

Lecture 3: Convolutional Neural Networks



Today:



Today:

- Expectations of visual recognition systems (10%)



Today:

- Expectations of visual recognition systems (10%)
- A Conv layer, stride, padding etc. (25%)



Today:

- Expectations of visual recognition systems (10%)
- A Conv layer, stride, padding etc. (25%)
- Conv2D (20%)

Today:

- Expectations of visual recognition systems (10%)
- A Conv layer, stride, padding etc. (25%)
- Conv2D (20%)
- Transposed Conv (5%)

Today:

- Expectations of visual recognition systems (10%)
- A Conv layer, stride, padding etc. (25%)
- Conv2D (20%)
- Transposed Conv (5%)
- Implementation and backprop (20%)

Today:

- Expectations of visual recognition systems (10%)
- A Conv layer, stride, padding etc. (25%)
- Conv2D (20%)
- Transposed Conv (5%)
- Implementation and backprop (20%)
- What is encoded in feature maps? (10%)

Today:

- Expectations of visual recognition systems (10%)
- A Conv layer, stride, padding etc. (25%)
- Conv2D (20%)
- Transposed Conv (5%)
- Implementation and backprop (20%)
- What is encoded in feature maps? (10%)
- Max-pooling, convnet and conv variants (10%)







vertical
line





vertical
line





vertical
line

horizontal
line






key







keyboard



key key key

key key key

key key


key key

key key





Monitor



keyboard

key key key key key key key key key
key key






Monitor



keyboard

Mouse





Monitor

Desk

keyboard

Mouse

key key key

key key key

key key

key key

key key



Monitor



Monitor





Monitor

Expectations of visual recognition network



Expectations of visual recognition network

1. Maintain 2D structure logic



Expectations of visual recognition network

1. Maintain 2D structure logic
2. Shift invariant (actually, equivariant)

Expectations of visual recognition network

1. Maintain 2D structure logic
2. Shift invariant (actually, equivariant)
3. Consider only local correlations

Expectations of visual recognition network

1. Maintain 2D structure logic
2. Shift invariant (actually, equivariant)
3. Consider only local correlations
4. Hierarchically growing field of view


Expectations of visual recognition network

1. Maintain 2D structure logic
2. Shift invariant (actually, equivariant)
3. Consider only local correlations
4. Hierarchically growing field of view
5. Hierarchically progressing complexity


Expectations of visual recognition network

1. Maintain 2D structure logic
2. Shift invariant (actually, equivariant)
3. Consider only local correlations
4. Hierarchically growing field of view
5. Hierarchically progressing complexity
6. Reasonable amount of params

Fully connected layer




Fully connected layer



$$\begin{bmatrix} w_{00}^l & w_{01}^l & \dots & \dots \\ w_{10}^l & w_{11}^l & \dots & \dots \\ \vdots & \vdots & \ddots & \ddots \\ \vdots & \vdots & \vdots & \vdots \end{bmatrix}$$


Fully connected layer



$$\begin{bmatrix} w_{00}^l & w_{01}^l & & & & & & & \\ w_{10}^l & w_{11}^l & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \end{bmatrix}$$

1. Maintain 2D structure logic
2. Shift invariant (actually, equivariant)
3. Consider only local correlations
4. Hierarchically growing field of view
5. Hierarchically progressing complexity
6. Reasonable amount of params


Fully connected layer



$$\begin{bmatrix} w_{00}^l & w_{01}^l & & & & & & & \\ w_{10}^l & w_{11}^l & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \end{bmatrix}$$

- ✗
1. Maintain 2D structure logic
 2. Shift invariant (actually, equivariant)
 3. Consider only local correlations
 4. Hierarchically growing field of view
 5. Hierarchically progressing complexity
 6. Reasonable amount of params


Fully connected layer



$$\begin{bmatrix} w_{00}^l & w_{01}^l & & & & & & & \\ w_{10}^l & w_{11}^l & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \end{bmatrix}$$

- ~~1.~~ 1. Maintain 2D structure logic
- ~~2.~~ 2. Shift invariant (actually, equivariant)
- 3. Consider only local correlations
- 4. Hierarchically growing field of view
- 5. Hierarchically progressing complexity
- 6. Reasonable amount of params


Fully connected layer



$$\begin{bmatrix} w_{00}^l & w_{01}^l & & & & & & & \\ w_{10}^l & w_{11}^l & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \end{bmatrix}$$

- ~~X~~ 1. Maintain 2D structure logic
- ~~X~~ 2. Shift invariant (actually, equivariant)
- ~~X~~ 3. Consider only local correlations
- 4. Hierarchically growing field of view
- 5. Hierarchically progressing complexity
- 6. Reasonable amount of params


Fully connected layer



$$\begin{bmatrix} w_{00}^l & w_{01}^l & & & & & & & \\ w_{10}^l & w_{11}^l & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \end{bmatrix}$$

- ~~X~~ 1. Maintain 2D structure logic
- ~~X~~ 2. Shift invariant (actually, equivariant)
- ~~X~~ 3. Consider only local correlations
- ~~X~~ 4. Hierarchically growing field of view
- 5. Hierarchically progressing complexity
- 6. Reasonable amount of params


Fully connected layer



$$\begin{bmatrix} w_{00}^l & w_{01}^l & & & & & & & \\ w_{10}^l & w_{11}^l & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \end{bmatrix}$$

- ~~1.~~ 1. Maintain 2D structure logic
- ~~2.~~ 2. Shift invariant (actually, equivariant)
- ~~3.~~ 3. Consider only local correlations
- ~~4.~~ 4. Hierarchically growing field of view
- ✓ 5. Hierarchically progressing complexity
- 6. Reasonable amount of params


Fully connected layer



$$\begin{bmatrix} w_{00}^l & w_{01}^l & & & & & & & \\ w_{10}^l & w_{11}^l & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \\ & & & & & & & & \end{bmatrix}$$

- ~~X~~ 1. Maintain 2D structure logic
- ~~X~~ 2. Shift invariant (actually, equivariant)
- ~~X~~ 3. Consider only local correlations
- ~~X~~ 4. Hierarchically growing field of view
- ✓ 5. Hierarchically progressing complexity
- ~~X~~ 6. Reasonable amount of params

Convolution



$$(f * g)(t) \stackrel{\text{def}}{=} \int_{-\infty}^{\infty} f(\tau) g(t - \tau) d\tau$$

Convolution layer

x_0^l

x_1^l

x_2^l

x_3^l

x_4^l

x_5^l

x_6^l

x_7^l

x_0^{l+1}


x_1^{l+1}

x_2^{l+1}



x_3^{l+1}

x_4^{l+1}



x_5^{l+1}





Convolution layer





Convolution layer



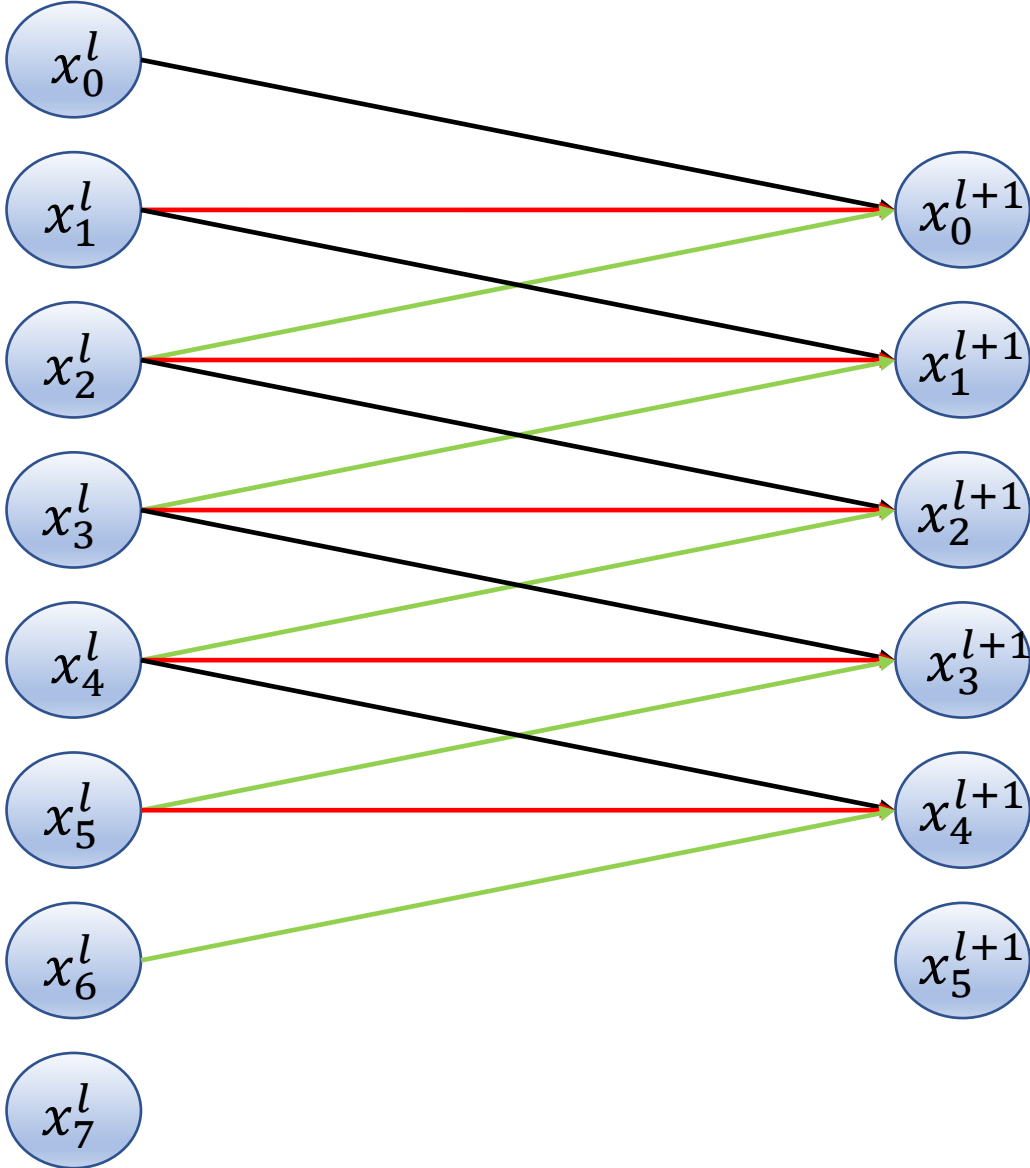

Convolution layer





Convolution layer




Convolution layer





Convolution layer





Convolution layer





Convolution layer





Convolution layer






Convolution layer





Convolution layer



Convolution layer





Convolution layer



$$\begin{bmatrix}
 \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 & 0 & 0 & 0 & 0 \\
 0 & & & & & & &
 \end{bmatrix}$$




Convolution layer



$$\begin{bmatrix}
 \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 & 0 & 0 & 0 & 0 \\
 0 & \mathbf{a} & & & & & &
 \end{bmatrix}$$




Convolution layer



\longrightarrow a
 \longrightarrow b
 \longrightarrow c

$$\begin{bmatrix}
 \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 & 0 & 0 & 0 & 0 \\
 0 & \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 & 0 & 0 & 0 \\
 & & & & & & & \\
 & & & & & & & \\
 & & & & & & & \\
 & & & & & & & \\
 & & & & & & & \\
 & & & & & & &
 \end{bmatrix}$$

Convolution layer




\longrightarrow a
 \longrightarrow b
 \longrightarrow c

$$\begin{bmatrix}
 \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 & 0 & 0 & 0 & 0 \\
 0 & \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 & 0 & 0 & 0 \\
 0 & 0 & \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 & 0 & 0 \\
 0 & 0 & 0 & \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 & 0 \\
 0 & 0 & 0 & 0 & \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 \\
 0 & 0 & 0 & 0 & 0 & \mathbf{a} & \mathbf{b} & \mathbf{c}
 \end{bmatrix}$$



Convolution layer

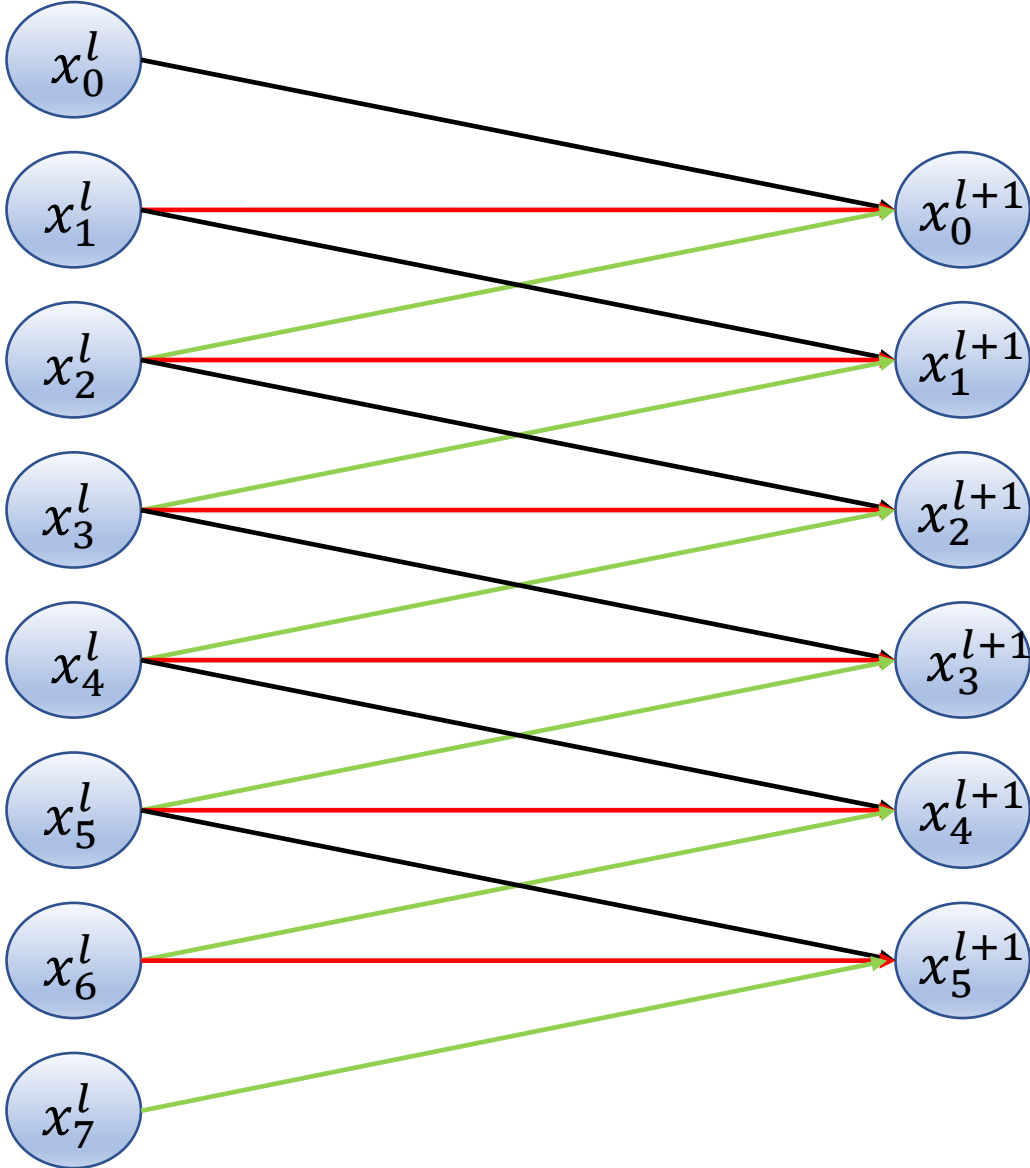


\longrightarrow a
 \longrightarrow b
 \longrightarrow c

$$\begin{bmatrix}
 \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 & 0 & 0 & 0 & 0 \\
 0 & \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 & 0 & 0 & 0 \\
 0 & 0 & \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 & 0 & 0 \\
 0 & 0 & 0 & \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 & 0 \\
 0 & 0 & 0 & 0 & \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 \\
 0 & 0 & 0 & 0 & 0 & \mathbf{a} & \mathbf{b} & \mathbf{c} \\
 0 & 0 & 0 & 0 & 0 & 0 & \mathbf{a} & \mathbf{b} & \mathbf{c}
 \end{bmatrix}$$

Toeplitz matrix

Convolution layer




\longrightarrow a
 \longrightarrow b
 \longrightarrow c

$$\begin{bmatrix}
 \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 & 0 & 0 & 0 & 0 \\
 0 & \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 & 0 & 0 & 0 \\
 0 & 0 & \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 & 0 & 0 \\
 0 & 0 & 0 & \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 & 0 \\
 0 & 0 & 0 & 0 & \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 \\
 0 & 0 & 0 & 0 & 0 & \mathbf{a} & \mathbf{b} & \mathbf{c}
 \end{bmatrix}$$


Special case of Linear Layer

Toeplitz matrix


Convolution Filter




Convolution Filter




Convolution Filter




Convolution Filter




Convolution Filter




Convolution Filter



Convolution Filter

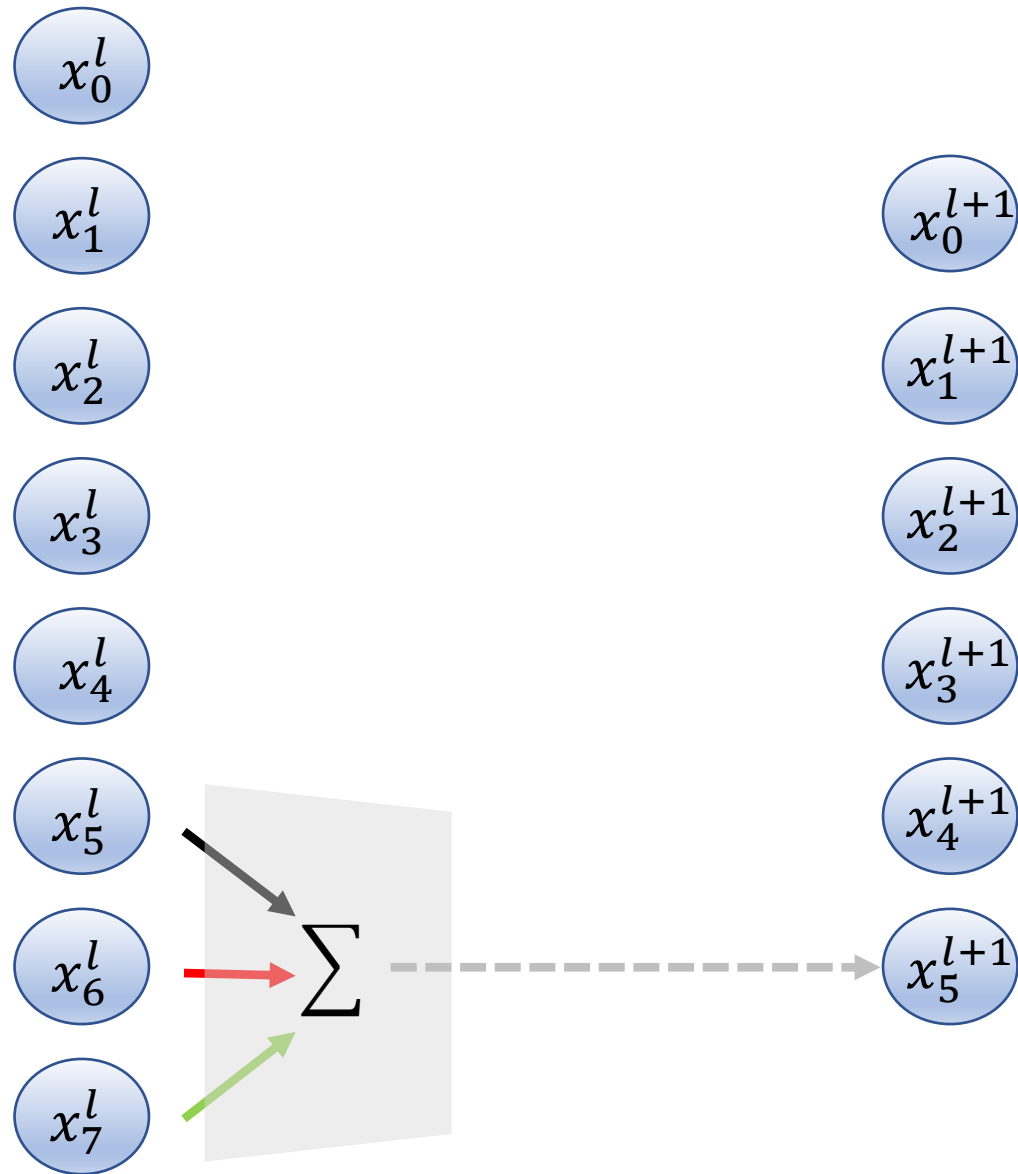


Convolution Filter



Q: Given input size N and filter size K , find output size.

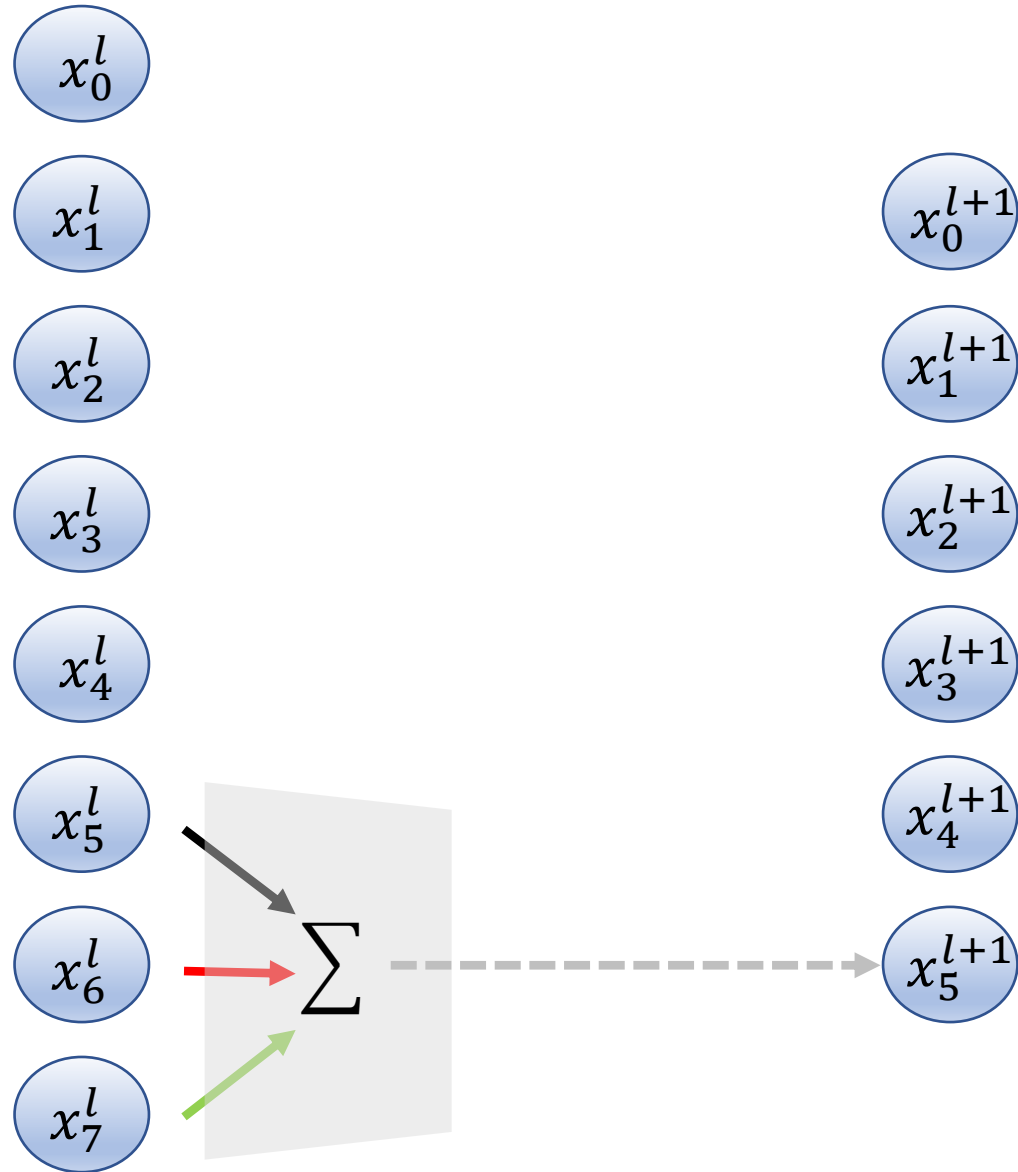
Convolution Filter



Q: Given input size N and filter size K , find output size.

A: $N - K + 1$

Convolution Filter




Q: Given input size N and filter size K , find output size.

A: $N - K + 1$

Q: is this a convolution?

Convolution Filter




Q: Given input size N and filter size K , find output size.


A: $N - K + 1$

Q: is this a convolution?


A: Yes, but with the flipped filter.
This is cross-correlation.




Stride




Stride



Stride




Stride




$$\begin{bmatrix} \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 & 0 & 0 & 0 \\ 0 & \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 & 0 & 0 \\ 0 & 0 & \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 & 0 \\ 0 & 0 & 0 & \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 \\ 0 & 0 & 0 & 0 & \mathbf{a} & \mathbf{b} & \mathbf{c} \\ 0 & 0 & 0 & 0 & 0 & \mathbf{a} & \mathbf{b} \end{bmatrix}$$

Stride




a	b	c	0	0	0	0
0	a	b	c	0	0	0
0	0	a	b	c	0	0
0	0	0	a	b	c	0
0	0	0	0	a	b	c
0	0	0	0	0	a	b

Stride



$$\begin{bmatrix} \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 & 0 & 0 & 0 \\ 0 & 0 & \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 & 0 \\ 0 & 0 & 0 & 0 & \mathbf{a} & \mathbf{b} & \mathbf{c} \end{bmatrix}$$


Stride



$$\begin{bmatrix} \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 & 0 & 0 & 0 \\ 0 & 0 & \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 & 0 \\ 0 & 0 & 0 & 0 & \mathbf{a} & \mathbf{b} & \mathbf{c} \end{bmatrix}$$

Q: Given input size N , filter size K and stride S , find output size.

