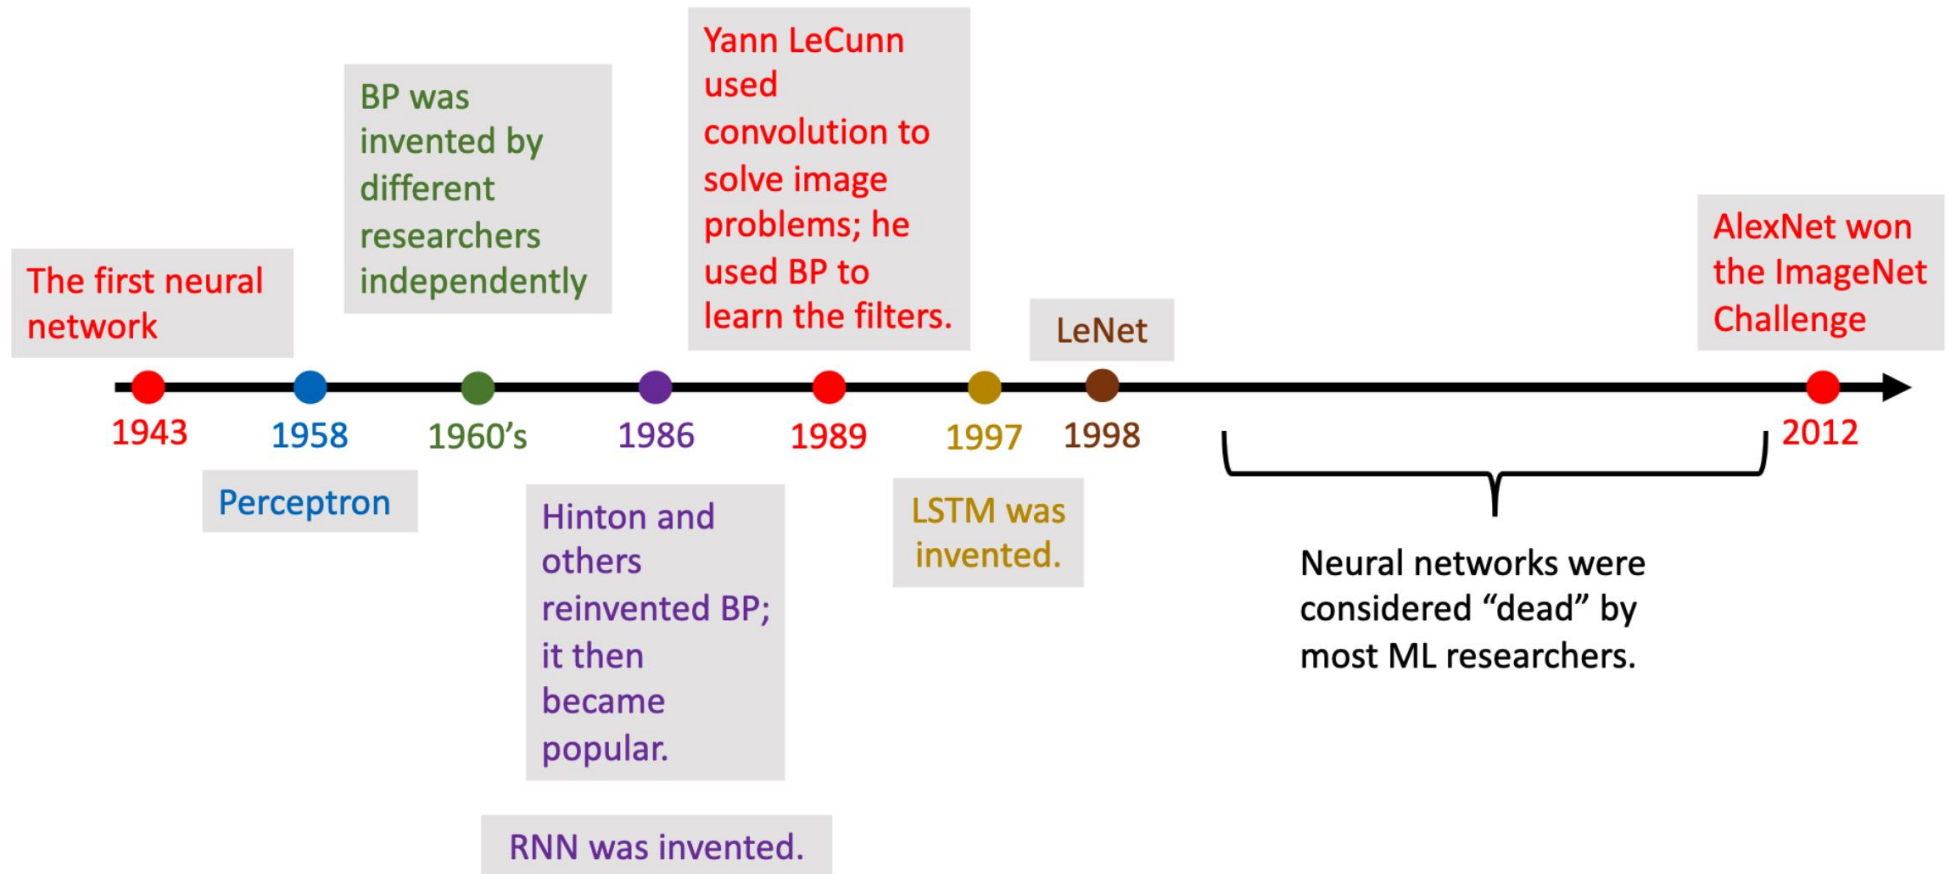


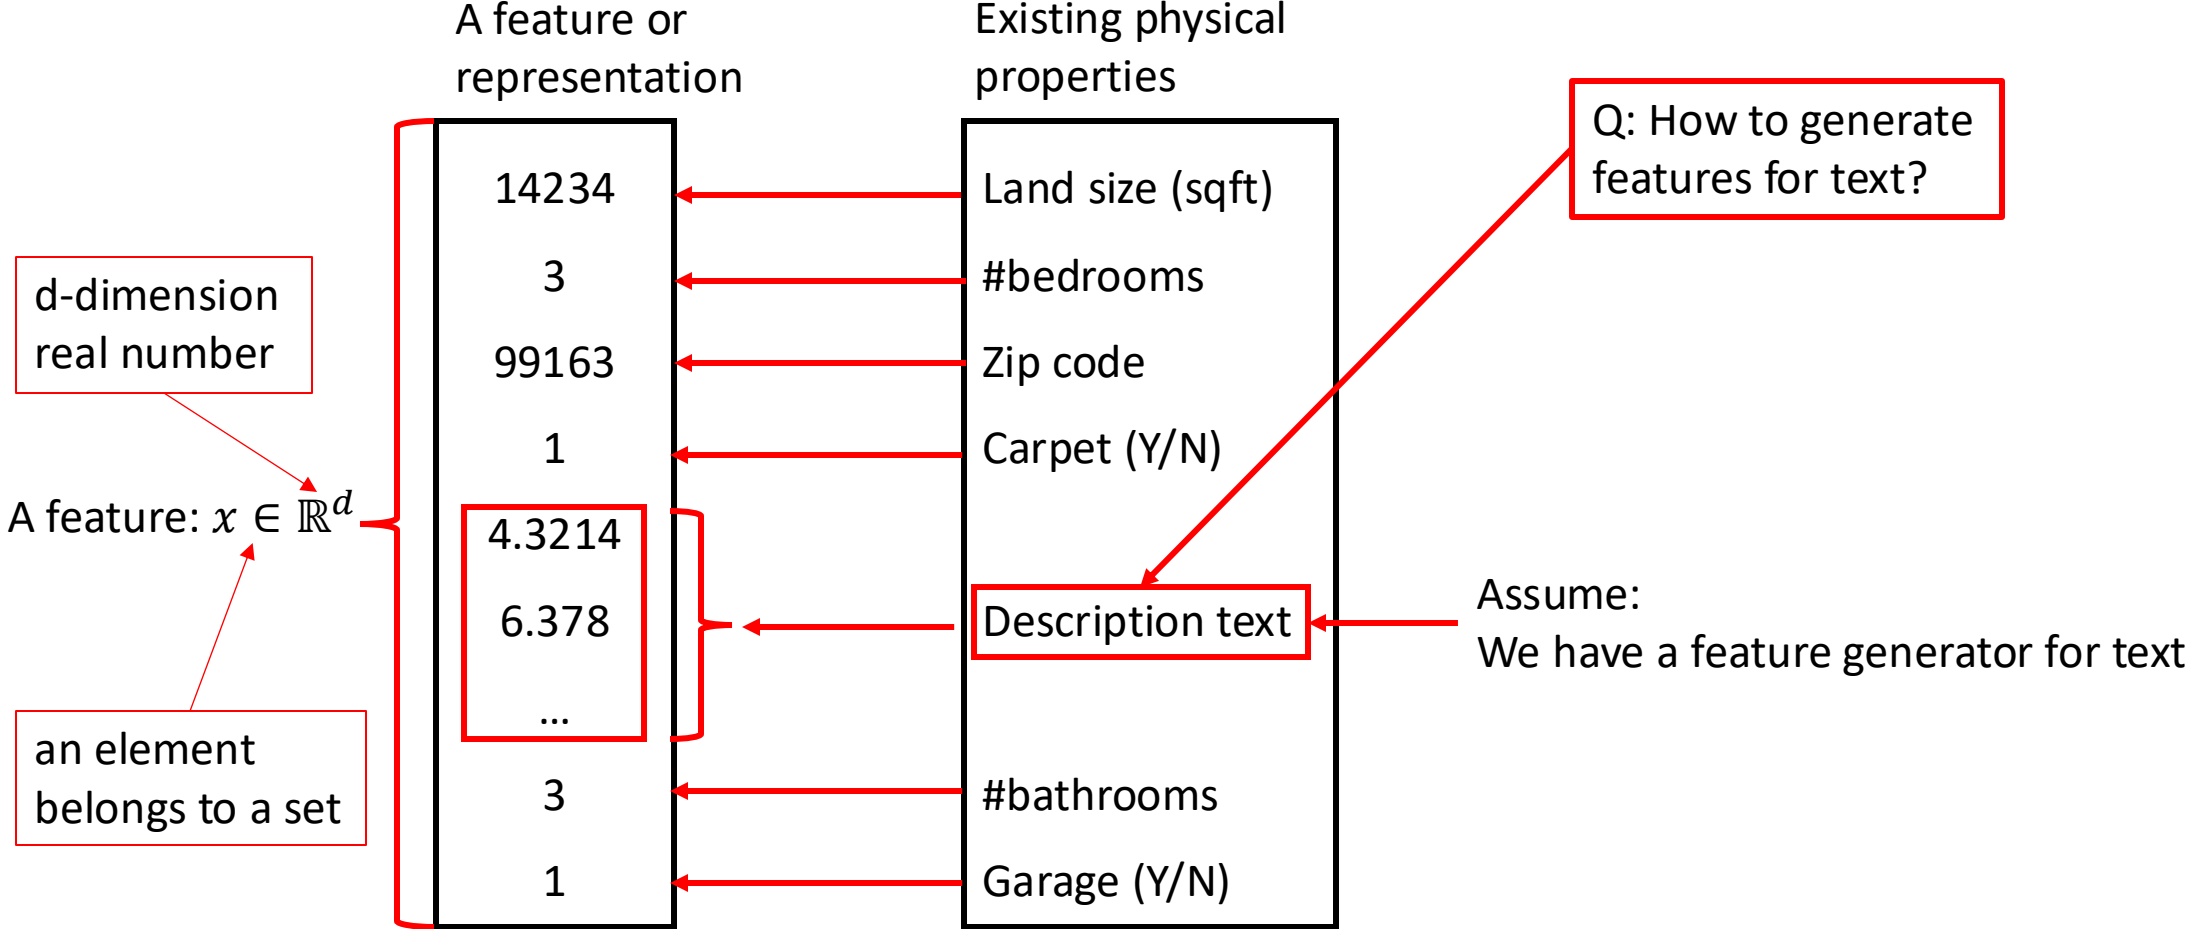
Convolutional Layer and Convolutional Neural Networks

Neural Networks Design And Application

History




Review: house price prediction



Review: histogram of oriented gradients

- Oriented gradients?
 - Gradients: changes in X and Y directions
 - Oriented:



121	10	78	96	125
48	152	68	125	111
145	78	85	89	65
154	214	56	200	66
214	87	45	102	45

X direction G_x

Subtract the value on the left from the pixel value on the right:

$$89 - 78 = 11$$

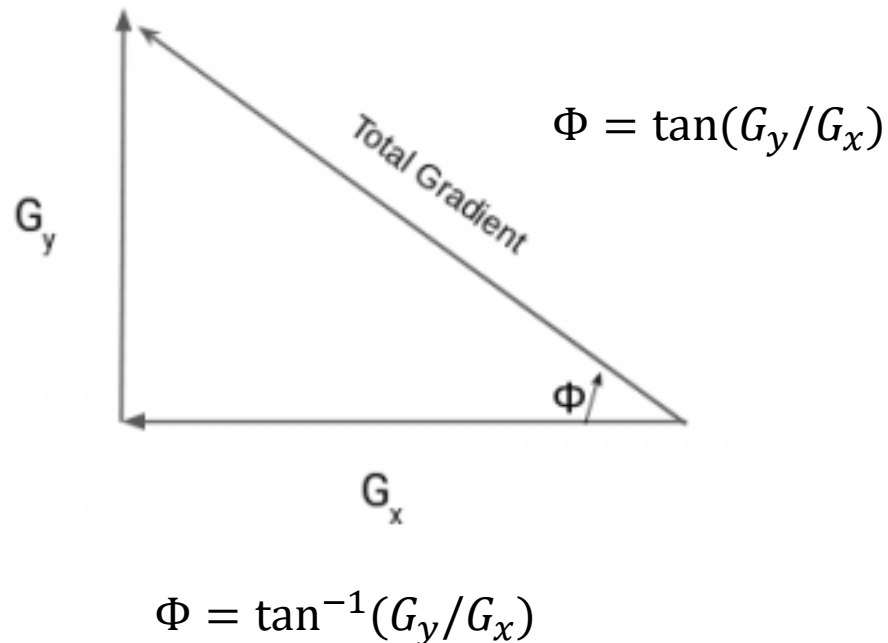
Y direction G_y

Subtract the pixel value below from the pixel value above the selected pixel:

$$68 - 56 = 12$$

Review: histogram of oriented gradients

- Oriented gradients?
 - Gradients: changes in X and Y directions
 - Oriented:



121	10	78	96	125
48	152	68	125	111
145	78	85	89	65
154	214	56	200	66
214	87	45	102	45

X direction G_x
Subtract the value on the left from the pixel value on the right:
 $89 - 78 = 11$

Y direction G_y
Subtract the pixel value below from the pixel value above the selected pixel:
 $68 - 56 = 8$

Review: histogram of oriented gradients

121	10	78	96	125
48	152	68	125	111
145	78	85	89	65
154	214	56	200	66
214	87	45	102	45

Frequency						1										
Angle	1	2	3	4 ...	35	36	37	38	39....		175	176	177	178	179	180

Review: ImageNet challenge 2012

Task 1

Team name	Filename	Error (5 guesses)	Description
SuperVision	test-preds-141-146.2009-131-137-145-146.2011-145f.	0.15315	Using extra training data from ImageNet Fall 2011 release
SuperVision	test-preds-131-137-145-135-145f.txt	0.16422	Using only supplied training data
ISI	pred_FVs_wLACs_weighted.txt	0.26172	Weighted sum of scores from each classifier with SIFT+FV, LBP+FV, GIST+FV, and CSIFT+FV, respectively.
ISI	pred_FVs_weighted.txt	0.26602	Weighted sum of scores from classifiers using each FV.
ISI	pred_FVs_summed.txt	0.26646	Naive sum of scores from classifiers using each FV.

AlexNet



Review: LeNet-5 in 1999





Fig. 1. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.


LeCun, Yann, Patrick Haffner, Léon Bottou, and Yoshua Bengio. "Object recognition with gradient-based learning." In *Shape, contour and grouping in computer vision*, pp. 319-345. Springer, Berlin, Heidelberg, 1999.

