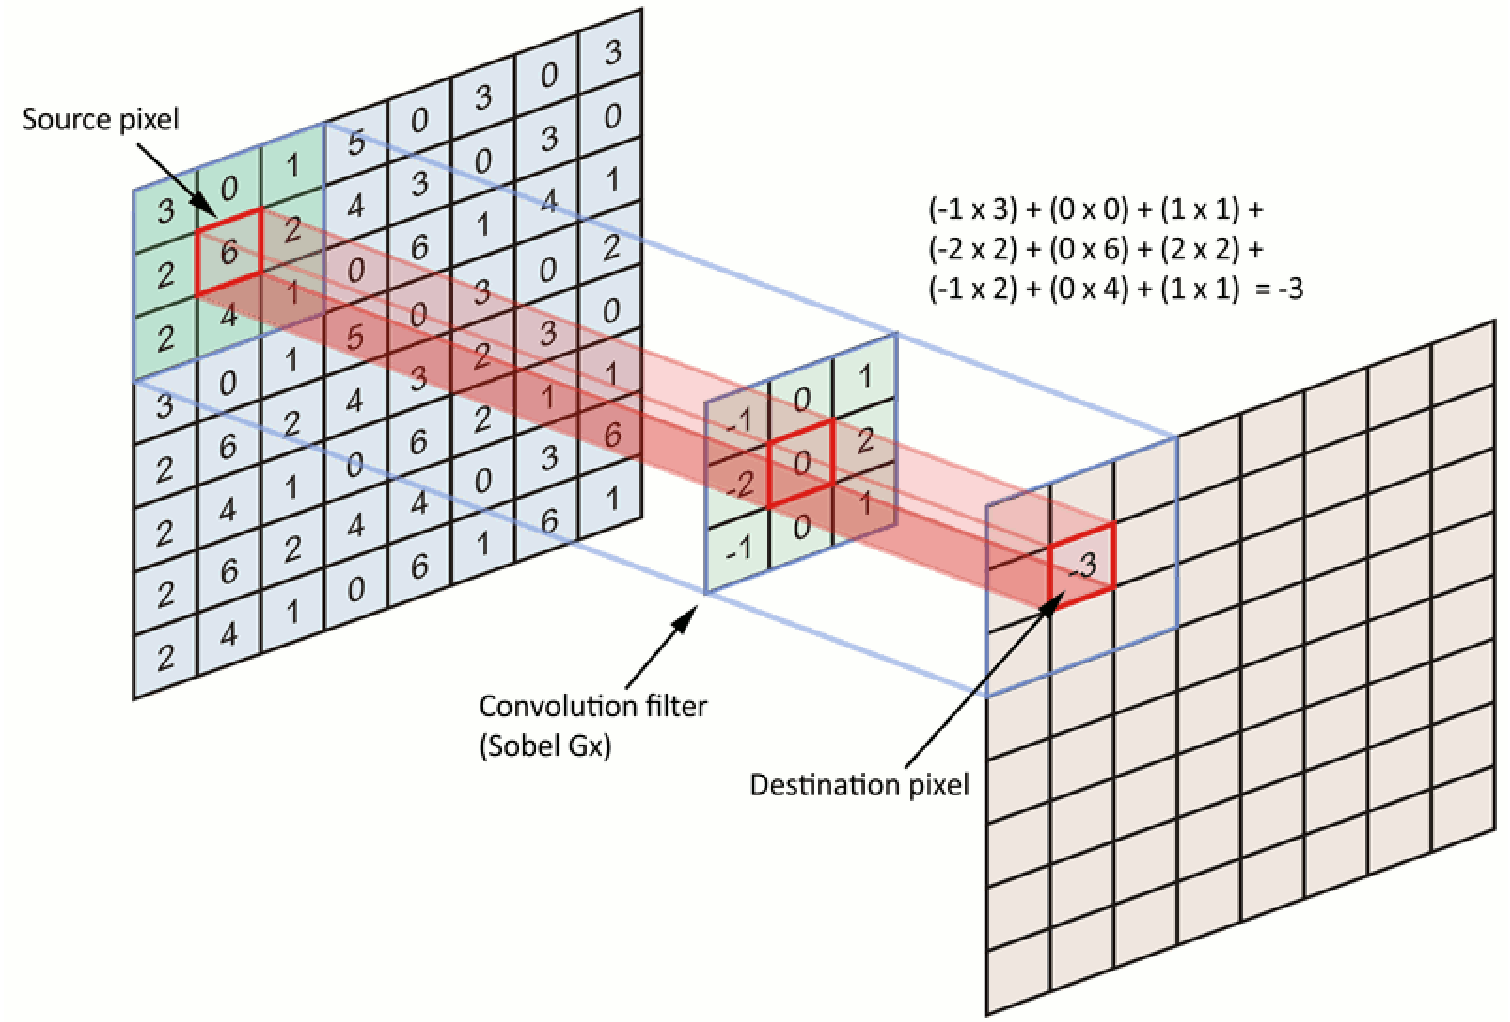

Image from https://i.stack.imgur.com/YDusp.png.

The K is usually called a *kernel* or a *filter*, and we generally apply several of them at the same time.

Consider an input image with C channels. The convolution layer with F filters of width W, height H and stride S produces an output with F channels, is parametrized by a kernel K of total size $W \times H \times C \times F$ and is computed as

$$(\mathsf{I} * \mathsf{K})_{i,j,k} = \sum_{m,n,o} \mathsf{I}_{i\cdot S+m, j\cdot S+n, o} \mathsf{K}_{m,n,o,k}.$$

We can consider the kernel to be composed of F independent kernels, one for every output channel.

Note that while only local interactions are performed in the image spacial dimensions (width and height), we combine input channels in a fully connected manner.

Convolution Layer

There are multiple padding schemes, most common are:

- `valid`: Only use valid pixels, which causes the result to be smaller than the input.
- `same`: Pad original image with zero pixels so that the result is exactly the size of the input.

There are two prevalent image formats (called `data_format` in TensorFlow):

- `channels_last`: The dimensions of the 4-dimensional image tensor are batch, height, width, and channels.

  The original TensorFlow format, faster on CPU.

- `channels_first`: The dimensions of the 4-dimensional image tensor are batch, channel, height, and width.

  Usual GPU format (used by CUDA and nearly all frameworks); on TensorFlow, not all CPU kernels are available with this layout.

TensorFlow has been implementing an approach that will convert data format to `channels_first` automatically depending on the backend.

Pooling

Pooling is an operation similar to convolution, but we perform a fixed operation instead of multiplying by a kernel.

- Max pooling (minor translation invariance)
- Average pooling

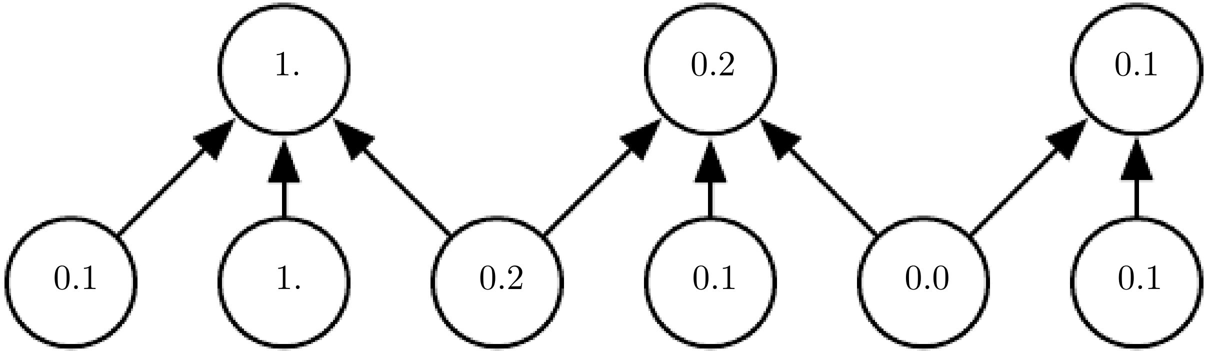

Figure 9.10, page 344 of Deep Learning Book, http://deeplearningbook.org

High-level CNN Architecture

We repeatedly use the following block:

1. Convolution operation
2. Non-linear activation (usually ReLU)
3. Pooling

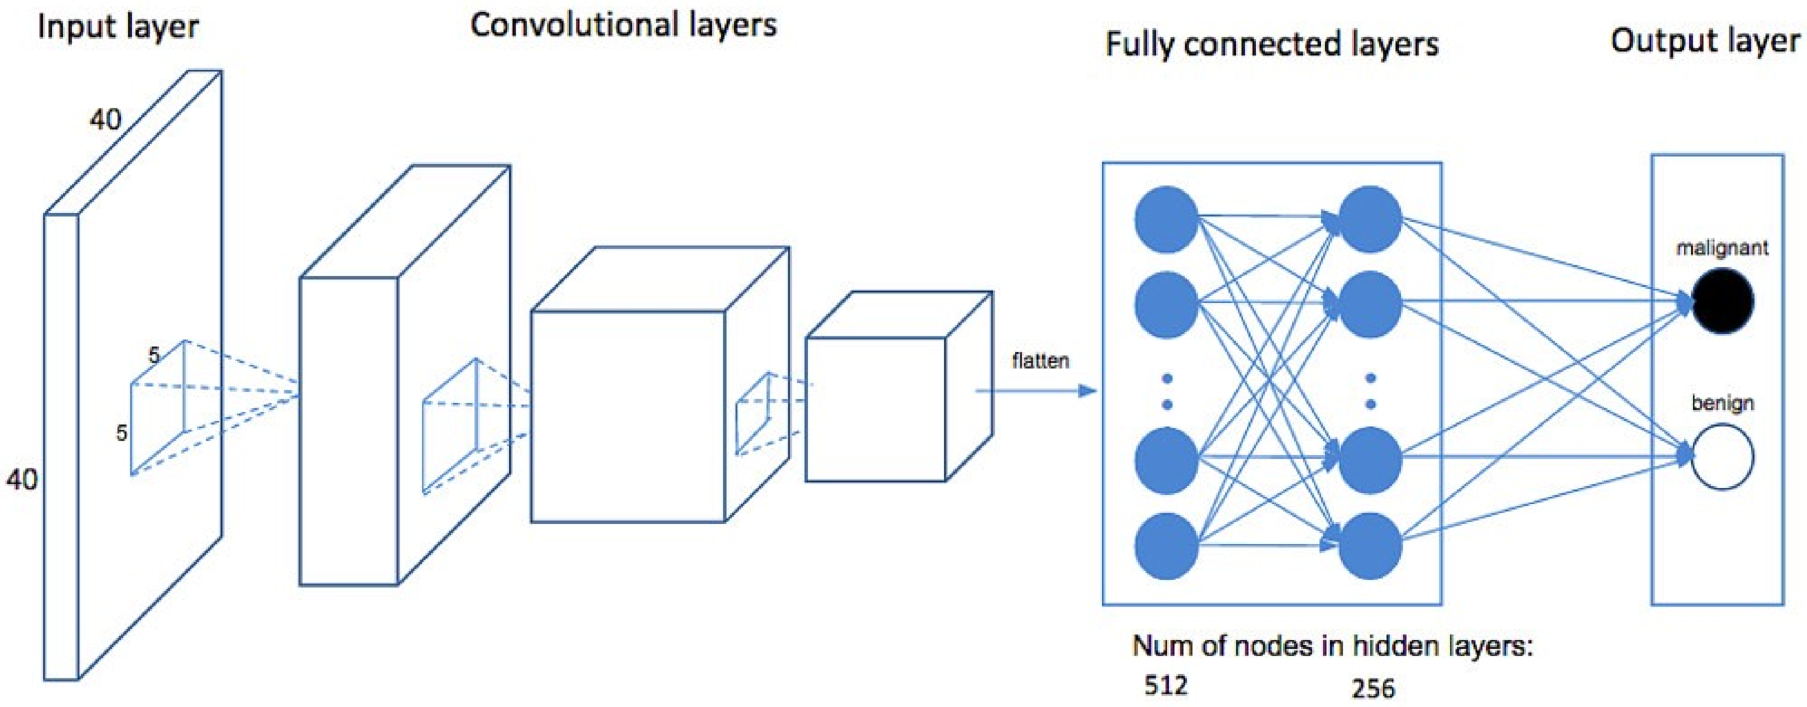

*Image from https://cdn-images-1.medium.com/max/1200/0*QyXSpqpm1wc_Dt6V. .*

AlexNet – 2012 (16.4% error)