Stride



$$\begin{bmatrix} \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 & 0 & 0 & 0 \\ 0 & 0 & \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 & 0 \\ 0 & 0 & 0 & 0 & \mathbf{a} & \mathbf{b} & \mathbf{c} \end{bmatrix}$$

Q: Given input size N , filter size K and stride S , find output size.

$$\text{A: } \left\lfloor \frac{N-K}{S} + 1 \right\rfloor$$

Padding

x_0^l

x_1^l

x_2^l

x_3^l

x_4^l

x_5^l

x_0^{l+1}

x_1^{l+1}


x_2^{l+1}

x_3^{l+1}


x_4^{l+1}

x_5^{l+1}


Padding




Padding




Padding




Padding



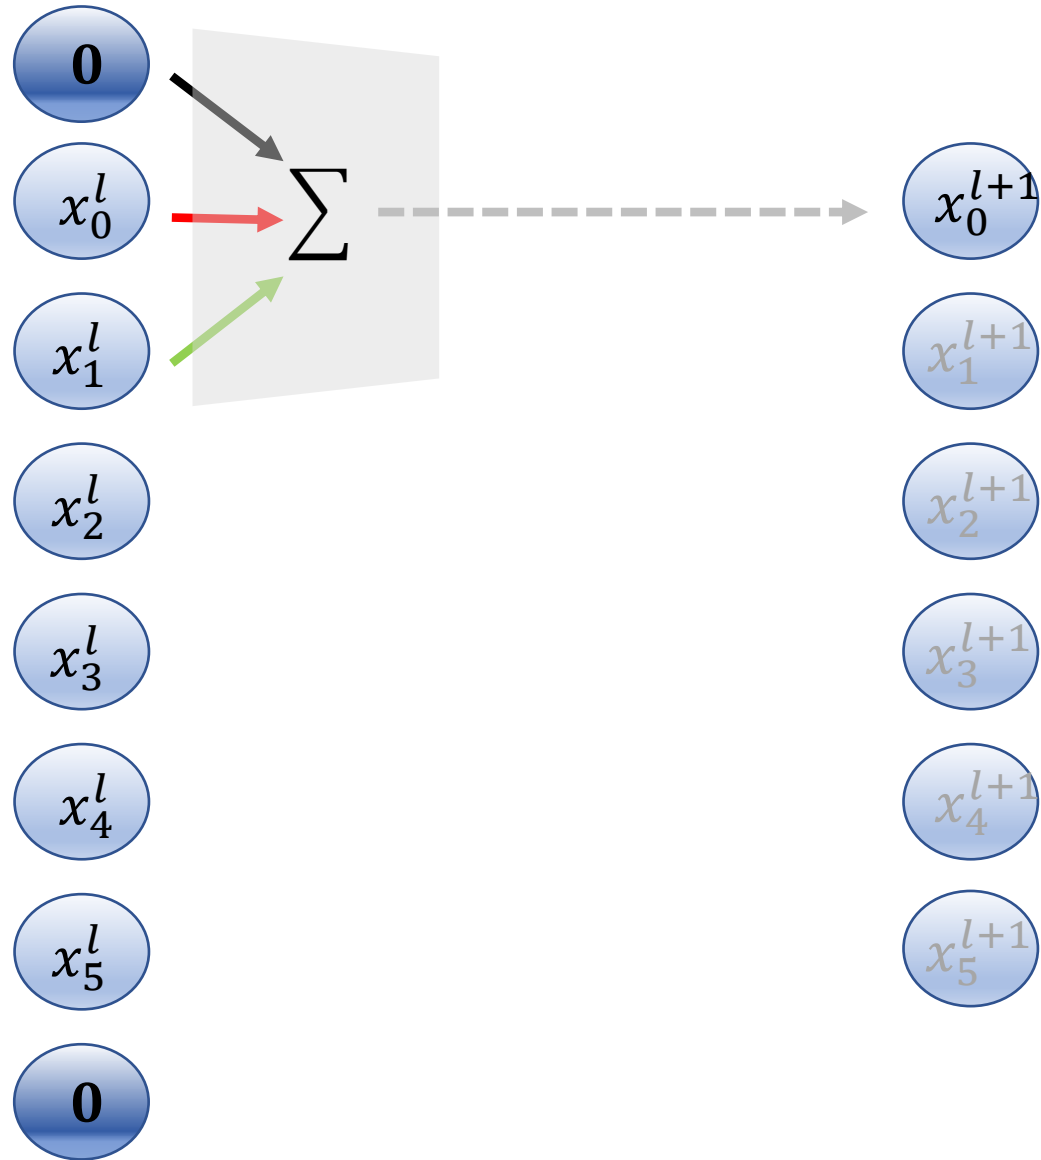
Padding



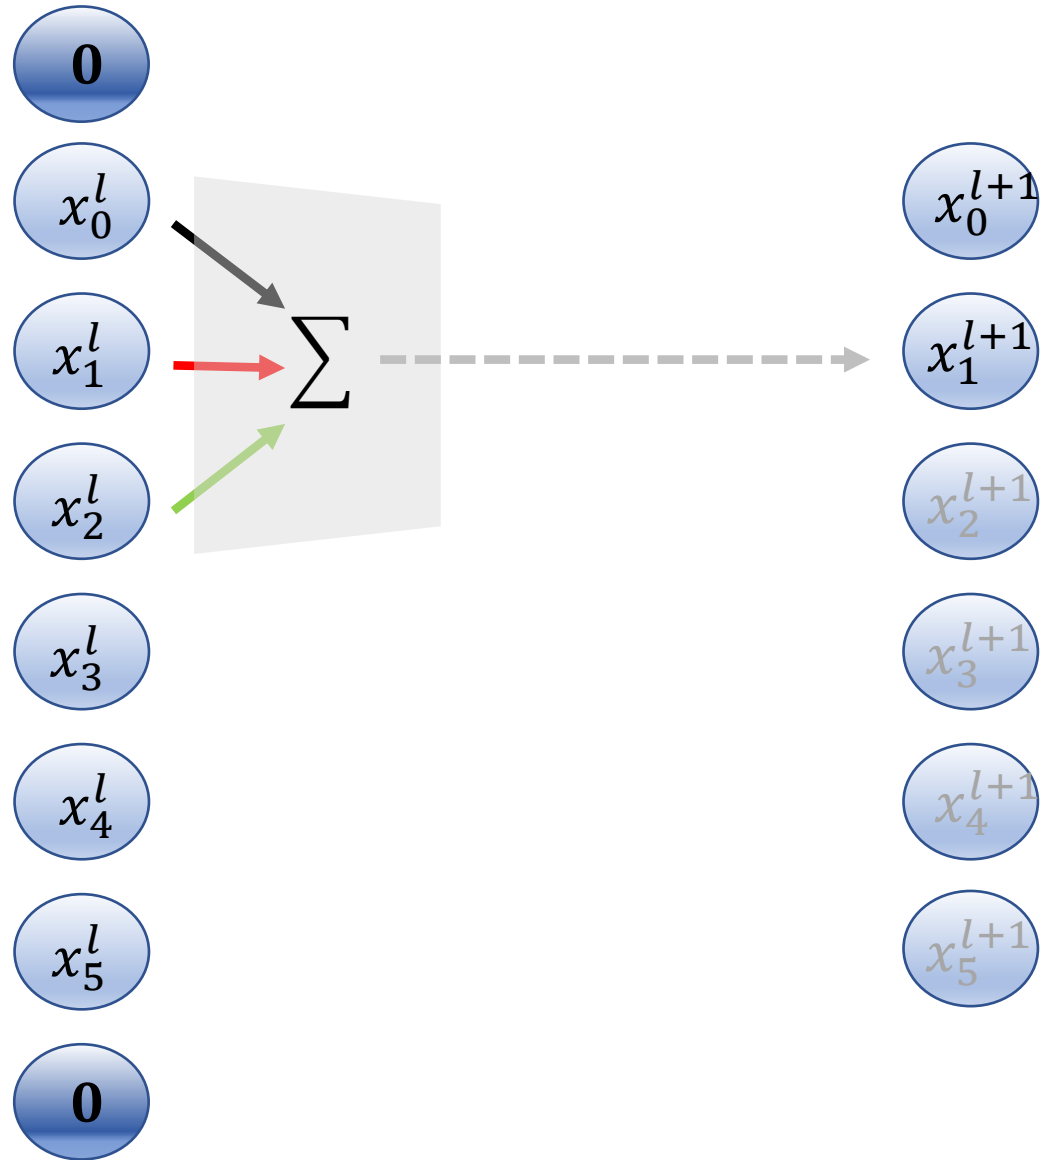
Padding




Padding




Padding




Padding




Padding



Padding




Padding




Conv2D



Conv2D





Conv2D



N, C, H, W





Conv2D



N, C, H, W


Conv2D



N, C, H, W


C, h, w

Conv2D



N, C, H, W


Conv2D



N, C, H, W





Conv2D



N, C, H, W




Conv2D



N, C, H, W




Conv2D



N, C, H, W




Conv2D




N, C, H, W



Conv2D





N, C, H, W





$[N, 1, \tilde{H}, \tilde{W}]$



Conv2D




Conv2D





Conv2D



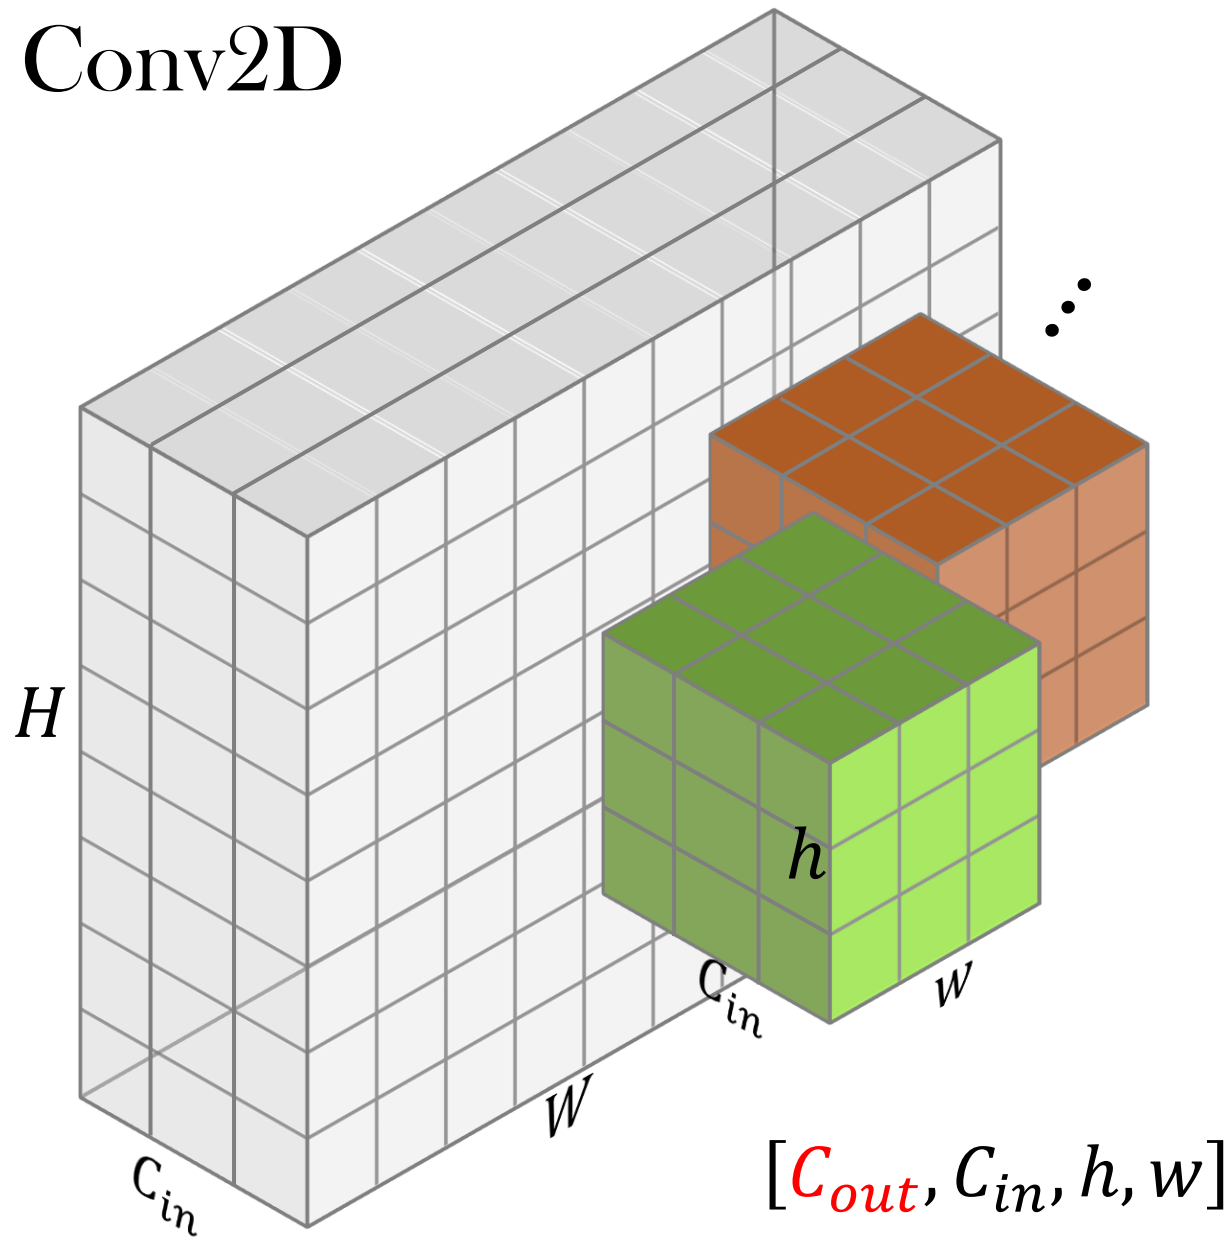
Conv2D



Conv2D




Conv2D



$$[C_{out}, C_{in}, h, w]$$

$$[N, C_{in}, H, W]$$



$$[N, C_{out}, \tilde{H}, \tilde{W}]$$

Receptive Field

[Filter size: 3]

x_0^0

x_1^0

x_2^0

x_3^0

x_4^0

x_5^0

x_6^0

x_7^0

x_0^1

x_1^1

x_2^1

x_3^1

x_4^1

x_5^1

x_6^1

x_7^1

x_0^2

x_1^2

x_2^2

x_3^2

x_4^2

x_5^2

x_6^2

x_7^2



Receptive Field

[Filter size: 3]

x_0^0

x_1^0

x_2^0

x_3^0

x_4^0

x_5^0

x_6^0

x_7^0

x_0^1

x_1^1

x_2^1

x_3^1

x_4^1

x_5^1

x_6^1

x_7^1

x_0^2

x_1^2

x_2^2

x_3^2

x_4^2

x_5^2

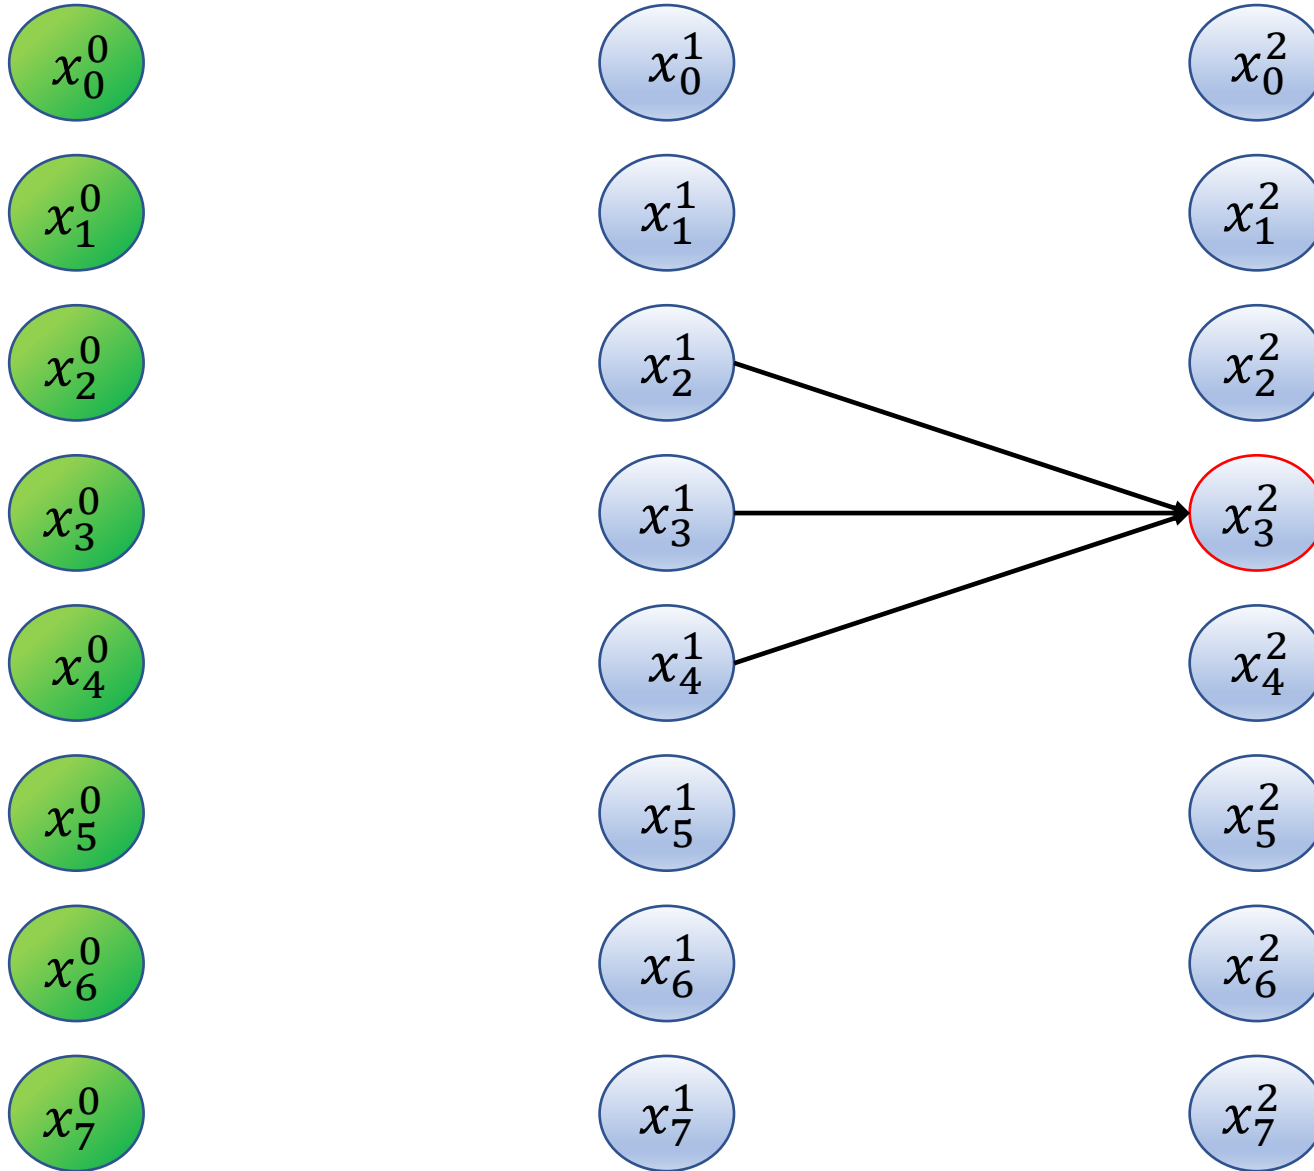
x_6^2

x_7^2



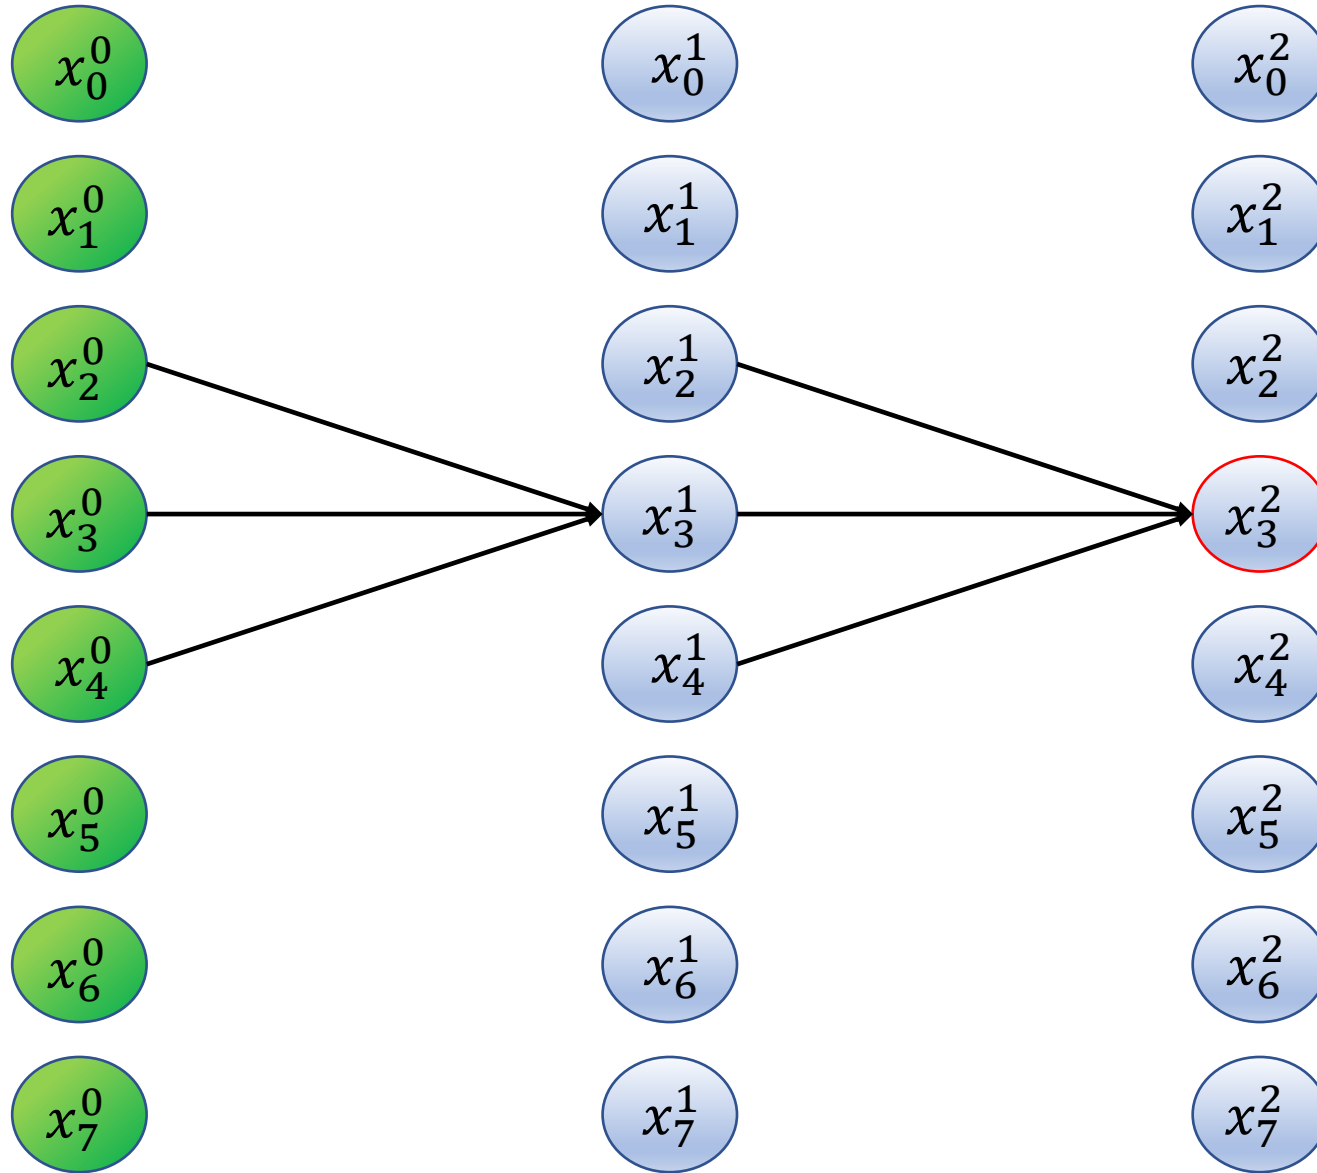
Receptive Field

[Filter size: 3]



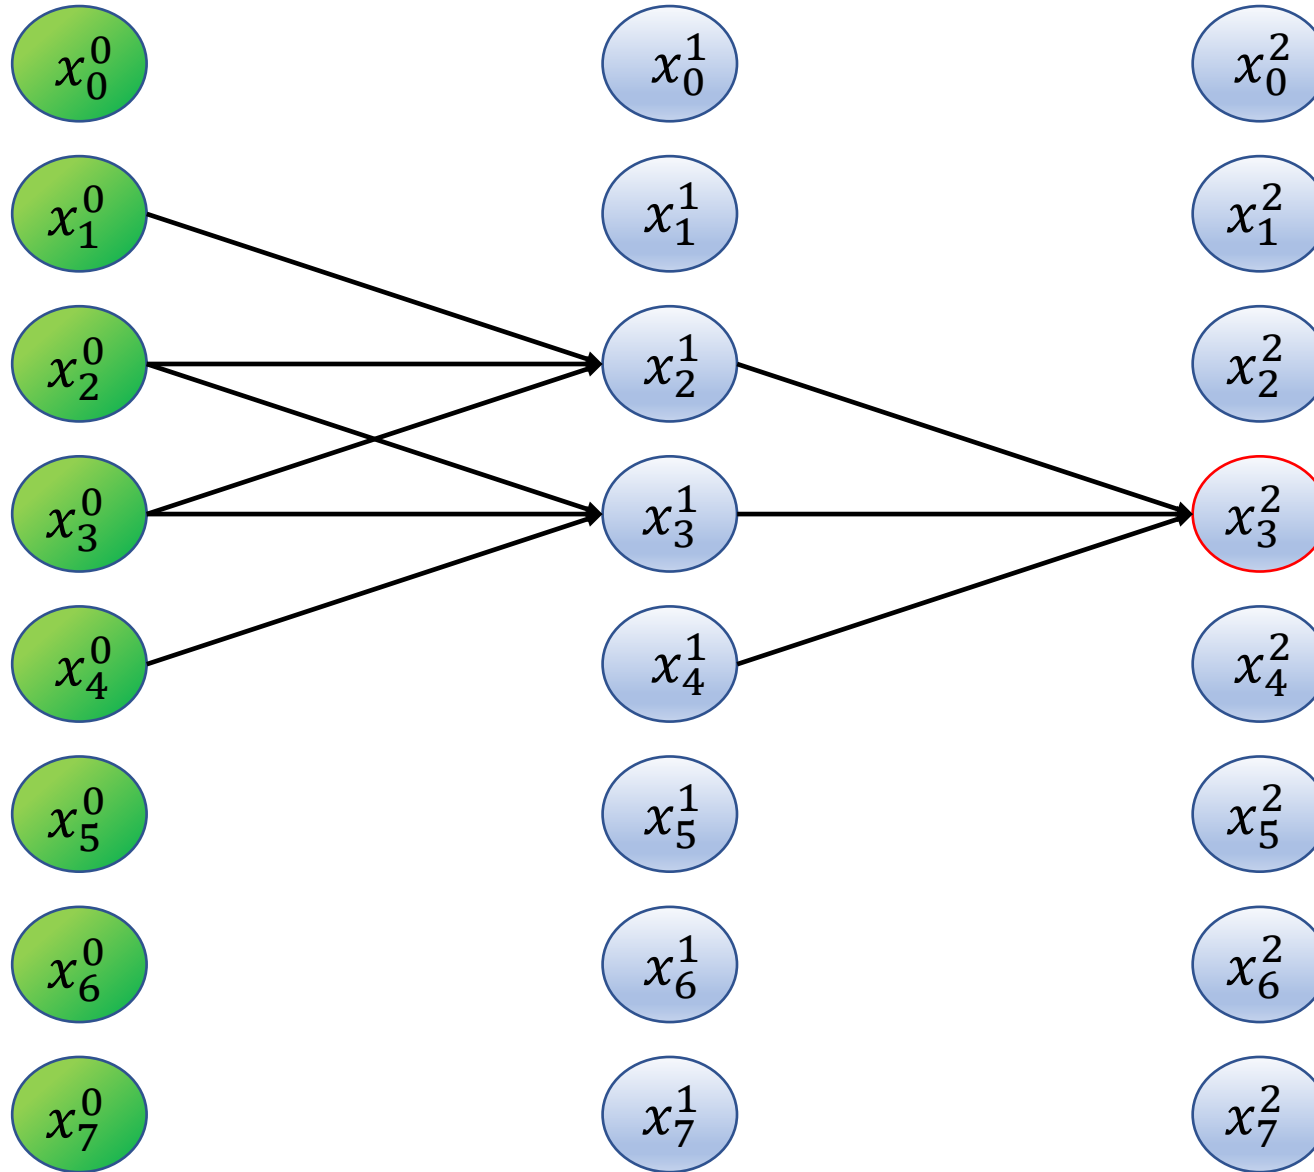
Receptive Field

[Filter size: 3]