What is **convolutional** neural network?



A grayscale image

What is **convolutional** neural network?




A grayscale image



An image \rightarrow a matrix

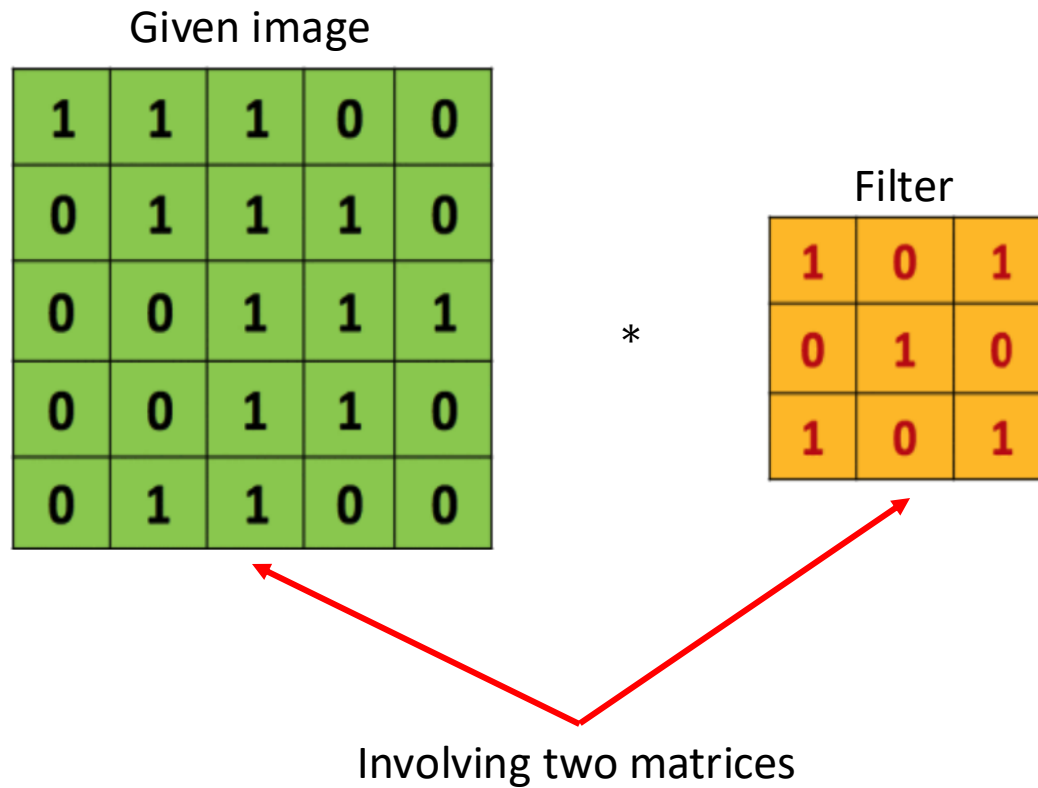
1	1	1	0	0
0	1	1	1	0
0	0	1	1	1
0	0	1	1	0
0	1	1	0	0

Convolution for images (matrices)

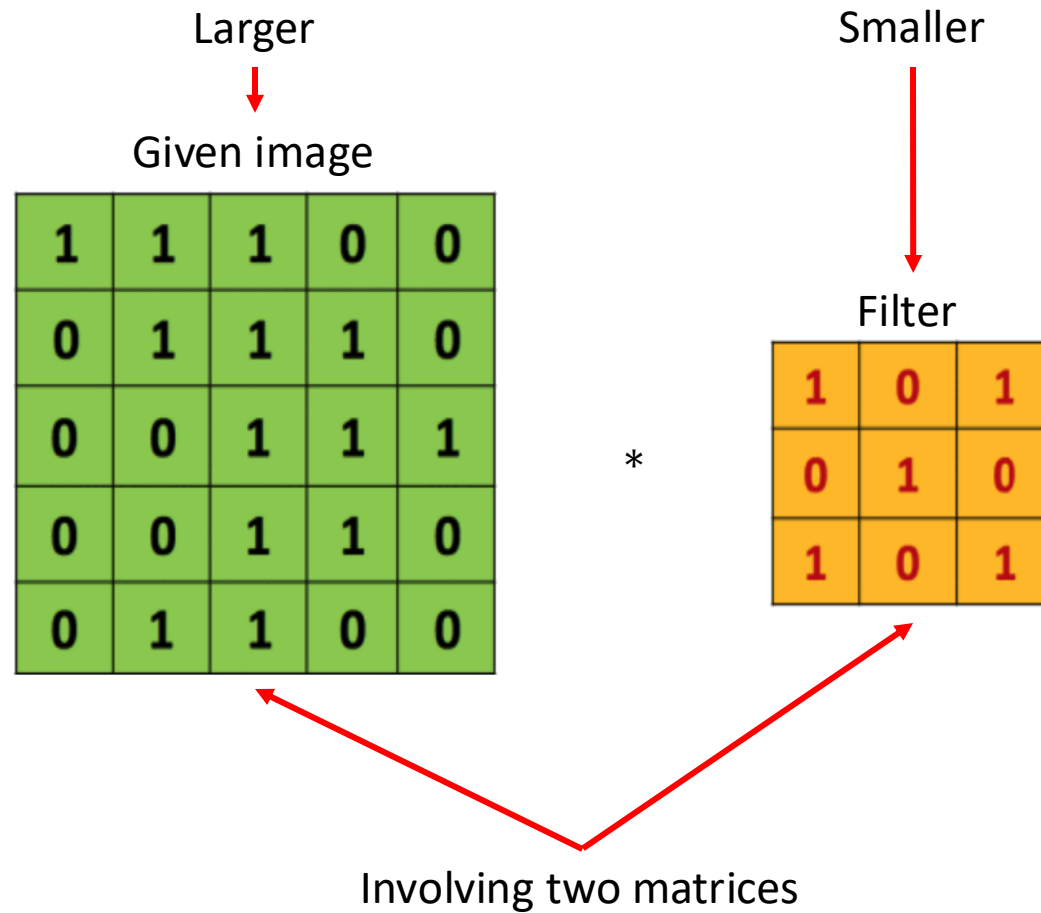


Involving two matrices


Convolution for images (matrices)



Convolution for images (matrices)




Convolution for images (matrices)




Finding pairs

Convolution for images (matrices)




Finding pairs

Convolution for images (matrices)




Finding pairs

Convolution for images (matrices)




Finding pairs

Convolution for images (matrices)



Finding pairs


Convolution for images (matrices)



Finding pairs

Q: how many pairs we have?

Convolution for images (matrices)




Finding pairs

Q: how many pairs we have?


$$(5-3+1) * (5-3+1)=9$$

Convolution for images (matrices)



Inner product of each pair


Convolution for images (matrices)



Inner product of each pair

Elementwise multiplication + summation

Convolution for images (matrices)




Inner product of each pair

Elementwise multiplication + summation

Q: what is your result for the first pair?


Convolution for images (matrices)



Inner product of each pair

Elementwise multiplication + summation


Convolution for images (matrices)



Inner product of each pair


Elementwise multiplication + summation

Convolution for images (matrices)




Q: the second pair?


Convolution for images (matrices)



Convolution for images (matrices)



Convolution for images (matrices)



We can repeat for each pair

Convolution for images (matrices)

1	1	1	0	0
0	1	1	1	0
0	0	1	1	1
0	0	1	1	0
0	1	1	0	0

*

1	0	1
0	1	0
1	0	1

→

4	3	4
2	4	3
2	3	4

Convolution for images (matrices)

1	1	1	0	0
0	1	1	1	0
0	0	1	1	1
0	0	1	1	0
0	1	1	0	0

*

1	0	1
0	1	0
1	0	1

→

4	3	4
2	4	3
2	3	4

Place each element
according to their positions

Convolution for images (matrices)

1	1	1	0	0
0	1	1	1	0
0	0	1	1	1
0	0	1	1	0
0	1	1	0	0

*

1	0	1
0	1	0
1	0	1

→

Row: 1
Column: 1

4	3	4
2	4	3
2	3	4

Place each element
according to their positions

Convolution for images (matrices)

1	1	1	0	0
0	1	1	1	0
0	0	1	1	1
0	0	1	1	0
0	1	1	0	0

*

1	0	1
0	1	0
1	0	1

→

Row: 1
Column: 2

4	3	4
2	4	3
2	3	4

Place each element
according to their positions

Convolution for images (matrices)

1	1	1	0	0
0	1	1	1	0
0	0	1	1	1
0	0	1	1	0
0	1	1	0	0

*

1	0	1
0	1	0
1	0	1

→

Row: 3
Column: 3

4	3	4
2	4	3
2	3	4

Place each element
according to their positions

Convolution for images (matrices)

1	1	1	0	0
0	1	1	1	0
0	0	1	1	1
0	0	1	1	0
0	1	1	0	0

n=5

*

1	0	1
0	1	0
1	0	1

m=3

→

4	3	4
2	4	3
2	3	4

Q: dimension?

Convolution for images (matrices)

1	1	1	0	0
0	1	1	1	0
0	0	1	1	1
0	0	1	1	0
0	1	1	0	0

$n=5$

*

1	0	1
0	1	0
1	0	1

$m=3$

→

4	3	4
2	4	3
2	3	4

$n-m+1=3$

Convolution for images (matrices)

1	1	1	0	0
0	1	1	1	0
0	0	1	1	1
0	0	1	1	0
0	1	1	0	0

$n=5$

One matrix

*

1	0	1
0	1	0
1	0	1

$m=3$

One matrix

→

4	3	4
2	4	3
2	3	4

$n-m+1=3$

One matrix

Convolution for images (matrices)

1	1	1	0	0
0	1	1	1	0
0	0	1	1	1
0	0	1	1	0
0	1	1	0	0

$n=5$

One matrix

*

1	0	1
0	1	0
1	0	1

$m=3$

One matrix

→


4	3	4
2	4	3
2	3	4

$n-m+1=3$

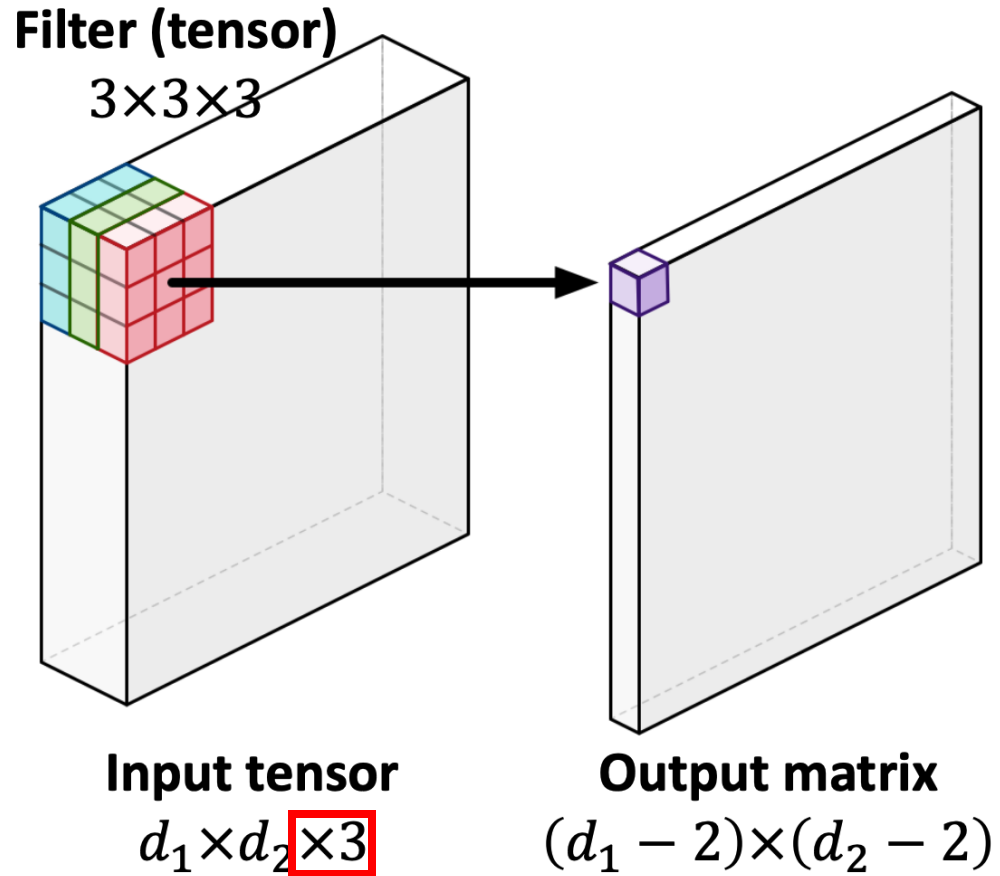
One matrix

One input matrix * one filter → one feature matrix


Convolution for images (**tensors**)



Convolution for images (**tensors**)




Convolution for images (**tensors**)




Q: why we care about tensors?


Convolution for images (**tensors**)




Q: why we care about tensors?



Convolution for images (**tensors**)




Q: why we care about tensors?




Reason 1:
RGB channels are more common

Convolution for images (**tensors**)



Q: why we care about tensors?



Reason 1:
RGB channels are more common
Each channel \rightarrow a matrix

LeNet-5 in 1999




Fig. 1. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

LeNet-5 in 1999




Fig. 1. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

LeCun, Yann, Patrick Haffner, Léon Bottou, and Yoshua Bengio. "Object recognition with gradient-based learning." In *Shape, contour and grouping in computer vision*, pp. 319-345. Springer, Berlin, Heidelberg, 1999.

LeNet-5 in 1999





Fig. 1. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

Subsampling operations

- Max pooling



Subsampling operations

- Max pooling

29	15	28	184
0	100	70	38
12	12	7	2
12	12	45	6

2 x 2
pool size

100	184
12	45

Q: what does max
Pooling really do?

Subsampling operations

- Max pooling
- Average pooling

29	15	28	184
0	100	70	38
12	12	7	2
12	12	45	6

2 x 2
pool size

100	184
12	45

31	15	28	184
0	100	70	38
12	12	7	2
12	12	45	6

2 x 2
pool size

36	80
12	15

Subsampling operations

- Max pooling
- Average pooling

Q: what does average Pooling really do?