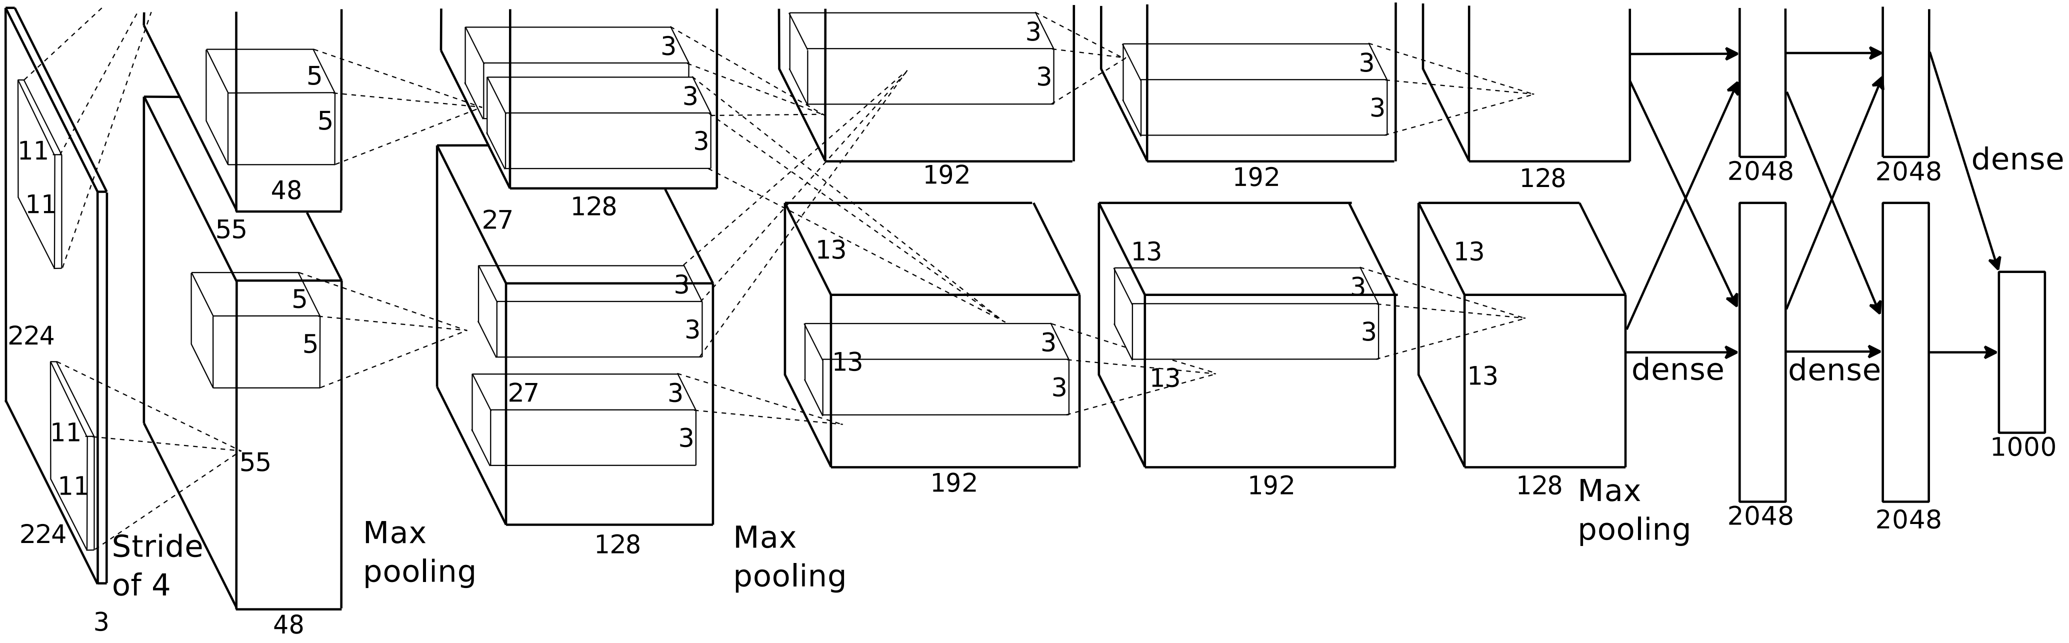

Figure 2: An illustration of the architecture of our CNN, explicitly showing the delineation of responsibilities between the two GPUs. One GPU runs the layer-parts at the top of the figure while the other runs the layer-parts at the bottom. The GPUs communicate only at certain layers. The network's input is 150,528-dimensional, and the number of neurons in the network's remaining layers is given by 253, 40–186,624–64,896–64,896–43,264–4096–4096–1000.

Figure 2 of paper "ImageNet Classification with Deep Convolutional Neural Networks" by Alex Krizhevsky et al.

Training details:

- 2 GPUs for 5-6 days
- SGD with batch size 128, momentum 0.9, weight decay 0.0005
- initial learning rate 0.01, manually divided by 10 when validation error rate stopped improving
- ReLU non-linearities
- dropout with rate 0.5 on fully-connected layers
- data augmentation using translations and horizontal reflections (choosing random 224×224 patches from 256×256 images)
  - during inference, 10 patches are used (four corner patches and a center patch, as well as their reflections)

AlexNet – ReLU vs tanh

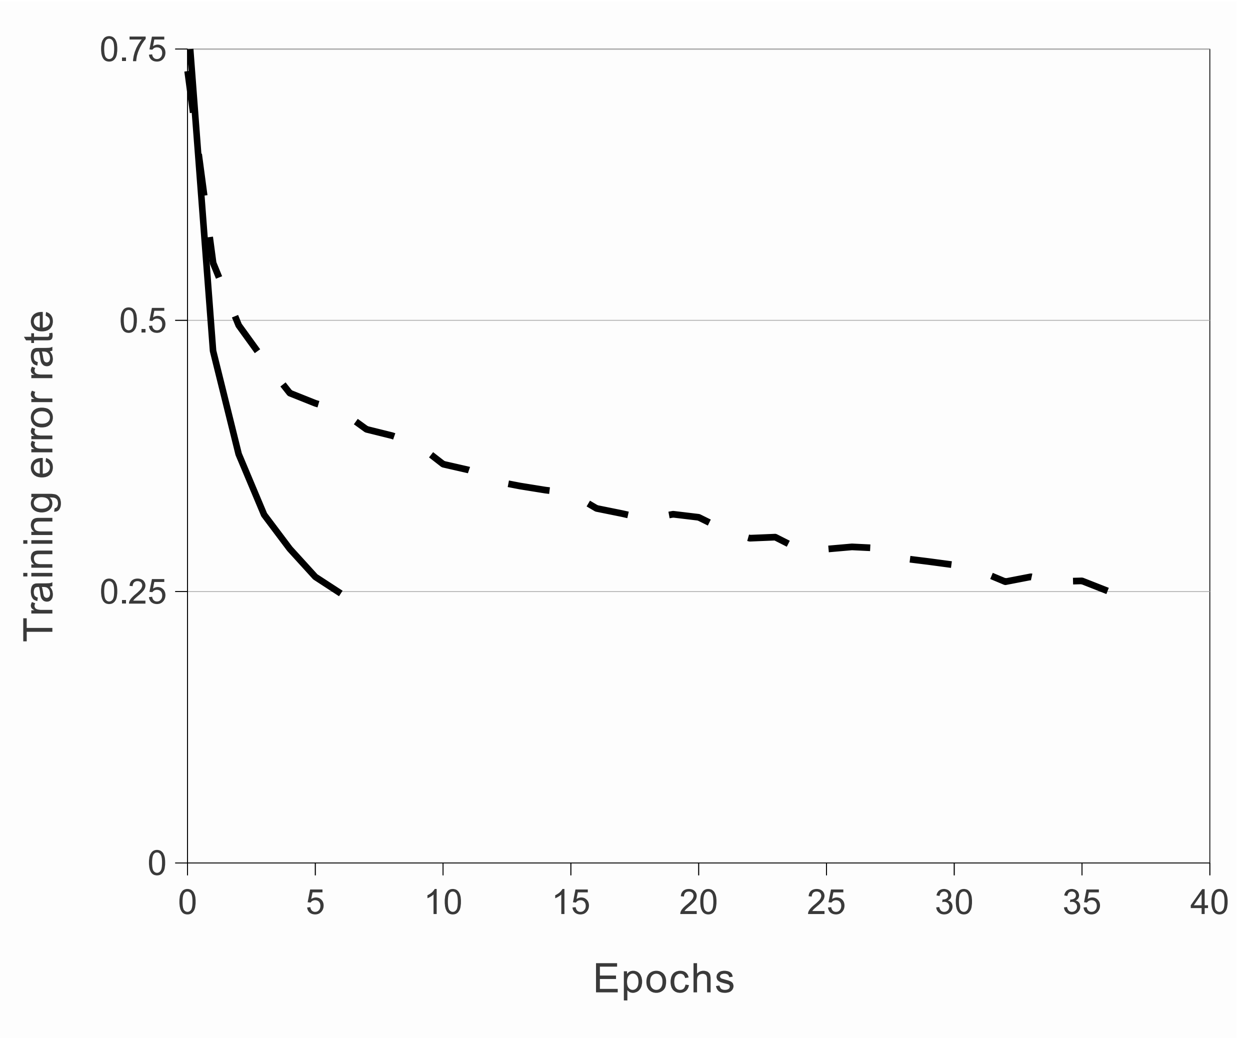

Figure 1 of paper "ImageNet Classification with Deep Convolutional Neural Networks" by Alex Krizhevsky et al.

AlexNet built on already existing CNN architectures, mostly on LeNet, which achieved 0.8% test error on MNIST.

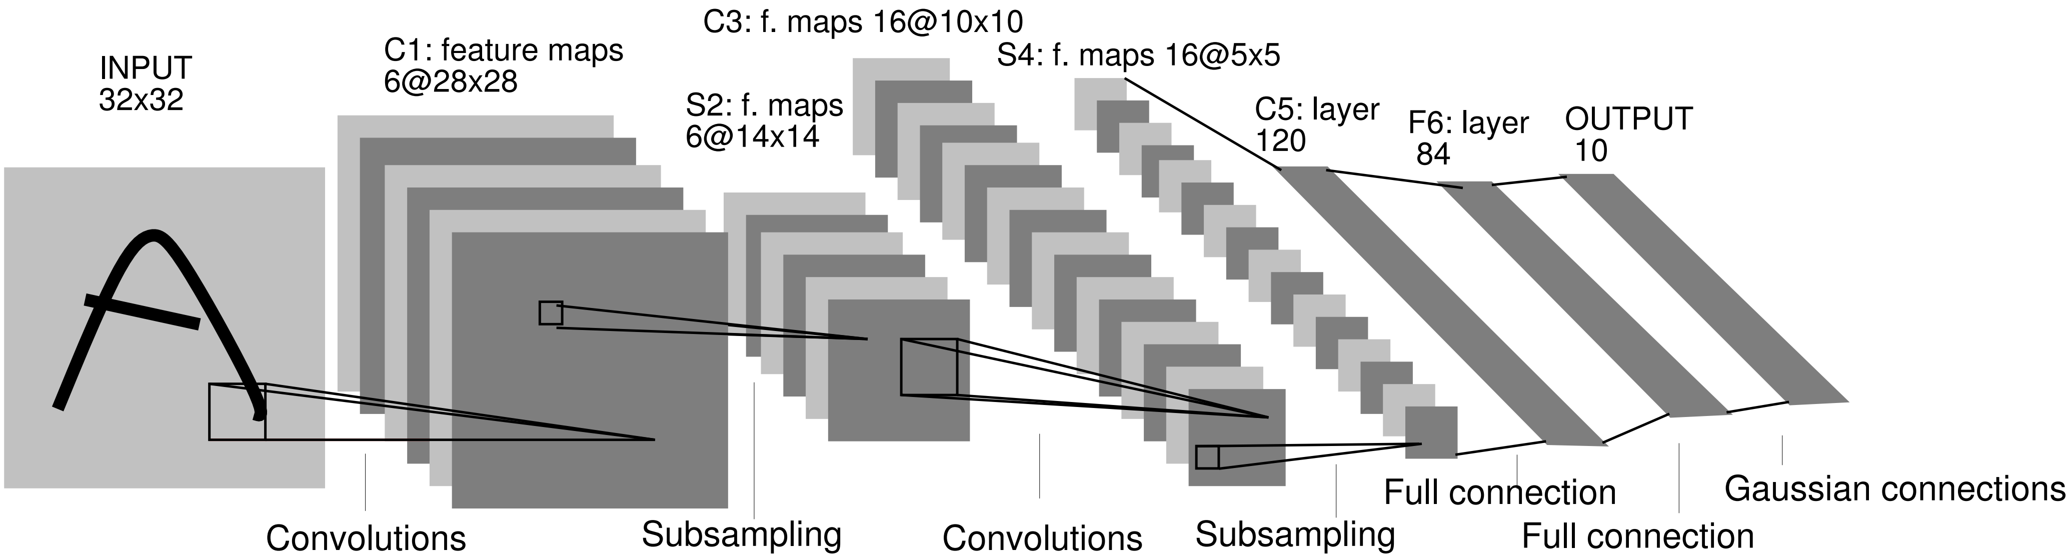

Figure 2 of paper "Gradient-Based Learning Applied to Document Recognition", http://yann.lecun.com/exdb/publis/pdf/lecun-01a.pdf.