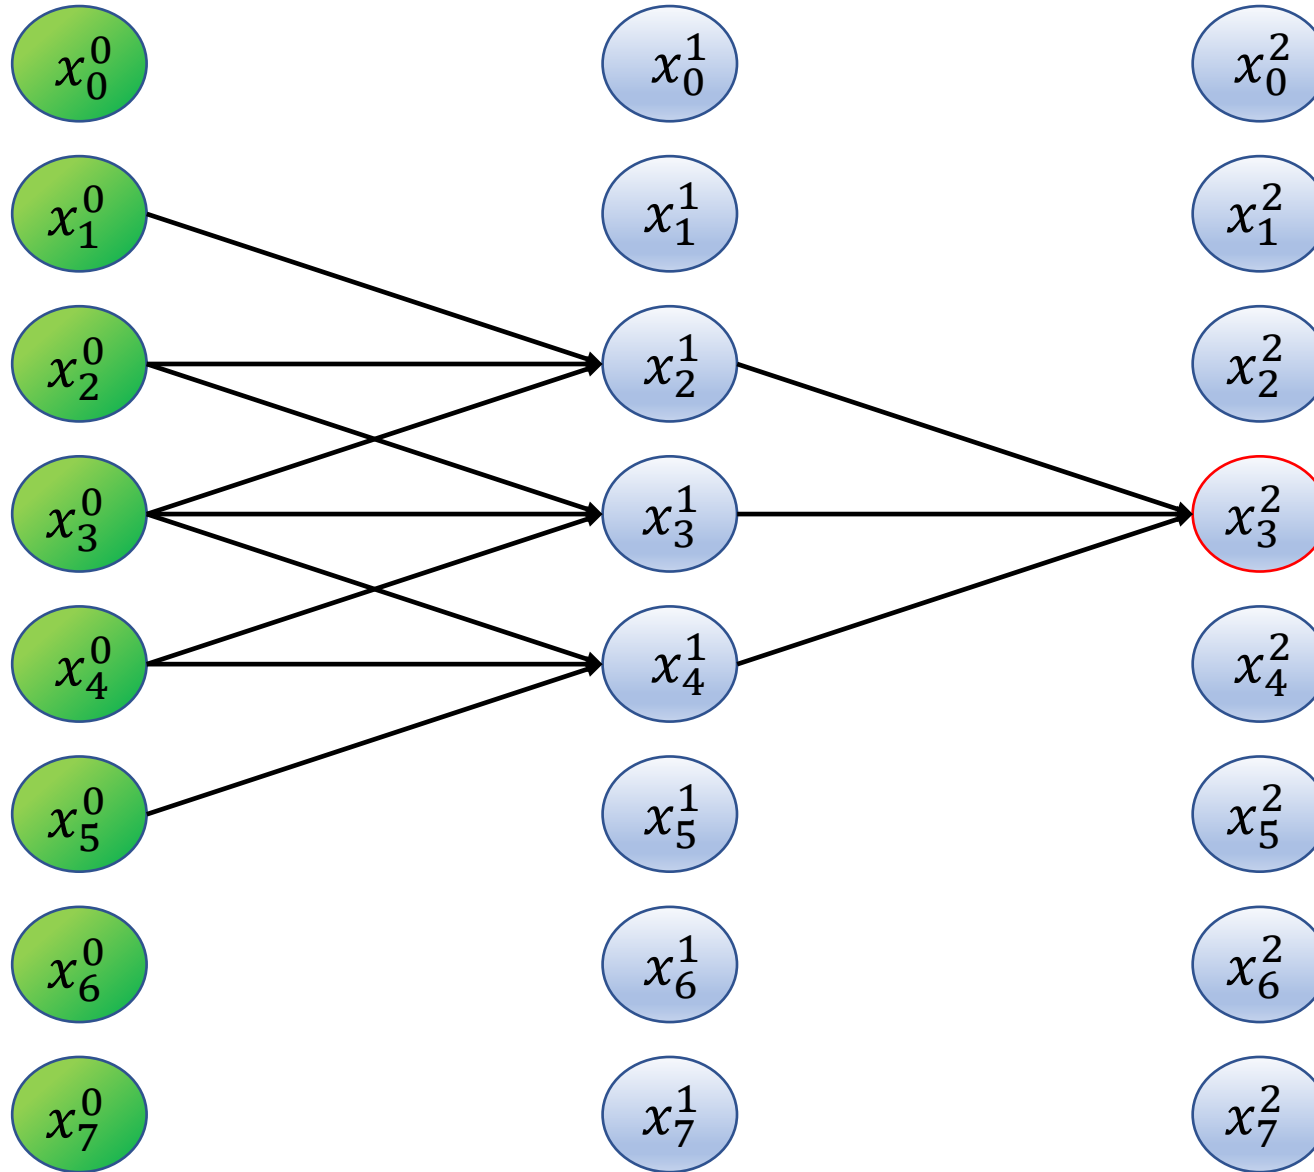
Receptive Field

[Filter size: 3]



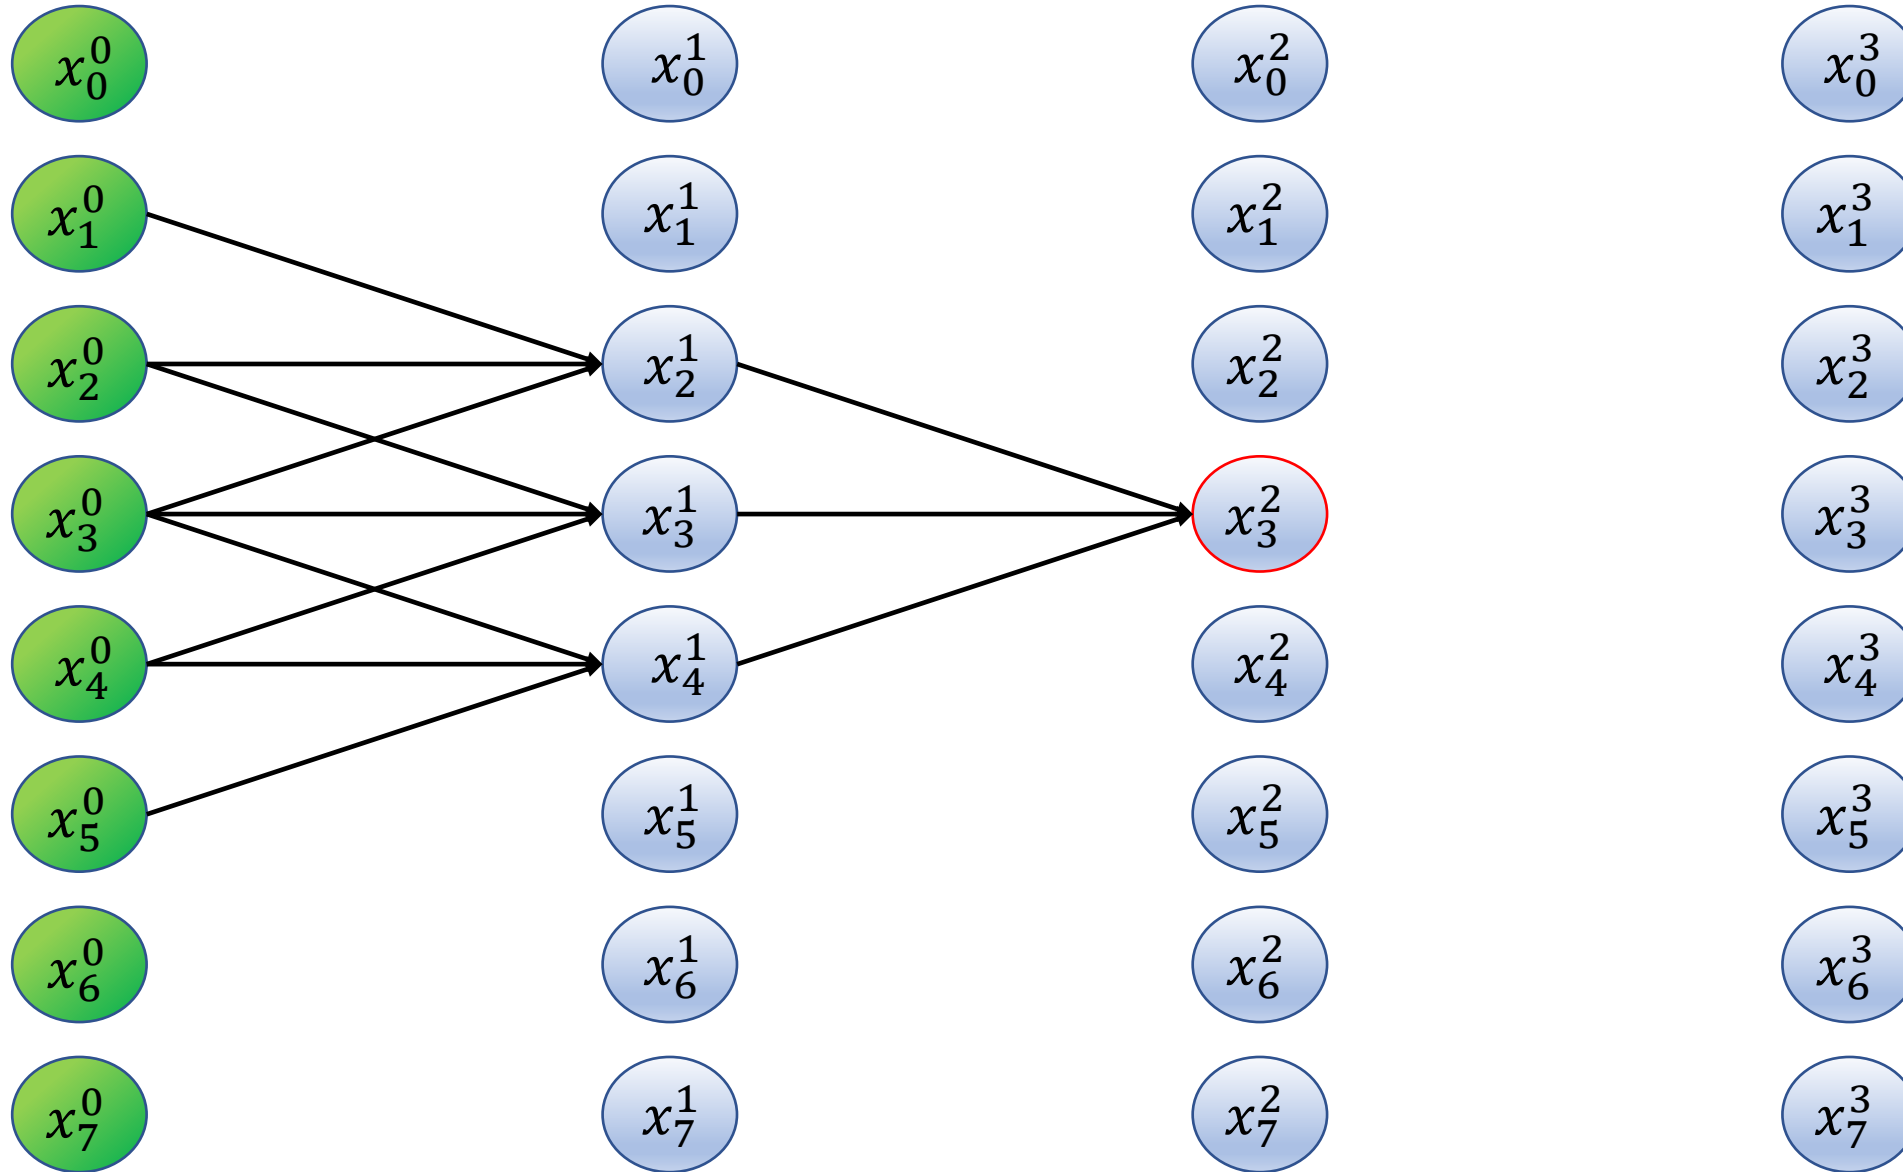
Receptive Field

[Filter size: 3]



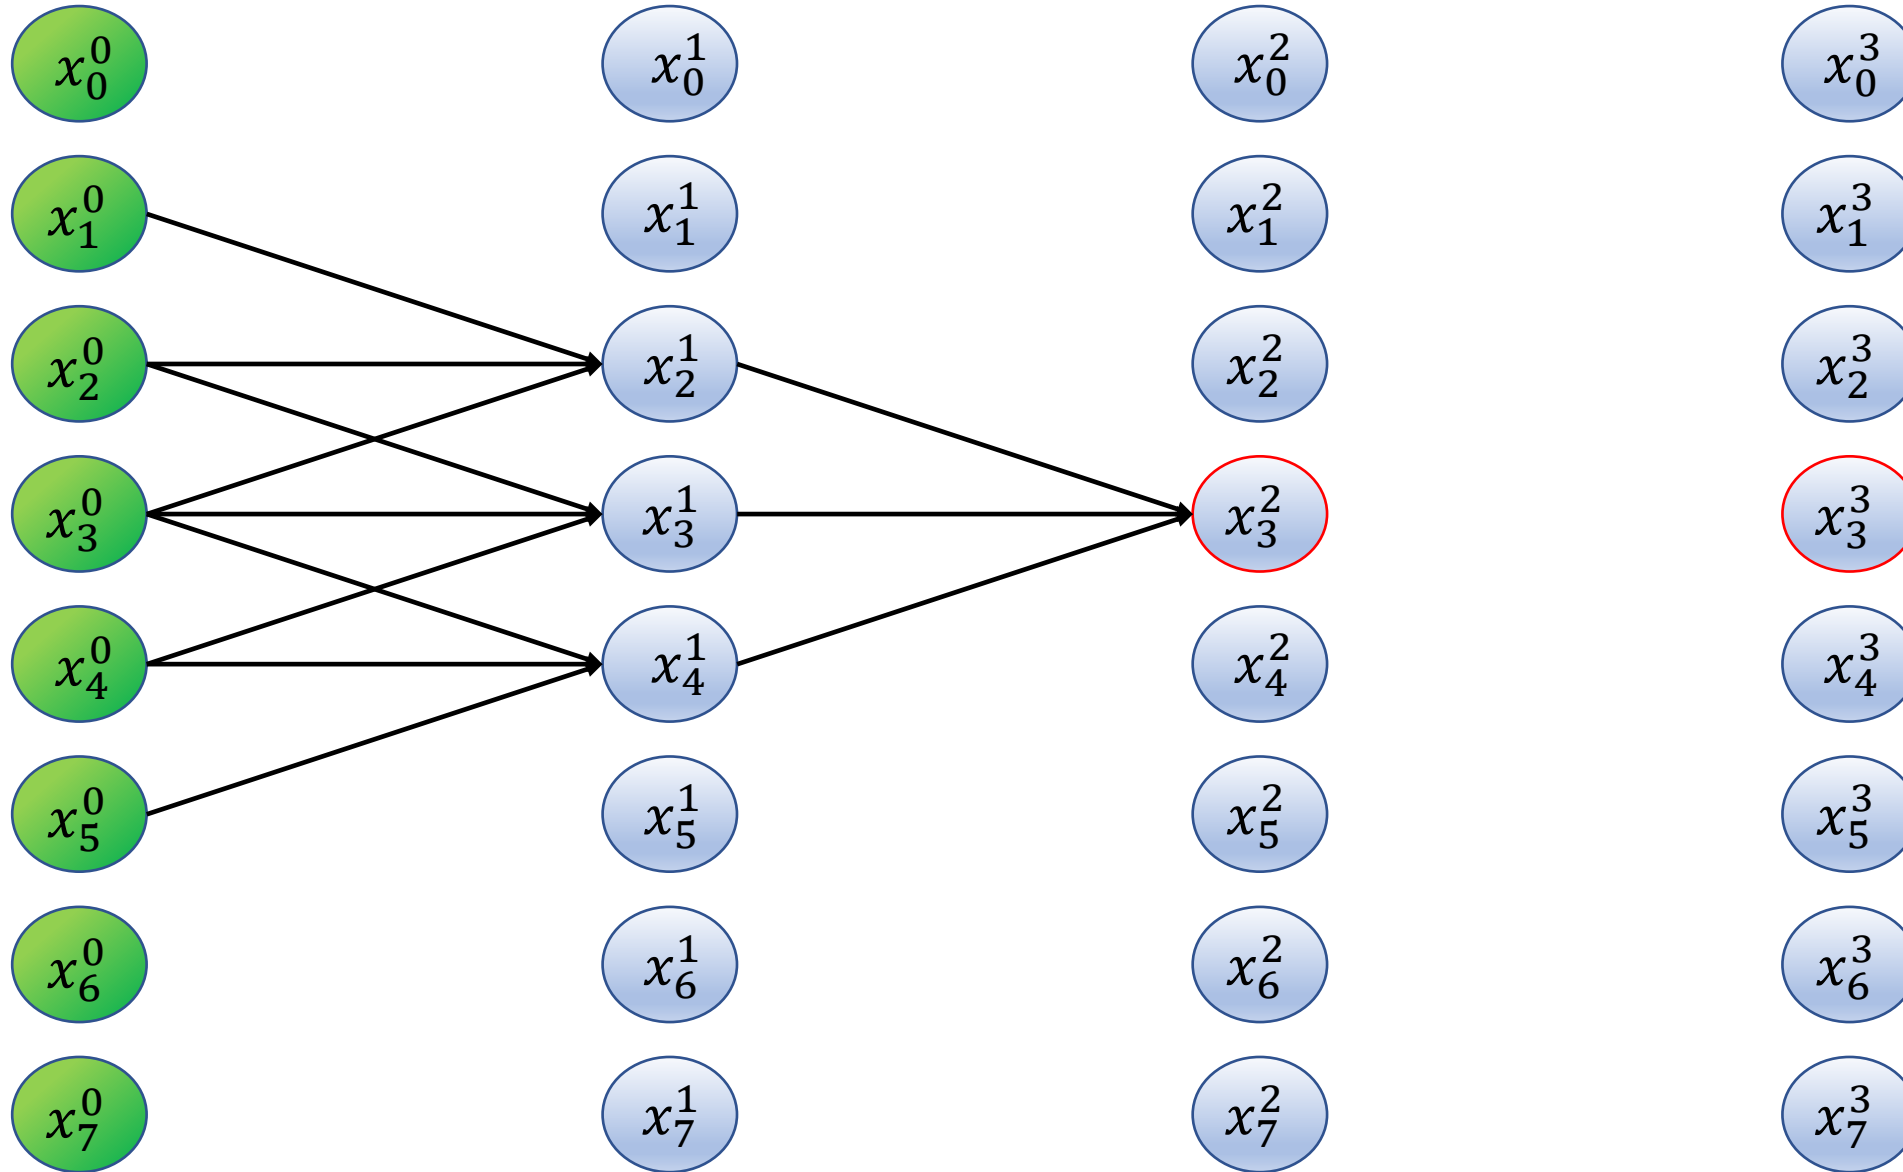
Receptive Field

[Filter size: 3]



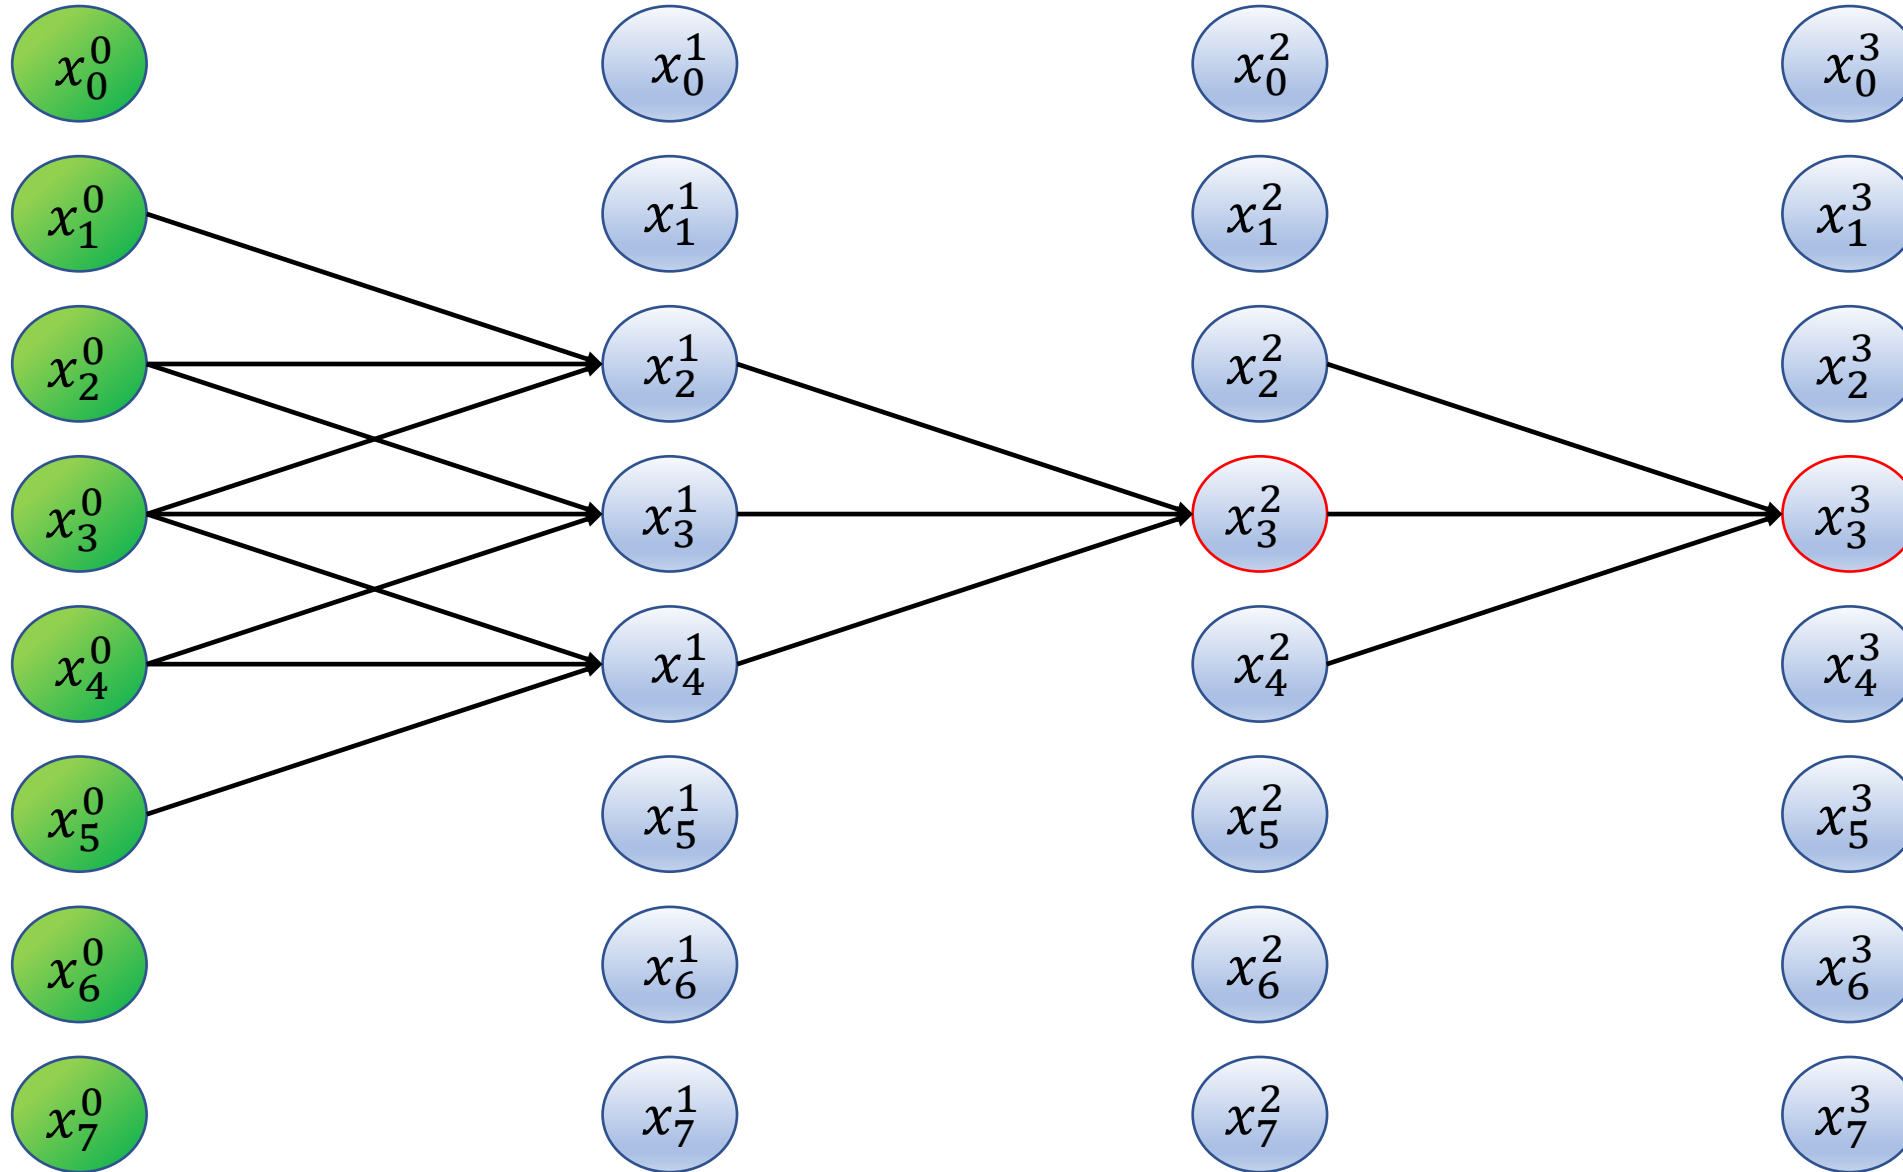
Receptive Field

[Filter size: 3]



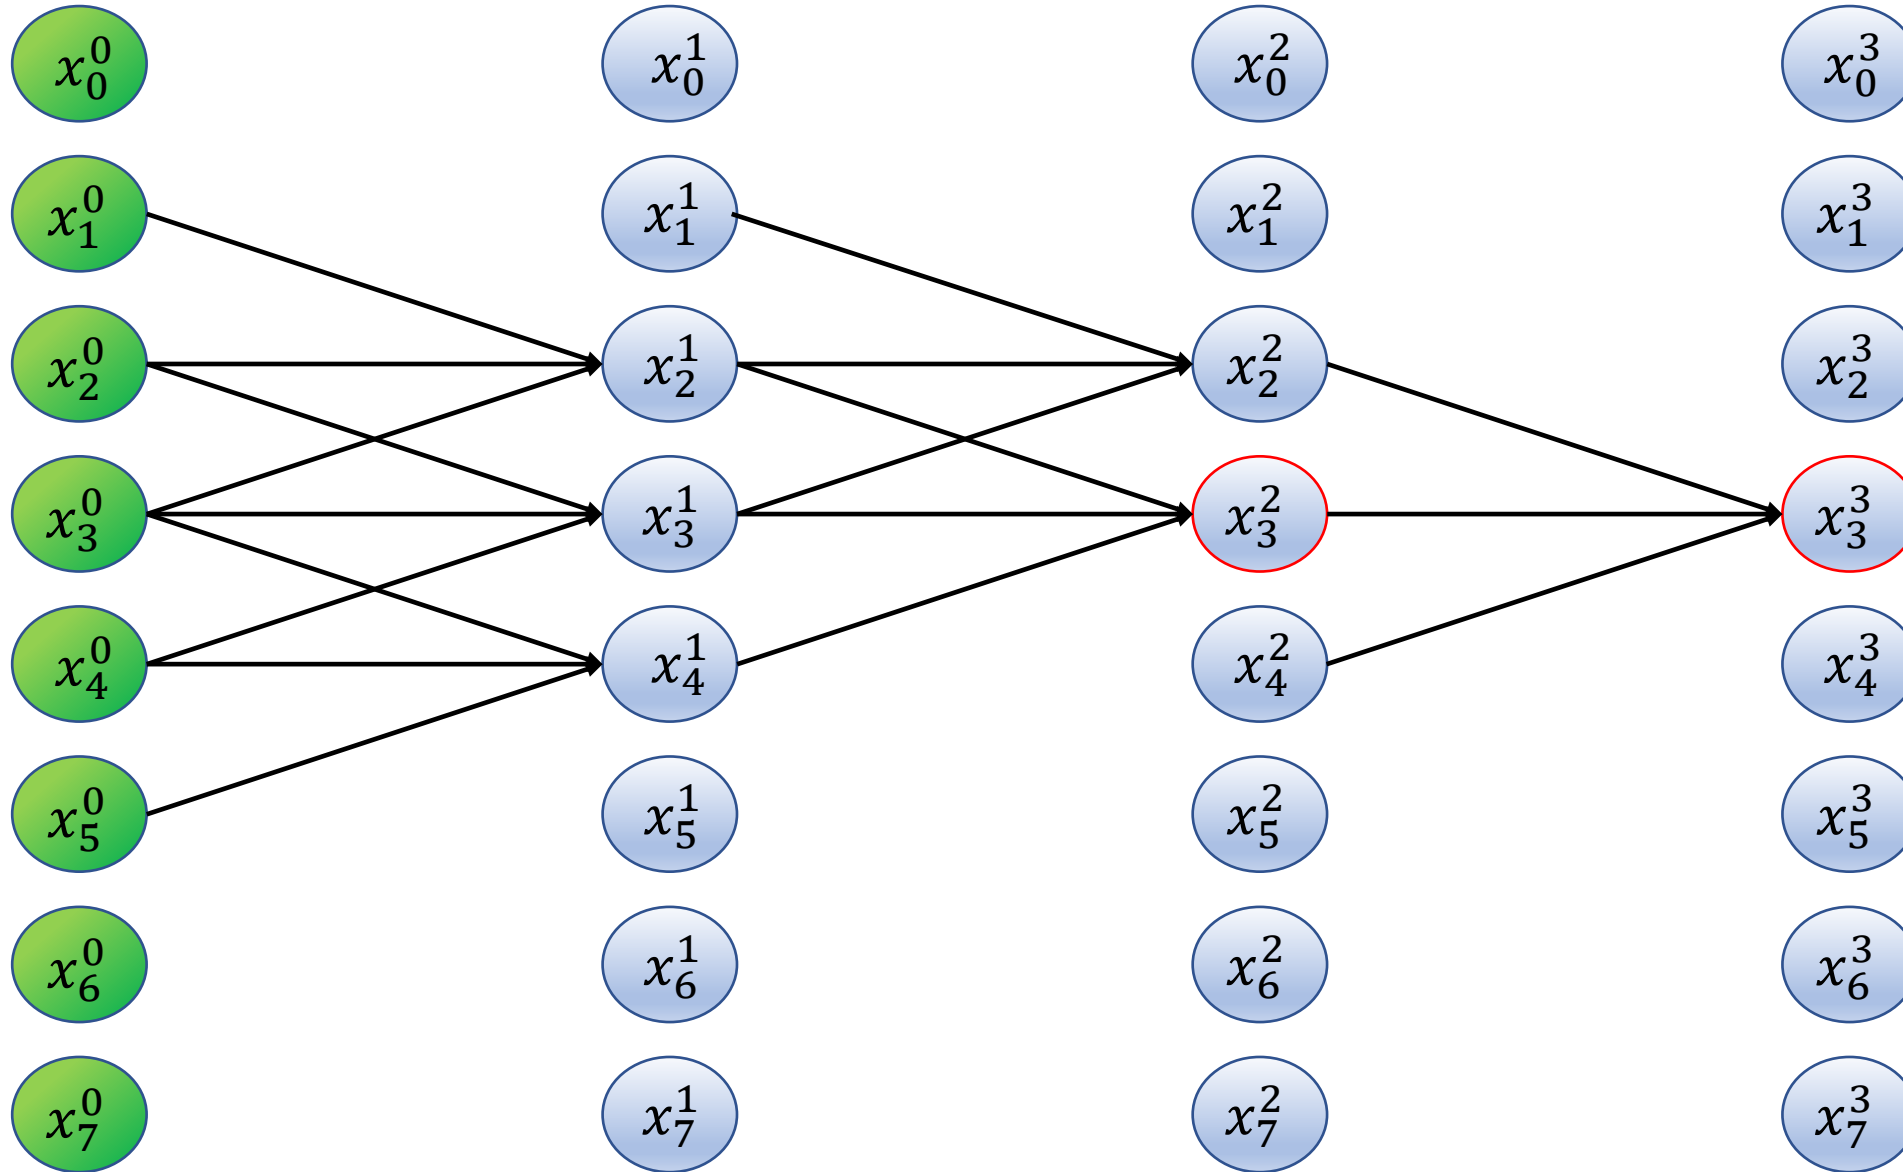
Receptive Field

[Filter size: 3]