29	15	28	184
0	100	70	38
12	12	7	2
12	12	45	6

2 x 2
pool size

100	184
12	45

31	15	28	184
0	100	70	38
12	12	7	2
12	12	45	6

2 x 2
pool size

36	80
12	15

Subsampling operations

- Max pooling
- Average pooling

29	15	28	184
0	100	70	38
12	12	7	2
12	12	45	6

2 x 2
pool size

100	184
12	45

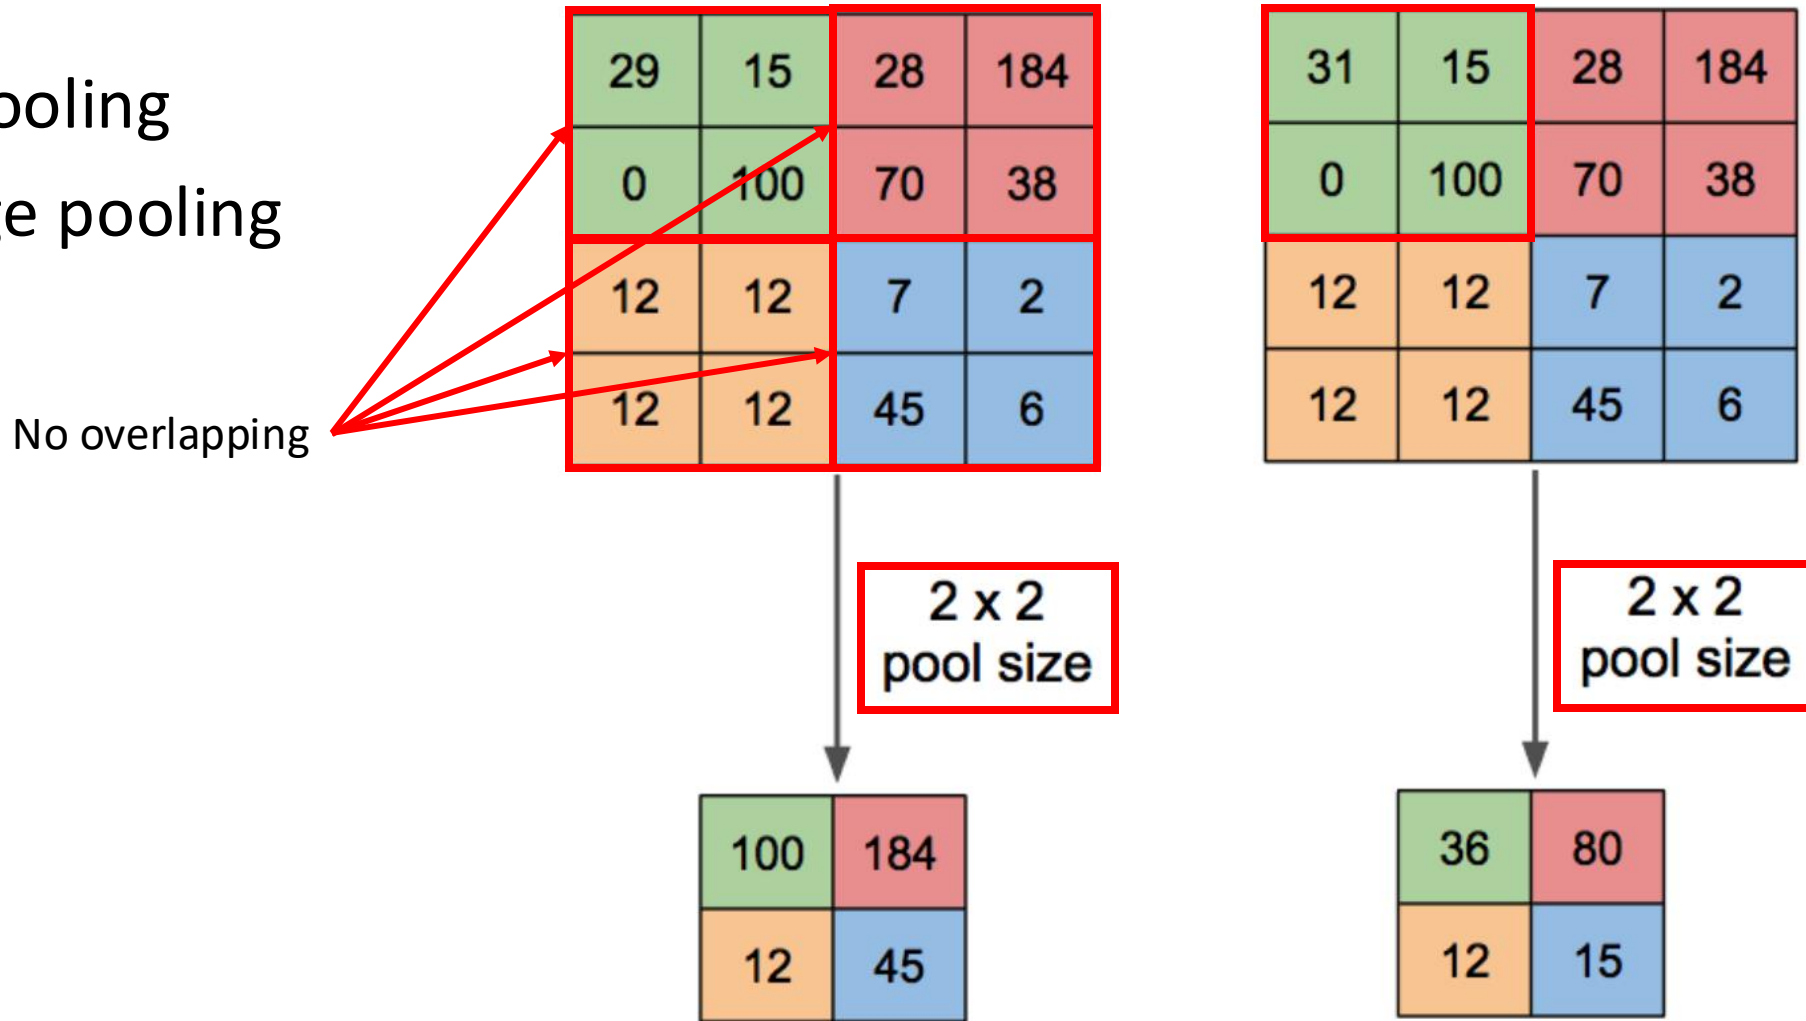
31	15	28	184
0	100	70	38
12	12	7	2
12	12	45	6

2 x 2
pool size

36	80
12	15

Subsampling operations

- Max pooling
- Average pooling



Subsampling operations

- Max pooling
- Average pooling

No overlapping
(stride=2*2)

29	15	28	184
0	100	70	38
12	12	7	2
12	12	45	6

2 x 2
pool size

100	184
12	45

31	15	28	184
0	100	70	38
12	12	7	2
12	12	45	6

2 x 2
pool size

36	80
12	15

Subsampling operations

- Max pooling
- Average pooling

No overlapping
(stride=2*2)

Row stride = 2
Column stride = 2

29	15	28	184
0	100	70	38
12	12	7	2
12	12	45	6

2 x 2
pool size

100	184
12	45

31	15	28	184
0	100	70	38
12	12	7	2
12	12	45	6

2 x 2
pool size

36	80
12	15

Subsampling operations

- Max pooling
- Average pooling

No overlapping
(stride=2*2)

Row stride = 2
Column stride = 2

Q: Why pooling?
Connection to subsampling?

29	15	28	184
0	100	70	38
12	12	7	2
12	12	45	6

2 x 2
pool size

100	184
12	45

31	15	28	184
0	100	70	38
12	12	7	2
12	12	45	6

2 x 2
pool size

36	80
12	15

Subsampling operations

- Max pooling
- Average pooling

No overlapping
(stride=2*2)

Row stride = 2
Column stride = 2

Q: Why pooling?
Connection to subsampling?

$4 \times 4 \rightarrow 2 \times 2$

29	15	28	184
0	100	70	38
12	12	7	2
12	12	45	6

2 x 2
pool size

100	184
12	45

31	15	28	184
0	100	70	38
12	12	7	2
12	12	45	6

2 x 2
pool size

36	80
12	15

Subsampling operations

- Max pooling
- Average pooling

No overlapping
(stride=2*2)

Row stride = 2
Column stride = 2

Q: Why pooling?
Connection to subsampling?

$4 \times 4 \rightarrow 2 \times 2$

Dimension reduced

29	15	28	184
0	100	70	38
12	12	7	2
12	12	45	6

2 x 2
pool size

100	184
12	45

31	15	28	184
0	100	70	38
12	12	7	2
12	12	45	6

2 x 2
pool size

36	80
12	15

Subsampling operations

- Max pooling
- Average pooling

No overlapping
(stride=2*2)

Row stride = 2
Column stride = 2

Q: Why pooling?
Connection to subsampling?

$4 \times 4 \rightarrow 2 \times 2$

Dimension reduced

29	15	28	184
0	100	70	38
12	12	7	2
12	12	45	6

2 x 2
pool size

100	184
12	45

Use one to represent all

31	15	28	184
0	100	70	38
12	12	7	2
12	12	45	6

2 x 2
pool size

36	80
12	15

LeNet-5 in 1999

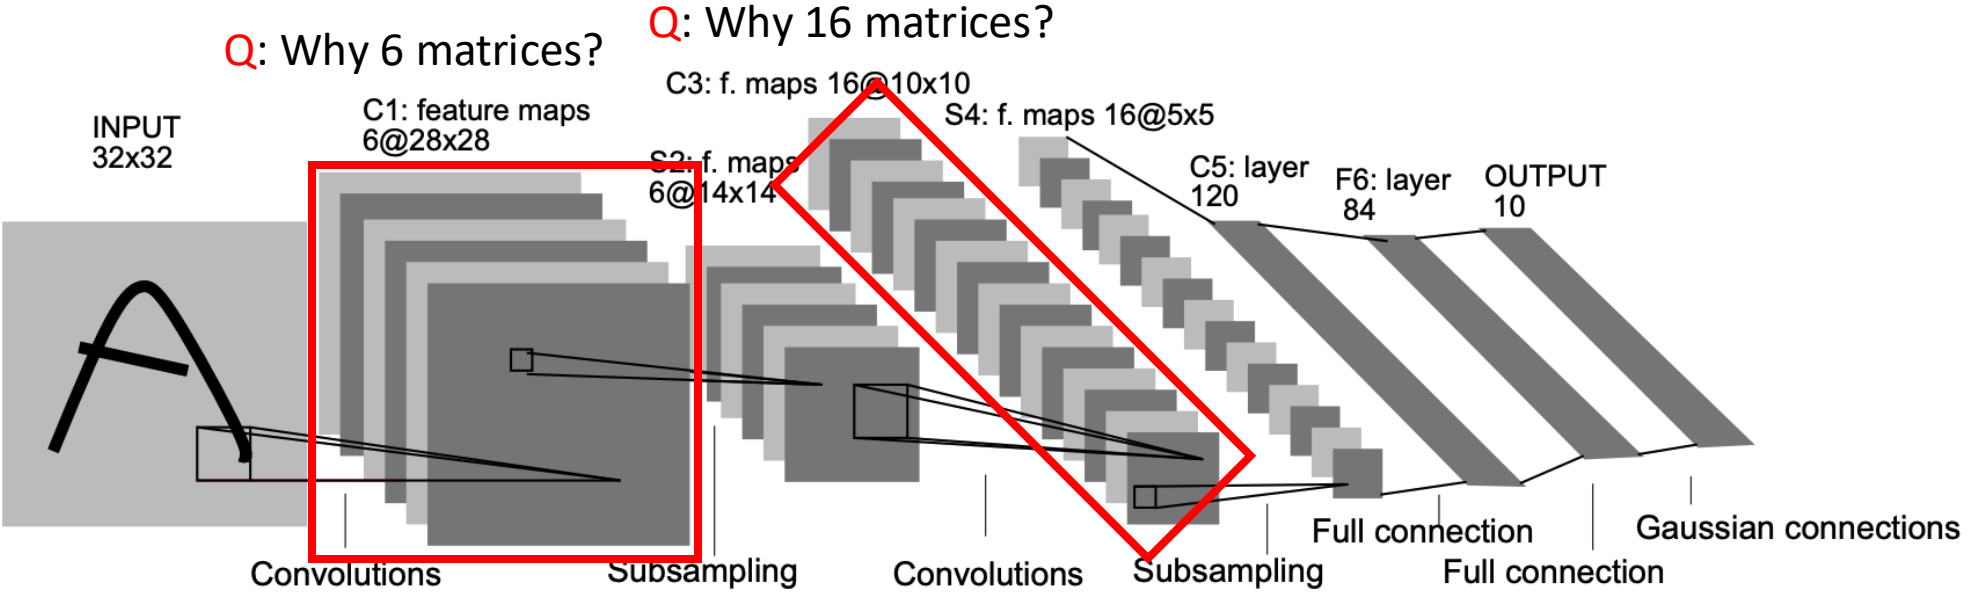


Fig. 1. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

LeCun, Yann, Patrick Haffner, Léon Bottou, and Yoshua Bengio. "Object recognition with gradient-based learning." In *Shape, contour and grouping in computer vision*, pp. 319-345. Springer, Berlin, Heidelberg, 1999.

LeNet-5 in 1999




Fig. 1. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

LeCun, Yann, Patrick Haffner, Léon Bottou, and Yoshua Bengio. "Object recognition with gradient-based learning." In *Shape, contour and grouping in computer vision*, pp. 319-345. Springer, Berlin, Heidelberg, 1999.

LeNet-5 in 1999




Fig. 1. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

LeCun, Yann, Patrick Haffner, Léon Bottou, and Yoshua Bengio. "Object recognition with gradient-based learning." In *Shape, contour and grouping in computer vision*, pp. 319-345. Springer, Berlin, Heidelberg, 1999.

LeNet-5 in 1999


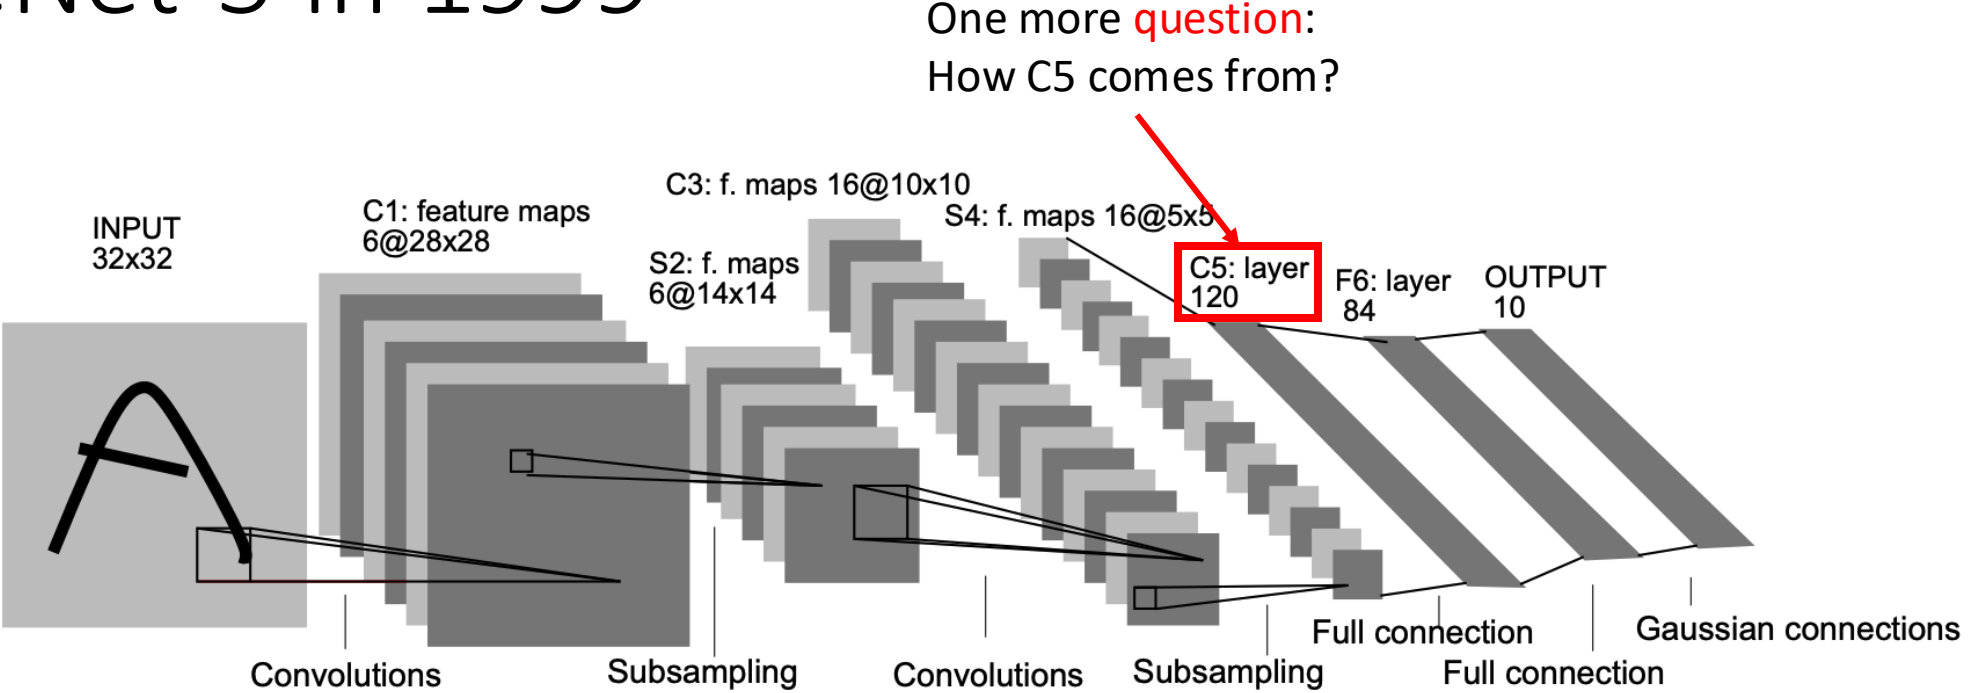


Fig. 1. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

Subsampling layer: max/average pooling

LeCun, Yann, Patrick Haffner, Léon Bottou, and Yoshua Bengio. "Object recognition with gradient-based learning." In *Shape, contour and grouping in computer vision*, pp. 319-345. Springer, Berlin, Heidelberg, 1999.