Similarities in V1 and CNNs

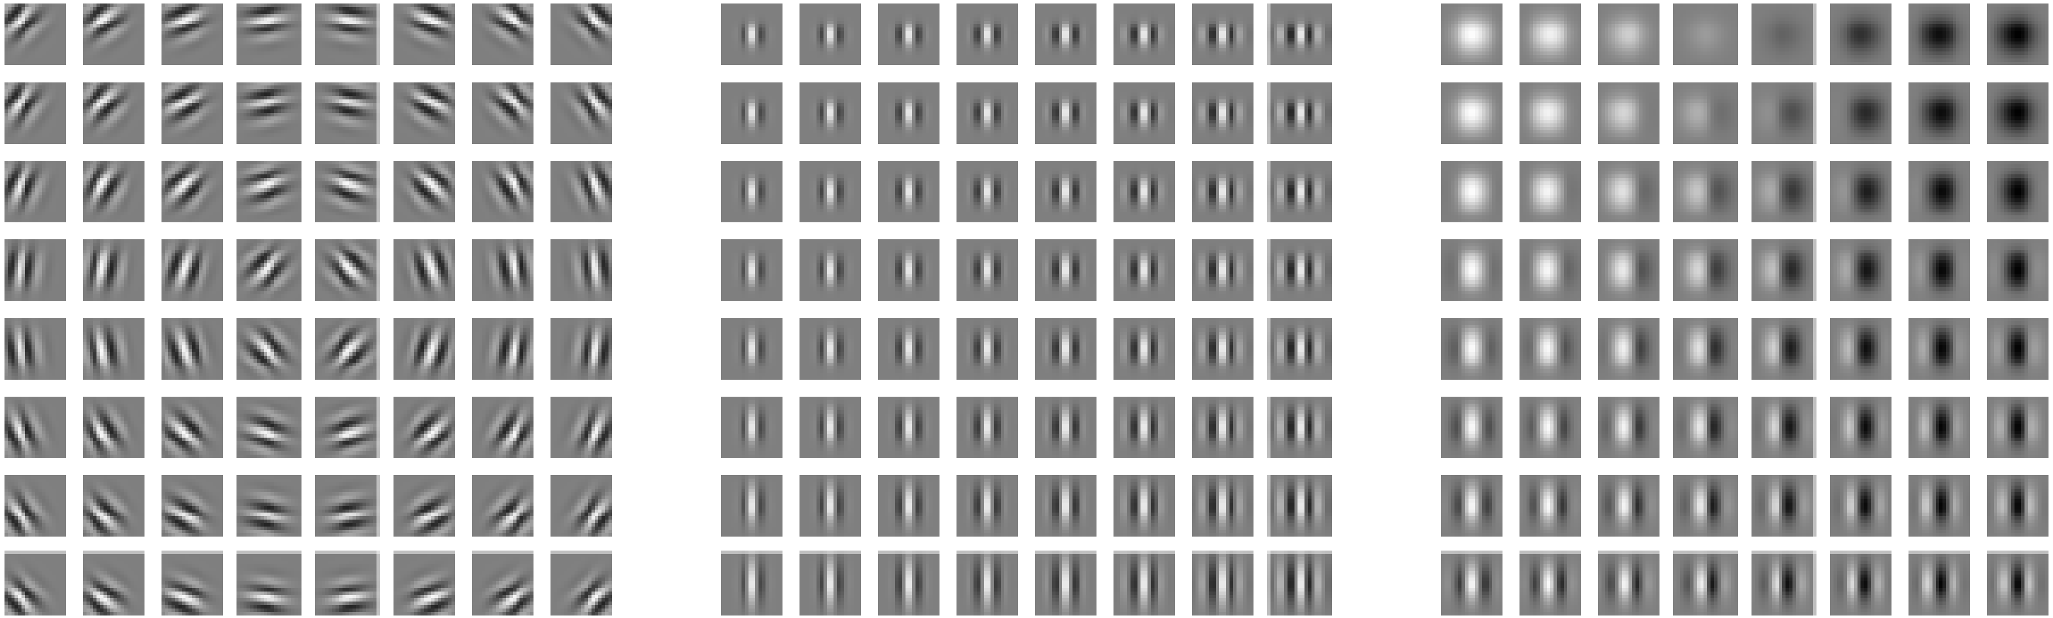

Figure 9.18, page 370 of Deep Learning Book, http://deeplearningbook.org

The primary visual cortex recognizes Gabor functions.

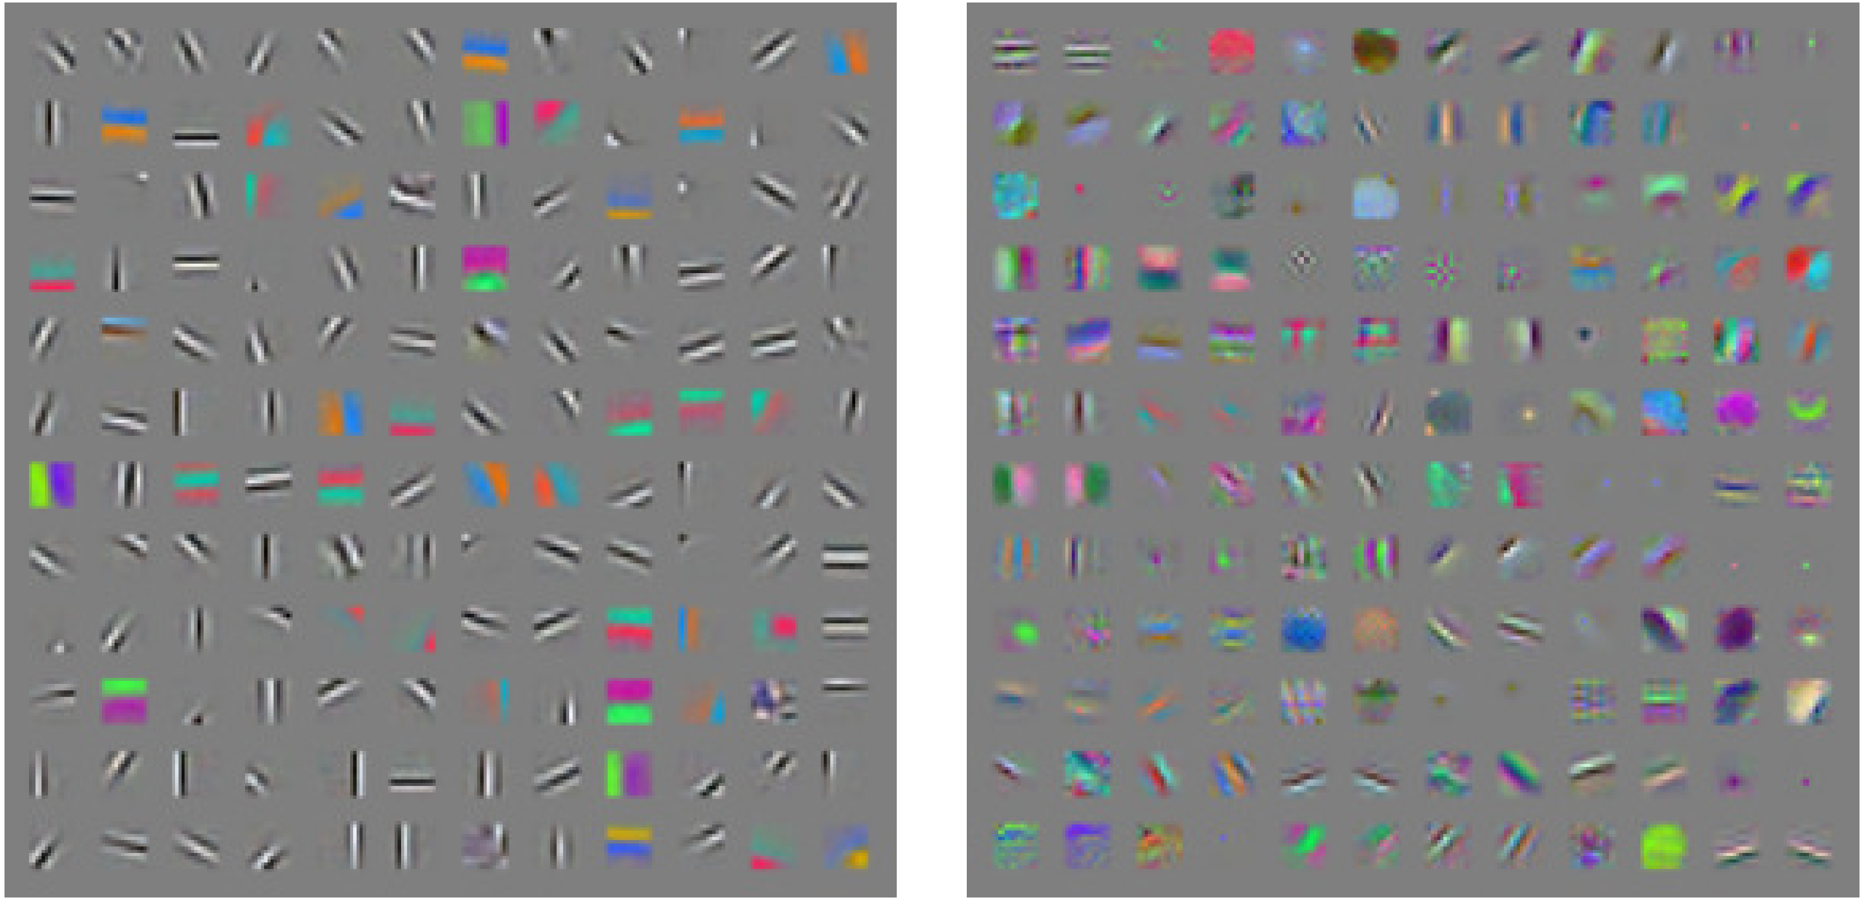

Figure 9.19, page 371 of Deep Learning Book, http://deeplearningbook.org

Similar functions are recognized in the first layer of a CNN.

(a) Ground truth

(b) SRResNet [18], **Trained**

(c) Bicubic, **Not trained**

(d) Deep prior, **Not trained**

Figure 1 of paper "Deep Prior", https://arxiv.org/abs/1712.05016

(a) Original image (b) Corrupted image (c) Shepard networks [26] (d) Deep Image Prior

(e) Original image (f) Corrupted image (g) [24], PSNR = 28.1 (h) Deep Img. Prior, PSNR = 30.9

Figure 7 of paper "Deep Prior", https://arxiv.org/abs/1712.05016

Figure 5: **Inpainting diversity.** Left: original image (black pixels indicate holes). The remaining four images show results obtained using deep prior corresponding to different input vector z.

Figure 5 of supplementary materials of paper "Deep Prior", https://arxiv.org/abs/1712.05016

CNNs as Regularizers – Deep Prior

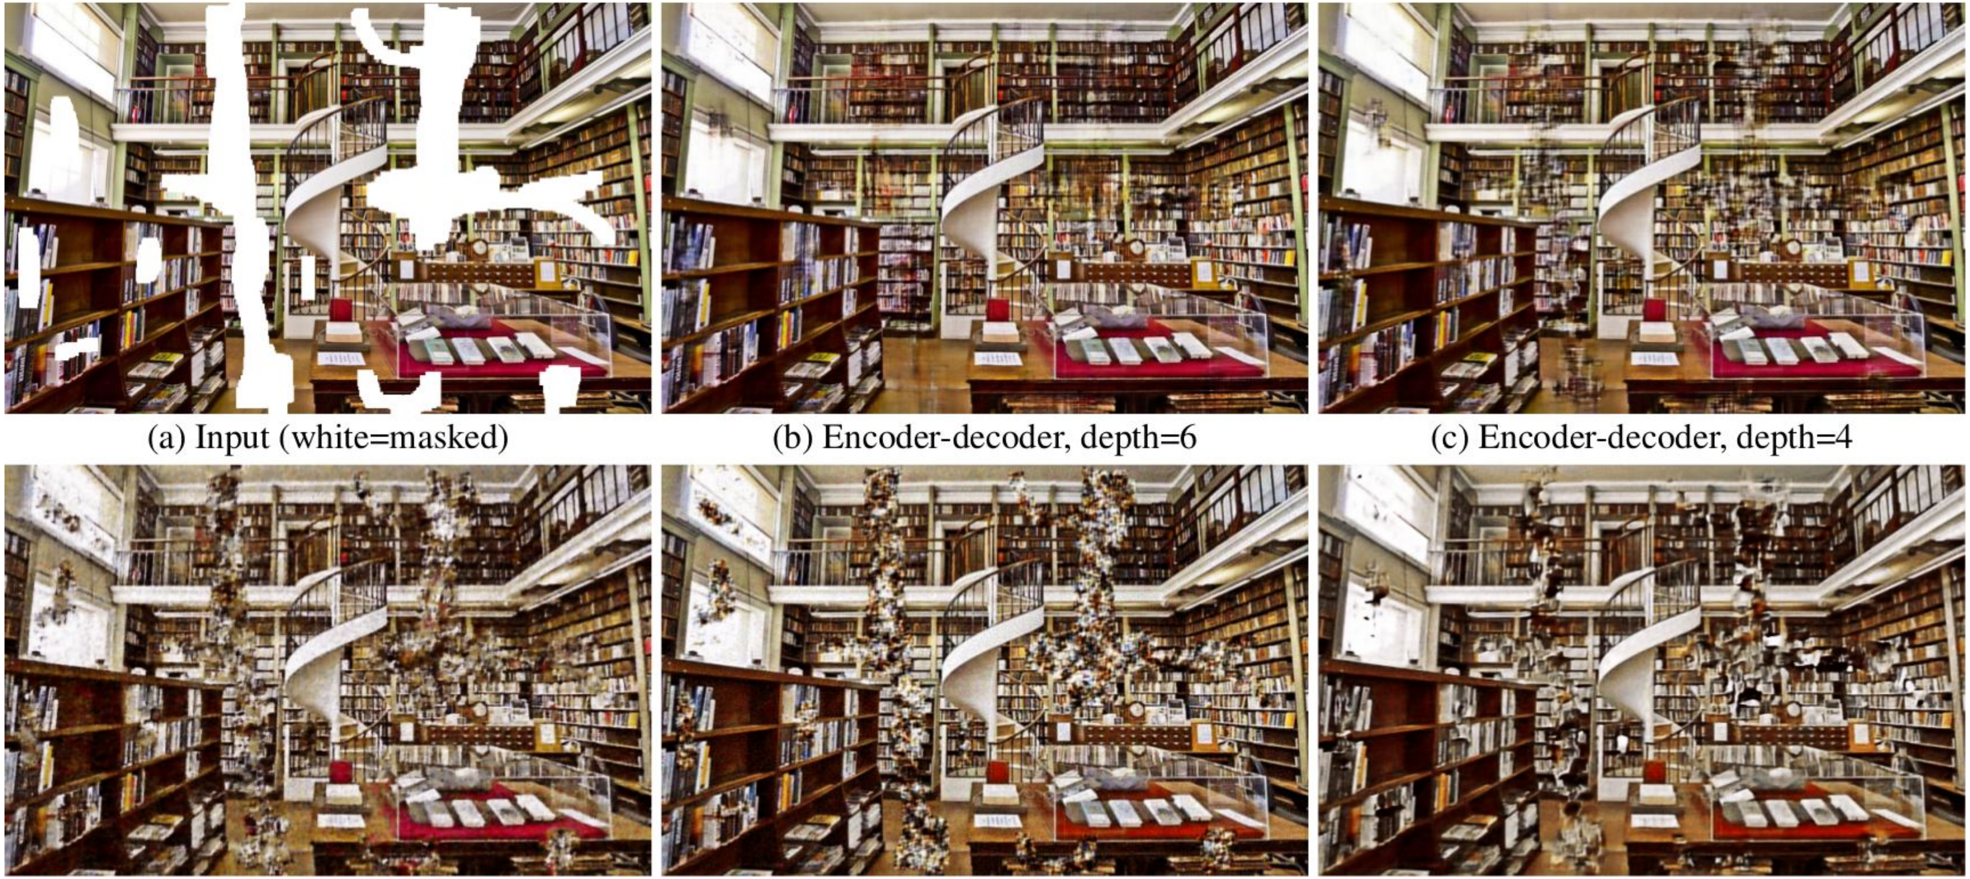

(a) Input (white=masked) (b) Encoder-decoder, depth=6 (c) Encoder-decoder, depth=4