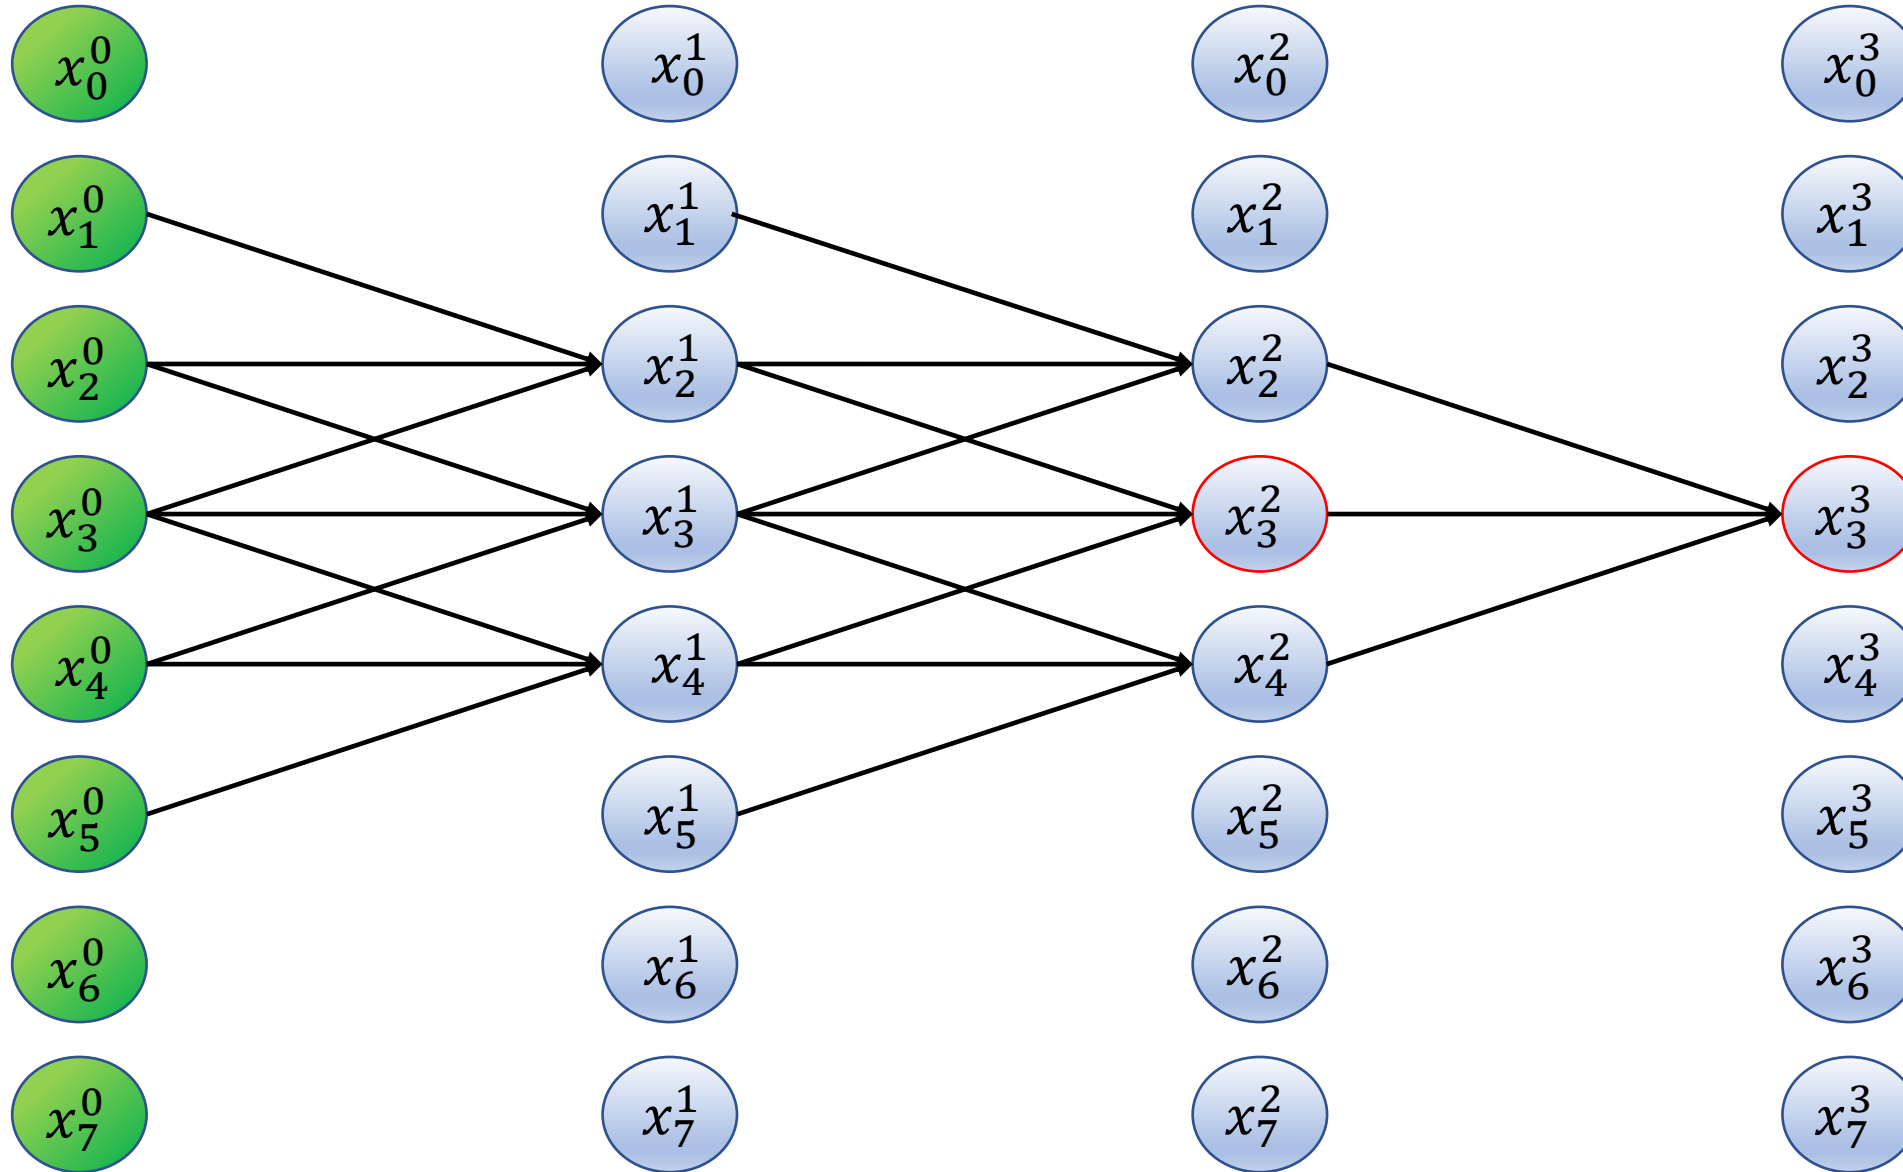
Receptive Field

[Filter size: 3]



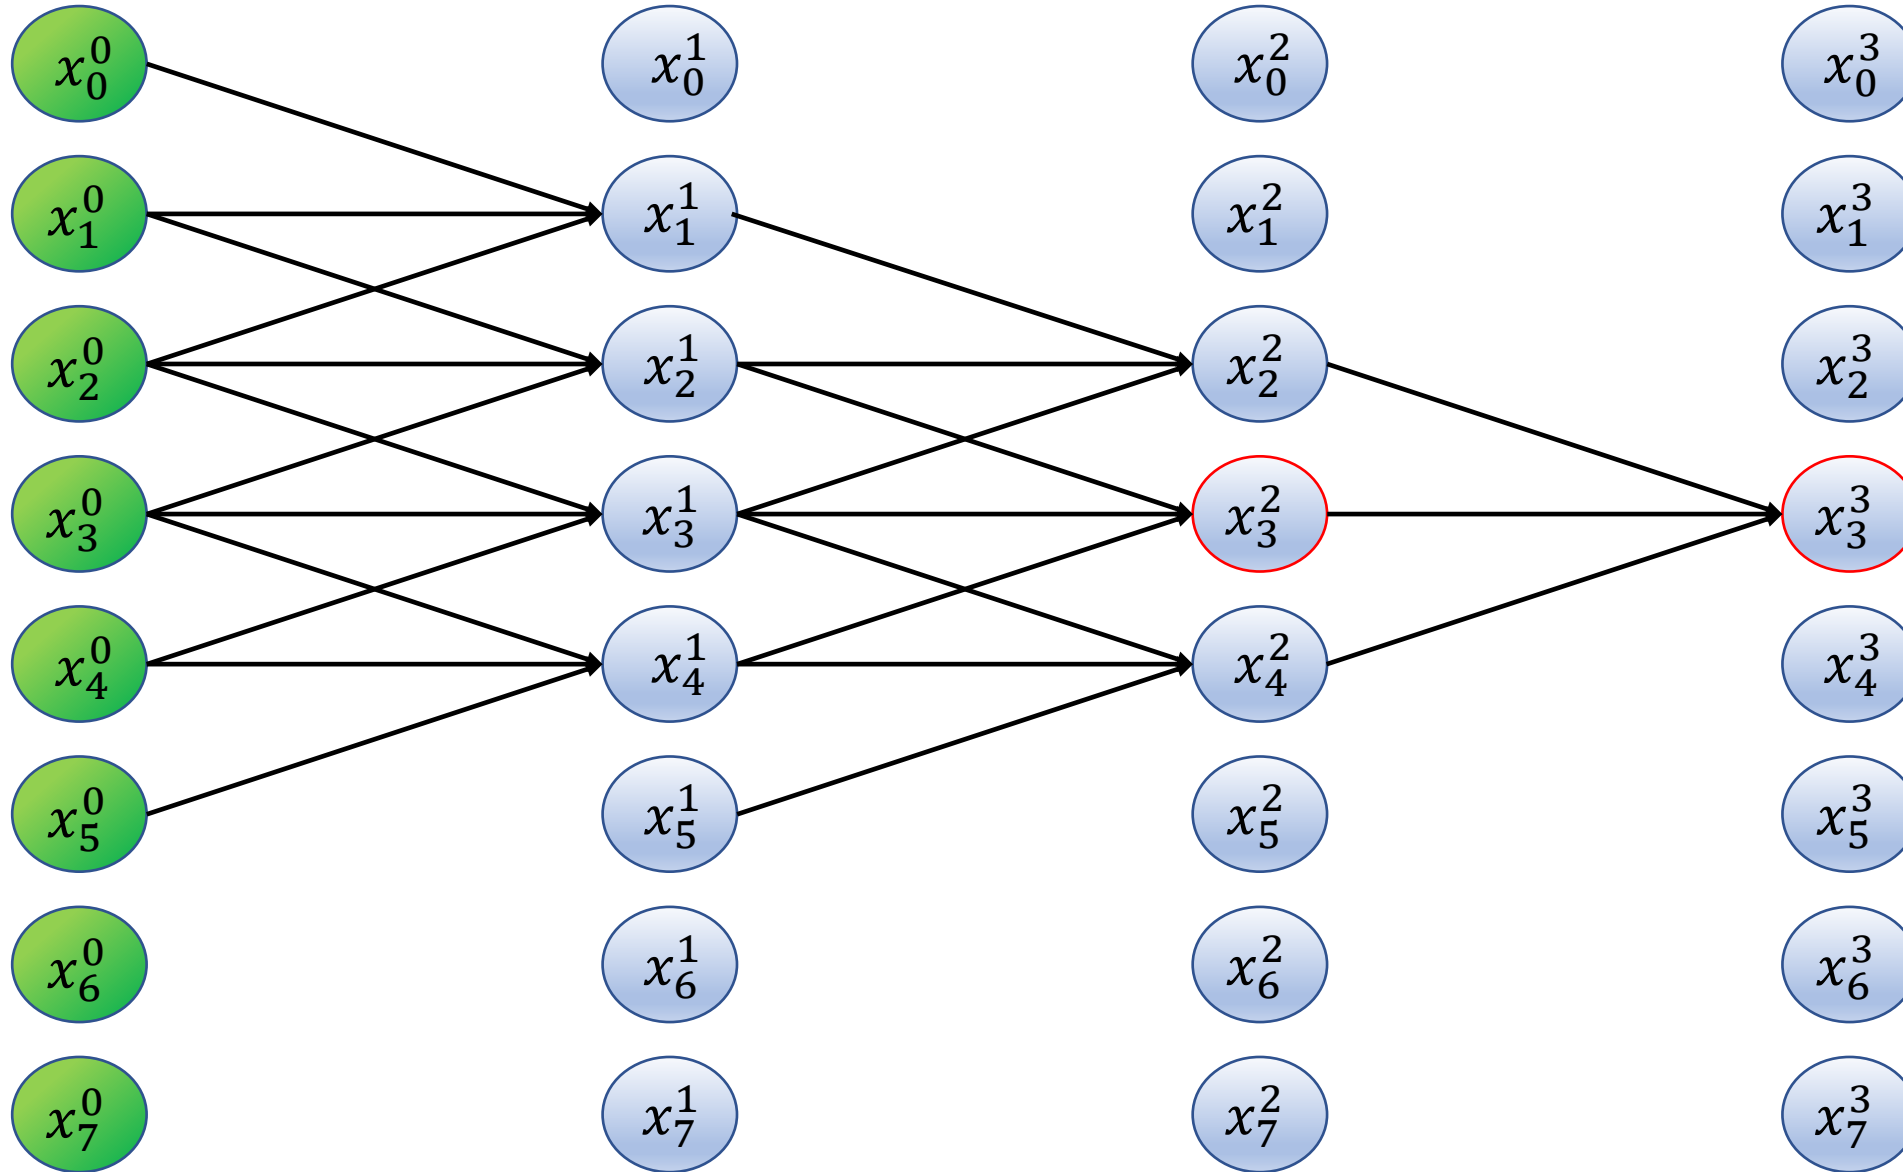
Receptive Field

[Filter size: 3]



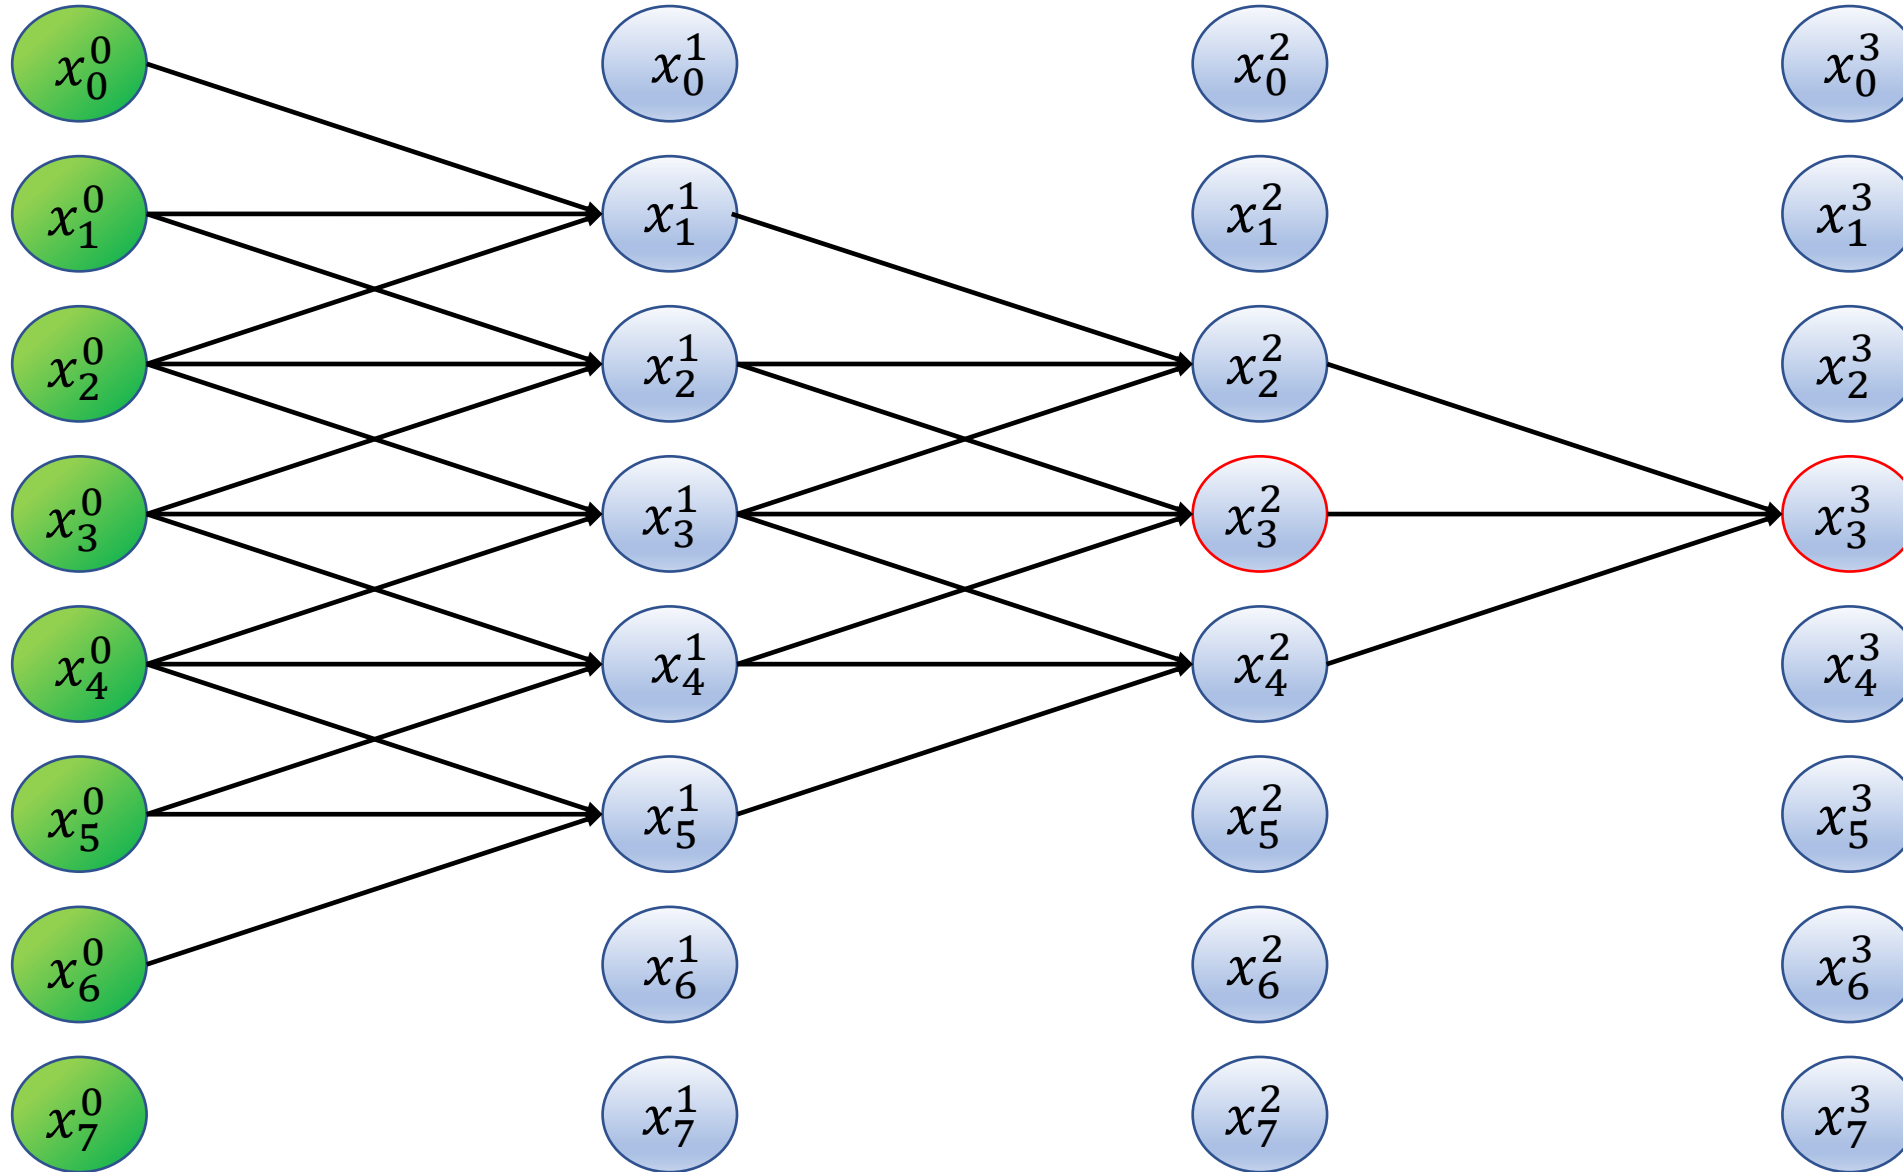
Receptive Field

[Filter size: 3]




Receptive Field

[Filter size: 3]




Receptive Field


[Filter size: 3]




Receptive Field



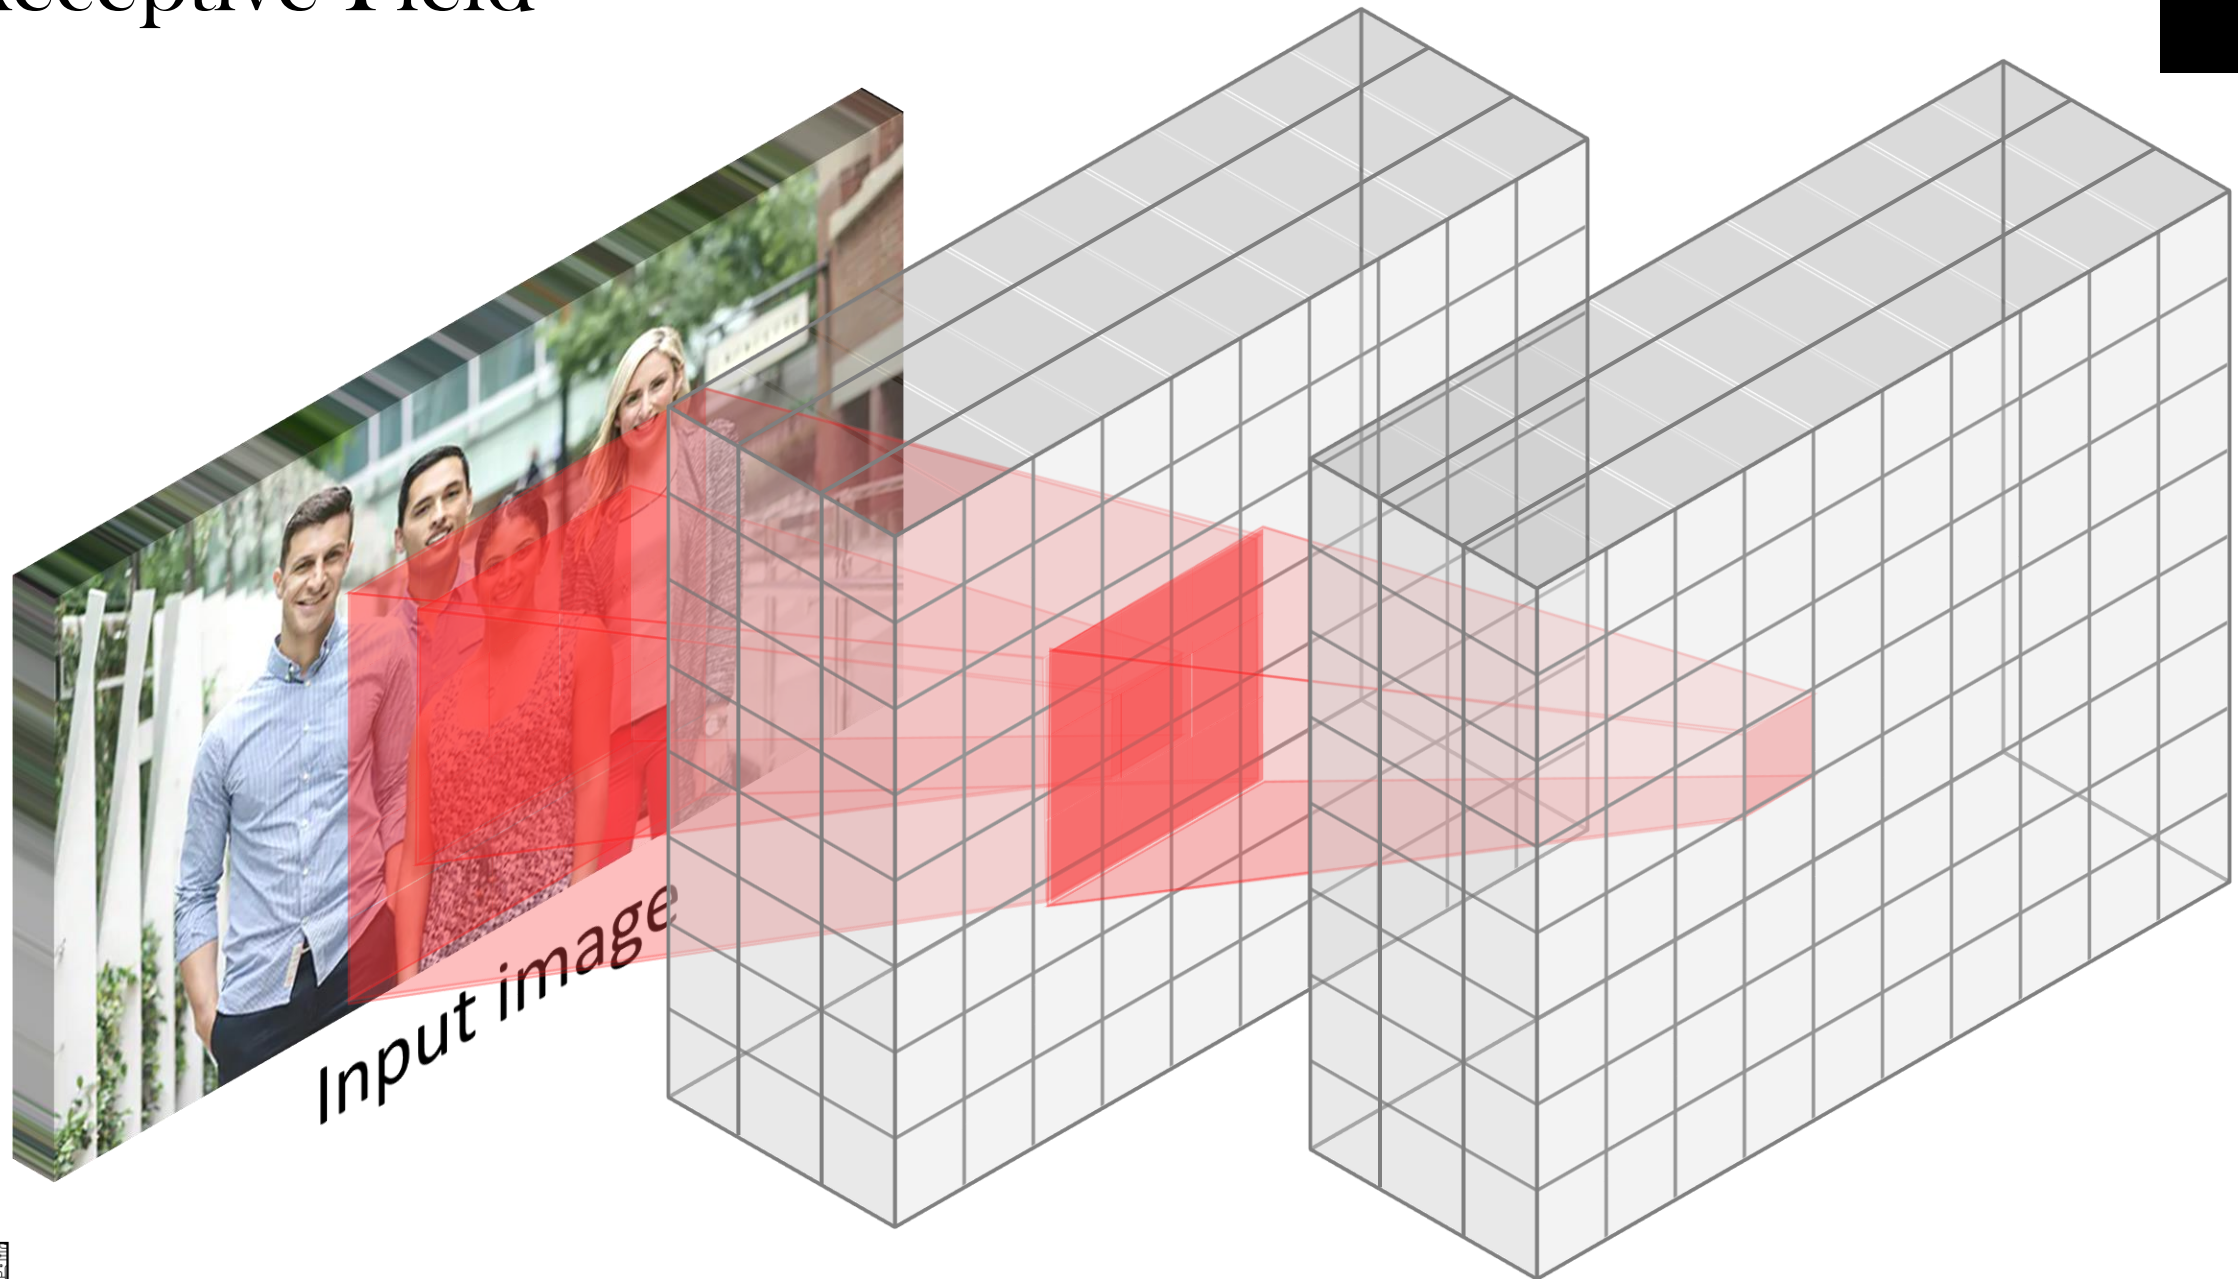
Receptive Field



Receptive Field



Receptive Field



Convs rock!