LeNet-5 in 1999



One more **question**:
How C5 comes from?

Fig. 1. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

LeCun, Yann, Patrick Haffner, Léon Bottou, and Yoshua Bengio. "Object recognition with gradient-based learning." In *Shape, contour and grouping in computer vision*, pp. 319-345. Springer, Berlin, Heidelberg, 1999.

LeNet-5 in 1999

One more **question**:
How C5 comes from? Matrices \rightarrow a vector?

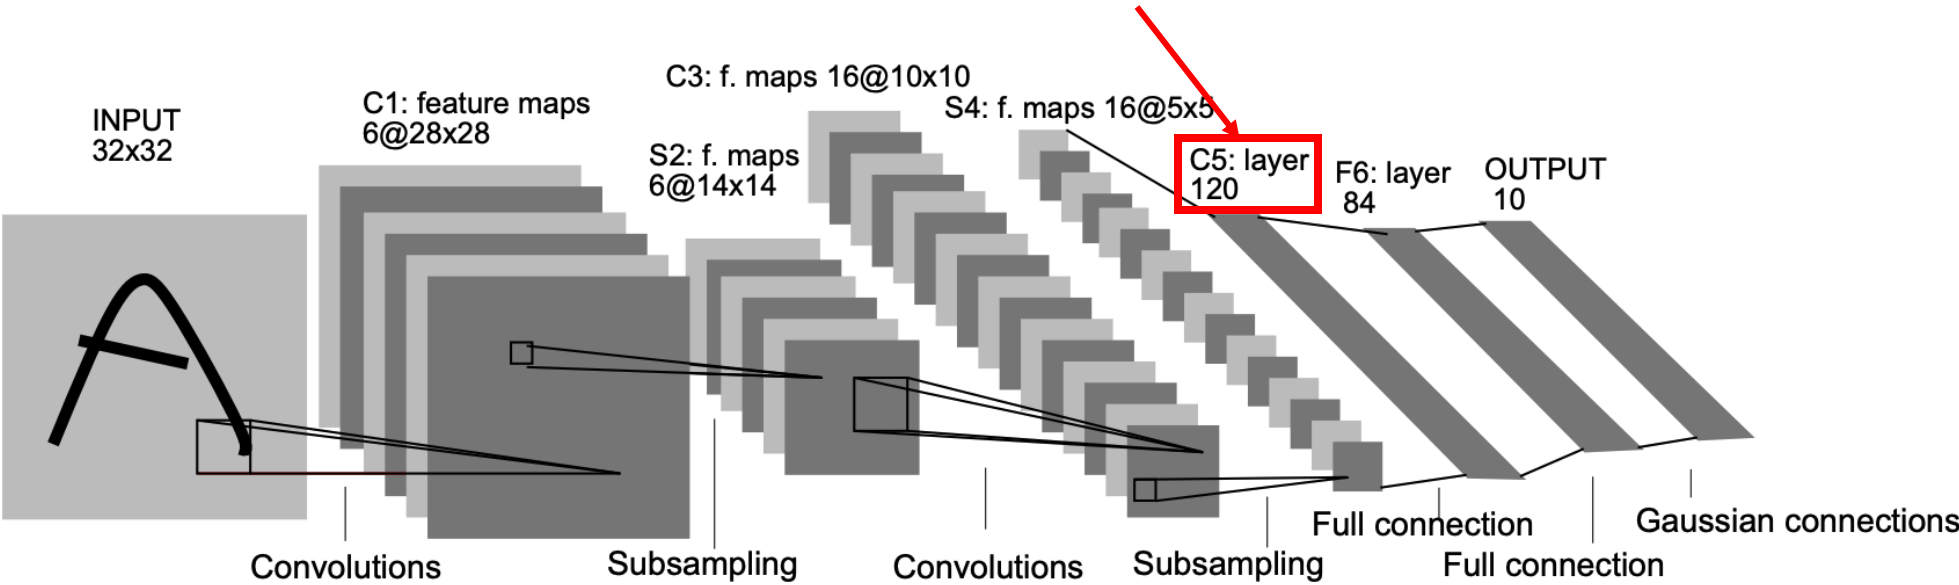


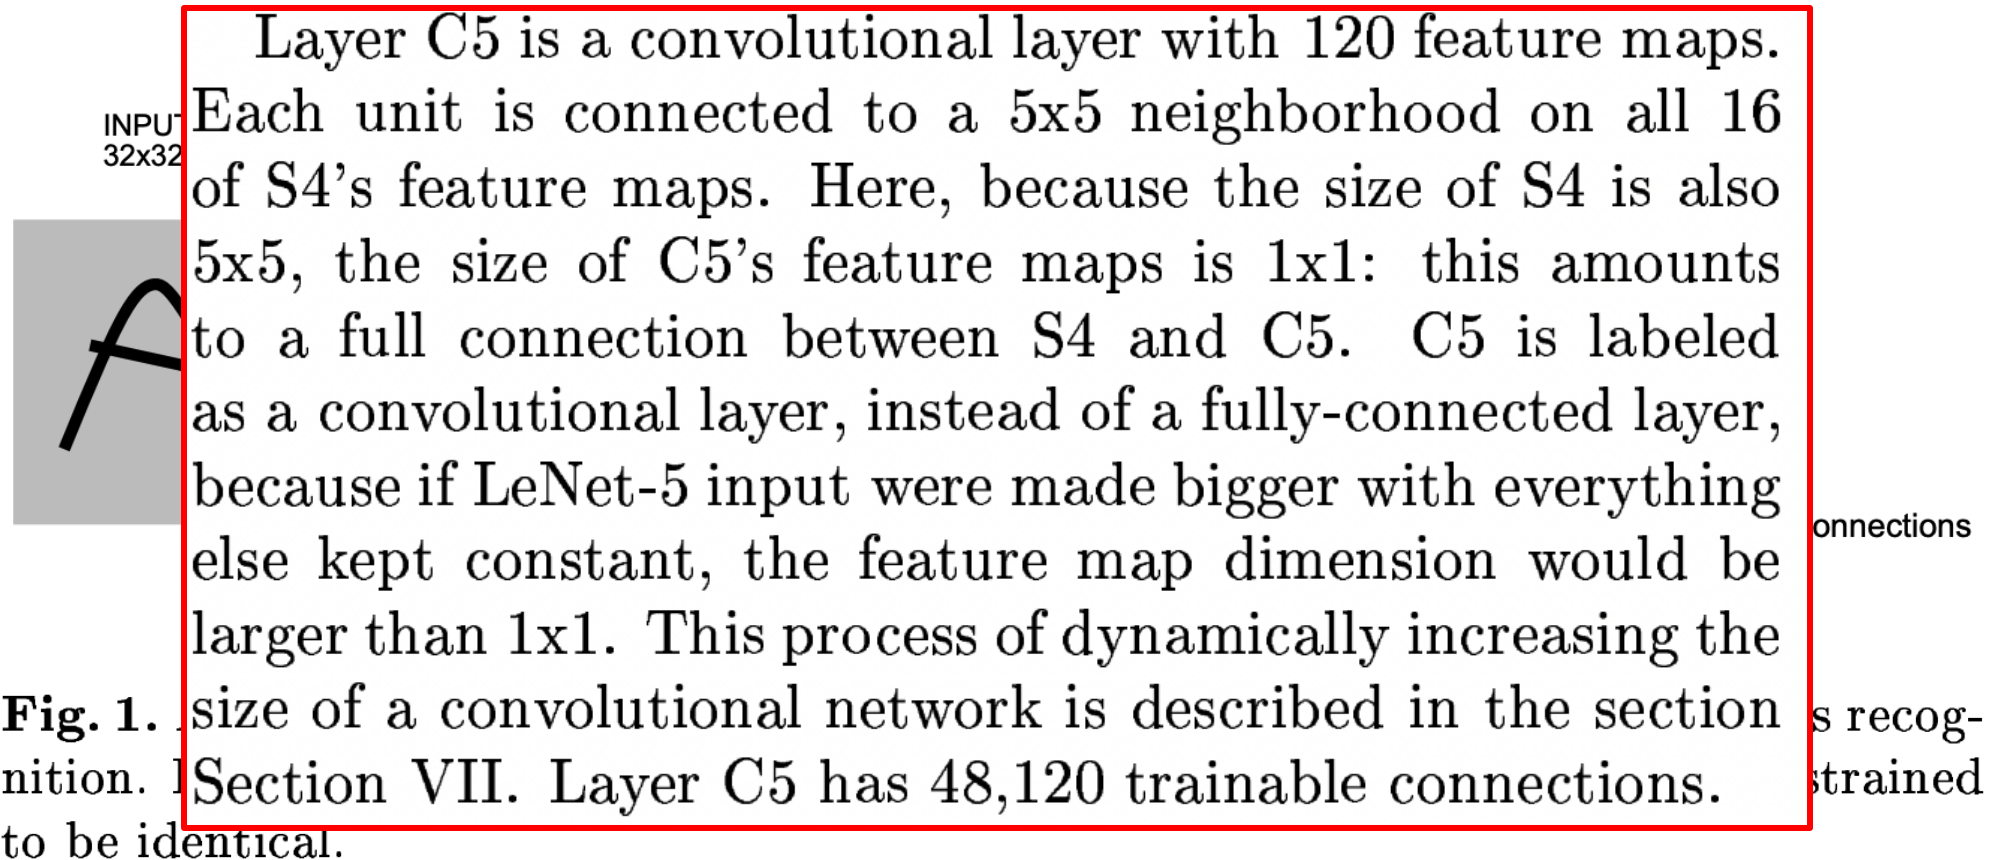
Fig. 1. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

LeCun, Yann, Patrick Haffner, Léon Bottou, and Yoshua Bengio. "Object recognition with gradient-based learning." In *Shape, contour and grouping in computer vision*, pp. 319-345. Springer, Berlin, Heidelberg, 1999.

LeNet-5 in 1999

One more **question**:

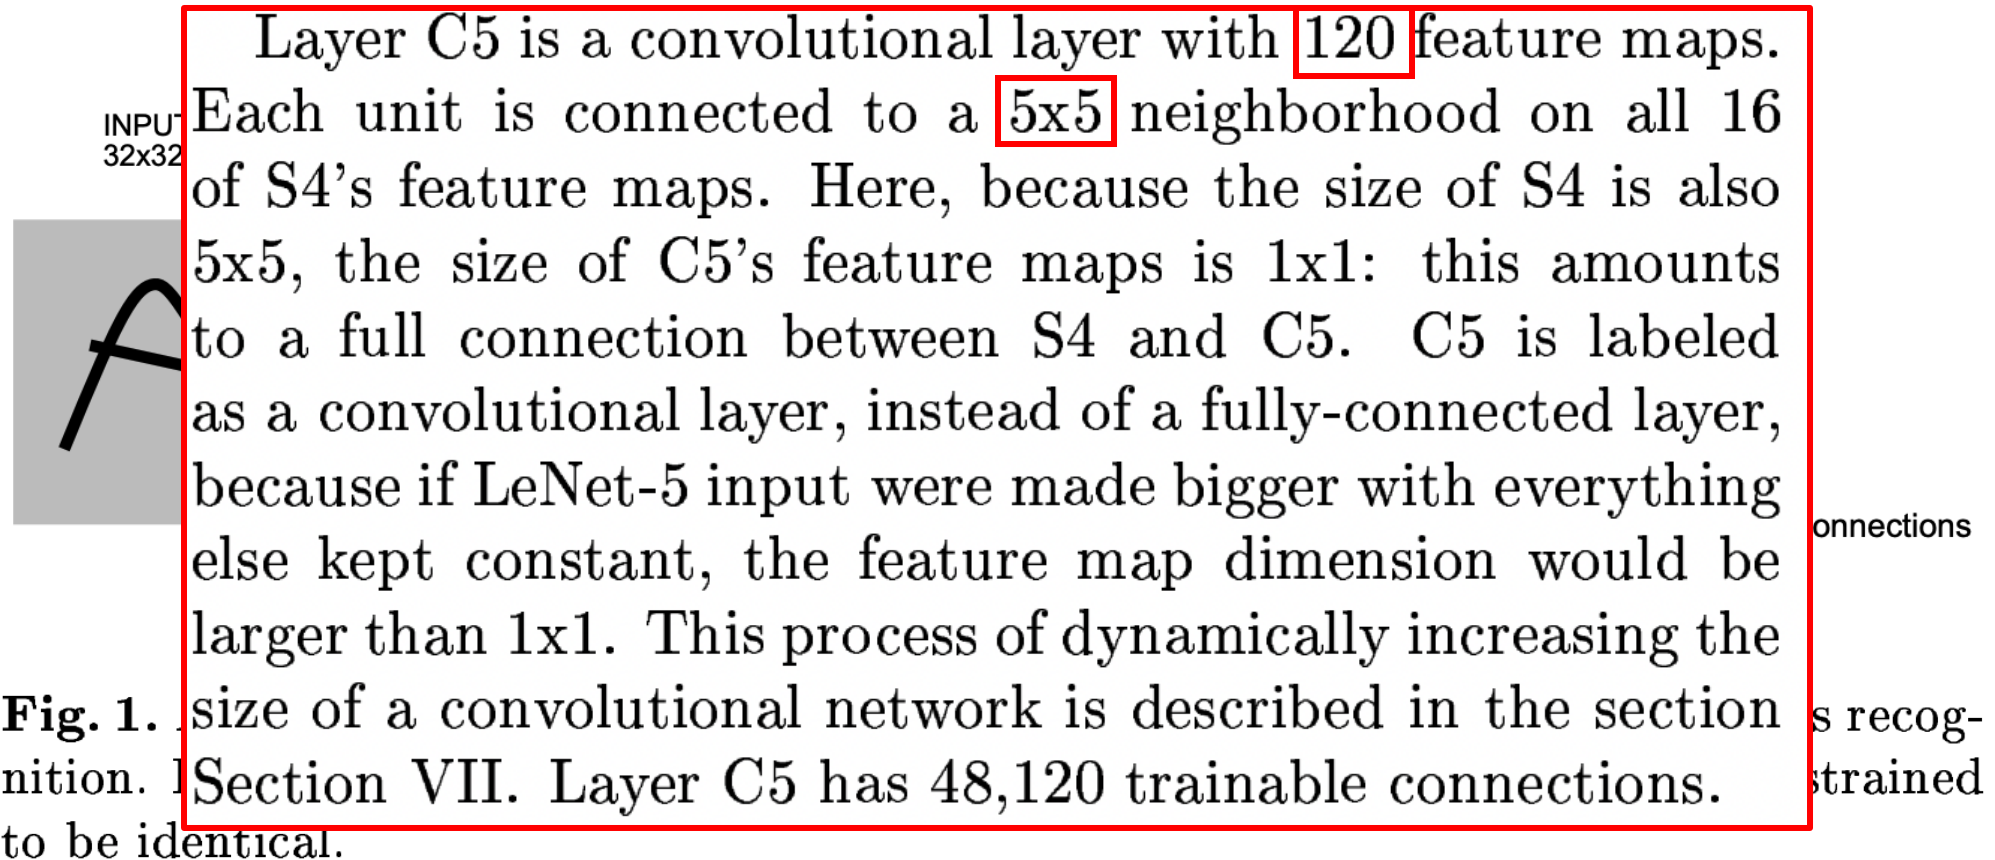
Where C5 comes from? 16 matrices \rightarrow a 120d vector?



LeNet-5 in 1999

One more **question**:

Where C5 comes from? 16 matrices \rightarrow a 120d vector?