(d) Encoder-decoder, depth=2 (e) ResNet, depth=8 (f) U-net, depth=5

Figure 8 of paper "Deep Prior", https://arxiv.org/abs/1712.05016

Deep Prior paper website with supplementary material

VGG – 2014 (6.8% error)

ConvNet Configuration					
A	A-LRN	B	C	D	E
11 weight layers	11 weight layers	13 weight layers	16 weight layers	16 weight layers	19 weight layers
input (224 × 224 RGB image)					
conv3-64	conv3-64 **LRN**	conv3-64 **conv3-64**	conv3-64 conv3-64	conv3-64 conv3-64	conv3-64 conv3-64
maxpool					
conv3-128	conv3-128	conv3-128 **conv3-128**	conv3-128 conv3-128	conv3-128 conv3-128	conv3-128 conv3-128
maxpool					
conv3-256 conv3-256	conv3-256 conv3-256	conv3-256 conv3-256	conv3-256 conv3-256 **conv1-256**	conv3-256 conv3-256 **conv3-256**	conv3-256 conv3-256 conv3-256 **conv3-256**
maxpool					
conv3-512 conv3-512	conv3-512 conv3-512	conv3-512 conv3-512	conv3-512 conv3-512 **conv1-512**	conv3-512 conv3-512 **conv3-512**	conv3-512 conv3-512 conv3-512 **conv3-512**
maxpool					
conv3-512 conv3-512	conv3-512 conv3-512	conv3-512 conv3-512	conv3-512 conv3-512 **conv1-512**	conv3-512 conv3-512 **conv3-512**	conv3-512 conv3-512 conv3-512 **conv3-512**
maxpool					
FC-4096					
FC-4096					
FC-1000					
soft-max					

Figure 1 of paper "Very Deep Convolutional Networks For Large-Scale Image Recognition", https://arxiv.org/abs/1409.1556.

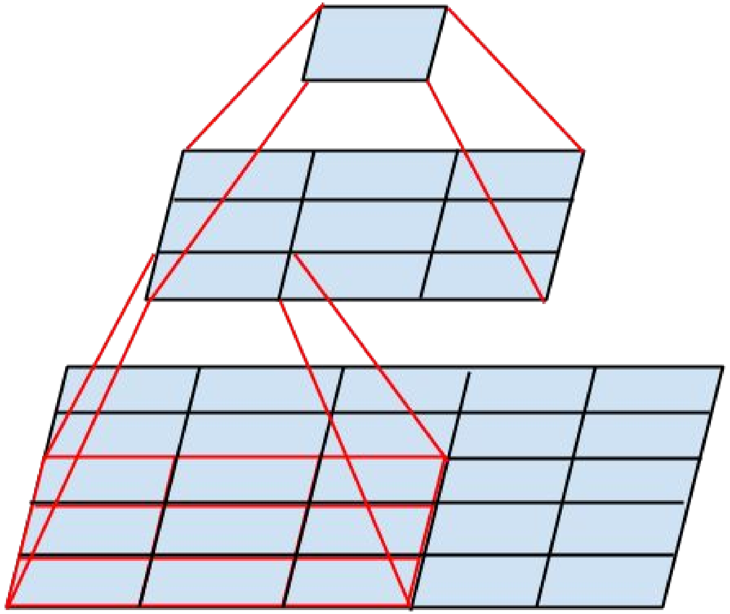

Figure 1 of paper "Rethinking the Inception Architecture for Computer Vision", https://arxiv.org/abs/1512.00567.

Table 2: **Number of parameters** (in millions).

Network	A,A-LRN	B	C	D	E
Number of parameters	133	133	134	138	144

Figure 2 of paper "Very Deep Convolutional Networks For Large-Scale Image Recognition", https://arxiv.org/abs/1409.1556.

VGG – 2014 (6.8% error)

Method	top-1 val. error (%)	top-5 val. error (%)	top-5 test error (%)
VGG (2 nets, multi-crop & dense eval.)	**23.7**	**6.8**	**6.8**
VGG (1 net, multi-crop & dense eval.)	24.4	7.1	7.0
VGG (ILSVRC submission, 7 nets, dense eval.)	24.7	7.5	7.3
GoogLeNet (Szegedy et al., 2014) (1 net)	-	7.9	
GoogLeNet (Szegedy et al., 2014) (7 nets)	-	**6.7**	
MSRA (He et al., 2014) (11 nets)	-	-	8.1
MSRA (He et al., 2014) (1 net)	27.9	9.1	9.1
Clarifai (Russakovsky et al., 2014) (multiple nets)	-	-	11.7
Clarifai (Russakovsky et al., 2014) (1 net)	-	-	12.5
Zeiler & Fergus (Zeiler & Fergus, 2013) (6 nets)	36.0	14.7	14.8
Zeiler & Fergus (Zeiler & Fergus, 2013) (1 net)	37.5	16.0	16.1
OverFeat (Sermanet et al., 2014) (7 nets)	34.0	13.2	13.6
OverFeat (Sermanet et al., 2014) (1 net)	35.7	14.2	-
Krizhevsky et al. (Krizhevsky et al., 2012) (5 nets)	38.1	16.4	16.4
Krizhevsky et al. (Krizhevsky et al., 2012) (1 net)	40.7	18.2	-

Figure 2 of paper "Very Deep Convolutional Networks For Large-Scale Image Recognition", https://arxiv.org/abs/1409.1556.

Inception (GoogLeNet) – 2014 (6.7% error)

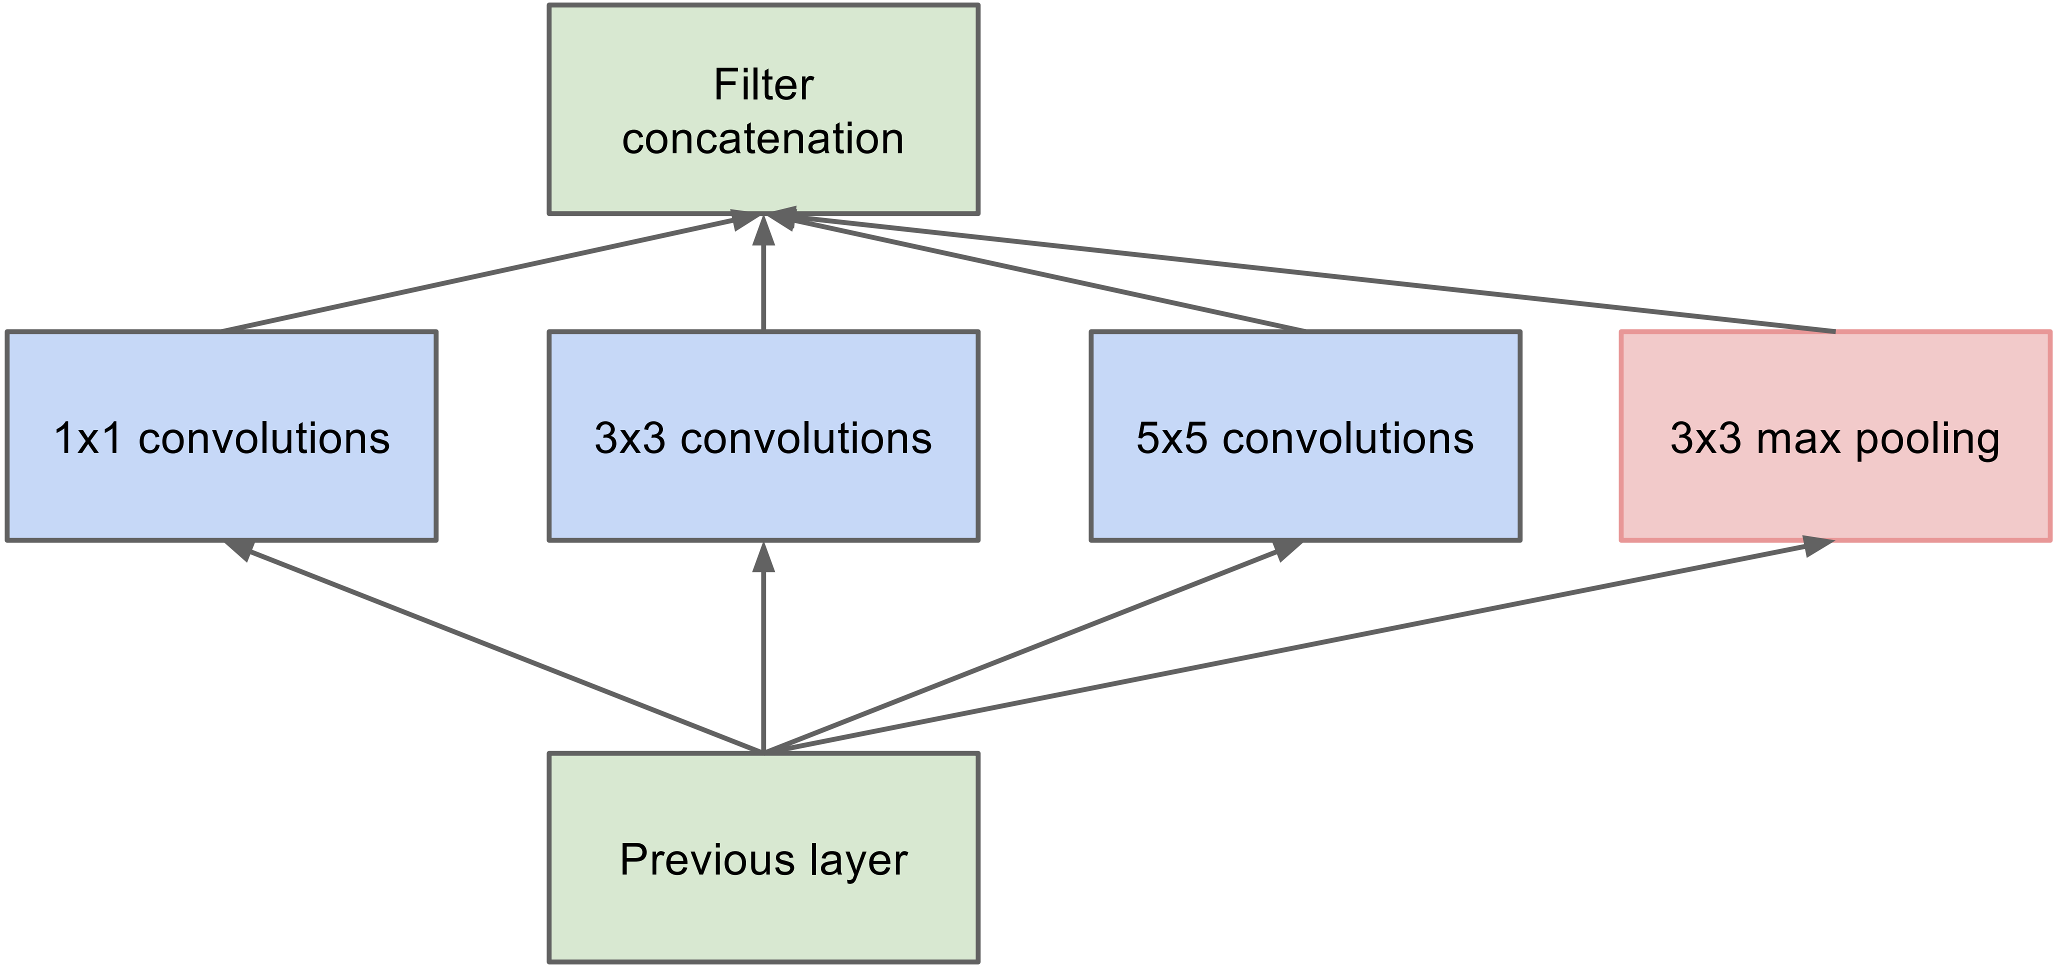

Figure 2 of paper "Going Deeper with Convolutions", https://arxiv.org/abs/1409.4842.