Convs rock!

1. Maintain 2D structure logic
2. Shift invariant (actually, equivariant)
3. Consider only local correlations
4. Hierarchically growing field of view
5. Hierarchically progressing complexity
6. Reasonable amount of params

Convs rock!

- ✓ 1. Maintain 2D structure logic
- 2. Shift invariant (actually, equivariant)
- 3. Consider only local correlations
- 4. Hierarchically growing field of view
- 5. Hierarchically progressing complexity
- 6. Reasonable amount of params

Convs rock!

- ✓ 1. Maintain 2D structure logic
- ✓ 2. Shift invariant (actually, equivariant)
3. Consider only local correlations
4. Hierarchically growing field of view
5. Hierarchically progressing complexity
6. Reasonable amount of params

Convs rock!

- ✓ 1. Maintain 2D structure logic
- ✓ 2. Shift invariant (actually, equivariant)
- ✓ 3. Consider only local correlations
- 4. Hierarchically growing field of view
- 5. Hierarchically progressing complexity
- 6. Reasonable amount of params

Convs rock!

- ✓ 1. Maintain 2D structure logic
- ✓ 2. Shift invariant (actually, equivariant)
- ✓ 3. Consider only local correlations
- ✓ 4. Hierarchically growing field of view
- 5. Hierarchically progressing complexity
- 6. Reasonable amount of params


Convs rock!

- ✓ 1. Maintain 2D structure logic
- ✓ 2. Shift invariant (actually, equivariant)
- ✓ 3. Consider only local correlations
- ✓ 4. Hierarchically growing field of view
- ✓ 5. Hierarchically progressing complexity
- 6. Reasonable amount of params

Convs rock!

- ✓ 1. Maintain 2D structure logic
- ✓ 2. Shift invariant (actually, equivariant)
- ✓ 3. Consider only local correlations
- ✓ 4. Hierarchically growing field of view
- ✓ 5. Hierarchically progressing complexity
- ✓ 6. Reasonable amount of params


Two important intuitions about feature maps



Two important intuitions about feature maps



H



C

W

Deep layer feature-map

Two important intuitions about feature maps



H

C

W



Deep layer feature-map




Two important intuitions about feature maps





H



Two important intuitions about feature maps





H




Deep layer feature-map

Two important intuitions about feature maps





Deep layer feature-map

Two important intuitions about feature maps





Deep layer feature-map

Two important intuitions about feature maps





Deep layer feature-map

Two important intuitions about feature maps



Two important intuitions about feature maps



Deep layer feature-map

Transposed Convolution

x_0^l

x_1^l

x_2^l

x_3^l

x_4^l

x_5^l

x_6^l

x_0^{l+1}

x_1^{l+1}


x_2^{l+1}

x_3^{l+1}


x_4^{l+1}

x_5^{l+1}



x_6^{l+1}





Transposed Convolution





Transposed Convolution





Transposed Convolution





Transposed Convolution





Transposed Convolution




Transposed Convolution



Transposed Convolution














Transposed Convolution



$\xrightarrow{\text{black}}$ a
 $\xrightarrow{\text{red}}$ b
 $\xrightarrow{\text{green}}$ c



$$\begin{bmatrix}
 \mathbf{b} & \mathbf{a} & 0 & 0 & 0 & 0 & 0 \\
 \mathbf{c} & \mathbf{b} & \mathbf{a} & 0 & 0 & 0 & 0 \\
 0 & \mathbf{c} & \mathbf{b} & \mathbf{a} & 0 & 0 & 0 \\
 0 & a & \mathbf{c} & \mathbf{b} & \mathbf{a} & 0 & 0 \\
 0 & 0 & a & \mathbf{c} & \mathbf{b} & \mathbf{a} & 0 \\
 0 & 0 & 0 & 0 & \mathbf{c} & \mathbf{b} & \mathbf{a} \\
 0 & 0 & 0 & 0 & 0 & \mathbf{c} & \mathbf{b}
 \end{bmatrix}$$

Transposed Convolution with stride



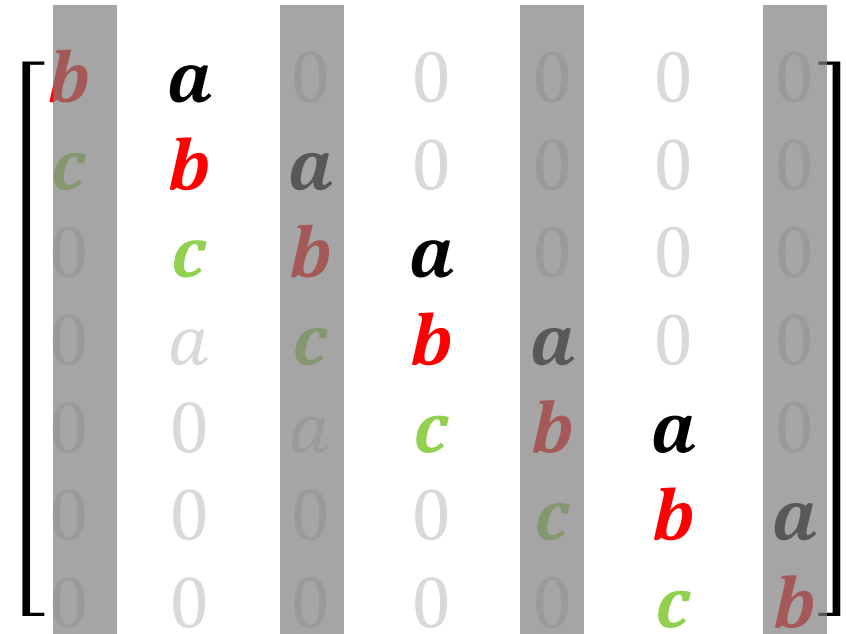


$$\begin{bmatrix}
 \mathbf{b} & \mathbf{a} & 0 & 0 & 0 & 0 & 0 \\
 \mathbf{c} & \mathbf{b} & \mathbf{a} & 0 & 0 & 0 & 0 \\
 0 & \mathbf{c} & \mathbf{b} & \mathbf{a} & 0 & 0 & 0 \\
 0 & a & \mathbf{c} & \mathbf{b} & \mathbf{a} & 0 & 0 \\
 0 & 0 & a & \mathbf{c} & \mathbf{b} & \mathbf{a} & 0 \\
 0 & 0 & 0 & 0 & \mathbf{c} & \mathbf{b} & \mathbf{a} \\
 0 & 0 & 0 & 0 & 0 & \mathbf{c} & \mathbf{b}
 \end{bmatrix}$$

Transposed Convolution with stride





$$\begin{bmatrix}
 \mathbf{b} & \mathbf{a} & 0 & 0 & 0 & 0 & 0 \\
 \mathbf{c} & \mathbf{b} & \mathbf{a} & 0 & 0 & 0 & 0 \\
 0 & \mathbf{c} & \mathbf{b} & \mathbf{a} & 0 & 0 & 0 \\
 0 & a & \mathbf{c} & \mathbf{b} & \mathbf{a} & 0 & 0 \\
 0 & 0 & a & \mathbf{c} & \mathbf{b} & \mathbf{a} & 0 \\
 0 & 0 & 0 & 0 & \mathbf{c} & \mathbf{b} & \mathbf{a} \\
 0 & 0 & 0 & 0 & 0 & \mathbf{c} & \mathbf{b}
 \end{bmatrix}$$

Transposed Convolution with stride




Transposed Convolution with stride



$$\begin{bmatrix}
 \mathbf{a} & 0 & 0 \\
 \mathbf{b} & 0 & 0 \\
 \mathbf{c} & \mathbf{a} & 0 \\
 a & \mathbf{b} & 0 \\
 0 & \mathbf{c} & \mathbf{a} \\
 0 & 0 & \mathbf{b} \\
 0 & 0 & \mathbf{c}
 \end{bmatrix}$$

Transposed Convolution with stride



$\xrightarrow{\text{black}}$ a
 $\xrightarrow{\text{red}}$ b
 $\xrightarrow{\text{green}}$ c


$$\begin{bmatrix} \mathbf{a} & 0 & 0 \\ \mathbf{b} & 0 & 0 \\ \mathbf{c} & \mathbf{a} & 0 \\ a & \mathbf{b} & 0 \\ 0 & \mathbf{c} & \mathbf{a} \\ 0 & 0 & \mathbf{b} \\ 0 & 0 & \mathbf{c} \end{bmatrix}$$

Recall- stride 2 conv:

$$\begin{bmatrix} \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 & 0 & 0 & 0 \\ 0 & 0 & \mathbf{a} & \mathbf{b} & \mathbf{c} & 0 & 0 \\ 0 & 0 & 0 & 0 & \mathbf{a} & \mathbf{b} & \mathbf{c} \end{bmatrix}$$



Transposed Convolution with stride



$$\begin{bmatrix} a & 0 & 0 \\ b & 0 & 0 \\ c & a & 0 \\ a & b & 0 \\ 0 & c & a \\ 0 & 0 & b \\ 0 & 0 & c \end{bmatrix}$$


Recall- stride 2 conv:

$$\begin{bmatrix} a & b & c & 0 & 0 & 0 & 0 \\ 0 & 0 & a & b & c & 0 & 0 \\ 0 & 0 & 0 & 0 & a & b & c \end{bmatrix}$$