Operations with convolution layers

- Padding
- Pooling layers for arbitrary input resolution

Operations with convolution layers

- Padding: convolution operation reduces the size of feature maps

Operations with convolution layers

- Padding: convolution operation reduces the size of feature maps

1	1	1	0	0
0	1	1	1	0
0	0	1	1	1
0	0	1	1	0
0	1	1	0	0

$n=5$

*

1	0	1
0	1	0
1	0	1

$m=3$


→

4	3	4
2	4	3
2	3	4

$n-m+1=3$

Operations with convolution layers


- Padding: convolution operation reduces the size of feature maps



If $m > 1 \rightarrow ??$

Operations with convolution layers


- Padding: convolution operation reduces the size of feature maps



If $m > 1 \rightarrow$ convolution will reduce the dimension

Operations with convolution layers


- Padding: convolution operation reduces the size of feature maps



If $m > 1 \rightarrow$ convolution will reduce the dimension
The input resolution introduces a limits of #convolution layers


Operations with convolution layers

- Padding: convolution operation reduces the size of feature maps



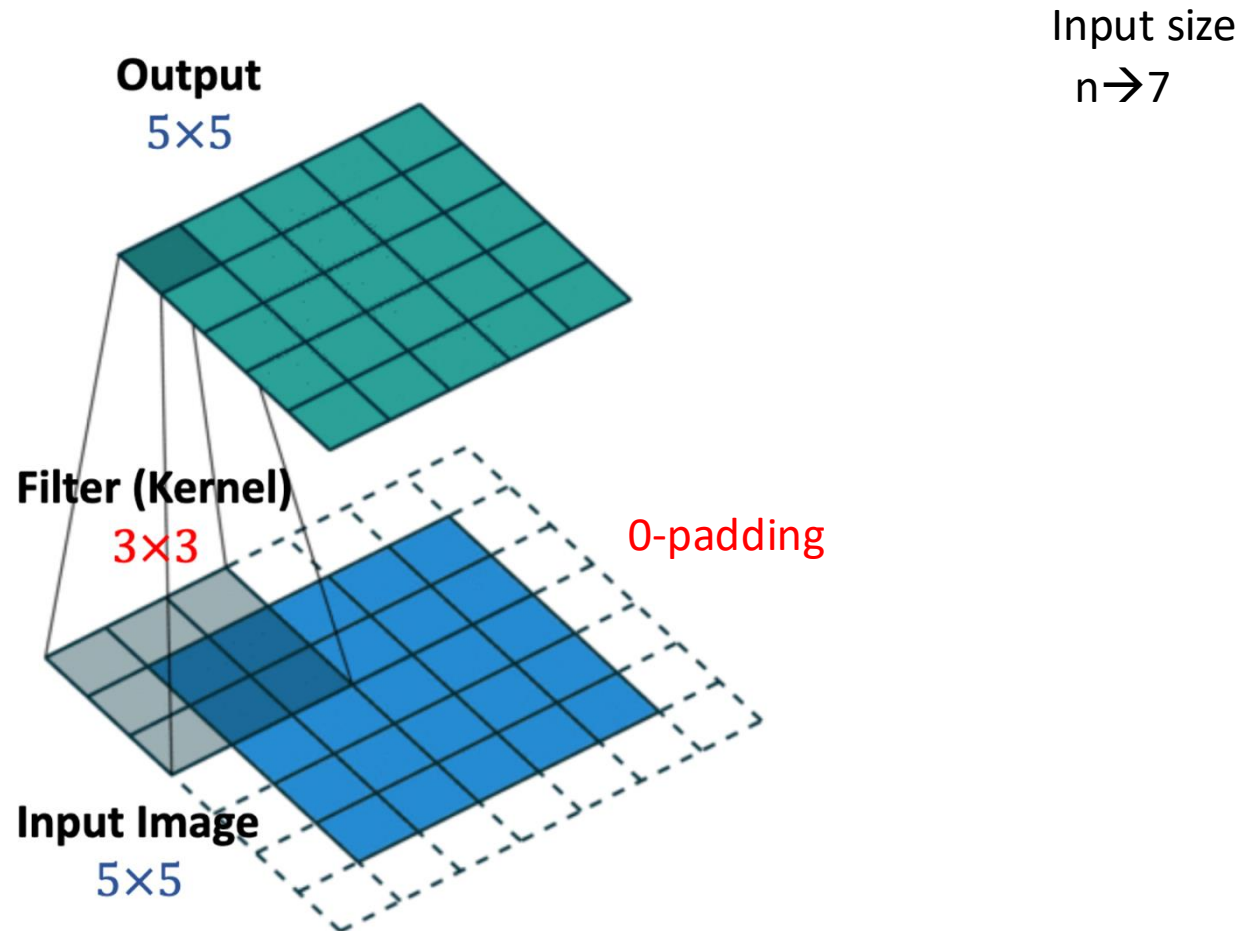
Operations with convolution layers

- Padding: convolution operation reduces the size of feature maps



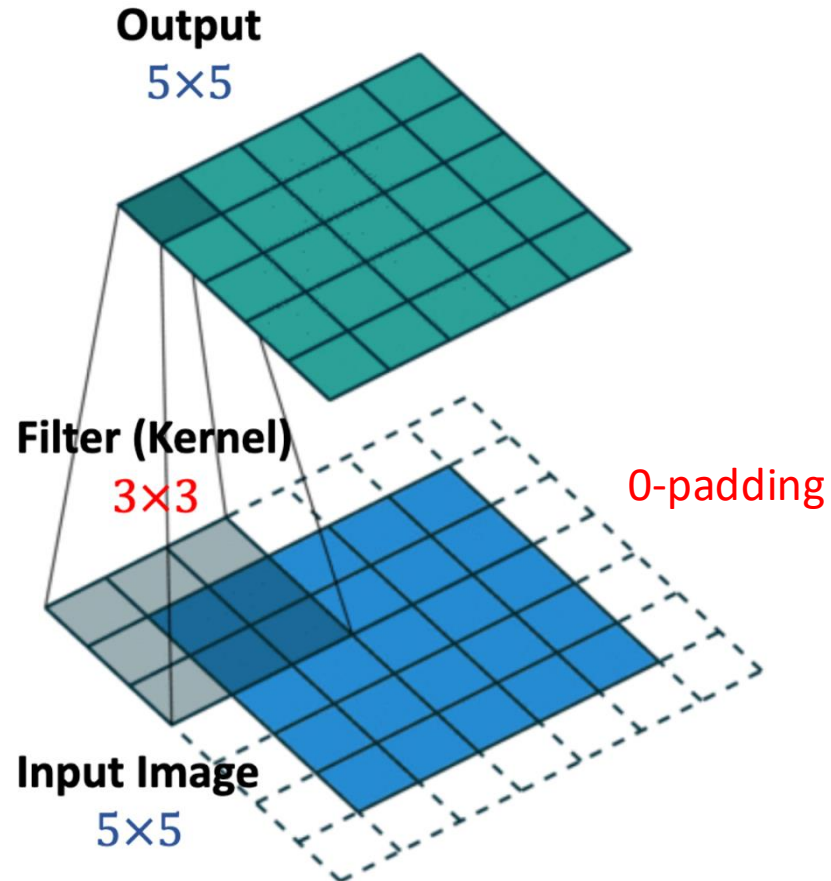
Operations with convolution layers

- Padding: convolution operation reduces the size of feature maps



Operations with convolution layers

- Padding: convolution operation reduces the size of feature maps



$$n \rightarrow 7 \rightarrow n - m + 1 = 7 - 3 + 1 = 5$$

Conclusion:
dimension of feature maps remains the same

Operations with convolution layers

- Padding: convolution operation reduces the size of feature maps
- Pooling layers for an arbitrary input resolution

Input resolution issue




Fig. 1. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

Input resolution issue




Fig. 1. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

Input resolution issue

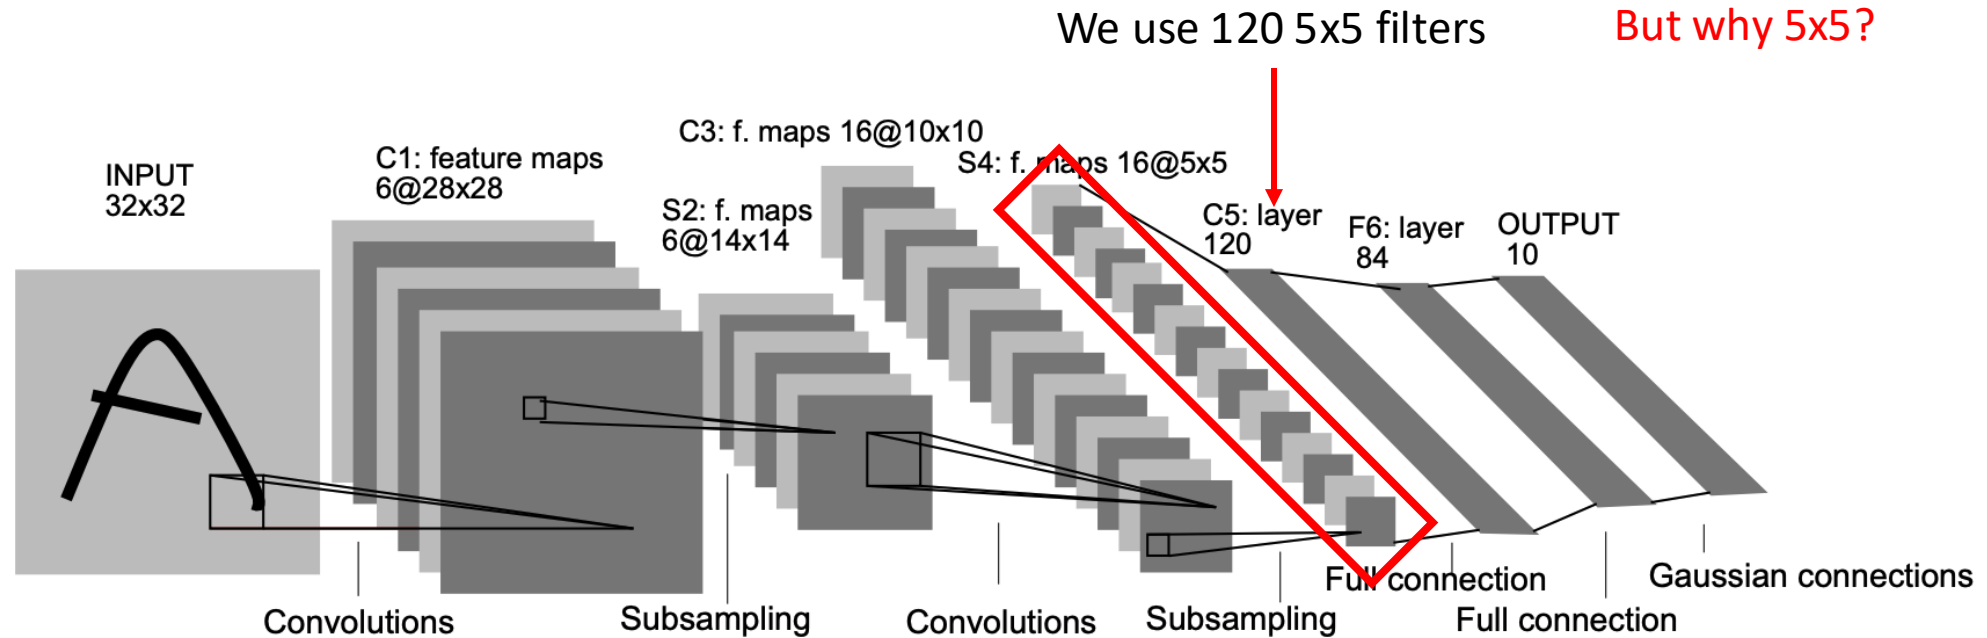


Fig. 1. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

Input resolution issue




Fig. 1. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

Input resolution issue




Fig. 1. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

LeCun, Yann, Patrick Haffner, Léon Bottou, and Yoshua Bengio. "Object recognition with gradient-based learning." In *Shape, contour and grouping in computer vision*, pp. 319-345. Springer, Berlin, Heidelberg, 1999.

Input resolution issue

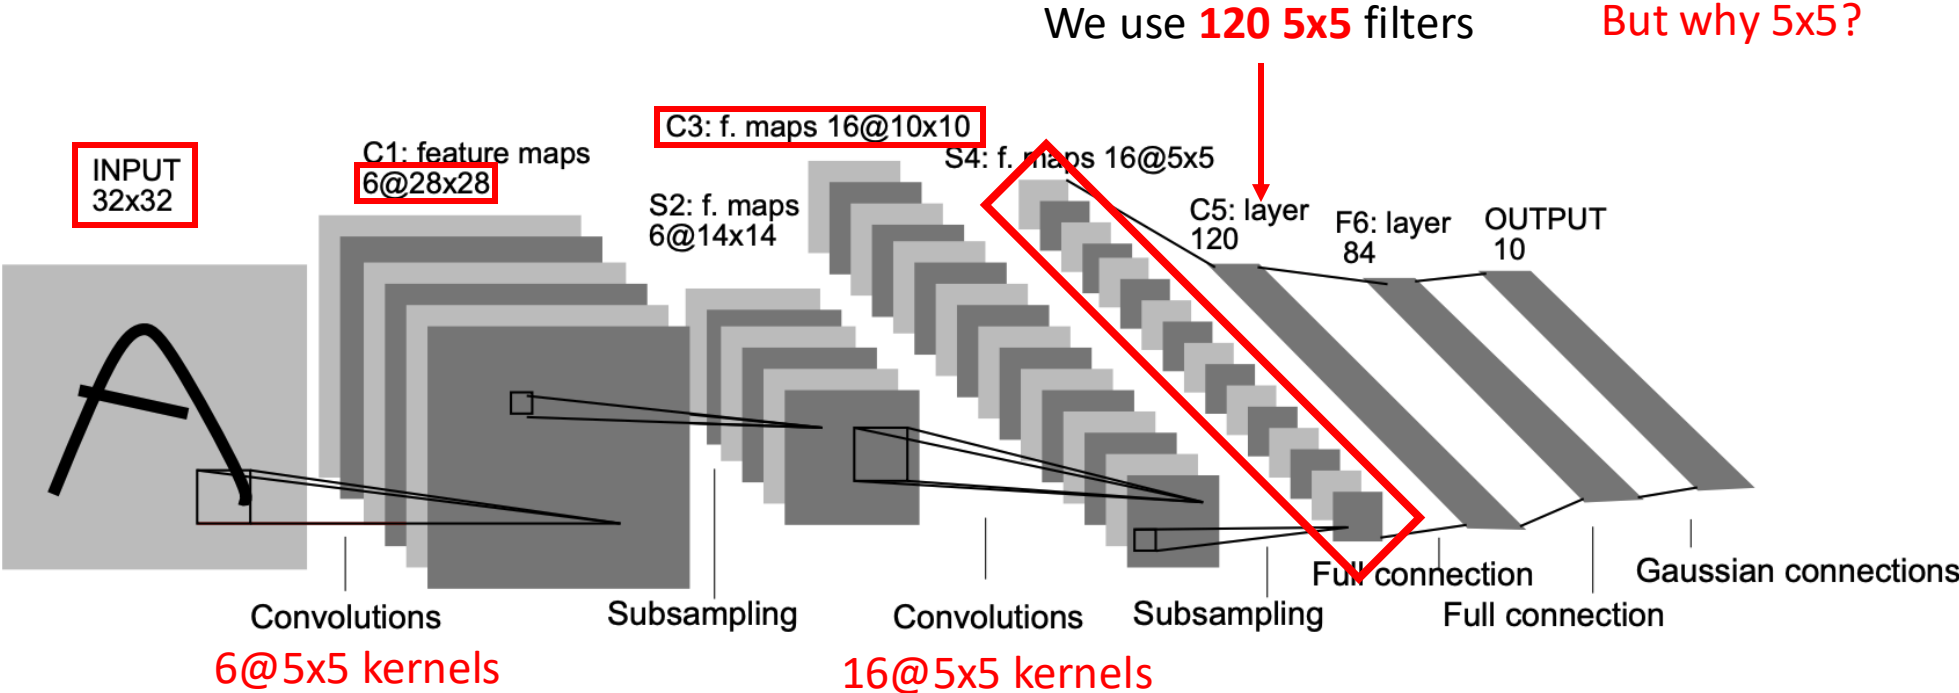


Fig. 1. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

LeCun, Yann, Patrick Haffner, Léon Bottou, and Yoshua Bengio. "Object recognition with gradient-based learning." In *Shape, contour and grouping in computer vision*, pp. 319-345. Springer, Berlin, Heidelberg, 1999.