Inception (GoogLeNet) – 2014 (6.7% error)

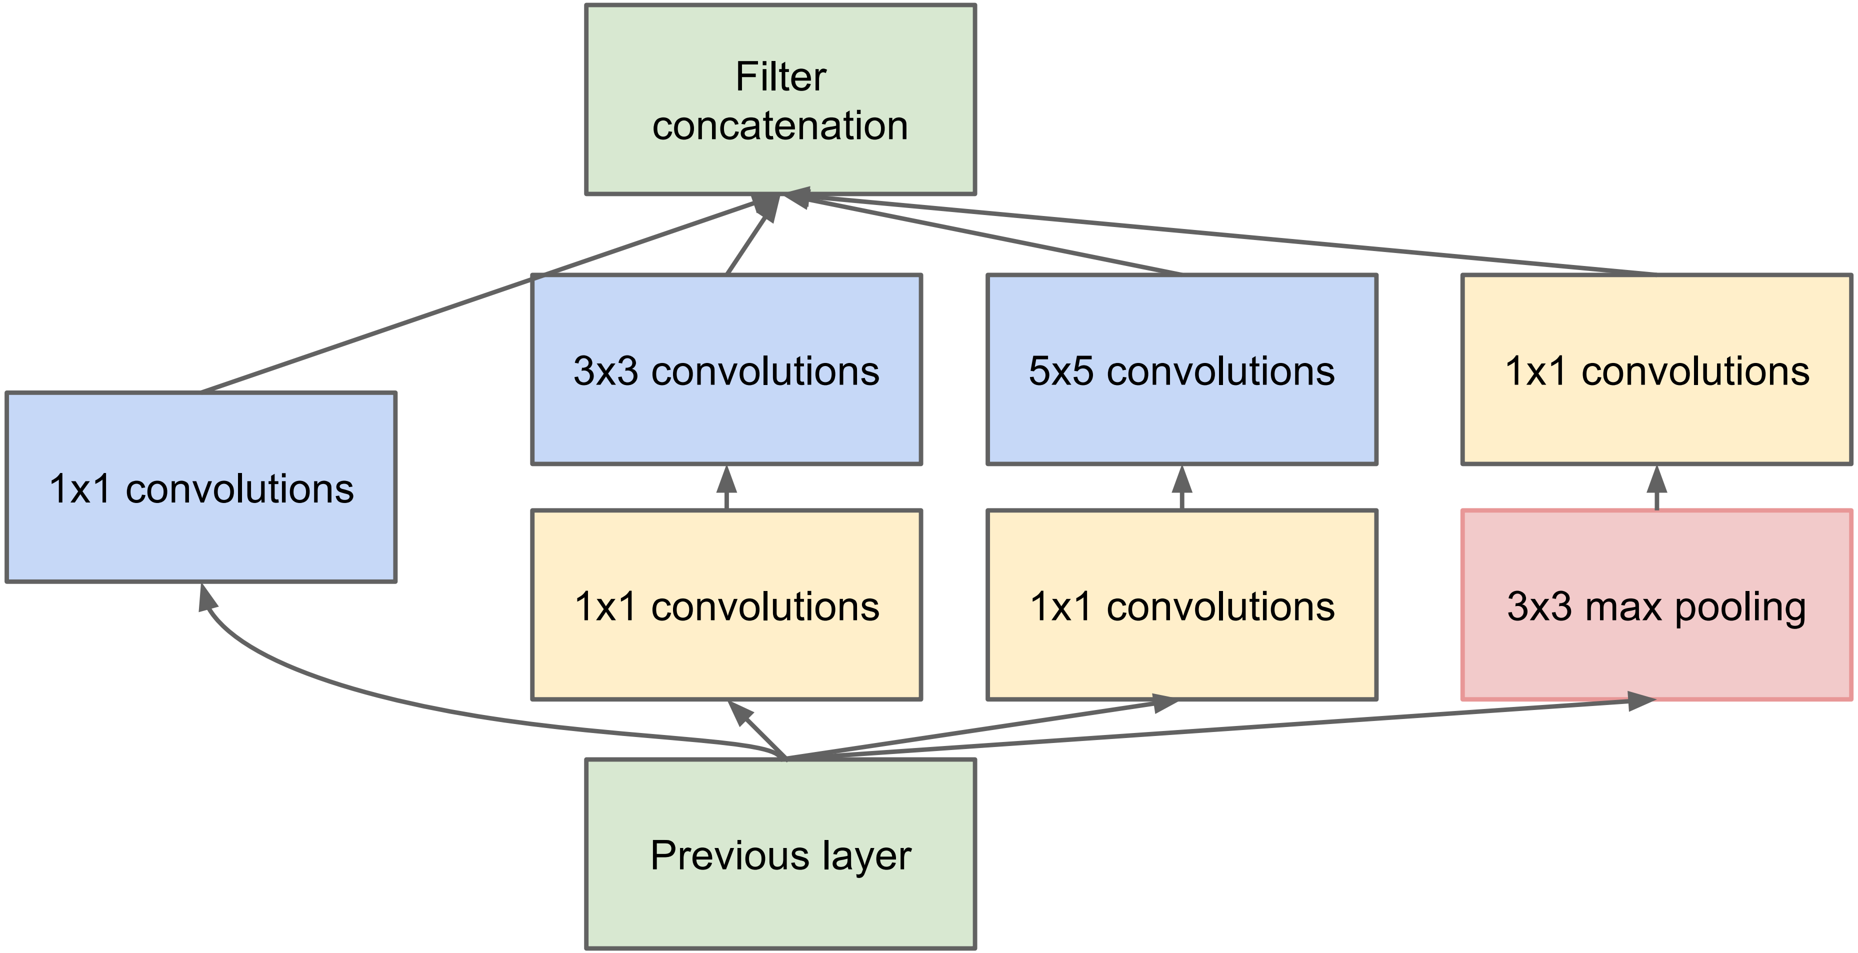

Figure 2 of paper "Going Deeper with Convolutions", https://arxiv.org/abs/1409.4842.

Inception (GoogLeNet) – 2014 (6.7% error)

type	patch size/ stride	output size	depth	#1×1	#3×3 reduce	#3×3	#5×5 reduce	#5×5	pool proj	params	ops
convolution	7×7/2	112×112×64	1							2.7K	34M
max pool	3×3/2	56×56×64	0								
convolution	3×3/1	56×56×192	2		64	192				112K	360M
max pool	3×3/2	28×28×192	0								
inception (3a)		28×28×256	2	64	96	128	16	32	32	159	128M
inception (3b)		28×28×480	2	128	128	192	32	96	64	380K	304M
max pool	3×3/2	14×14×480	0								
inception (4a)		14×14×512	2	192	96	208	16	48	64	364K	73M
inception (4b)		14×14×512	2	160	112	224	24	64	64	437K	88M
inception (4c)		14×14×512	2	128	128	256	24	64	64	463K	100M
inception (4d)		14×14×528	2	112	144	288	32	64	64	580K	119M
inception (4e)		14×14×832	2	256	160	320	32	128	128	840K	170M
max pool	3×3/2	7×7×832	0								
inception (5a)		7×7×832	2	256	160	320	32	128	128	1072K	54M
inception (5b)		7×7×1024	2	384	192	384	48	128	128	1388K	71M
avg pool	7×7/1	1×1×1024	0								
dropout (40%)		1×1×1024	0								
linear		1×1×1000	1							1000K	1M
softmax		1×1×1000	0								

Table 1 of paper "Going Deeper with Convolutions", https://arxiv.org/abs/1409.4842.

Inception (GoogLeNet) – 2014 (6.7% error)

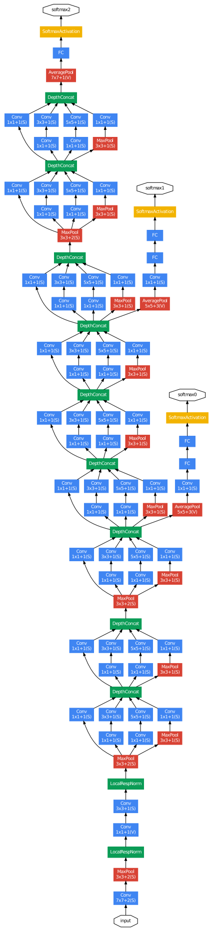

Figure 3 of paper "Going Deeper with Convolutions", https://arxiv.org/abs/1409.4842.

Also note the two auxiliary classifiers (they have weight 0.3).

Batch Normalization

Internal covariate shift refers to the change in the distributions of hidden node activations due to the updates of network parameters during training.

Let $\boldsymbol{x} = (x_1, \ldots, x_d)$ be d-dimensional input. We would like to normalize each dimension as

$$\hat{x}_i = \frac{x_i - \mathbb{E}[x_i]}{\sqrt{\operatorname{Var}[x_i]}}.$$

Furthermore, it may be advantageous to learn suitable scale γ_i and shift β_i to produce normalized value

$$y_i = \gamma_i \hat{x}_i + \beta_i.$$

Consider a mini-batch of m examples $(\boldsymbol{x}^{(1)}, \ldots, \boldsymbol{x}^{(m)})$.

Batch normalizing transform of the mini-batch is the following transformation.

Inputs: Mini-batch $(\boldsymbol{x}^{(1)}, \ldots, \boldsymbol{x}^{(m)})$, $\varepsilon \in \mathbb{R}$
Outputs: Normalized batch $(\boldsymbol{y}^{(1)}, \ldots, \boldsymbol{y}^{(m)})$

- $\boldsymbol{\mu} \leftarrow \frac{1}{m} \sum_{i=1}^{m} \boldsymbol{x}^{(i)}$
- $\boldsymbol{\sigma}^2 \leftarrow \frac{1}{m} \sum_{i=1}^{m} (\boldsymbol{x}^{(i)} - \mu)^2$
- $\hat{\boldsymbol{x}}^{(i)} \leftarrow (\boldsymbol{x}^{(i)} - \boldsymbol{\mu}) / \sqrt{\sigma^2 + \varepsilon}$
- $\boldsymbol{y}^{(i)} \leftarrow \boldsymbol{\gamma} \hat{\boldsymbol{x}}^{(i)} + \boldsymbol{\beta}$

Batch normalization is commonly added just before a nonlinearity. Therefore, we replace $f(\boldsymbol{W}\boldsymbol{x} + \boldsymbol{b})$ by $f(BN(\boldsymbol{W}\boldsymbol{x}))$.

During inference, $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}^2$ are fixed. They are either precomputed after training on the whole training data, or an exponential moving average is updated during training.

When a batch normalization is used on a fully connected layer, each neuron is normalized individually across the minibatch.

However, for convolutional networks we would like the normalization to honour their properties, most notably the shift invariance. We therefore normalize each channel across not only the minibatch, but also across all corresponding spacial/temporal locations.

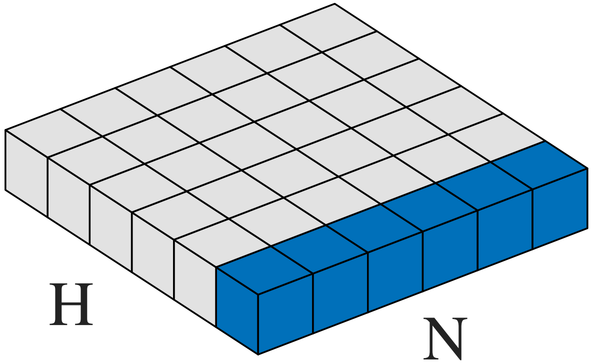

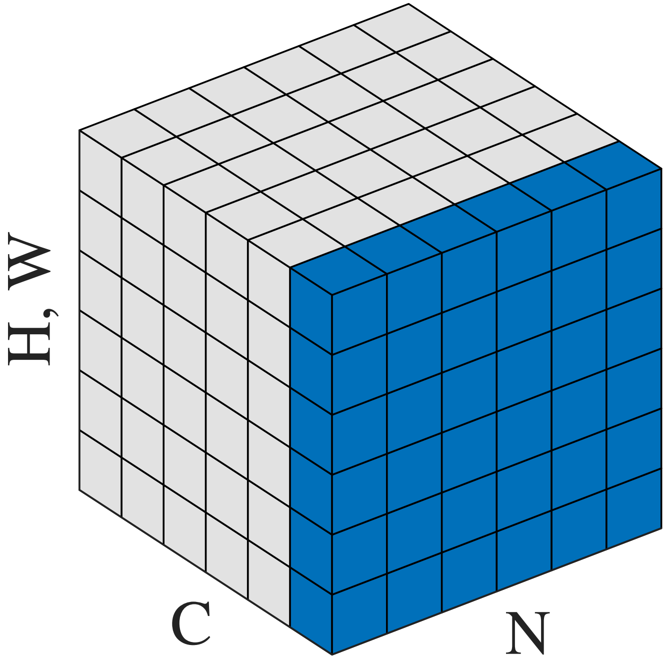

Adapted from Figure 2 of paper "Group Normalization", https://arxiv.org/abs/1803.08494.

Inception with BatchNorm (4.8% error)

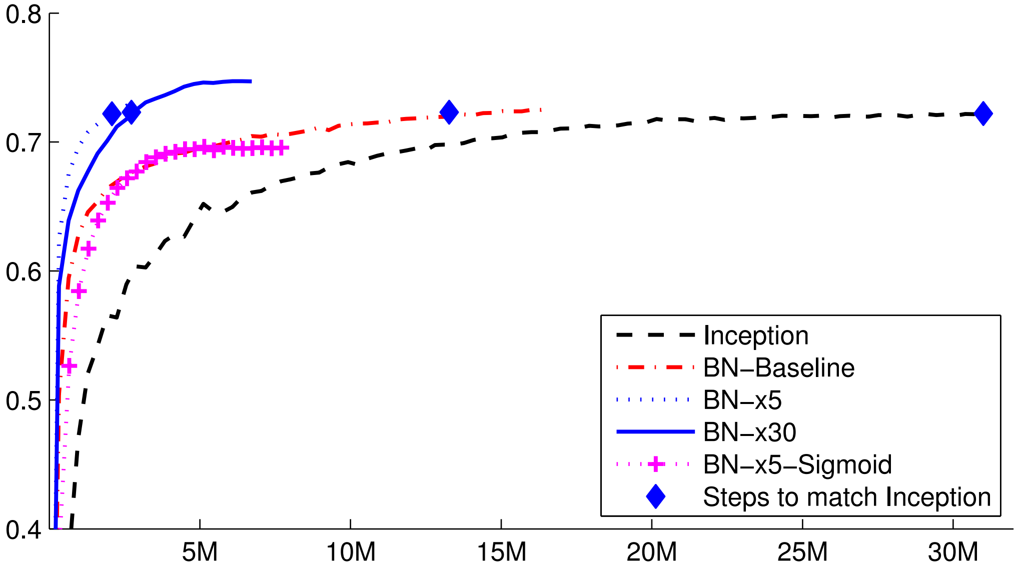

Figure 2: *Single crop validation accuracy of Inception and its batch-normalized variants, vs. the number of training steps.*

Model	Steps to 72.2%	Max accuracy
Inception	$31.0 \cdot 10^6$	72.2%
BN-Baseline	$13.3 \cdot 10^6$	72.7%
BN-x5	$2.1 \cdot 10^6$	73.0%
BN-x30	$2.7 \cdot 10^6$	74.8%
BN-x5-Sigmoid		69.8%

F gure 3: *For Inception and the batch-normalized variants, the number of training steps required to reach the maximum accuracy of Inception (72.2%), and the maximum accuracy achieved by the network.*

Figures 2 and 3 of paper "Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift", https://arxiv.org/abs/1502.03167.