A note about the implementation of conv




A note about the implementation of conv




$[N, C_{in}, H, W]$

A note about the implementation of conv

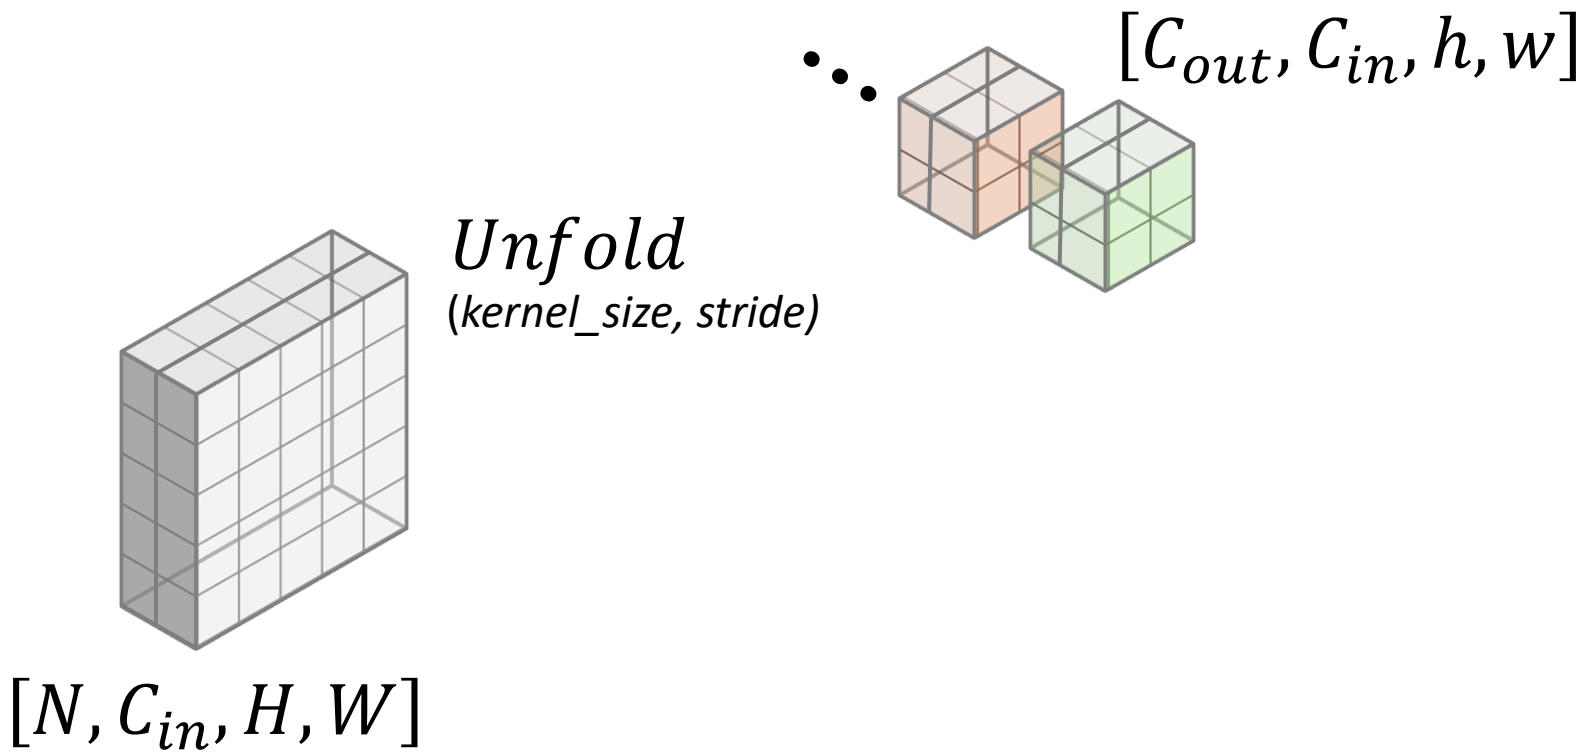


$[N, C_{in}, H, W]$




$[C_{out}, C_{in}, h, w]$


A note about the implementation of conv



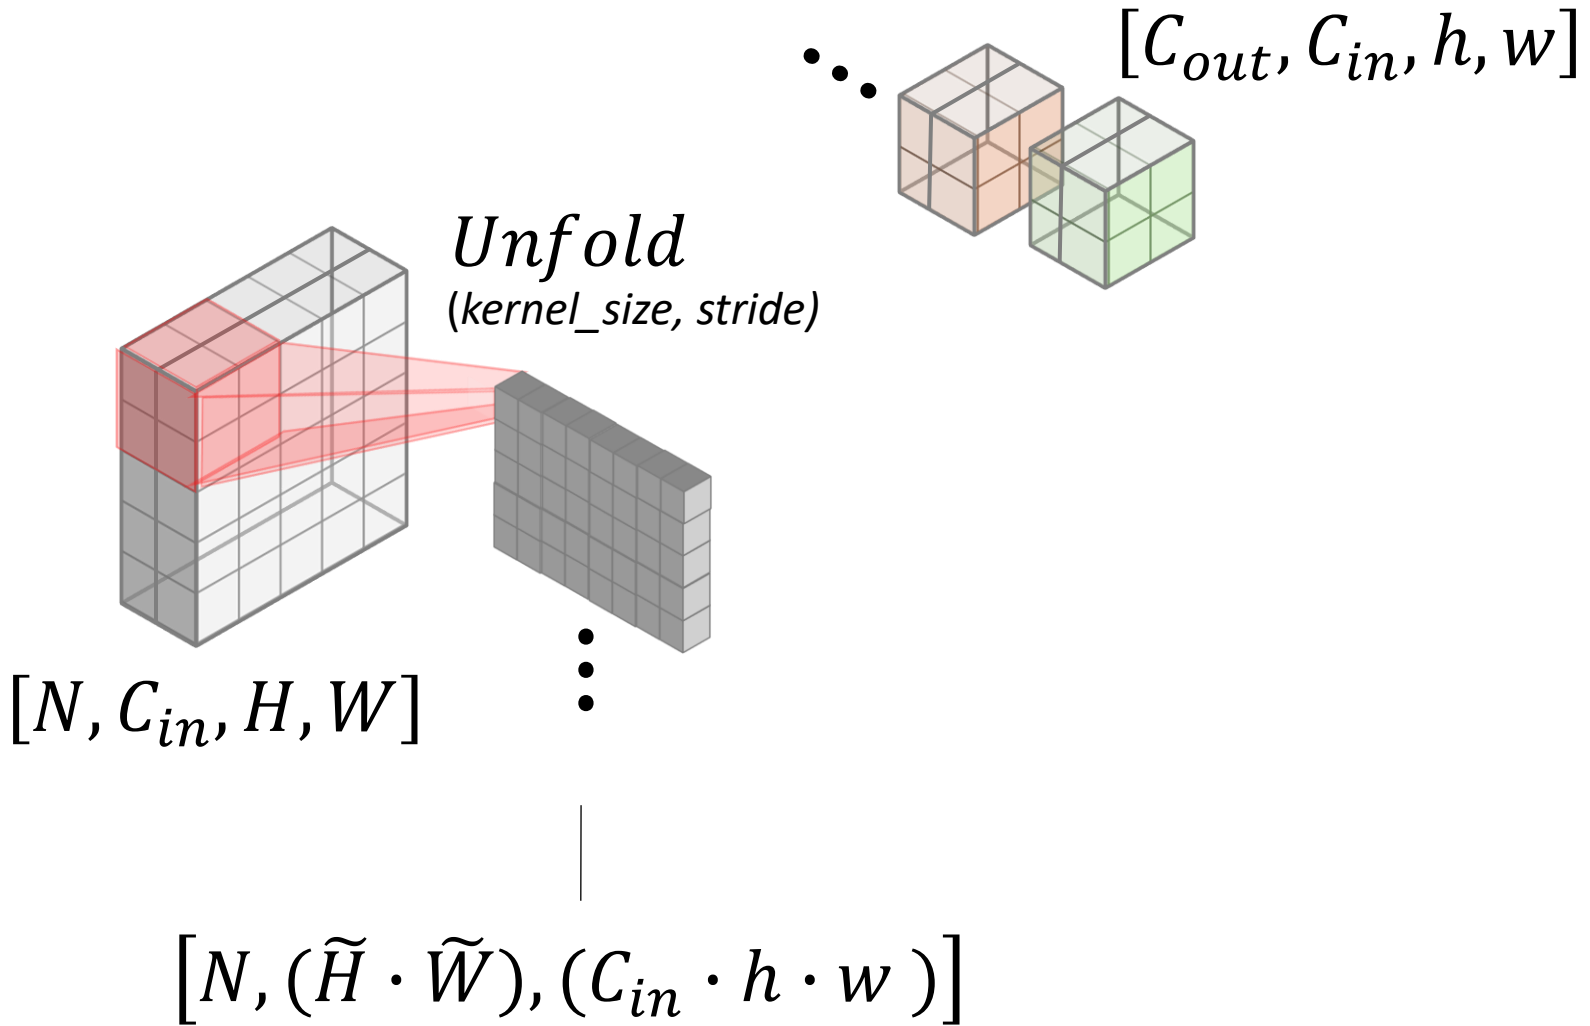
A note about the implementation of conv



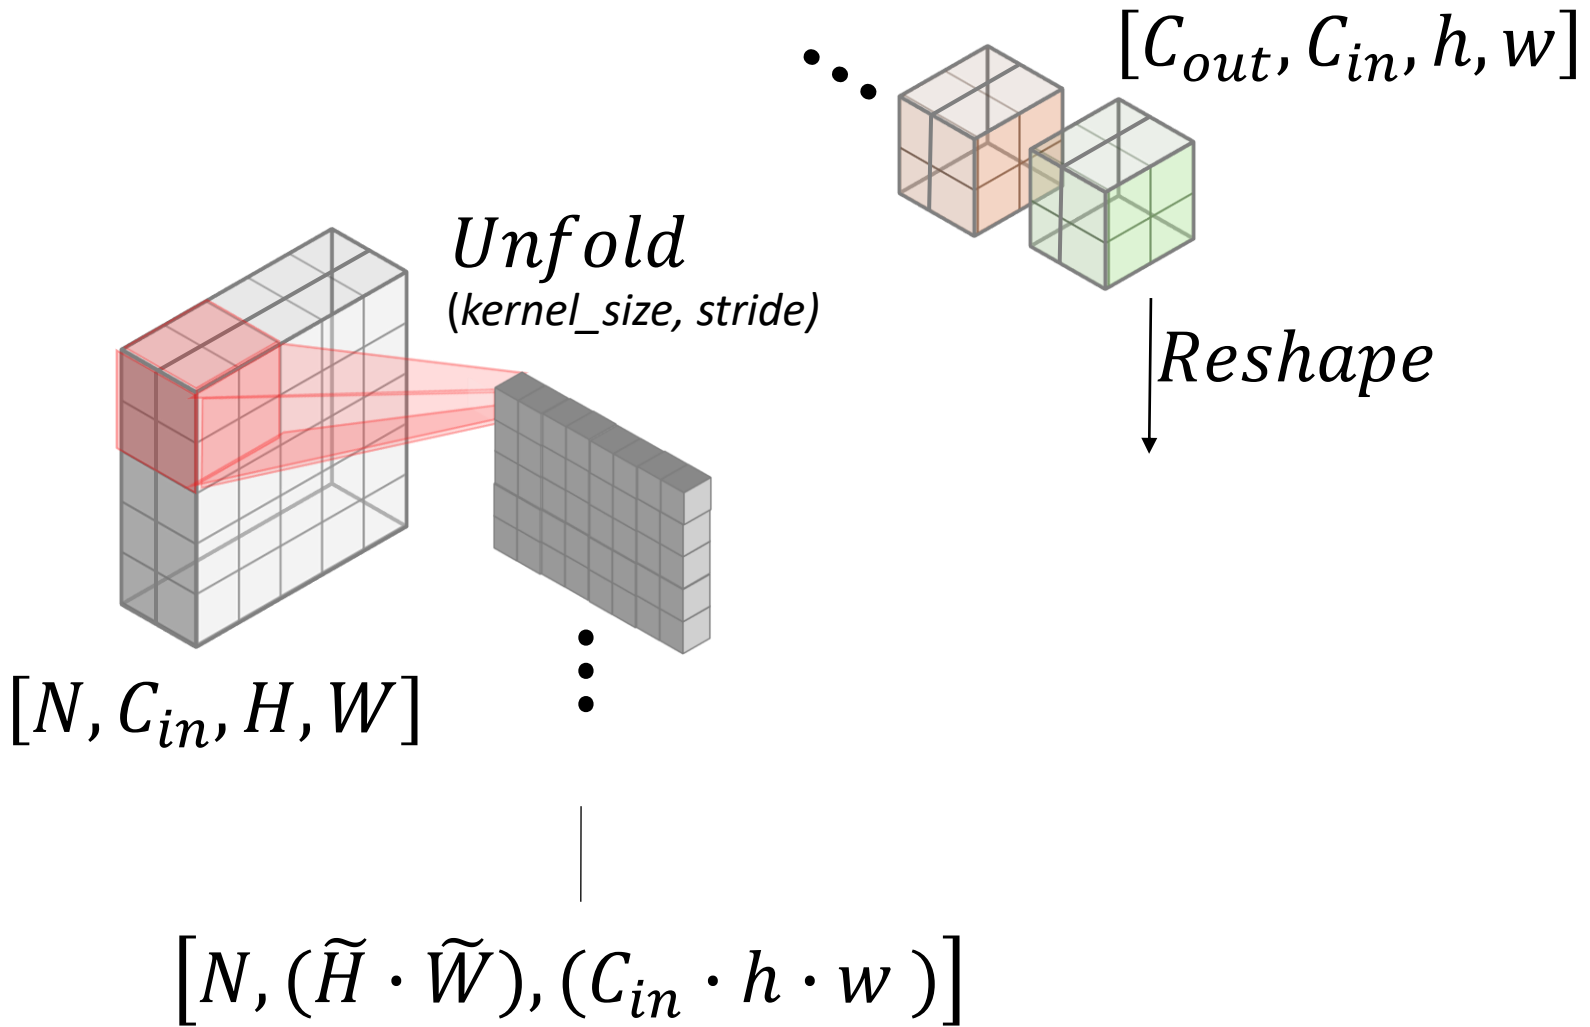
A note about the implementation of conv



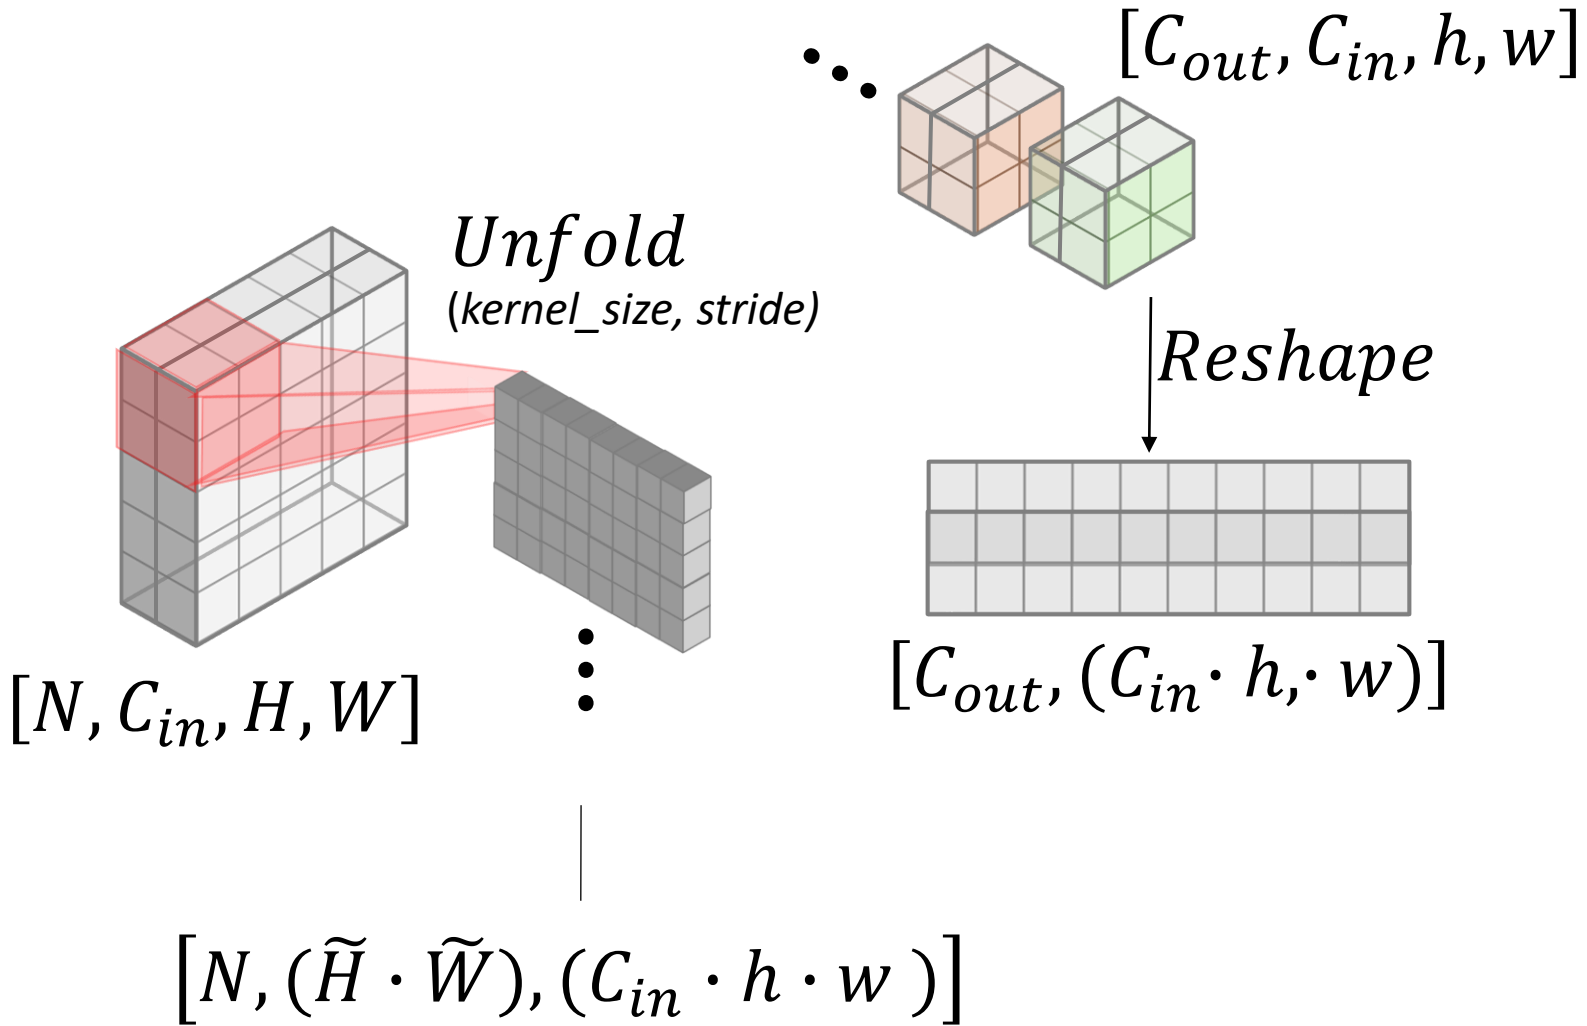
A note about the implementation of conv



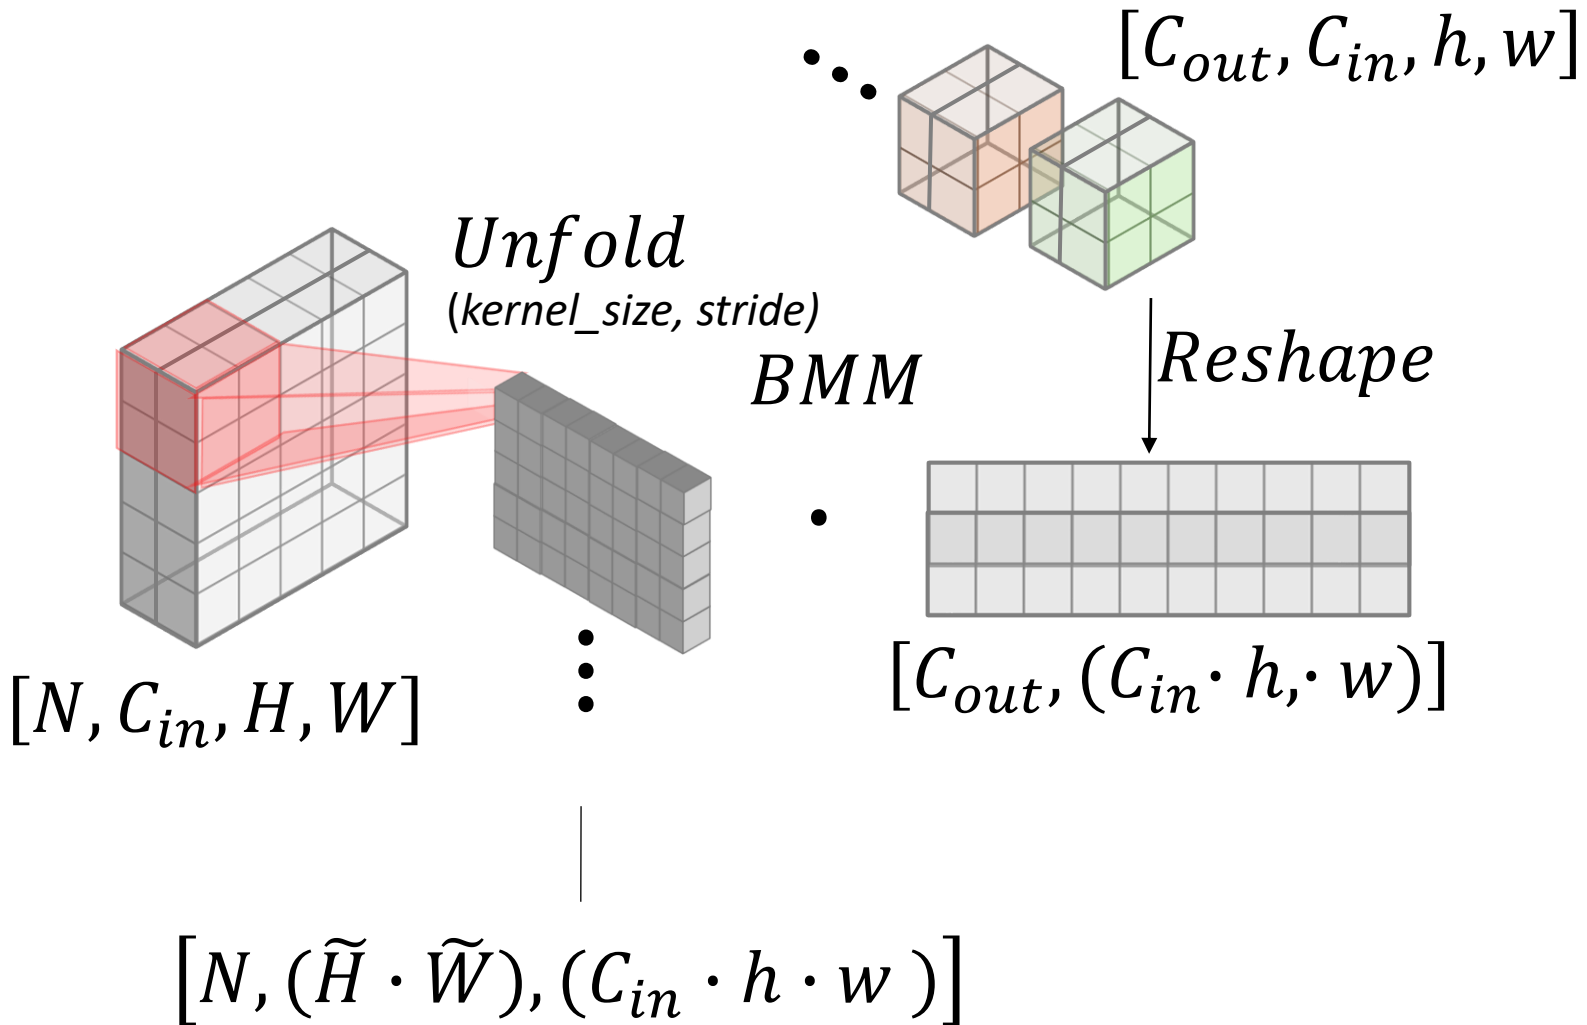
A note about the implementation of conv



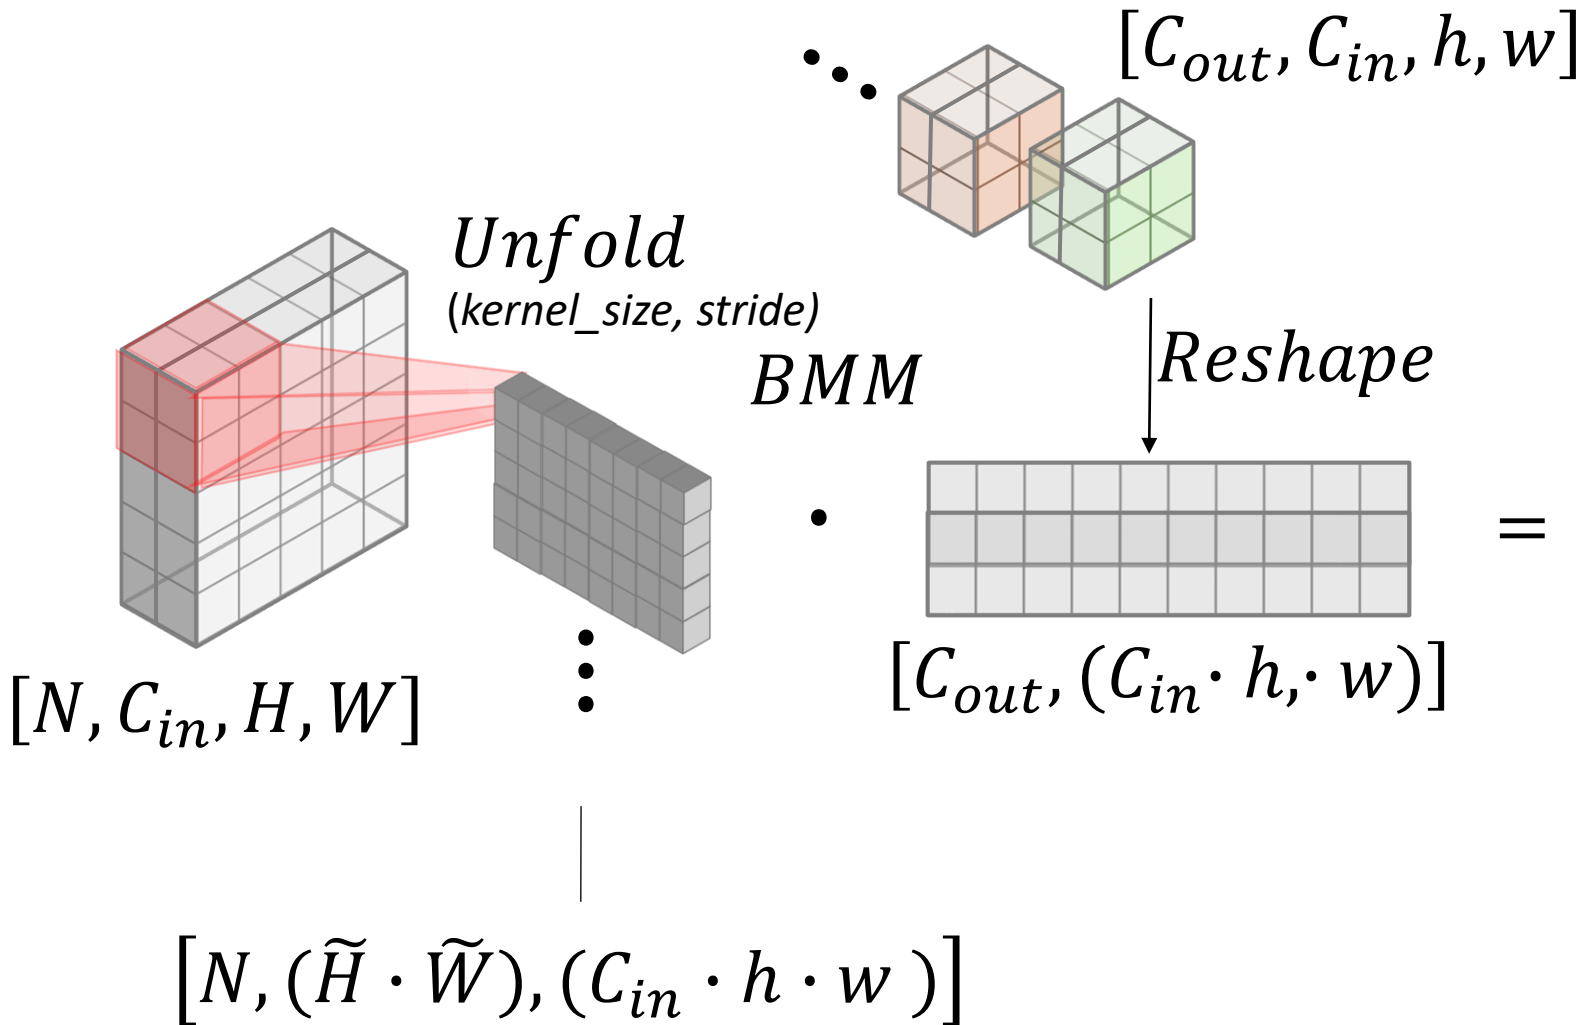
A note about the implementation of conv



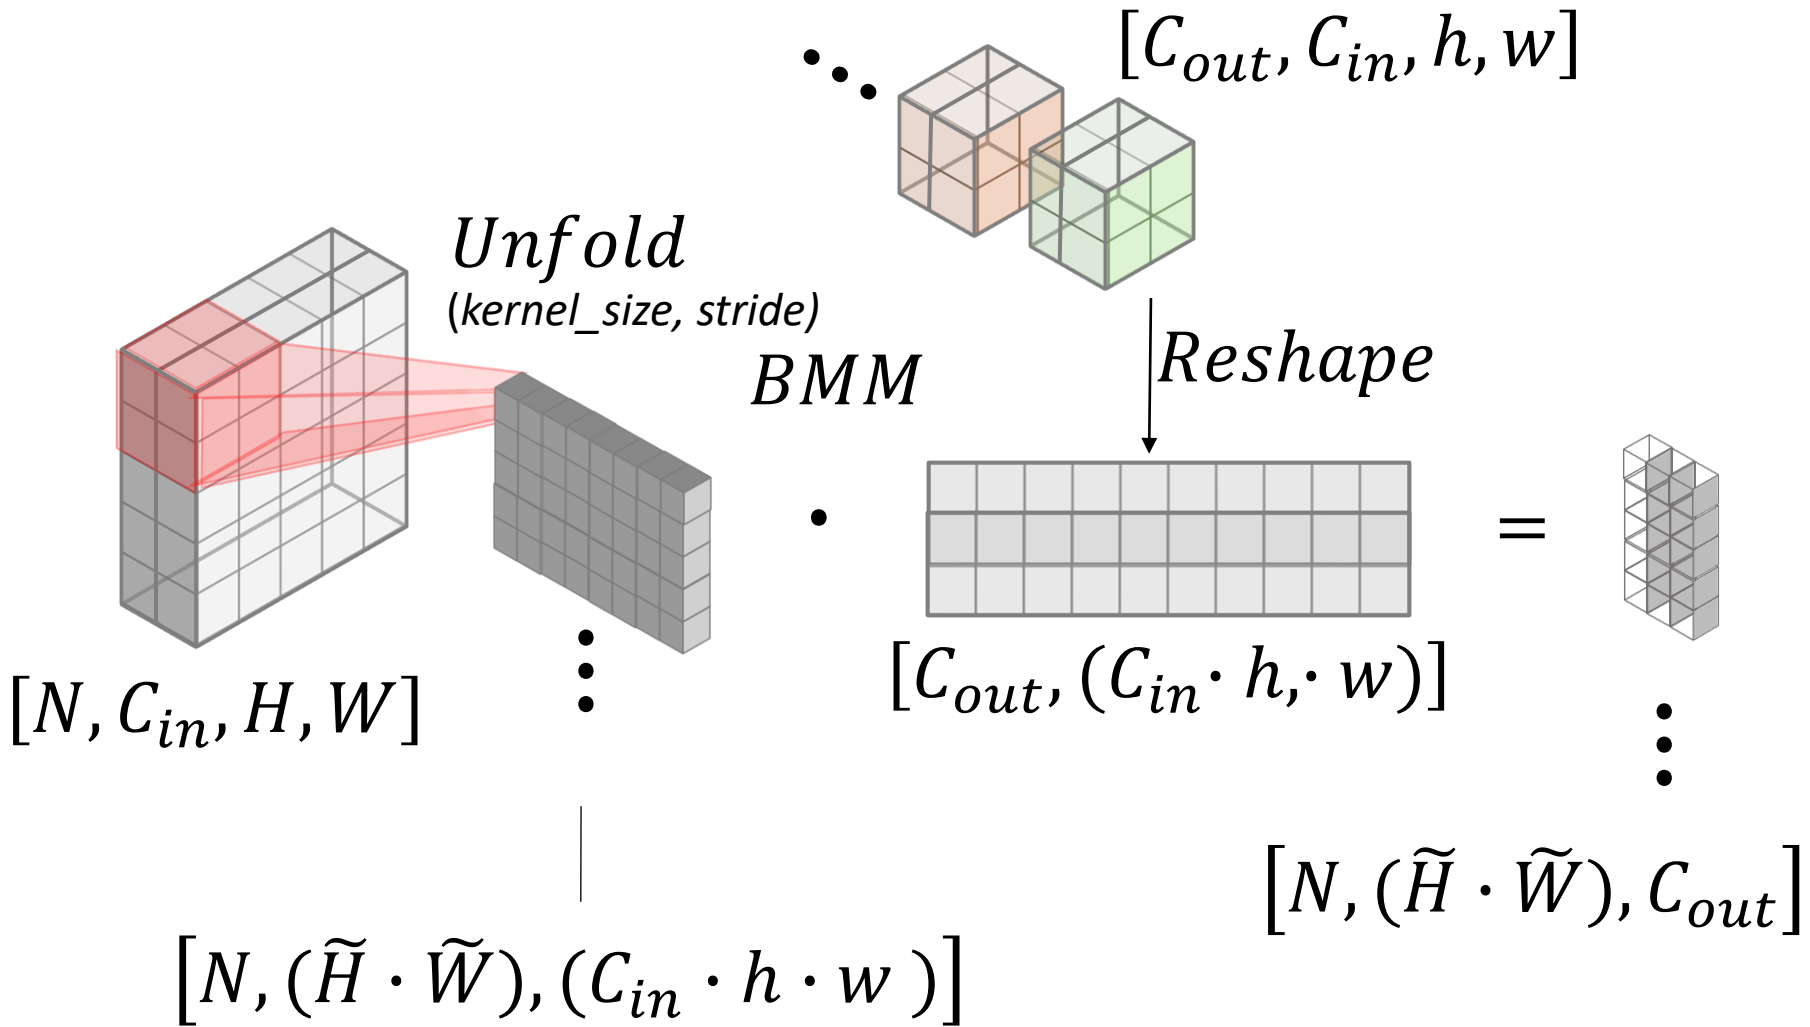
A note about the implementation of conv



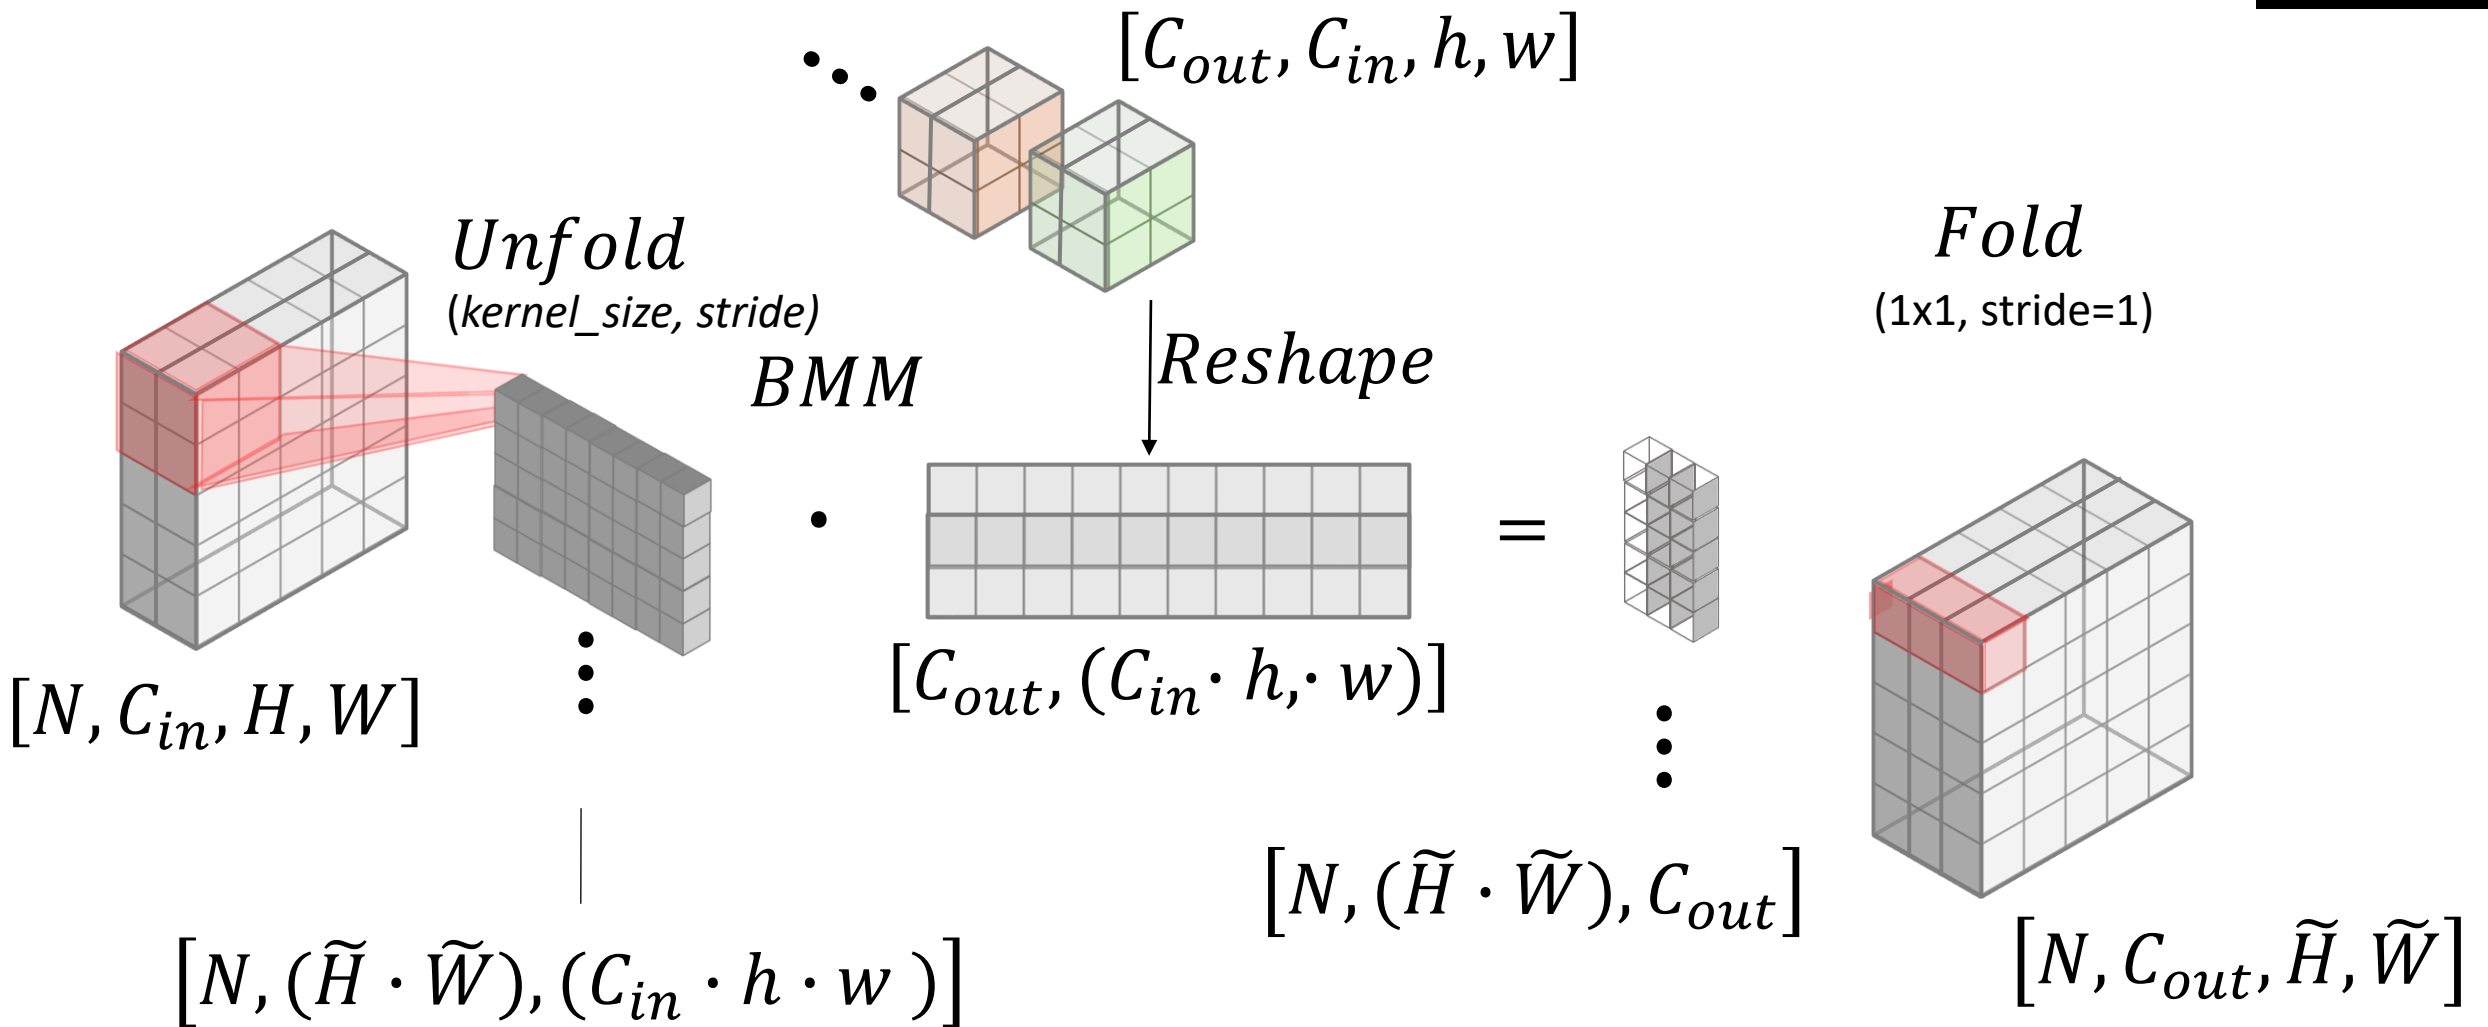
A note about the implementation of conv




A note about the implementation of conv




A note about the implementation of conv




A note about the implementation of conv




A note about the implementation of conv



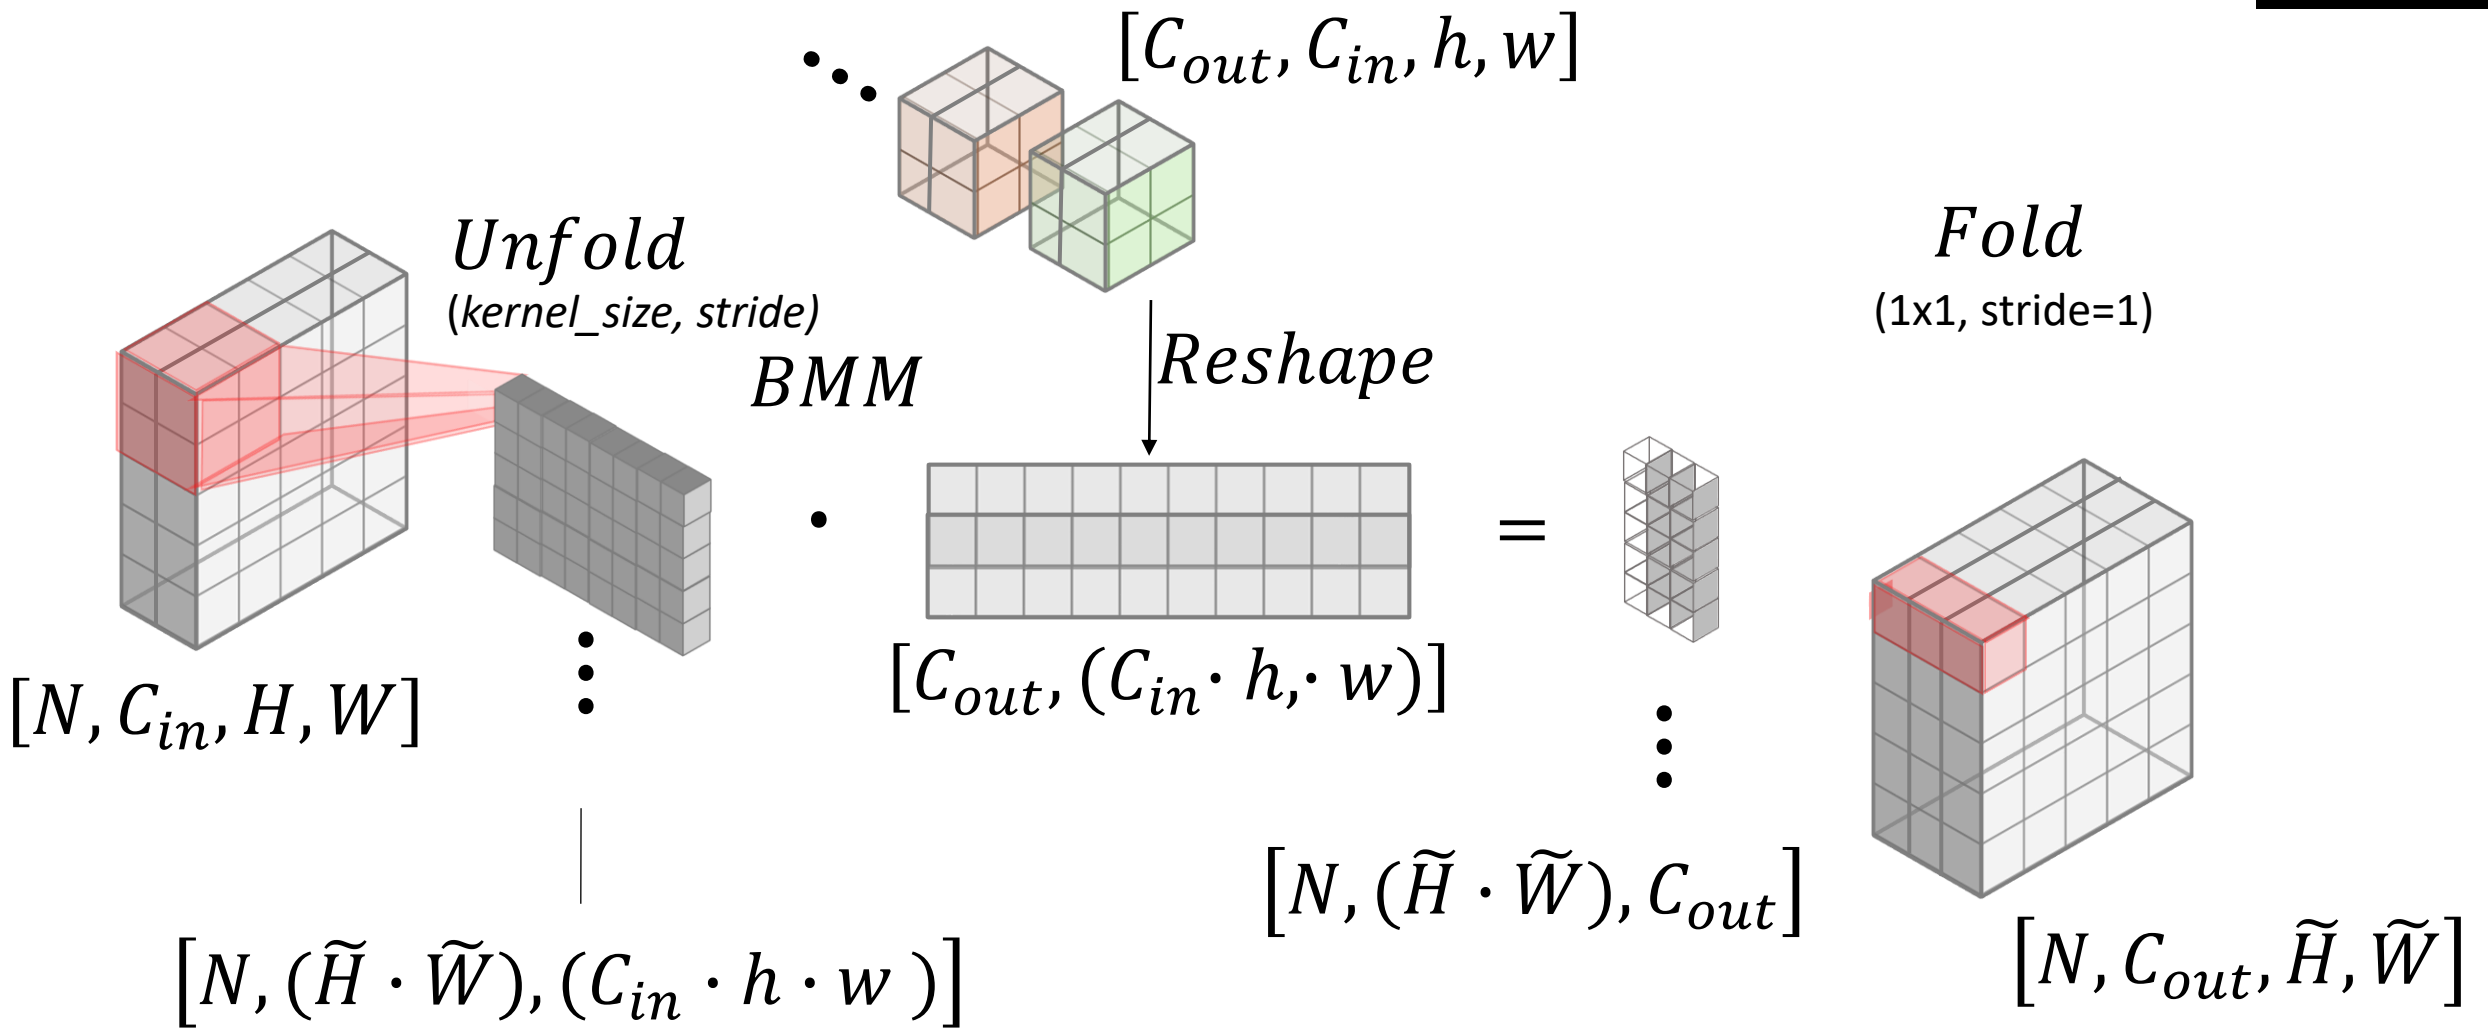
A note about the implementation of conv