Input resolution issue

Q: What if 64x64?

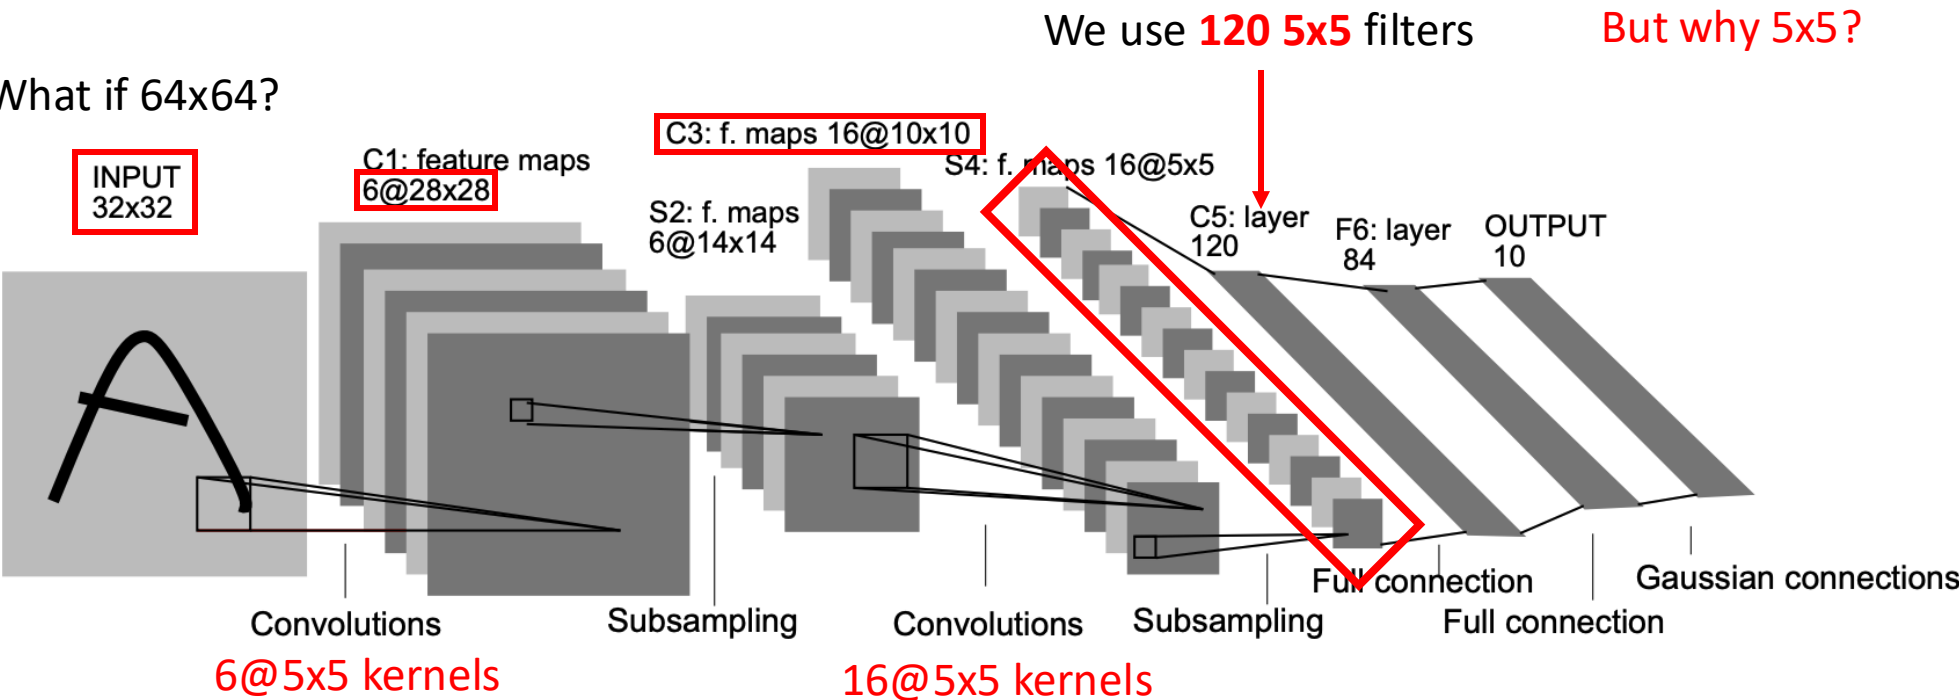



Fig. 1. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.


LeCun, Yann, Patrick Haffner, Léon Bottou, and Yoshua Bengio. "Object recognition with gradient-based learning." In *Shape, contour and grouping in computer vision*, pp. 319-345. Springer, Berlin, Heidelberg, 1999.

Input resolution issue

Original Image




Cropped Image




Input resolution issue


Original Image



Resized Image




Input resolution issue




Input resolution issue

Q: can you understand the following architecture?




Input resolution issue



Any input image must be 224x224

Input resolution issue



Any input image must be 224x224

Q: how to handle an arbitrary resolution?

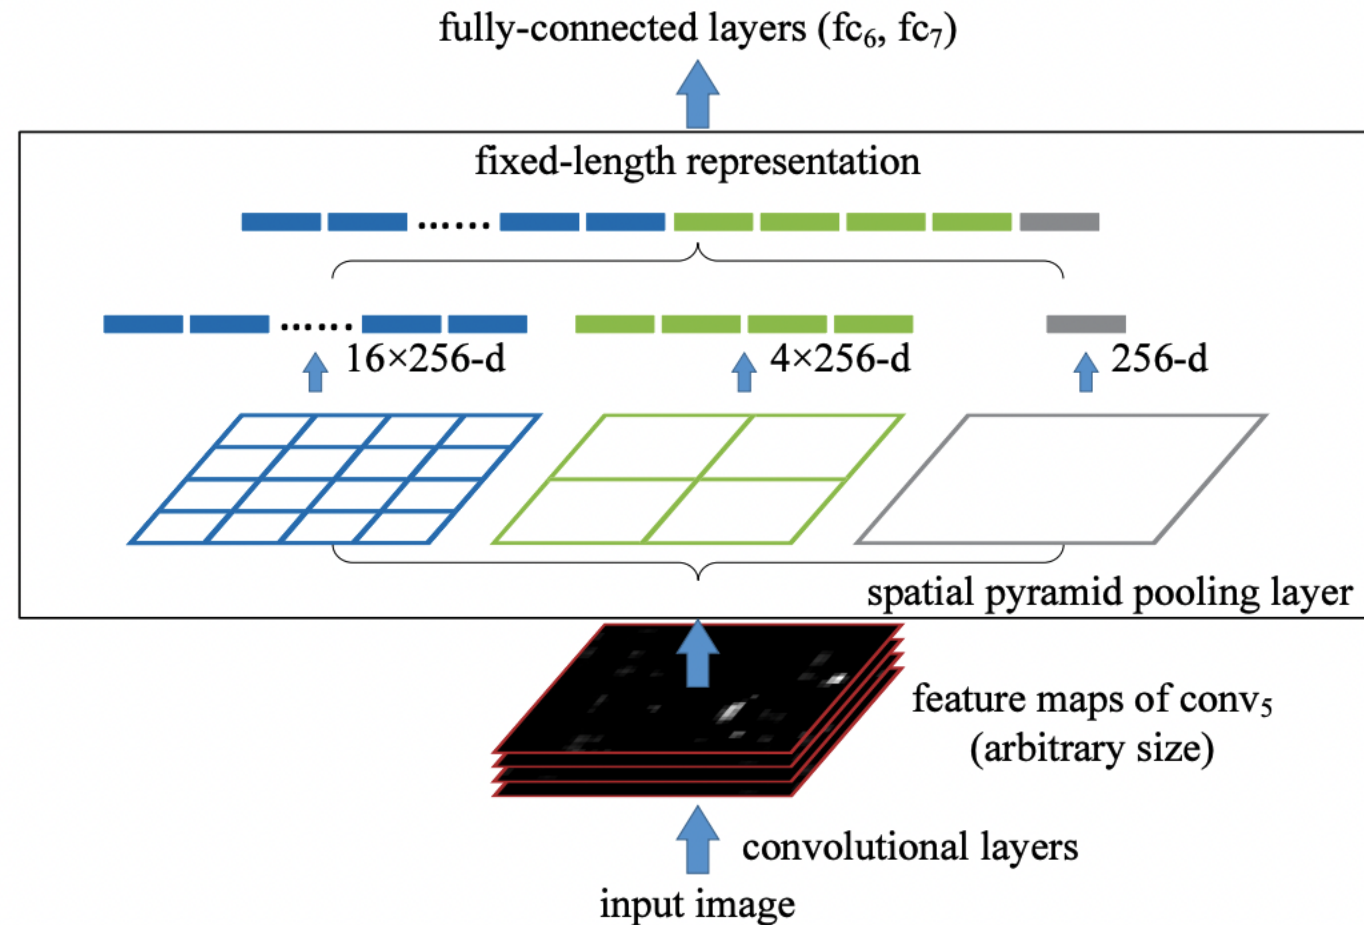
[Alexnet]

Input resolution issue

- Spatial pyramid pooling [pyramid]
- Global average pooling [NIN]
-

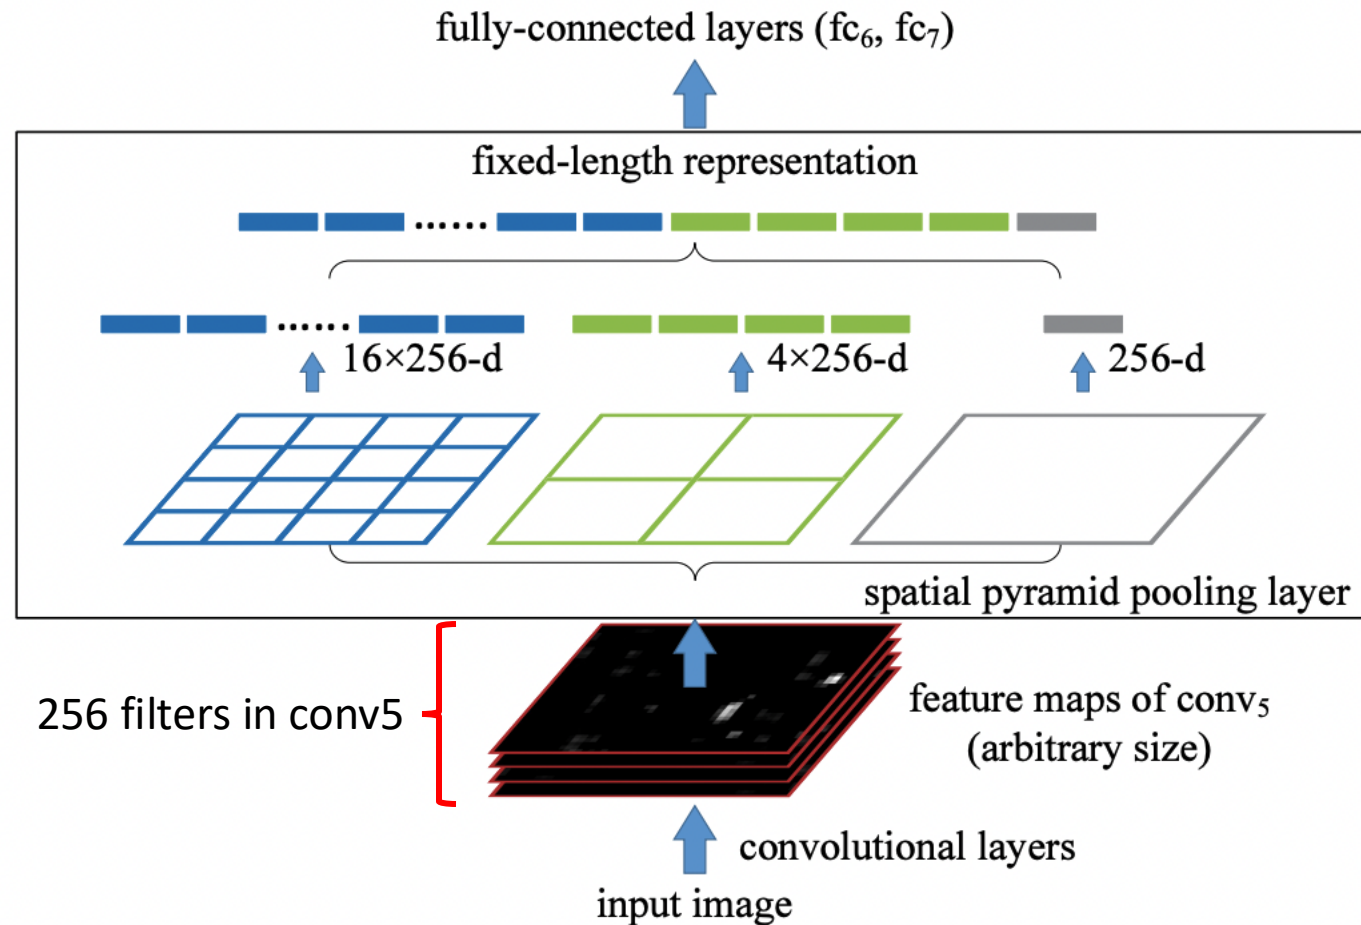
Input resolution issue

- Spatial pyramid pooling



Input resolution issue

- Spatial pyramid pooling




Input resolution issue


- Spatial pyramid pooling

fully-connected layers (fc_6, fc_7)

fixed-length representation




256 filters in conv5
256 feature maps
(matrices)



feature maps of conv5
(arbitrary size)



convolutional layers
input image



Some pooling (max/average)