Inception v2 and v3 – 2015 (3.6% error)

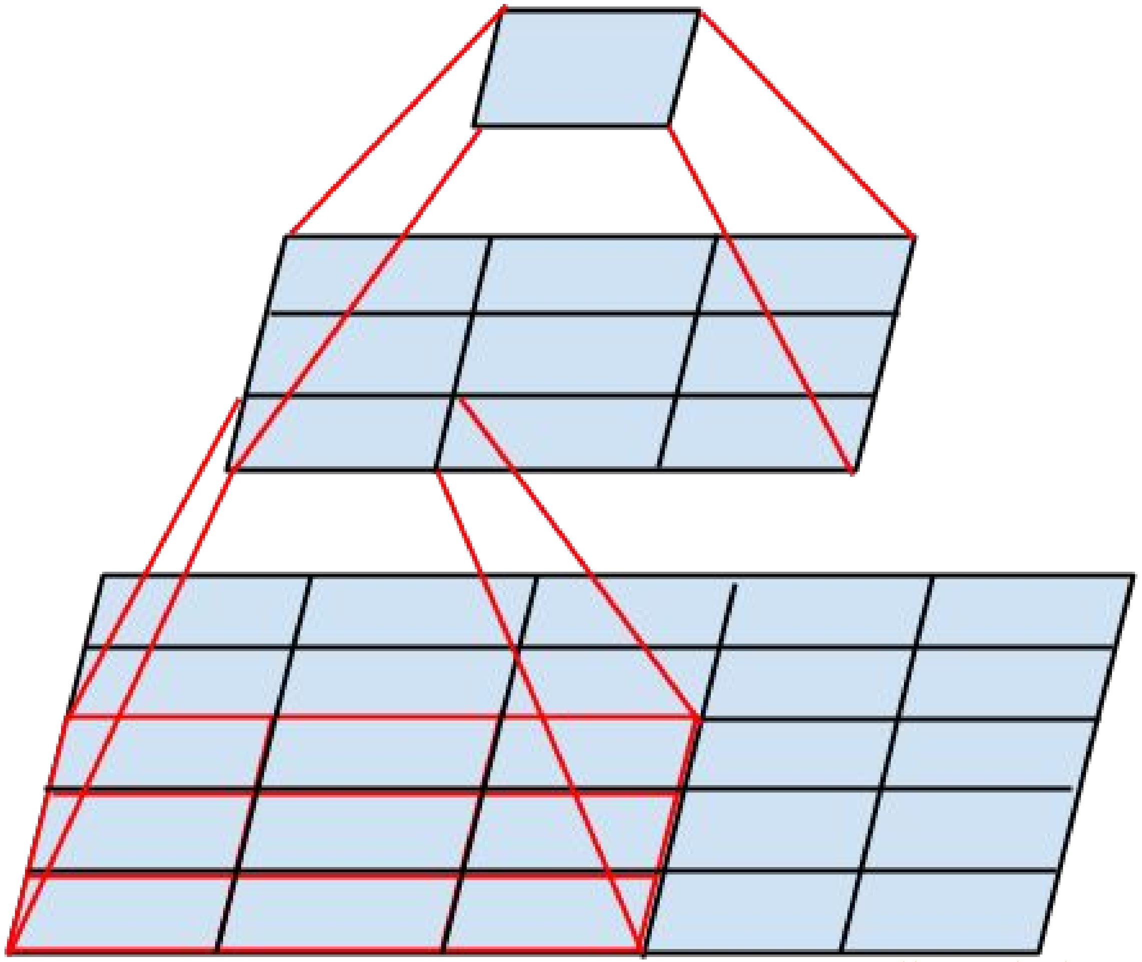

Figure 1 of paper "Rethinking the Inception Architecture for Computer Vision", https://arxiv.org/abs/1512.00567.

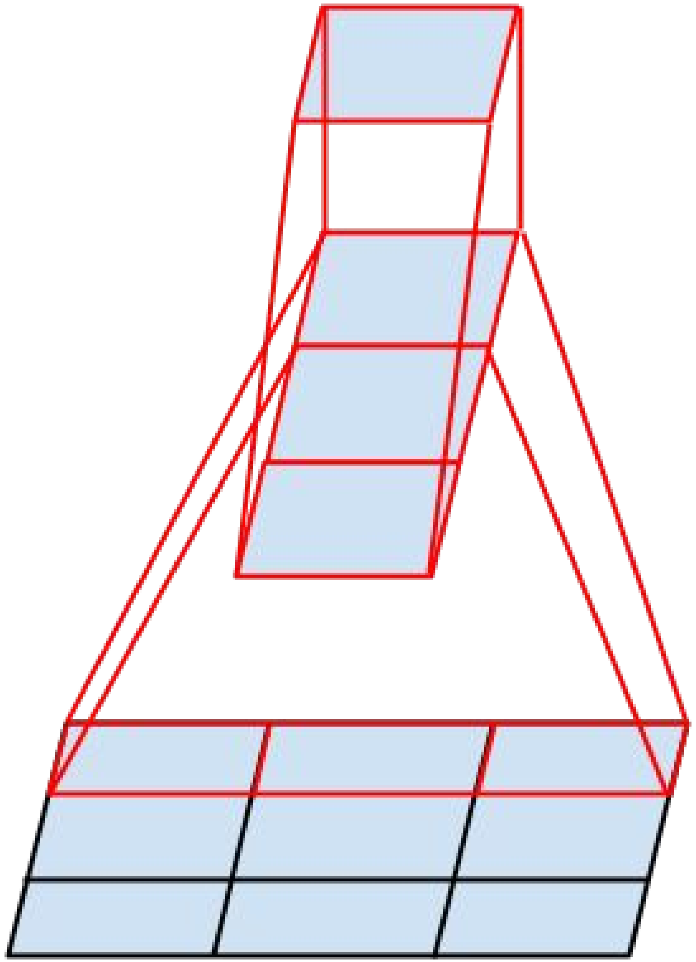

Figure 3 of paper "Rethinking the Inception Architecture for Computer Vision", https://arxiv.org/abs/1512.00567.

Inception v2 and v3 – 2015 (3.6% error)

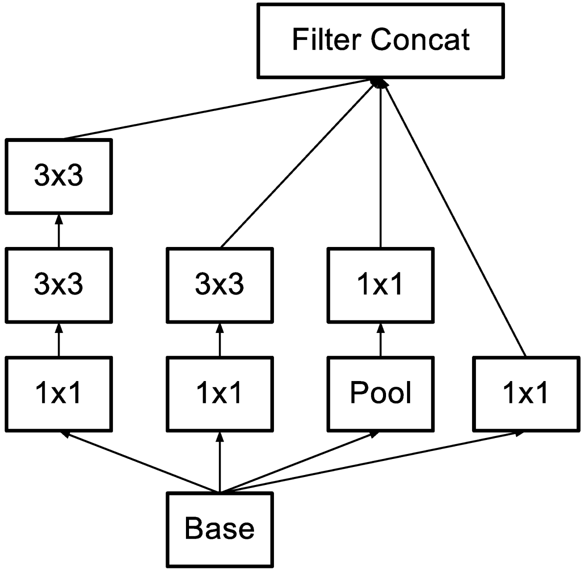

Figure 5. Inception modules where each 5×5 convolution is replaced by two 3×3 convolution, as suggested by principle 3 of Section 2.

Figure 5 of paper "Rethinking the Inception Architecture for Computer Vision", https://arxiv.org/abs/1512.00567.

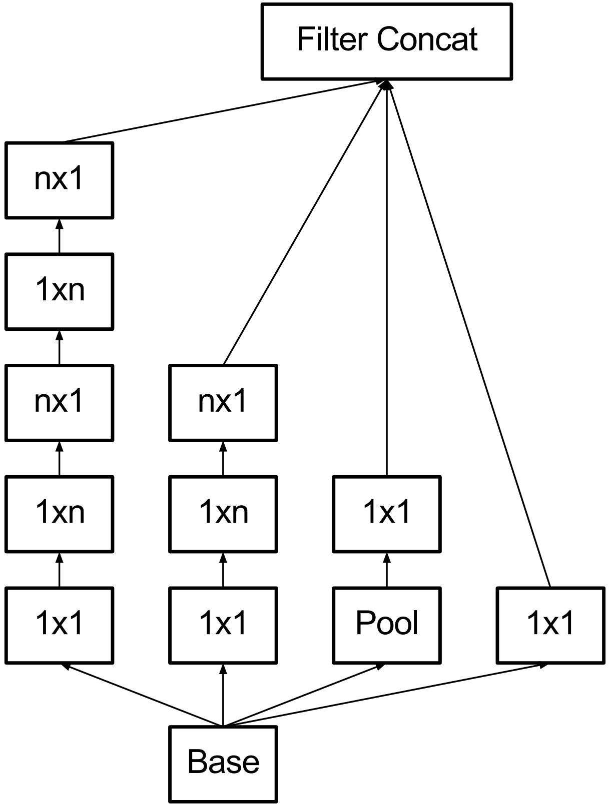

Figure 6. Inception modules after the factorization of the $n \times n$ convolutions. In our proposed architecture, we chose $n = 7$ for the 17×17 grid. (The filter sizes are picked using principle 3)

.

Figure 6 of paper "Rethinking the Inception Architecture for Computer Vision", https://arxiv.org/abs/1512.00567.

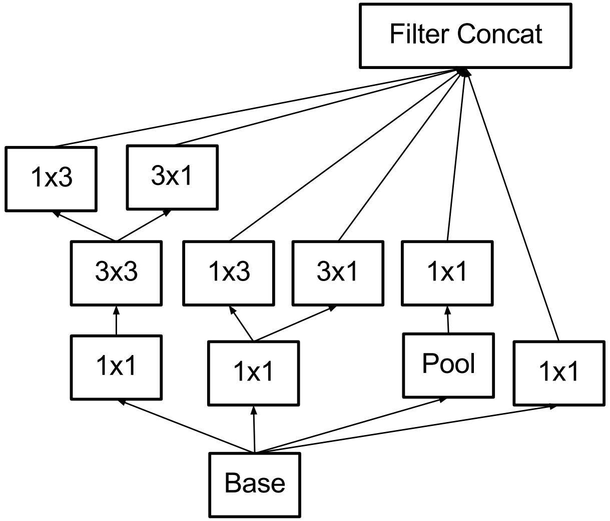

Figure 7. Inception modules with expanded the filter bank outputs. This architecture is used on the coarsest (8×8) grids to promote high dimensional representations, as suggested by principle 2 of Section 2. We are using this solution only on the coarsest grid, since that is the place where producing high dimensional sparse representation is the most critical as the ratio of local processing (by 1×1 convolutions) is increased compared to the spatial aggregation.

Figure 7 of paper "Rethinking the Inception Architecture for Computer Vision", https://arxiv.org/abs/1512.00567.

Inception v2 and v3 – 2015 (3.6% error)

type	patch size/stride or remarks	input size
conv	$3\times3/2$	$299\times299\times3$
conv	$3\times3/1$	$149\times149\times32$
conv padded	$3\times3/1$	$147\times147\times32$
pool	$3\times3/2$	$147\times147\times64$
conv	$3\times3/1$	$73\times73\times64$
conv	$3\times3/2$	$71\times71\times80$
conv	$3\times3/1$	$35\times35\times192$
$3\times$Inception	As in figure 5	$35\times35\times288$
$5\times$Inception	As in figure 6	$17\times17\times768$
$2\times$Inception	As in figure 7	$8\times8\times1280$
pool	8×8	$8\times8\times2048$
linear	logits	$1\times1\times2048$
softmax	classifier	$1\times1\times1000$

Table 1 of paper "Rethinking the Inception Architecture for Computer Vision", https://arxiv.org/abs/1512.00567.

Inception v2 and v3 – 2015 (3.6% error)

Network	Top-1 Error	Top-5 Error	Cost Bn Ops
GoogLeNet [20]	29%	9.2%	1.5
BN-GoogLeNet	26.8%	-	**1.5**
BN-Inception [7]	25.2%	7.8	2.0
Inception-v2	23.4%	-	3.8
Inception-v2 RMSProp	23.1%	6.3	3.8
Inception-v2 Label Smoothing	22.8%	6.1	3.8
Inception-v2 Factorized 7×7	21.6%	5.8	4.8
Inception-v2 BN-auxiliary	**21.2%**	**5.6%**	4.8

Table 3 of paper "Rethinking the Inception Architecture for Computer Vision", https://arxiv.org/abs/1512.00567.