A note about the implementation of conv




A note about the implementation of conv



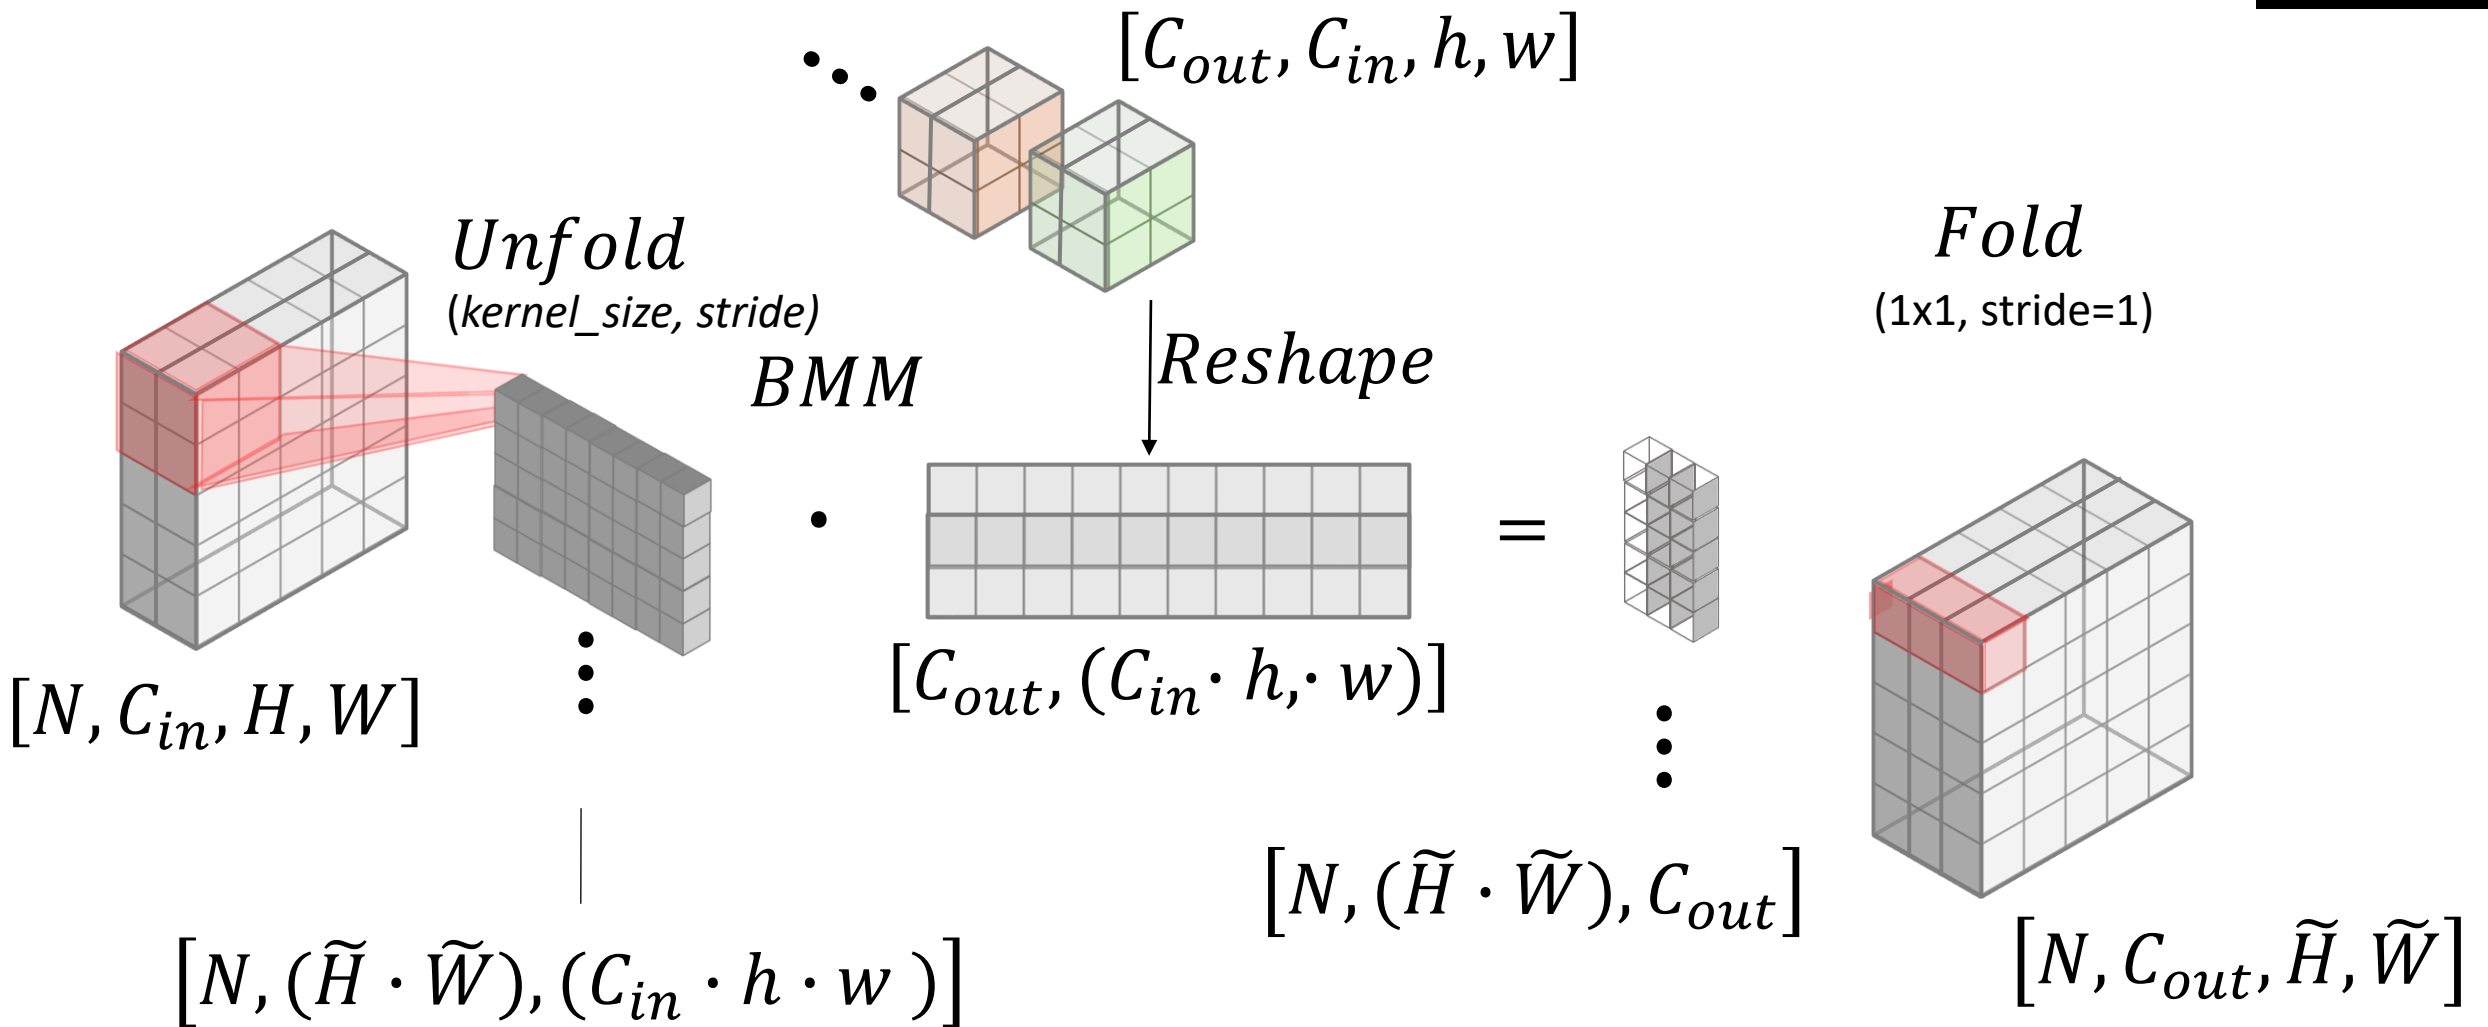
A note about the implementation of conv



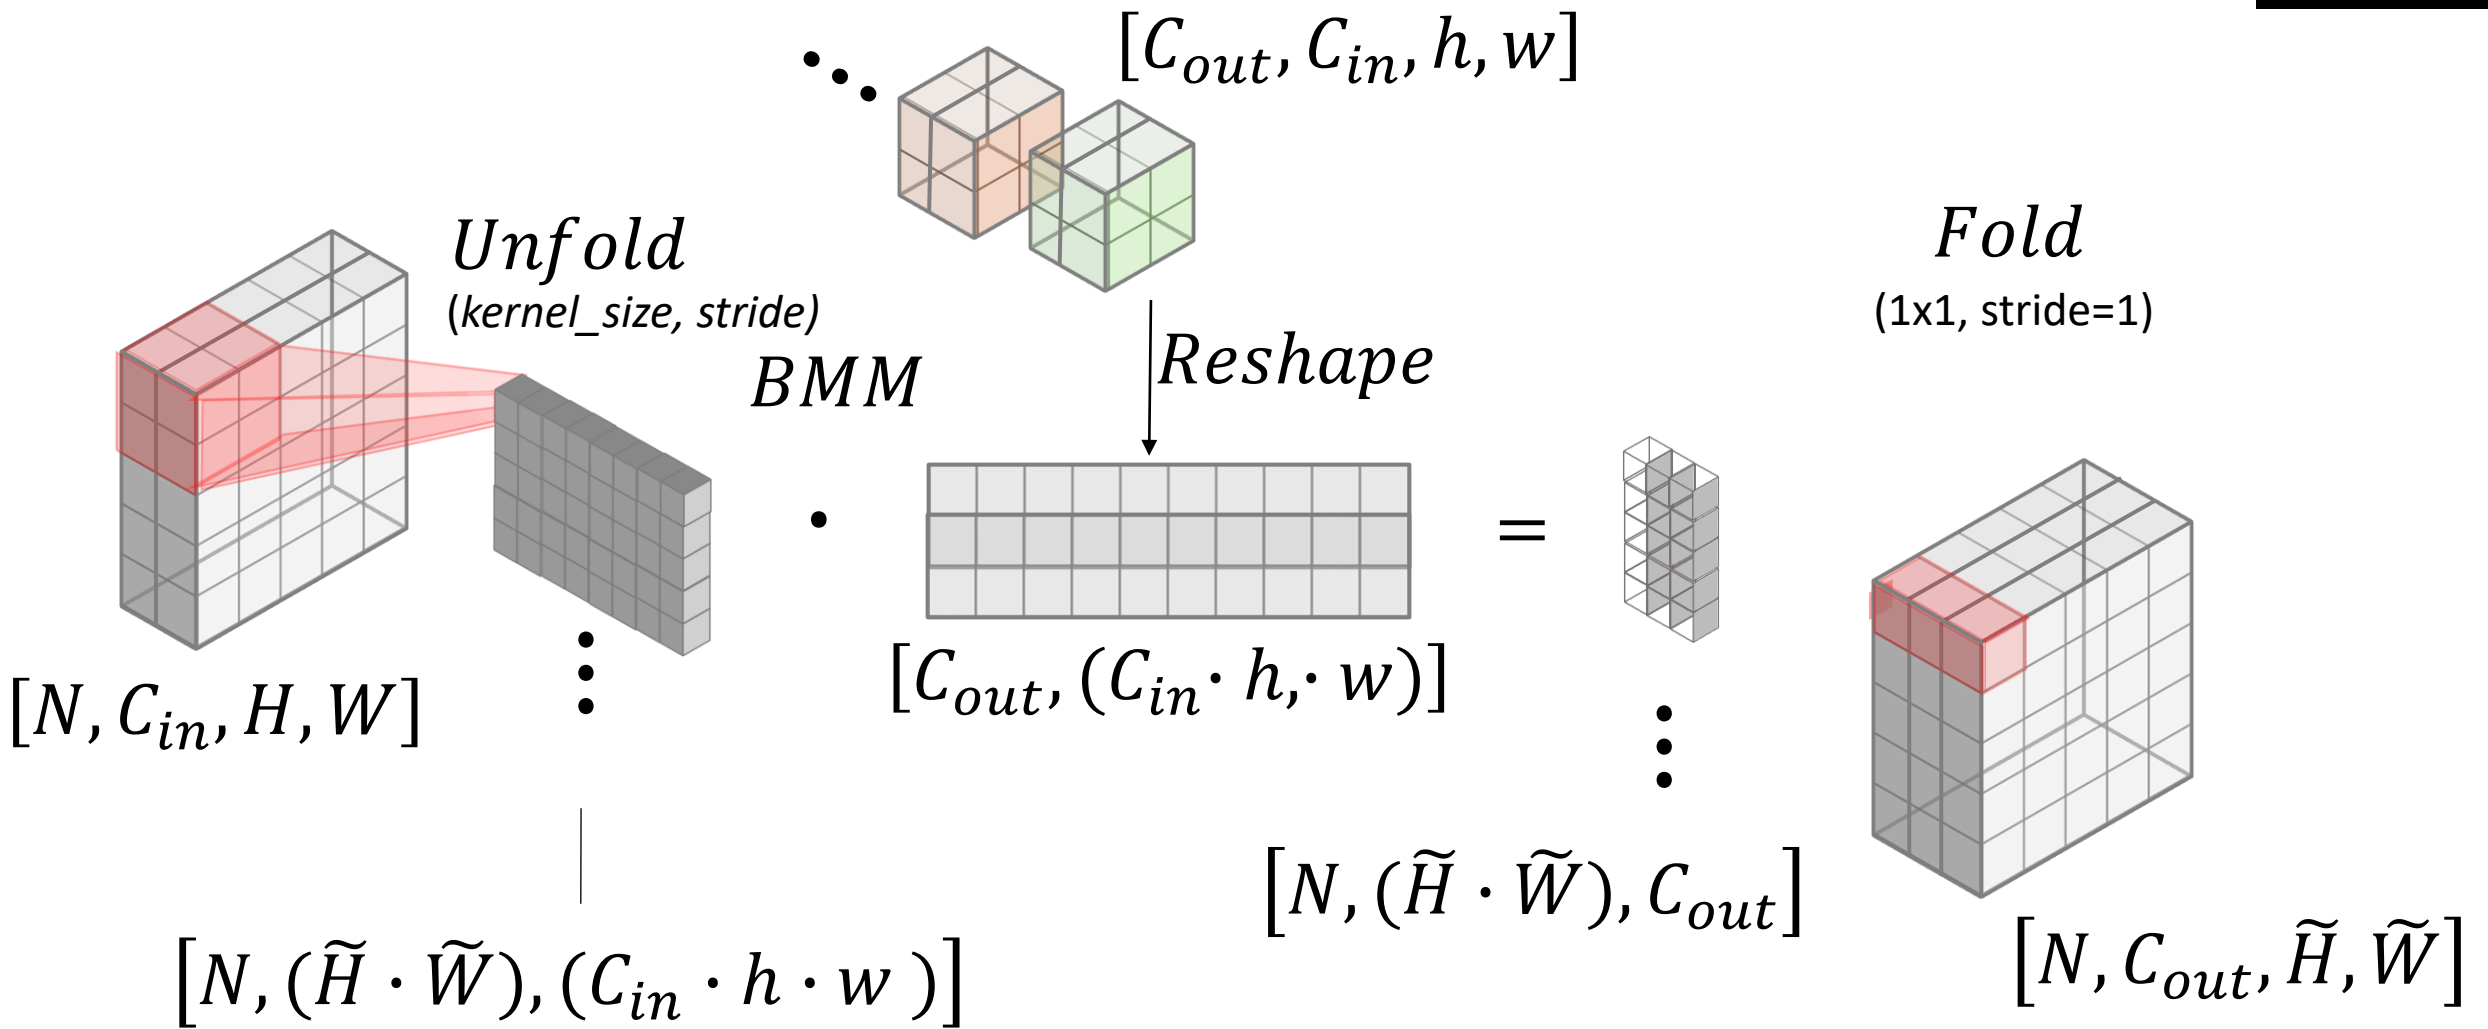
A note about the implementation of conv




A note about the implementation of conv



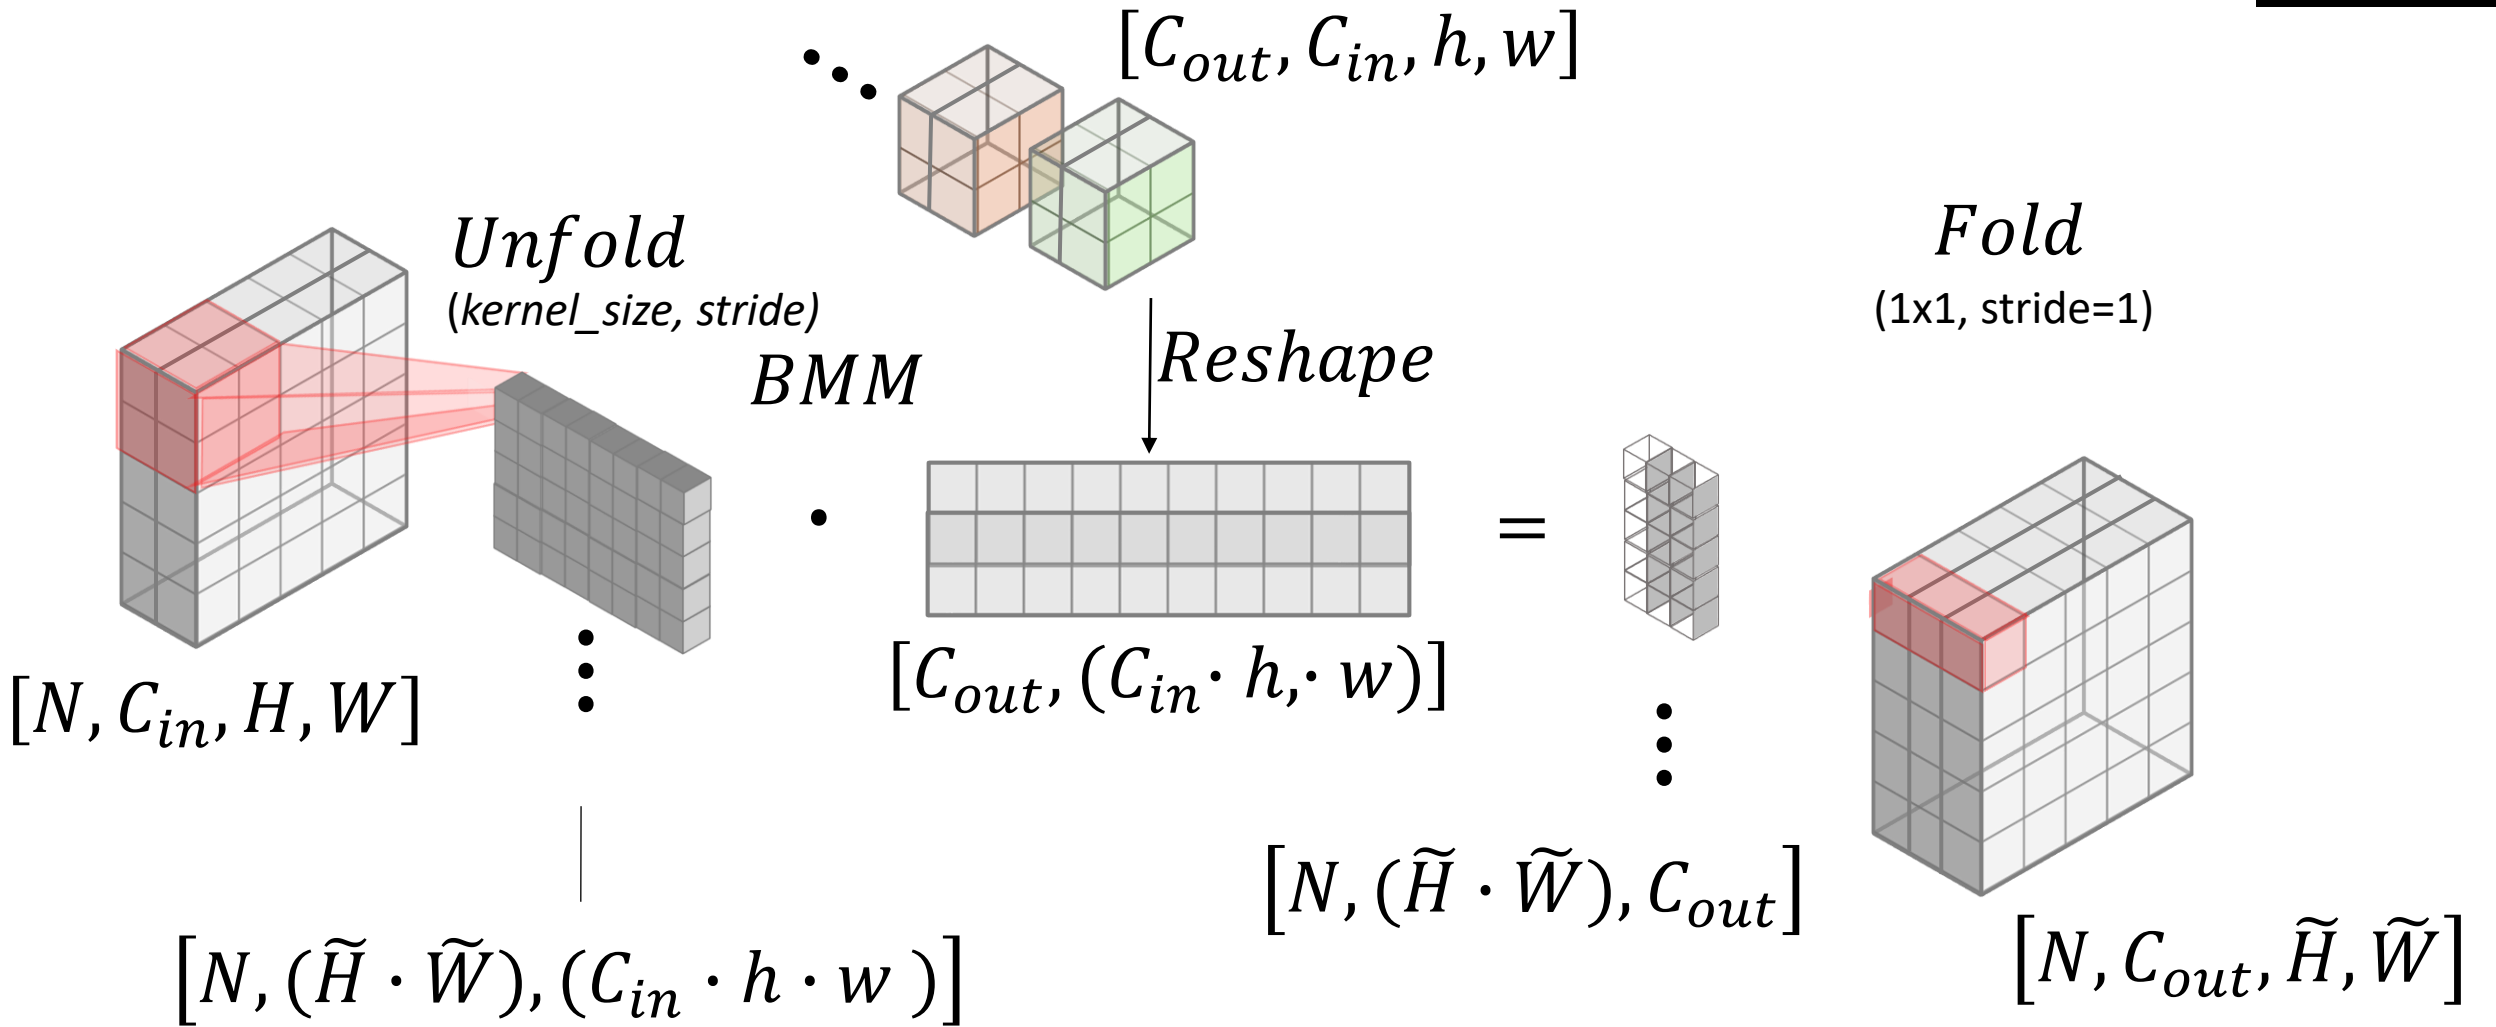
A note about the implementation of conv



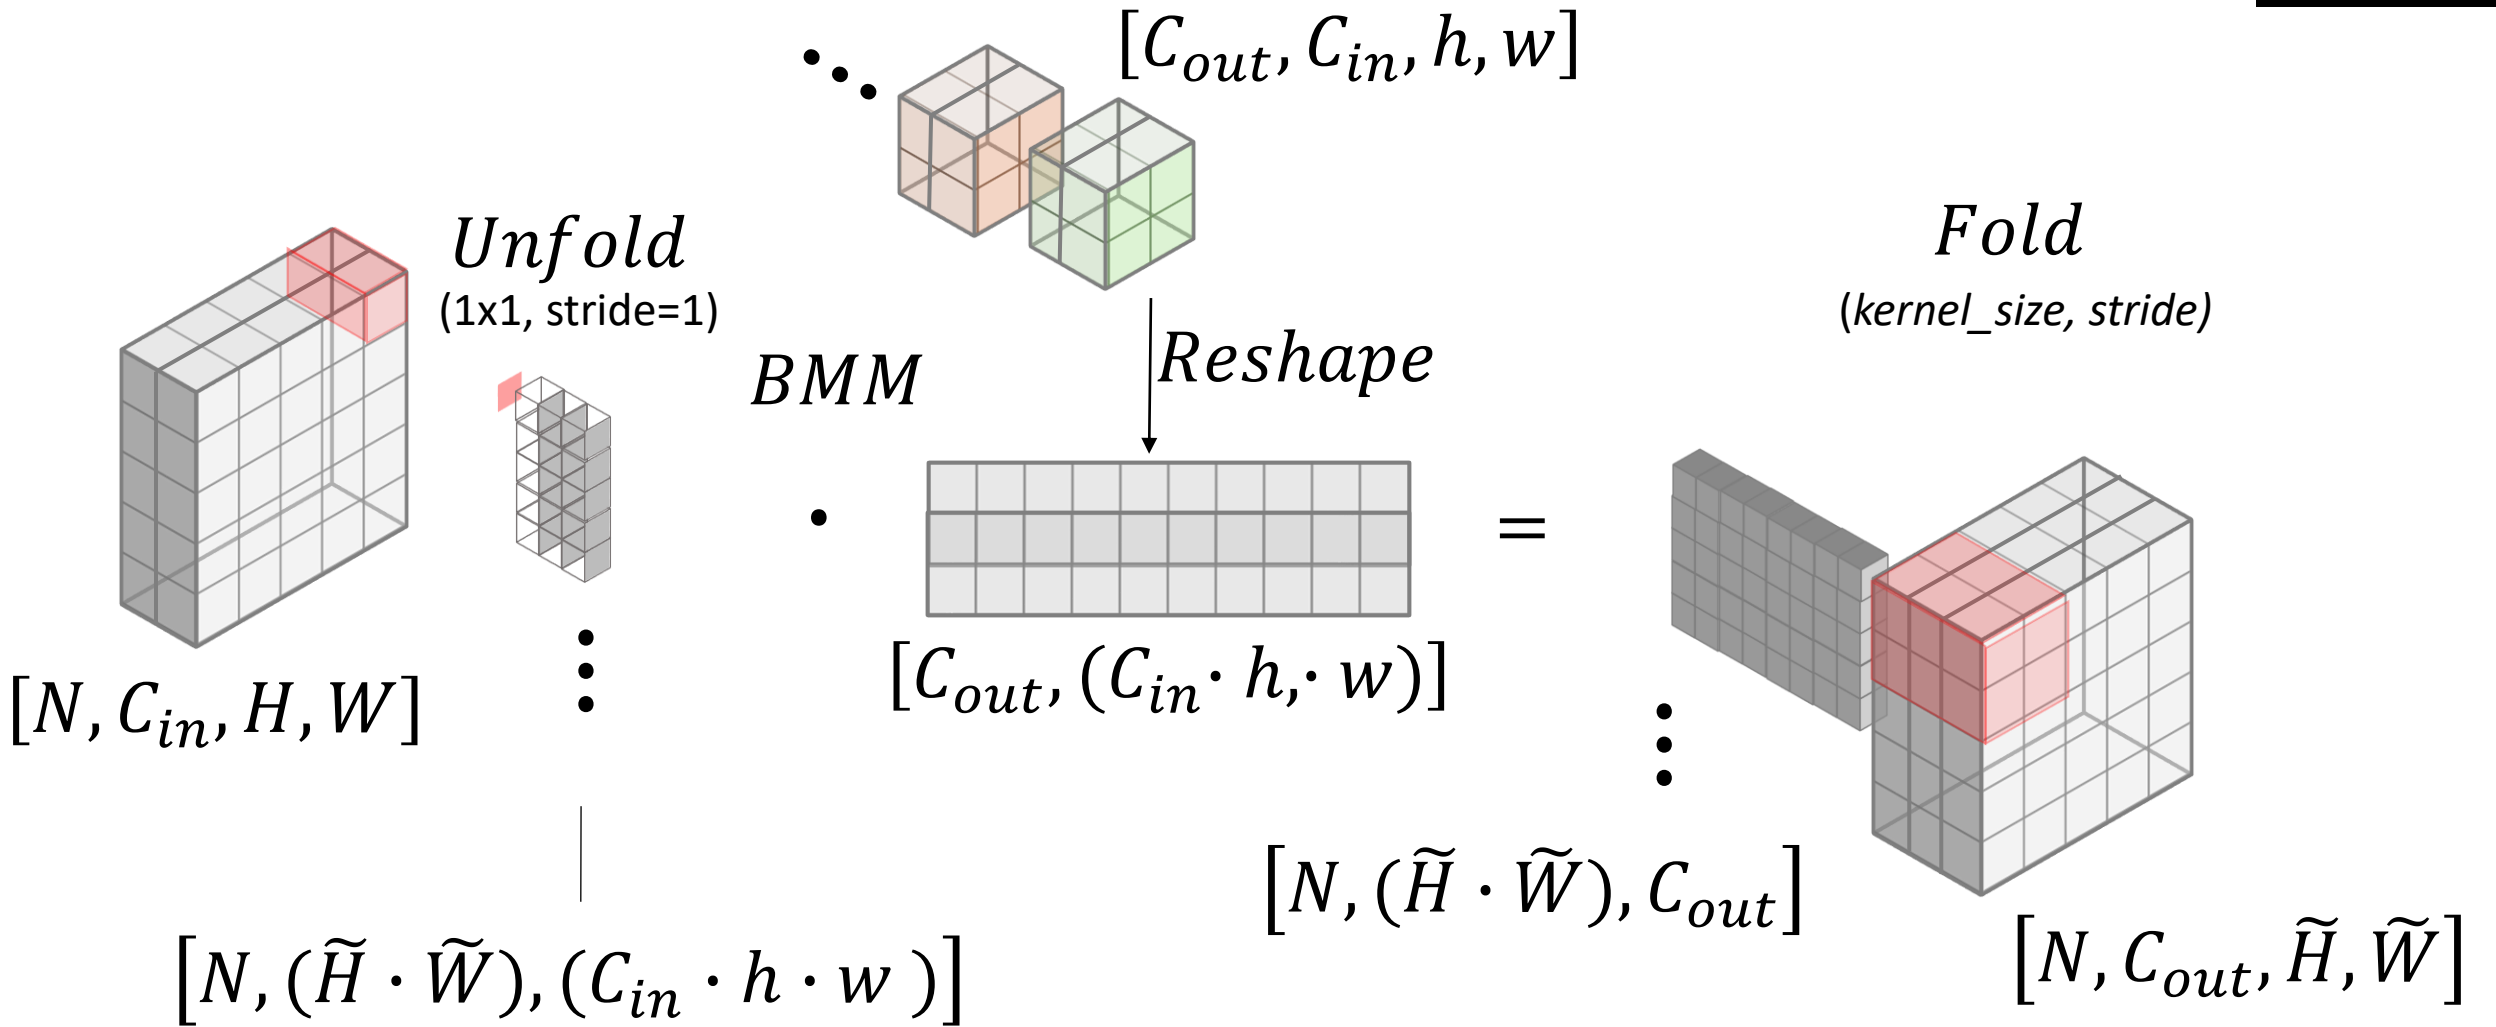
A note about the implementation of conv



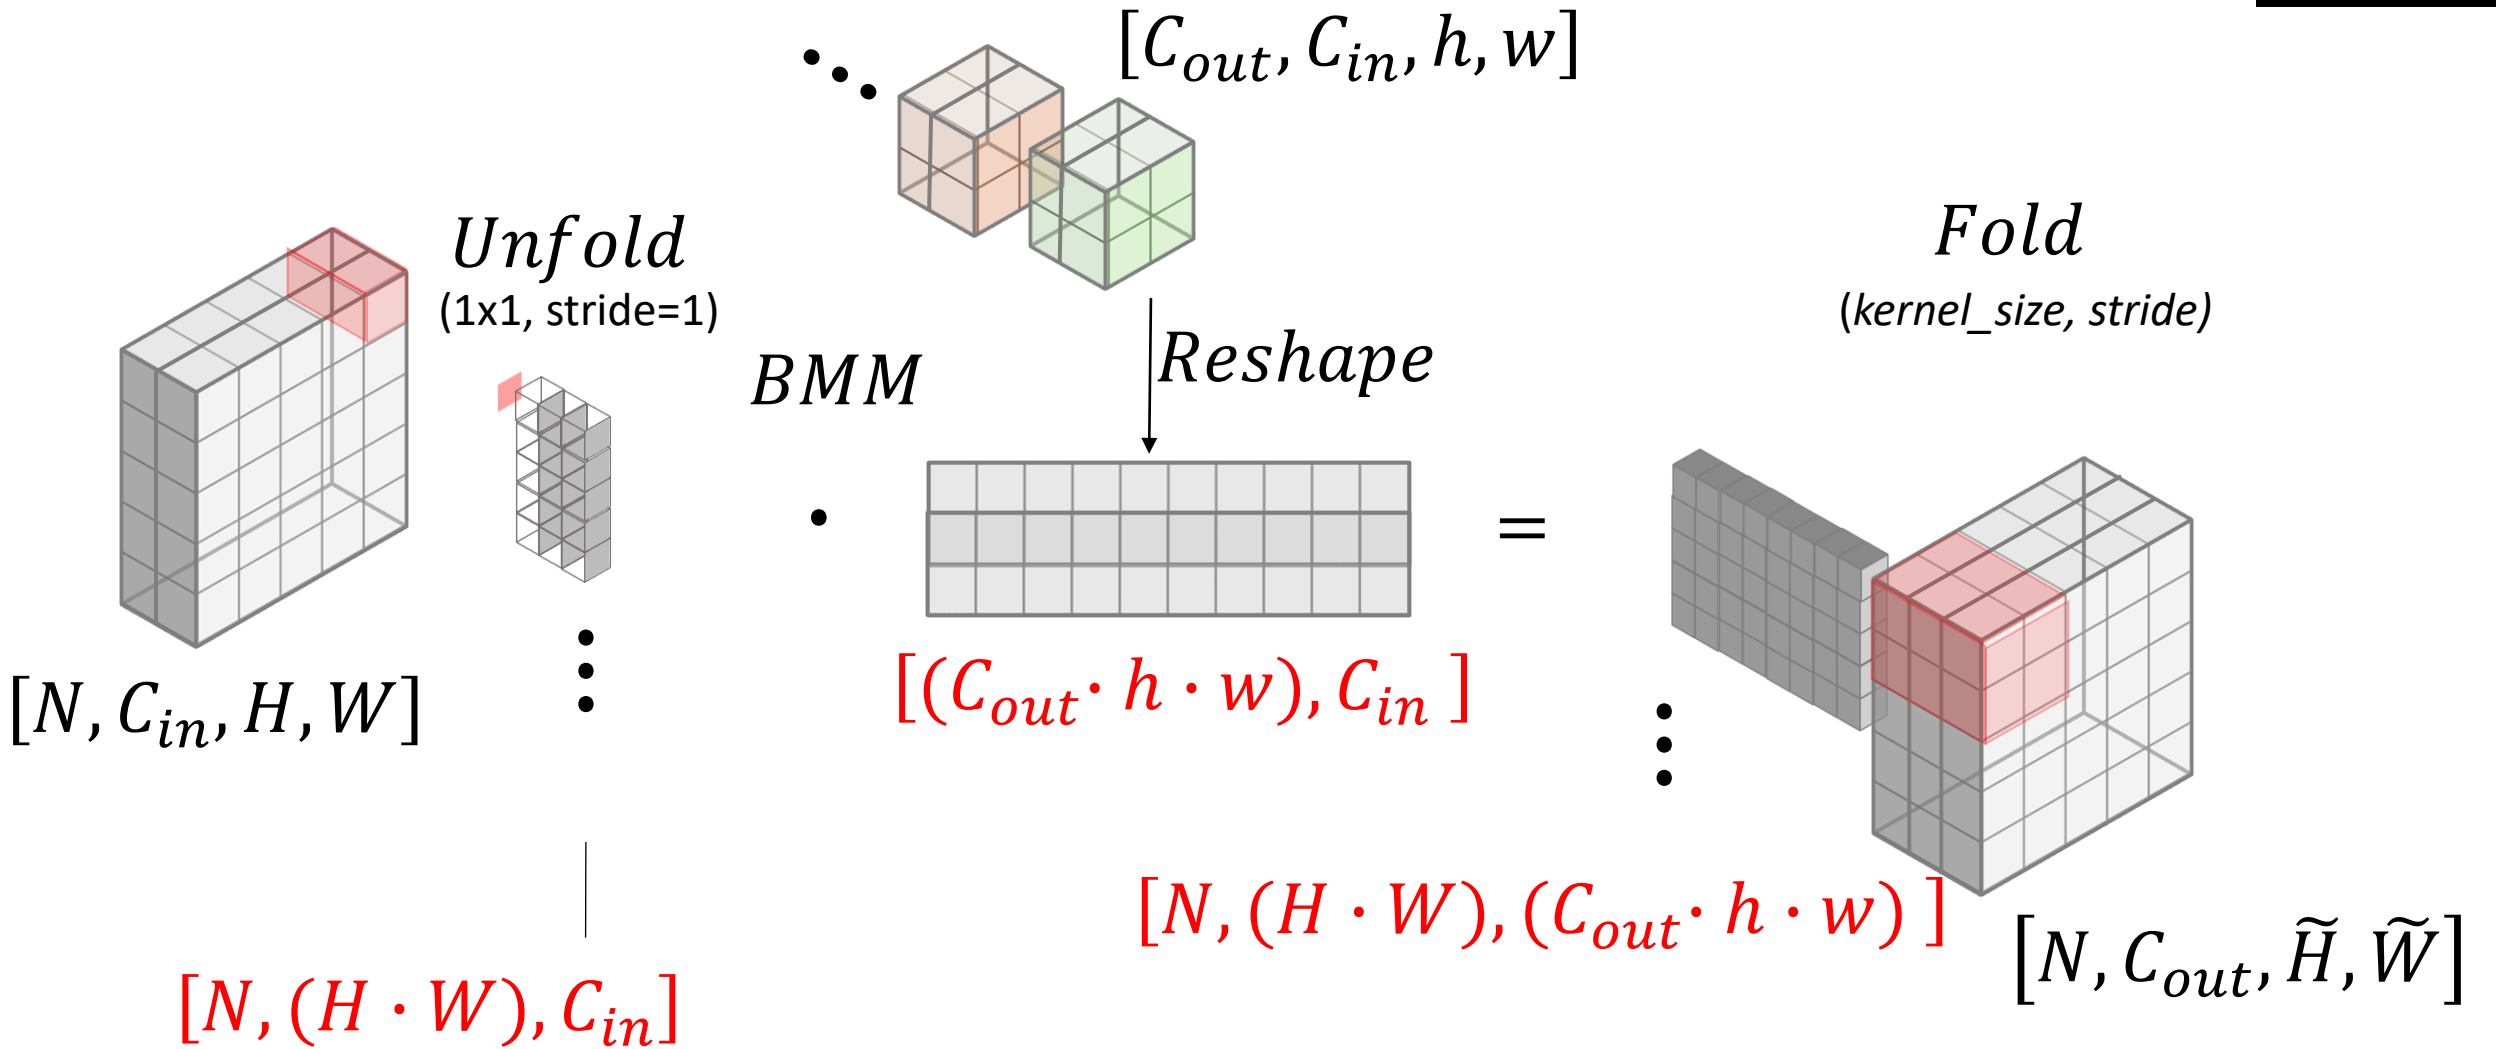
A note about the implementation of conv *TRANSPPOSED*



A note about the implementation of conv *TRANSPPOSED*



A note about the implementation of conv *TRANSPPOSED*



Q: How do you backprop a Conv? (1D, single channel)



Q: How do you backprop a Conv? (1D, single channel)

For derivative of loss w.r.t all neurons:
Transposed Conv with the same filter!

(Conv is a special case of linear)

Q: How do you backprop a Conv? (1D, single channel)

For derivative of loss w.r.t all neurons:
Transposed Conv with the same filter!

(Conv is a special case of linear)

Q: How about Conv2D?

Q: How do you backprop a Conv? (1D, single channel)