Input resolution issue

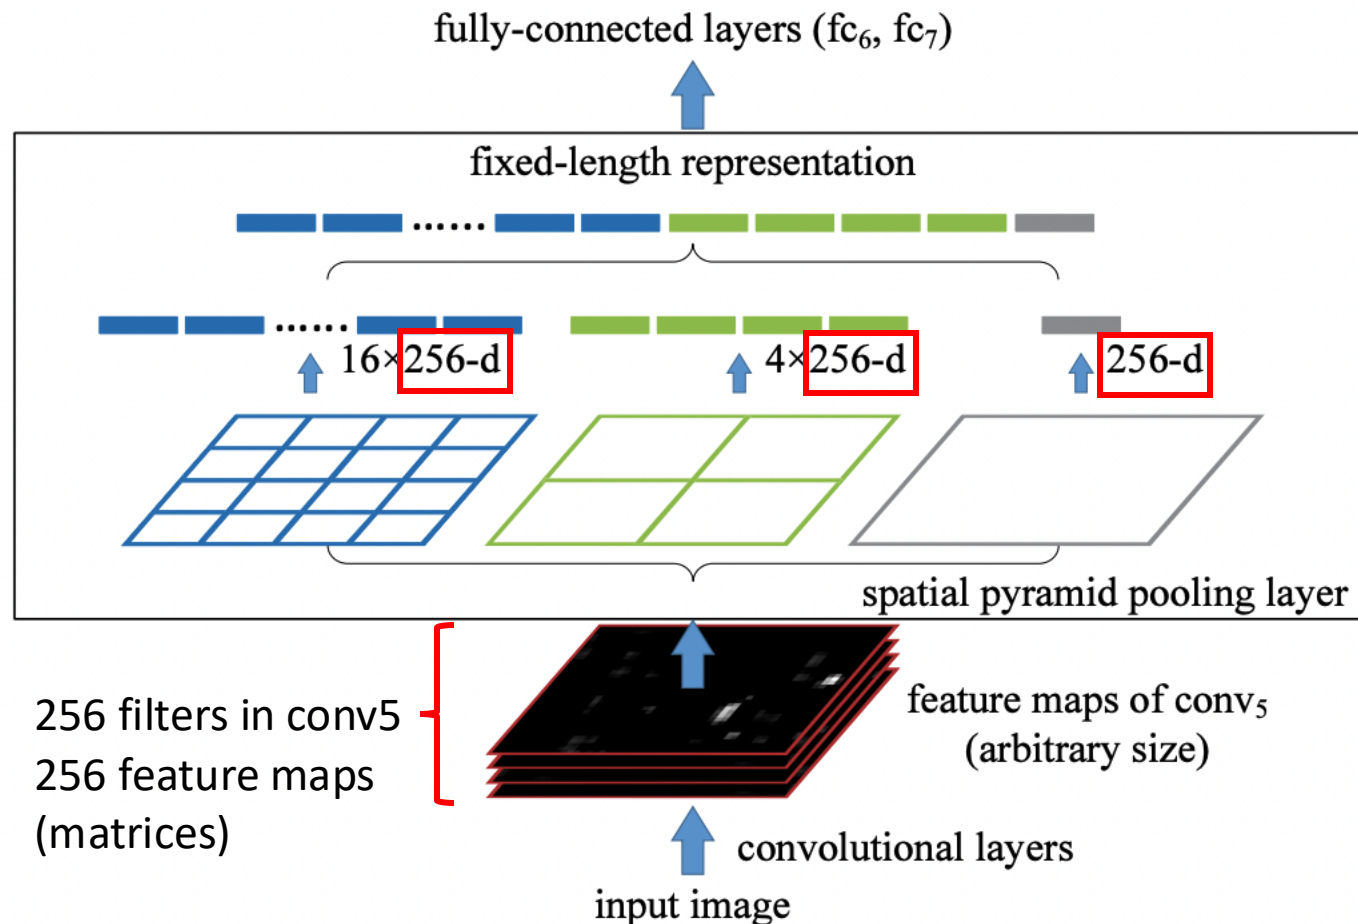
- Spatial pyramid pooling




Some pooling (max/average)

Input resolution issue

- Spatial pyramid pooling



16 numbers



Some pooling (max/average)

Input resolution issue

- Spatial pyramid pooling

fully-connected layers (fc_6, fc_7)

fixed-length representation

Concatenation:
(1+4+16) x 256 numbers

$16 \times 256-d$ $4 \times 256-d$ $256-d$

spatial pyramid pooling layer

256 filters in conv5
256 feature maps
(matrices)

feature maps of conv5
(arbitrary size)

convolutional layers

input image

Input resolution issue

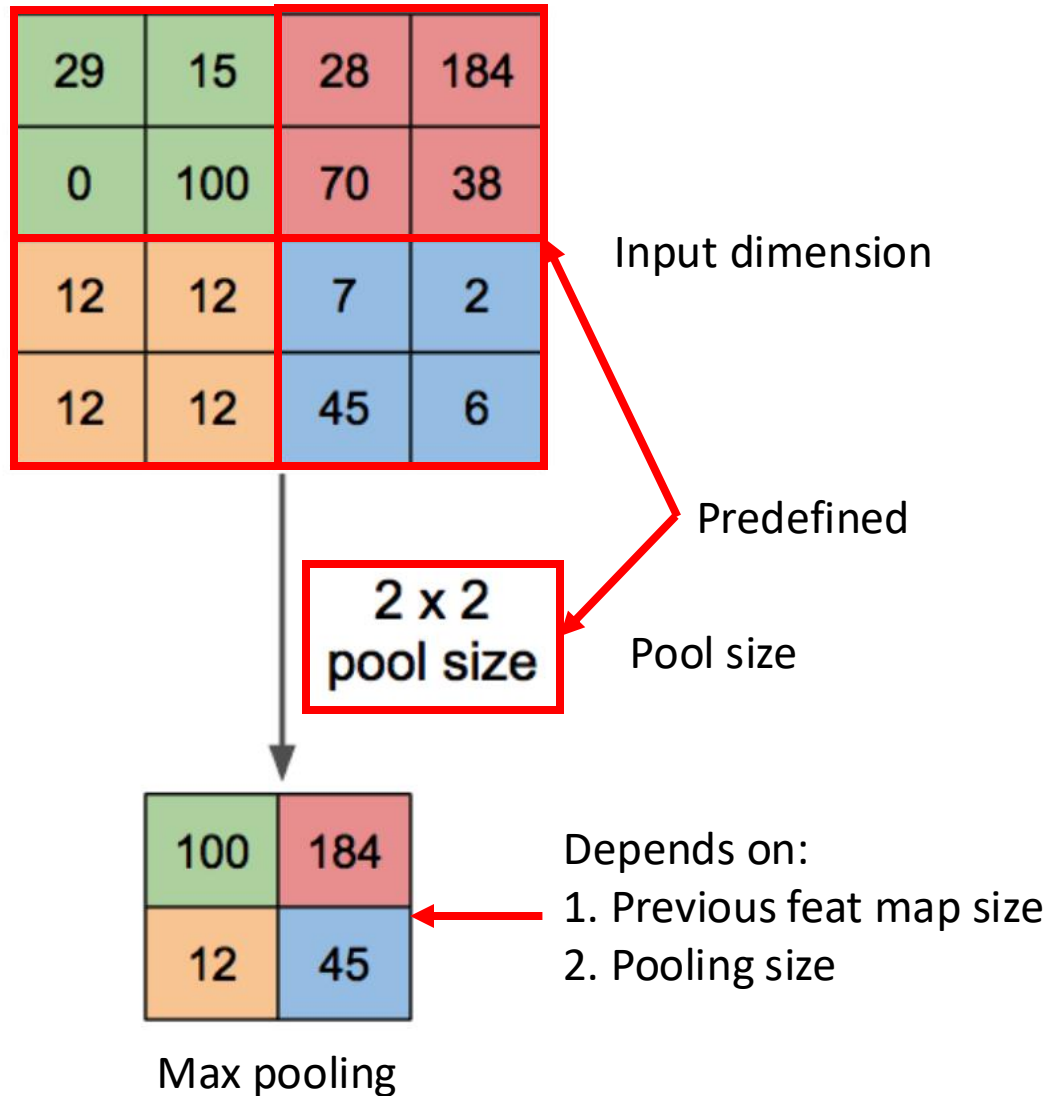
29	15	28	184
0	100	70	38
12	12	7	2
12	12	45	6

2 x 2
pool size

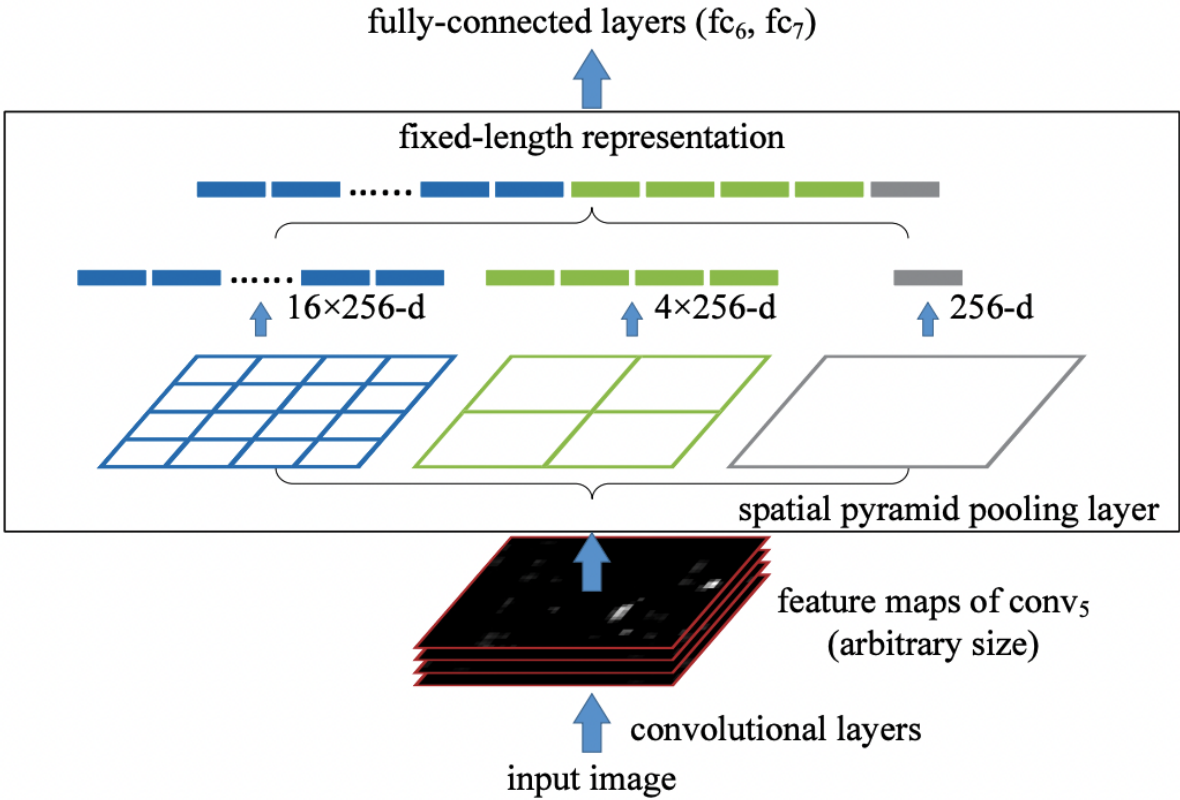

100	184
12	45

Max pooling

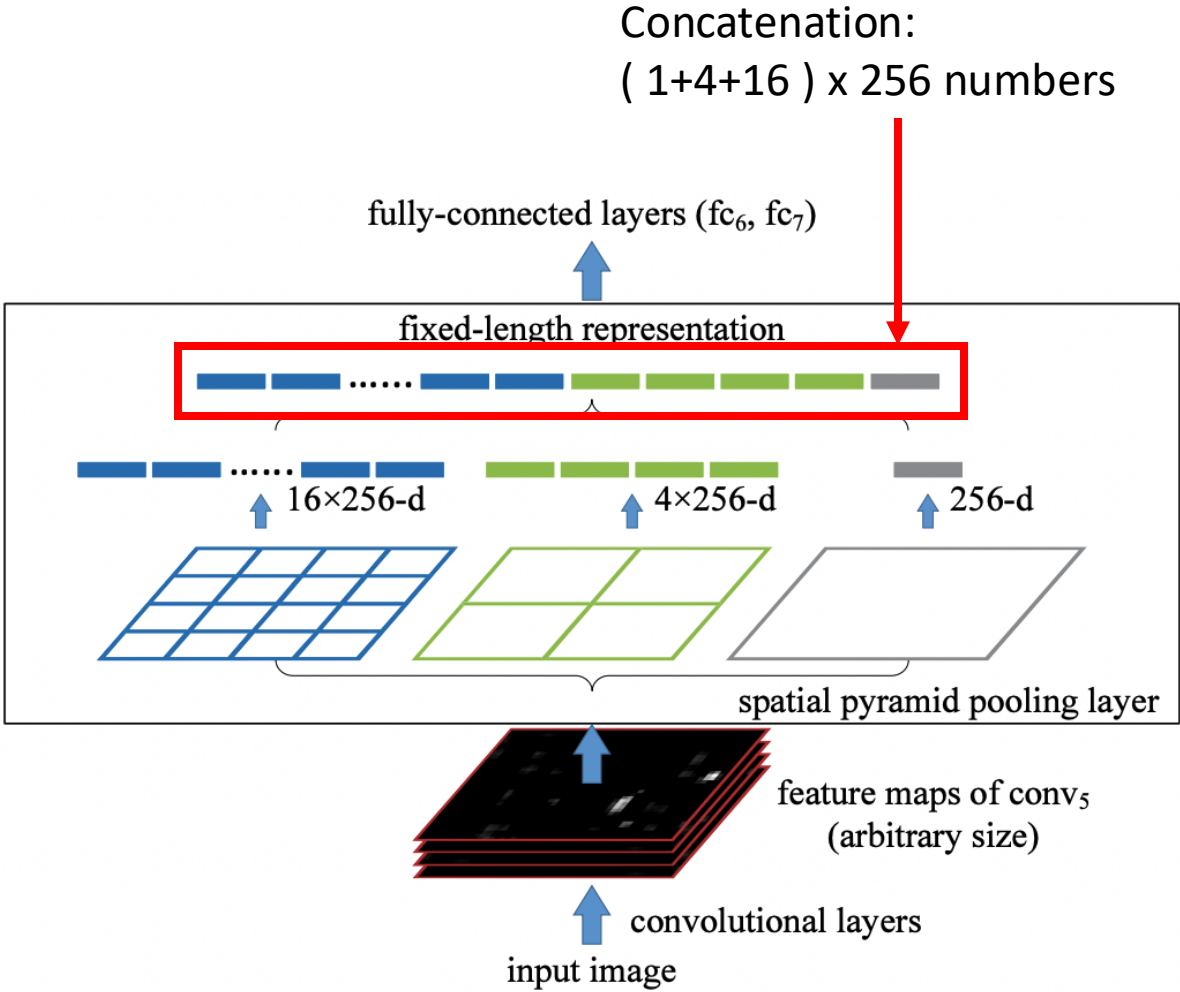
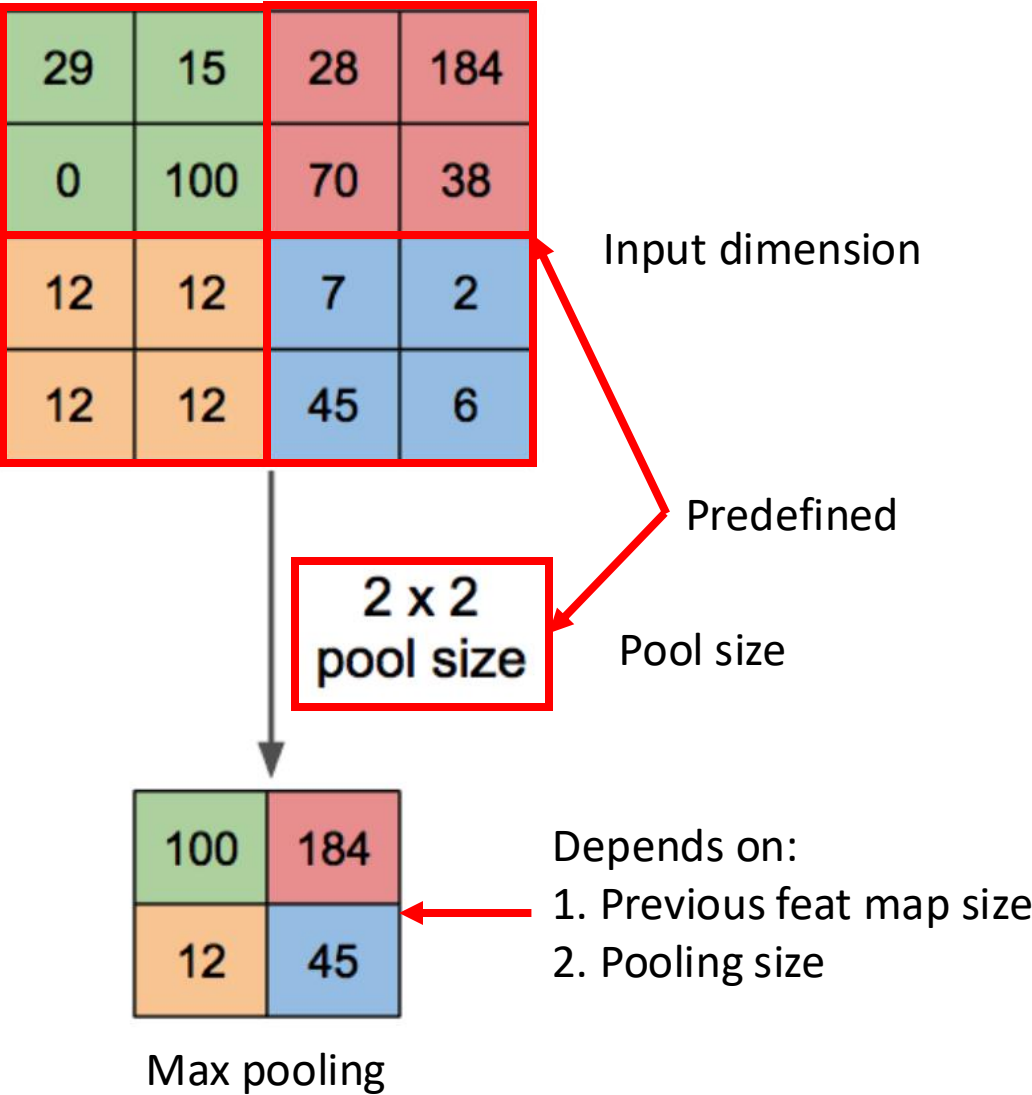
Input resolution issue