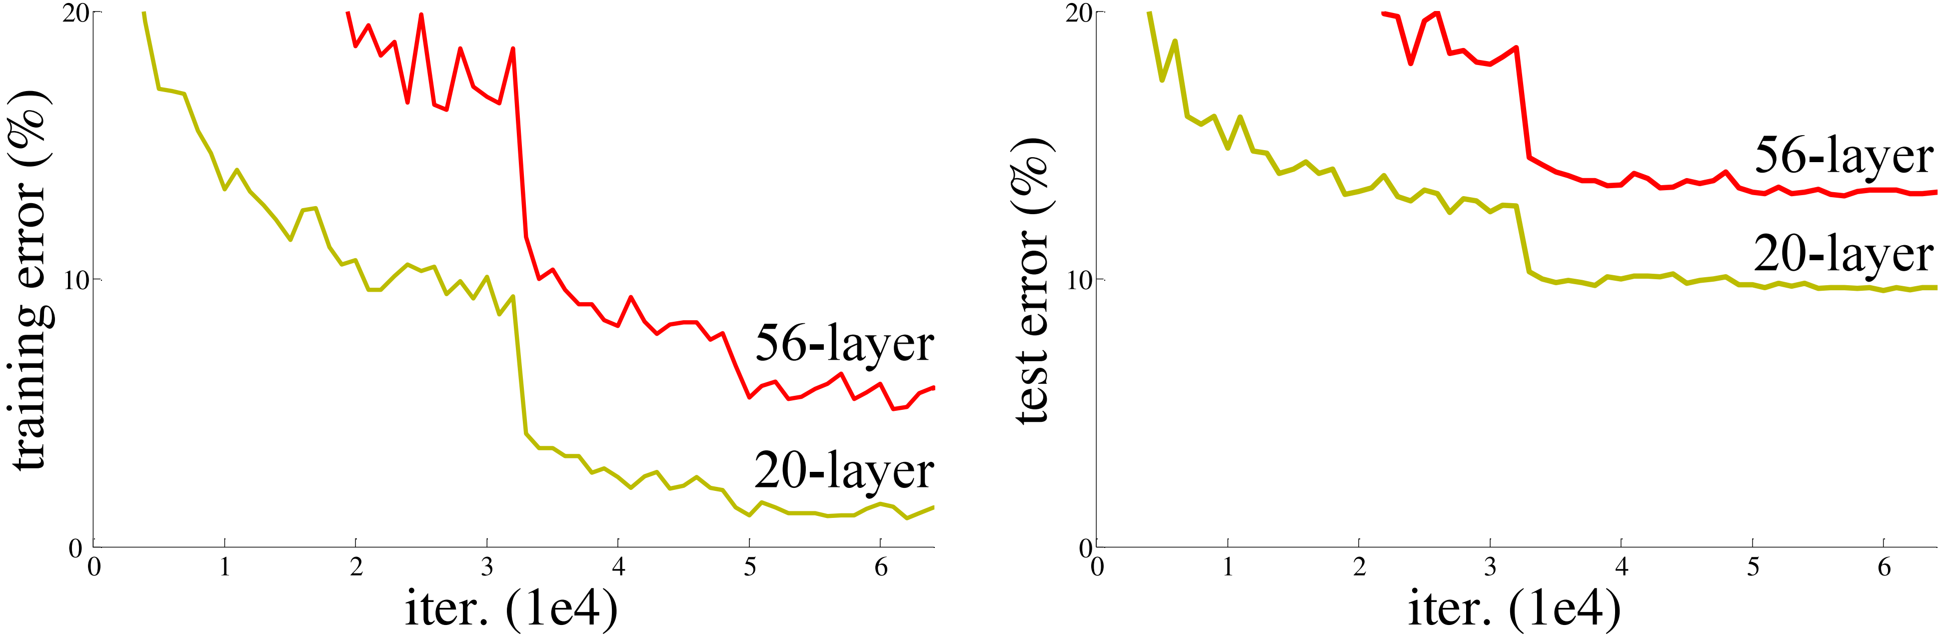

Figure 1. Training error (left) and test error (right) on CIFAR-10 with 20-layer and 56-layer "plain" networks. The deeper network has higher training error, and thus test error. Similar phenomena on ImageNet is presented in Fig. 4.

Figure 1 of paper "Deep Residual Learning for Image Recognition", https://arxiv.org/abs/1512.03385.

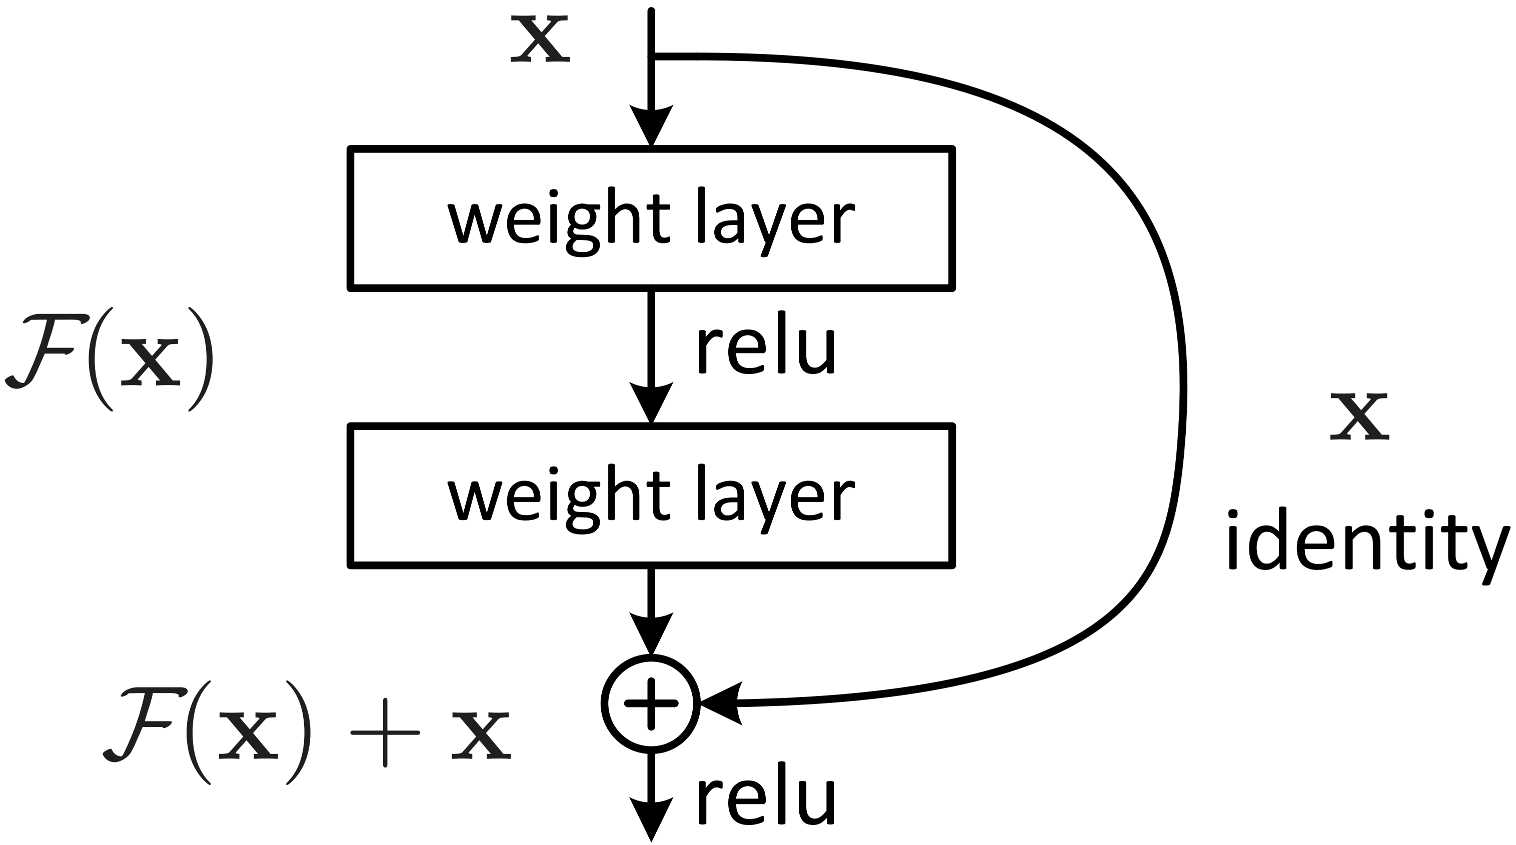

Figure 2. Residual learning: a building block.

Figure 2 of paper "Deep Residual Learning for Image Recognition", https://arxiv.org/abs/1512.03385.

ResNet – 2015 (3.6% error)

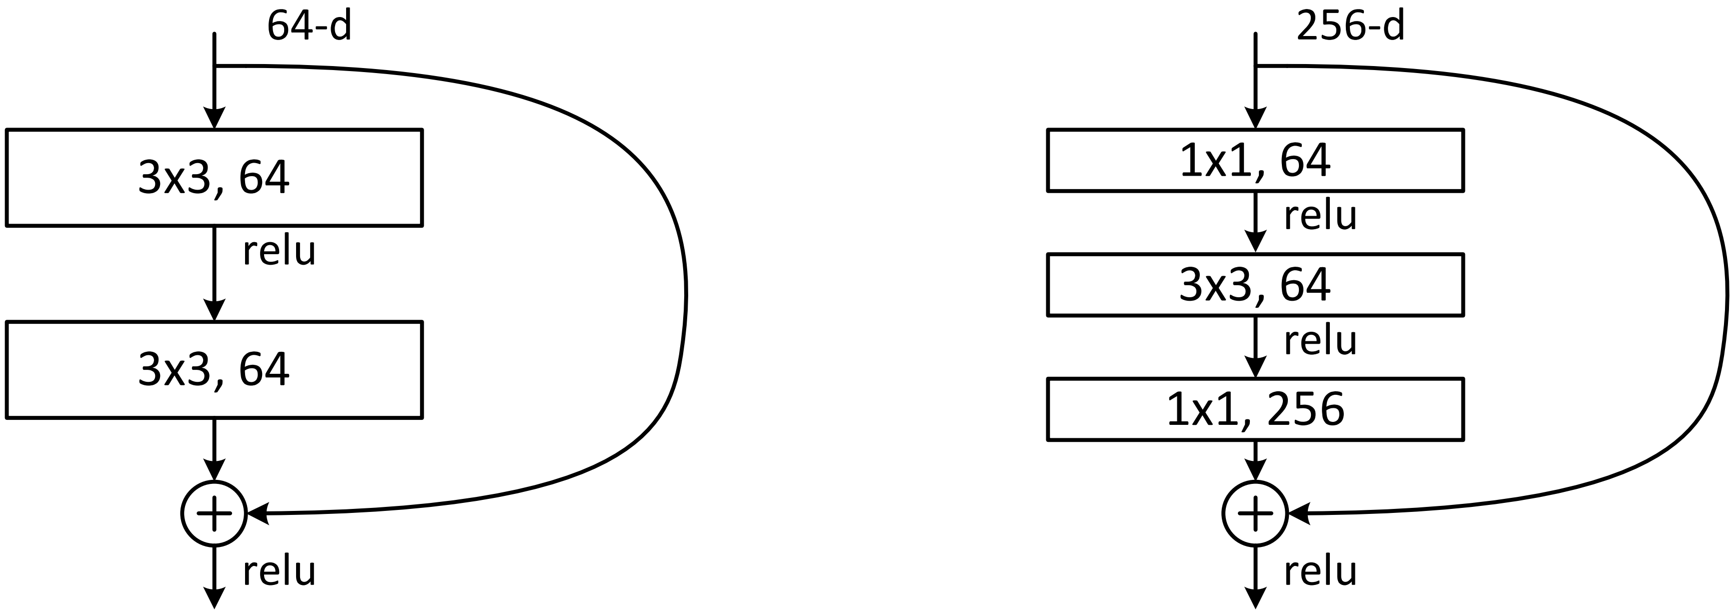

Figure 5. A deeper residual function $\mathcal{F}$ for ImageNet. Left: a building block (on 56×56 feature maps) as in Fig. 3 for ResNet-34. Right: a "bottleneck" building block for ResNet-50/101/152.

Figure 5 of paper "Deep Residual Learning for Image Recognition", https://arxiv.org/abs/1512.03385.

ResNet – 2015 (3.6% error)

layer name	output size	18-layer	34-layer	50-layer	101-layer	152-layer
conv1	112×112	7×7, 64, stride 2				
conv2_x	56×56	3×3 max pool, stride 2				
		$\begin{bmatrix} 3\times3, 64 \\ 3\times3, 64 \end{bmatrix}\times2$	$\begin{bmatrix} 3\times3, 64 \\ 3\times3, 64 \end{bmatrix}\times3$	$\begin{bmatrix} 1\times1, 64 \\ 3\times3, 64 \\ 1\times1, 256 \end{bmatrix}\times3$	$\begin{bmatrix} 1\times1, 64 \\ 3\times3, 64 \\ 1\times1, 256 \end{bmatrix}\times3$	$\begin{bmatrix} 1\times1, 64 \\ 3\times3, 64 \\ 1\times1, 256 \end{bmatrix}\times3$
conv3_x	28×28	$\begin{bmatrix} 3\times3, 128 \\ 3\times3, 128 \end{bmatrix}\times2$	$\begin{bmatrix} 3\times3, 128 \\ 3\times3, 128 \end{bmatrix}\times4$	$\begin{bmatrix} 1\times1, 128 \\ 3\times3, 128 \\ 1\times1, 512 \end{bmatrix}\times4$	$\begin{bmatrix} 1\times1, 128 \\ 3\times3, 128 \\ 1\times1, 512 \end{bmatrix}\times4$	$\begin{bmatrix} 1\times1, 128 \\ 3\times3, 128 \\ 1\times1, 512 \end{bmatrix}\times8$
conv4_x	14×14	$\begin{bmatrix} 3\times3, 256 \\ 3\times3, 256 \end{bmatrix}\times2$	$\begin{bmatrix} 3\times3, 256 \\ 3\times3, 256 \end{bmatrix}\times6$	$\begin{bmatrix} 1\times1, 256 \\ 3\times3, 256 \\ 1\times1, 1024 \end{bmatrix}\times6$	$\begin{bmatrix} 1\times1, 256 \\ 3\times3, 256 \\ 1\times1, 1024 \end{bmatrix}\times23$	$\begin{bmatrix} 1\times1, 256 \\ 3\times3, 256 \\ 1\times1, 1024 \end{bmatrix}\times36$
conv5_x	7×7	$\begin{bmatrix} 3\times3, 512 \\ 3\times3, 512 \end{bmatrix}\times2$	$\begin{bmatrix} 3\times3, 512 \\ 3\times3, 512 \end{bmatrix}\times3$	$\begin{bmatrix} 1\times1, 512 \\ 3\times3, 512 \\ 1\times1, 2048 \end{bmatrix}\times3$	$\begin{bmatrix} 1\times1, 512 \\ 3\times3, 512 \\ 1\times1, 2048 \end{bmatrix}\times3$	$\begin{bmatrix} 1\times1, 512 \\ 3\times3, 512 \\ 1\times1, 2048 \end{bmatrix}\times3$
	1×1	average pool, 1000-d fc, softmax				
FLOPs		1.8×10^9	3.6×10^9	3.8×10^9	7.6×10^9	11.3×10^9

Table 1 of paper "Deep Residual Learning for Image Recognition", https://arxiv.org/abs/1512.03385.

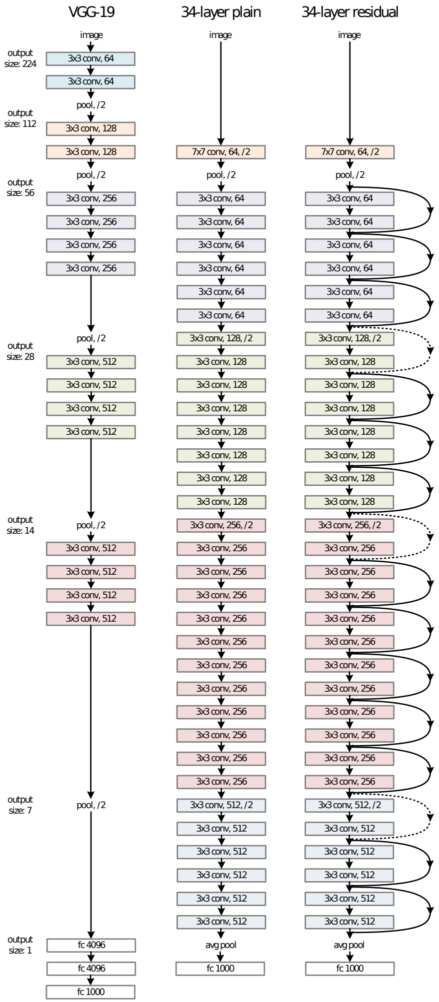

Figure 3 of paper "Deep Residual Learning for Image Recognition", https://arxiv.org/abs/1512.03385.

The residual connections cannot be applied directly when number of channels increase.

The authors considered several alternatives, and chose the one where in case of channels increase a 1×1 convolution is used on the projections to match the required number of channels.

ResNet – 2015 (3.6% error)

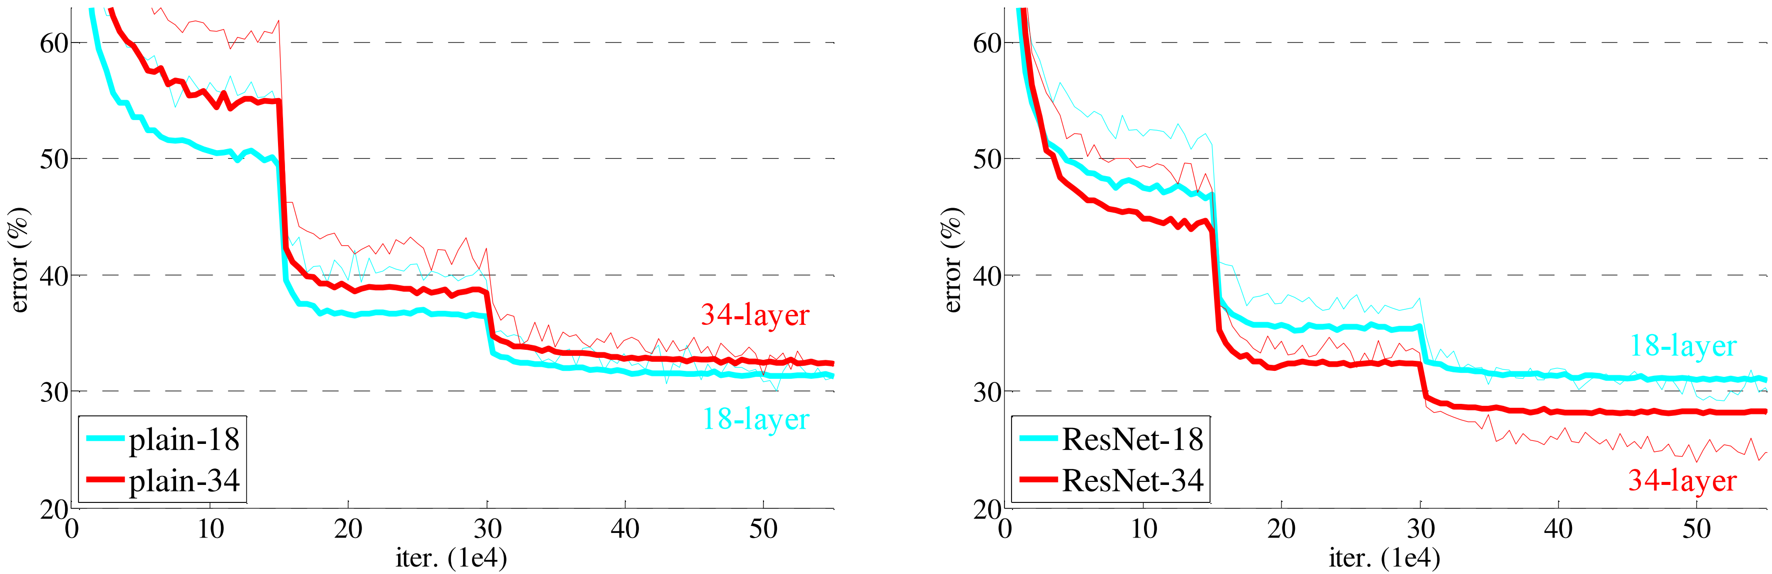

Figure 4. Training on **ImageNet**. Thin curves denote training error, and bold curves denote validation error of the center crops. Left: plain networks of 18 and 34 laye s. Right: ResNets of 18 and 34 layers. In this plot, the residual networks have no extra parameter compared to their plain counterparts.

Figure 4 of paper "Deep Residual Learning for Image Recognition", https://arxiv.org/abs/1512.03385.

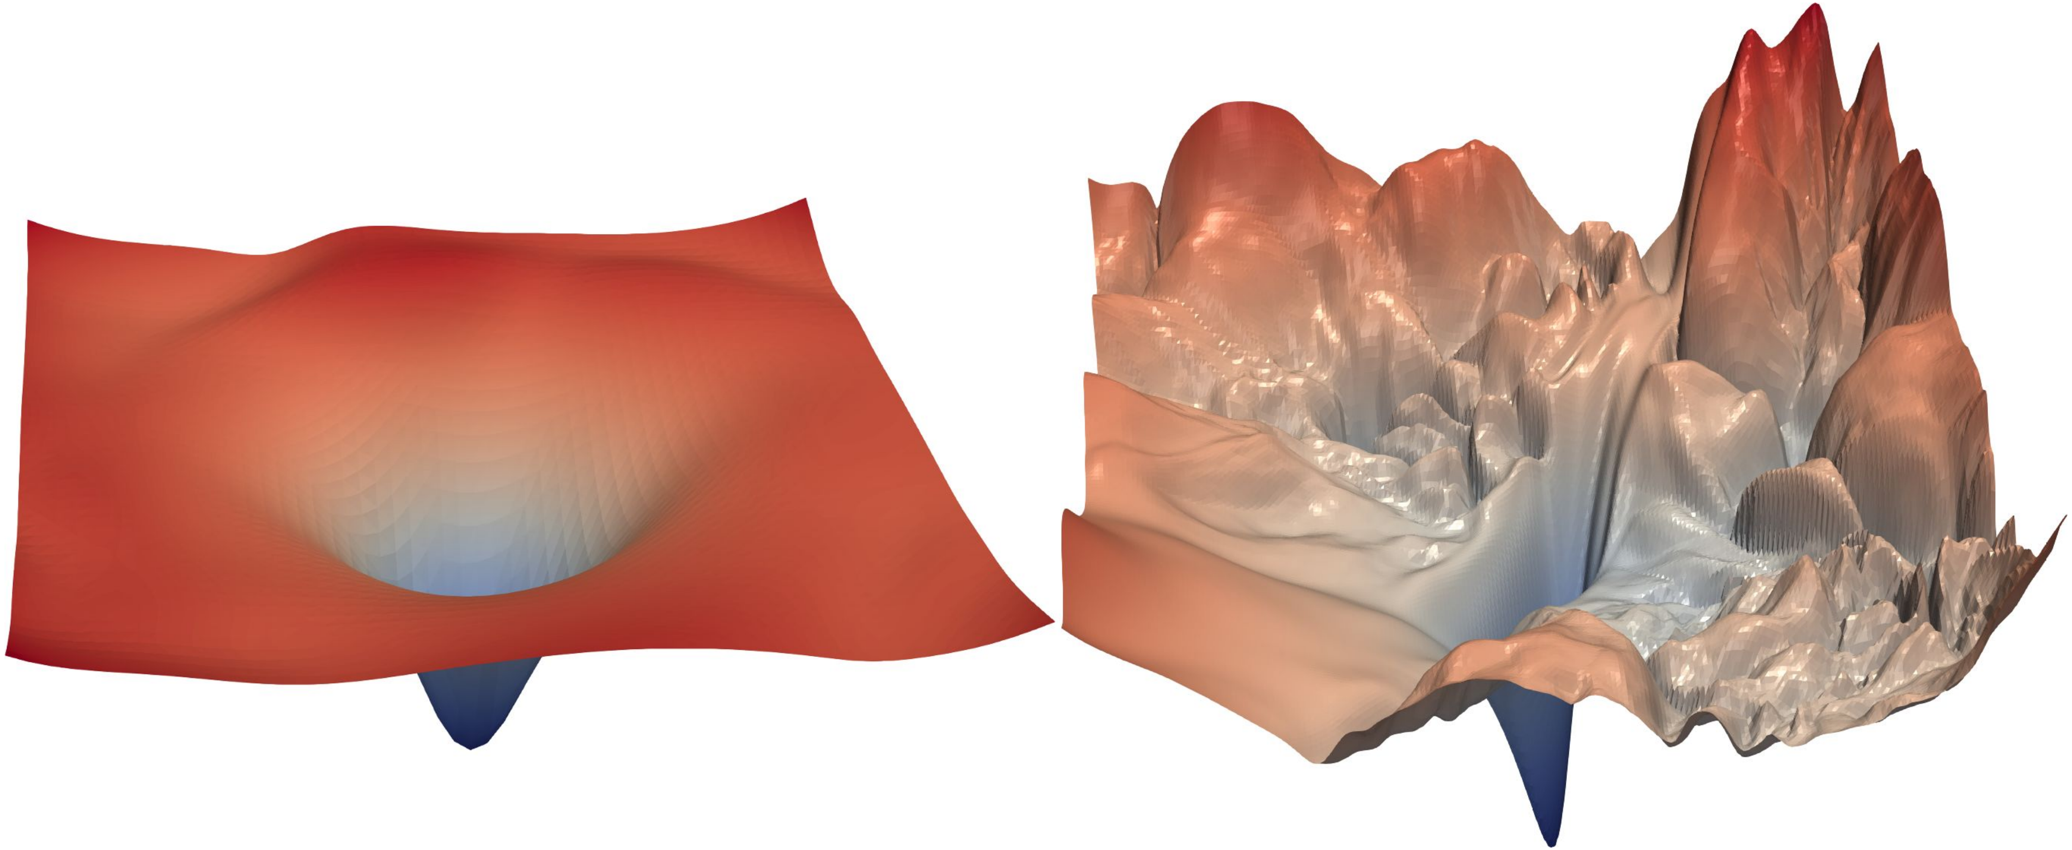

Figure 1 of paper "Visualizing the Loss Landscape of Neural Nets", https://arxiv.org/abs/1712.09913.

ResNet – 2015 (3.6% error)

method	top-1 err.	top-5 err.
VGG [41] (ILSVRC'14)	-	8.43†
GoogLeNet [44] (ILSVRC'14)	-	7.89
VGG [41] (v5)	24.4	7.1
PReLU-net [13]	21.59	5.71
BN-inception [16]	21.99	5.81
ResNet-34 B	21.84	5.71
ResNet-34 C	21.53	5.60
ResNet-50	20.74	5.25
ResNet-101	19.87	4.60
ResNet-152	**19.38**	**4.49**

Table 4. Error rates (%) of **single-model** results on the ImageNet validation set (except † reported on the test set).

Table 4 of paper "Deep Residual Learning for Image Recognition", https://arxiv.org/abs/1512.03385.

method	top-5 err. (**test**)
VGG [41] (ILSVRC'14)	7.32
GoogLeNet [44] (ILSVRC'14)	6.66
VGG [41] (v5)	6.8
PReLU-net [13]	4.94
BN-inception [16]	4.82
ResNet (ILSVRC'15)	**3.57**

Table 5. Error rates (%) of **ensembles**. The top-5 error is on the test set of ImageNet and reported by the test server.

Table 5 of paper "Deep Residual Learning for Image Recognition", https://arxiv.org/abs/1512.03385.

Main Takeaways

- Convolutions can provide
  - local interactions in spacial/temporal dimensions
  - shift invariance
  - *much* less parameters than a fully connected layer
- Usually repeated 3×3 convolutions are enough, no need for larger filter sizes.
- When pooling is performed, double number of channels.
- Final fully connected layers are not needed, global average pooling is usually enough.
- Batch normalization is a great regularization method for CNNs.