Input resolution issue



Input resolution issue



Input resolution issue



Input resolution issue

- Global average pooling




Fig. 1. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

Input resolution issue

- Global average pooling

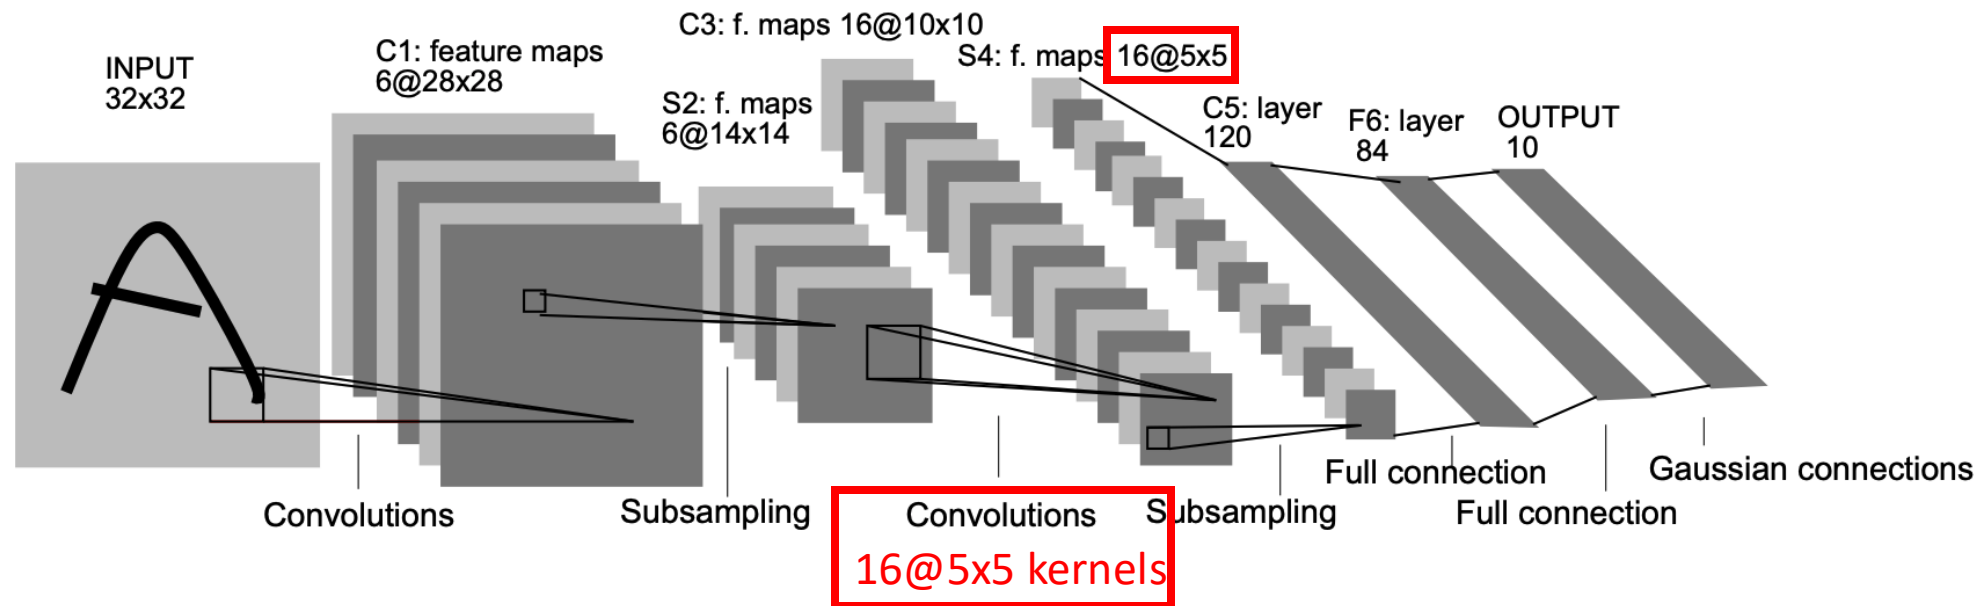


Fig. 1. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

Input resolution issue

- Global average pooling

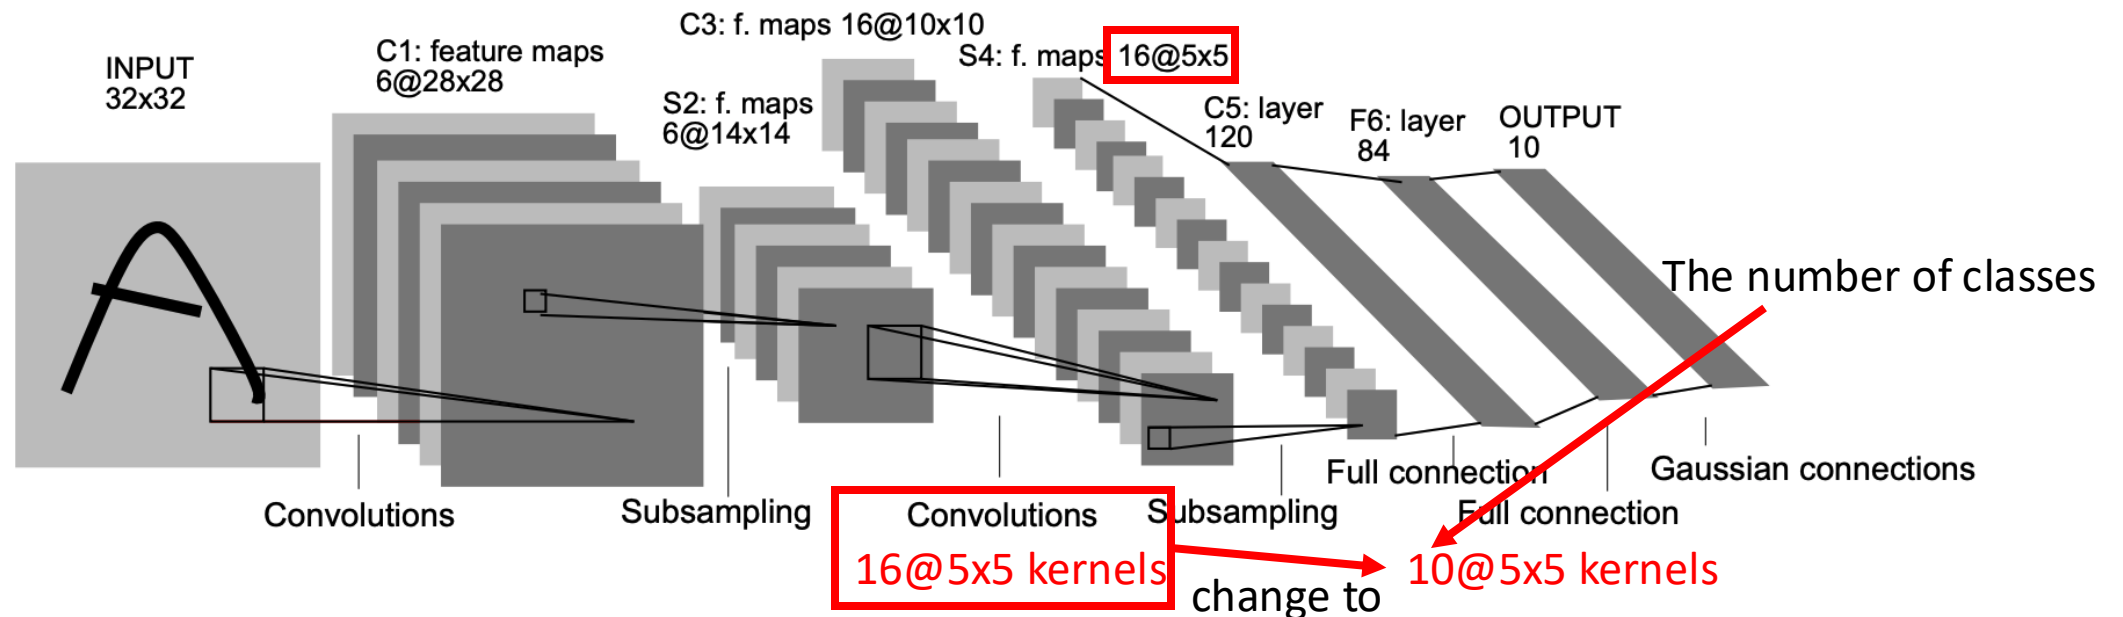


Fig. 1. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

Input resolution issue

- Global average pooling

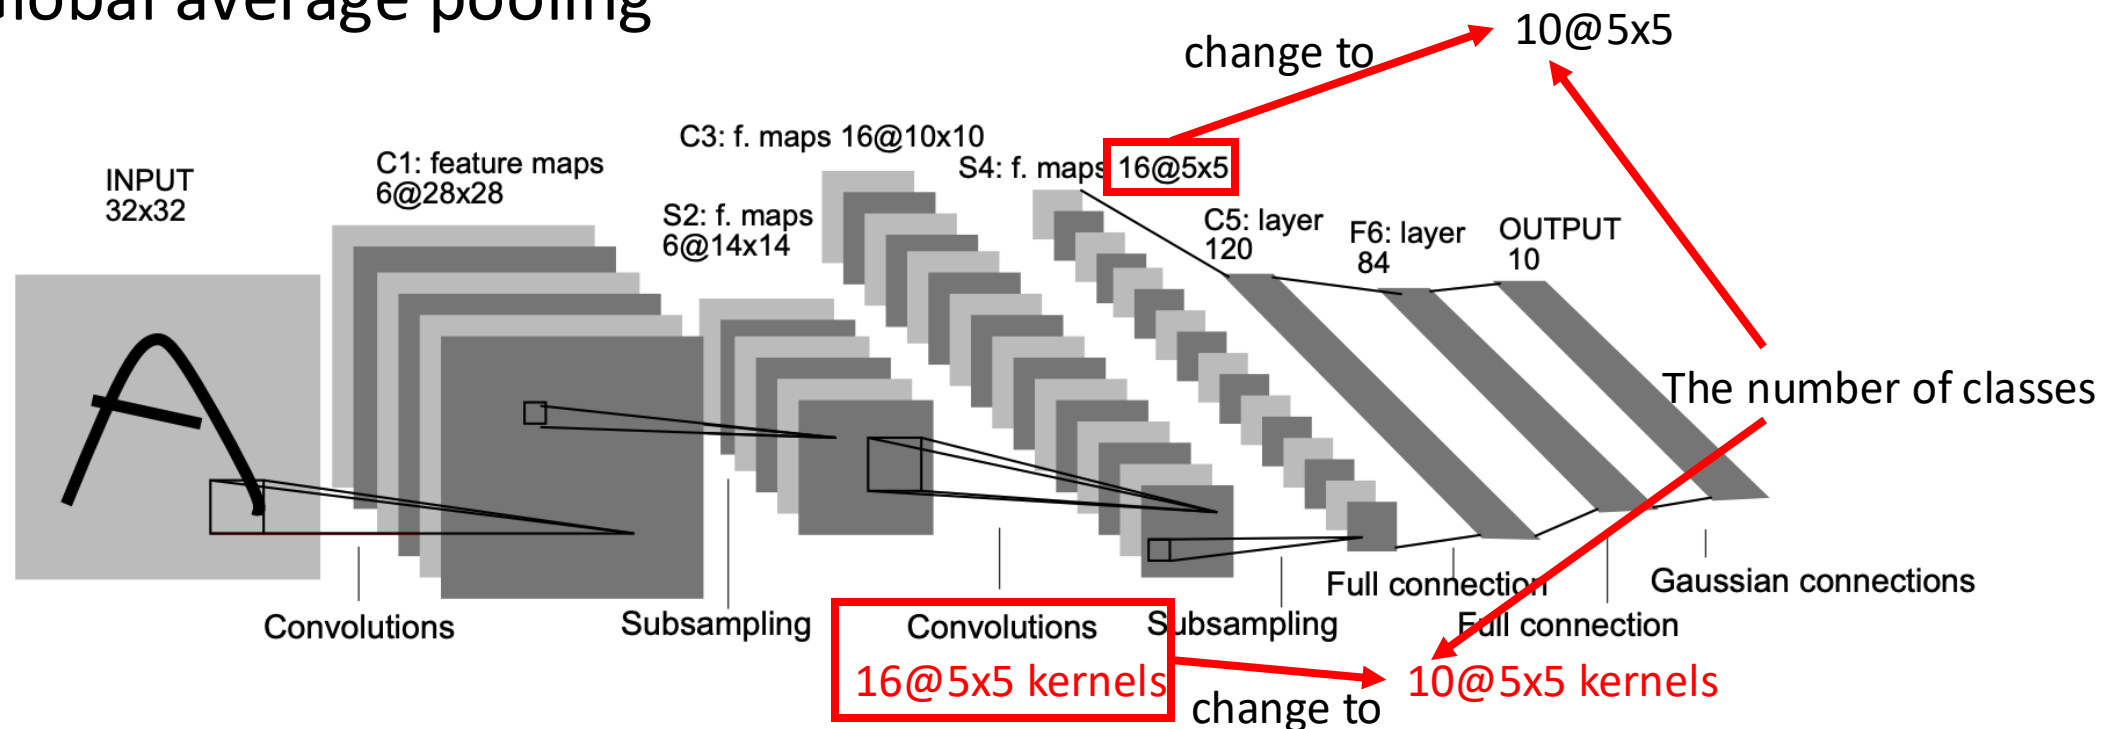


Fig. 1. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

Input resolution issue

- Global average pooling




Fig. 1. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

Input resolution issue

- Global average pooling




Fig. 1. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

Input resolution issue

- Global average pooling




Fig. 1. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

References

- LeCun, Yann, Patrick Haffner, Léon Bottou, and Yoshua Bengio. "Object recognition with gradient-based learning." In *Shape, contour and grouping in computer vision*, pp. 319-345. Springer, Berlin, Heidelberg, 1999.
 - Online at <http://yann.lecun.com/exdb/publis/pdf/lecun-99.pdf>
 - Section 2.2
 - Understand architecture of LeNet-5
- LeCun, Yann, Léon Bottou, Yoshua Bengio, and Patrick Haffner. "Gradient-based learning applied to document recognition." *Proceedings of the IEEE* 86, no. 11 (1998): 2278-2324.
 - Online at http://vision.stanford.edu/cs598_spring07/papers/Lecun98.pdf
 - Section II.B

References

- [Alexnet] Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. "Imagenet classification with deep convolutional neural networks." *Advances in neural information processing systems* 25 (2012): 1097-1105. Conference proceeding version at <https://papers.nips.cc/paper/2012/hash/c399862d3b9d6b76c8436e924a68c45b-Abstract.html> or <https://papers.nips.cc/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf> (Section 3.5)
- [pyramid] He, Kaiming, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. "Spatial pyramid pooling in deep convolutional networks for visual recognition." *IEEE transactions on pattern analysis and machine intelligence* 37, no. 9 (2015): 1904-1916. ArXiv version at <https://arxiv.org/abs/1406.4729> (Section 2.2)
- [NIN] Lin, Min, Qiang Chen, and Shuicheng Yan. "Network in network." *arXiv preprint arXiv:1312.4400* (2013). ArXiv version at <https://arxiv.org/abs/1312.4400> (Section 3.2)