For derivative of loss w.r.t all neurons:
Transposed Conv with the same filter!

(Conv is a special case of linear)


Q: How about Conv2D?

For derivative of loss w.r.t all neurons:


1. Transpose dims of filter c_{in}, c_{out} (as in FC)
2. Transposed Conv2D with the modified filter!

(Think of a conv as a linear layer applied to each patch separately,
going from c_{in} to c_{out})


A note about the implementation of conv - BACKWARD




A note about the implementation of conv - BACKWARD



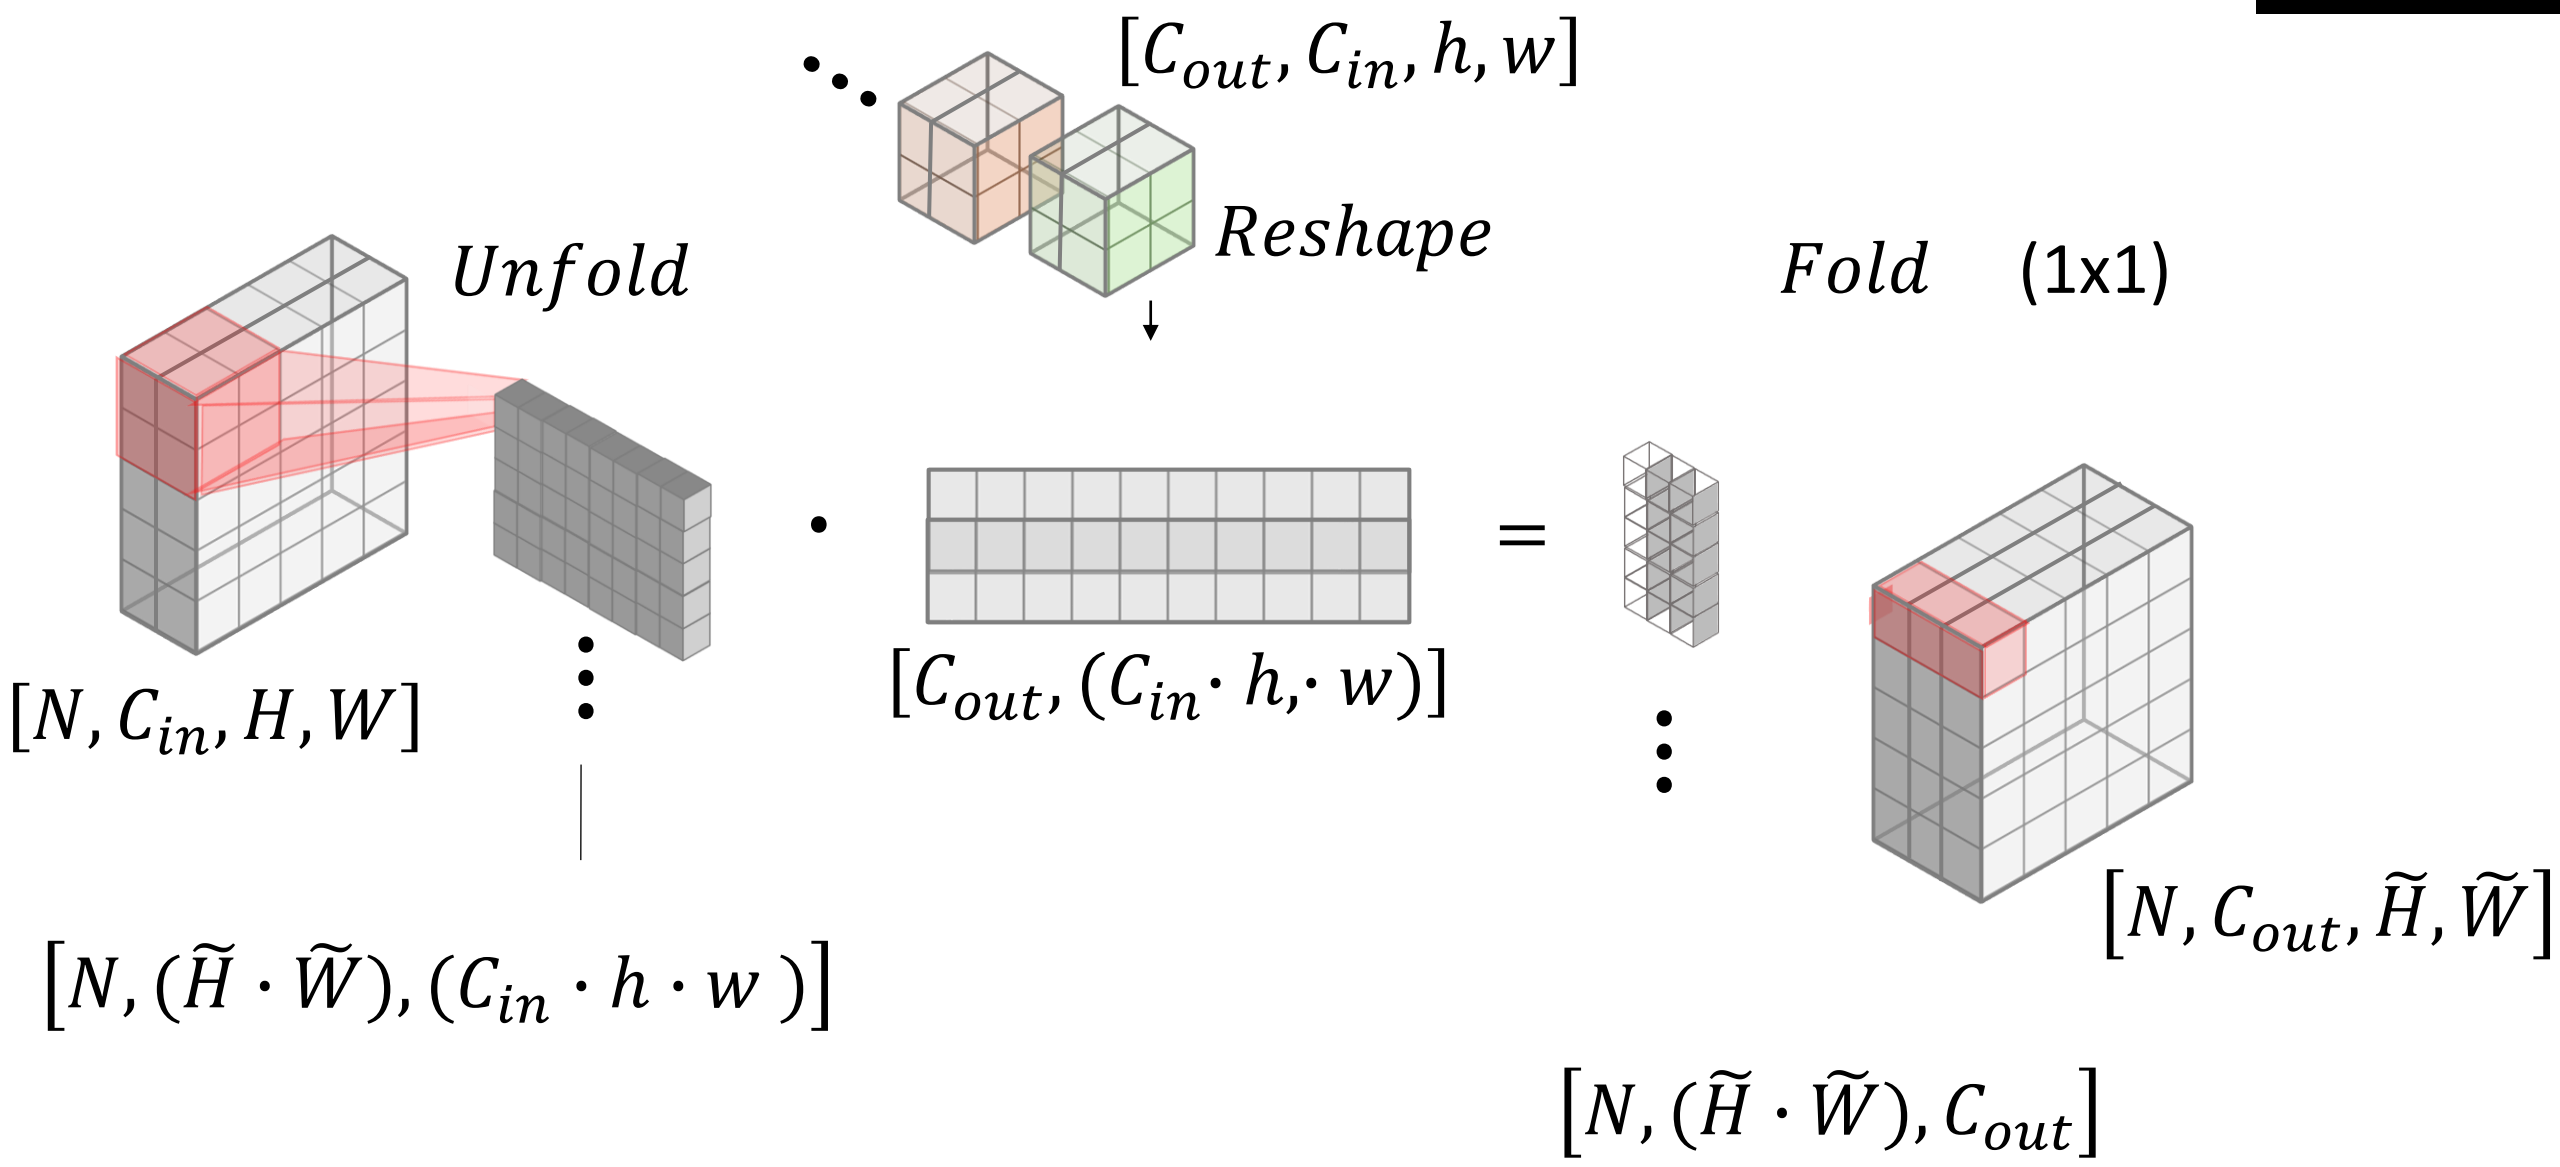
A note about the implementation of conv - BACKWARD




A note about the implementation of conv - BACKWARD




A note about the implementation of conv - BACKWARD




A note about the implementation of conv - BACKWARD




A note about the implementation of conv - BACKWARD




A note about the implementation of conv - BACKWARD




A note about the implementation of conv - BACKWARD




A note about the implementation of conv - BACKWARD




A note about the implementation of conv - BACKWARD




A note about the implementation of conv - BACKWARD




A note about the implementation of conv - BACKWARD



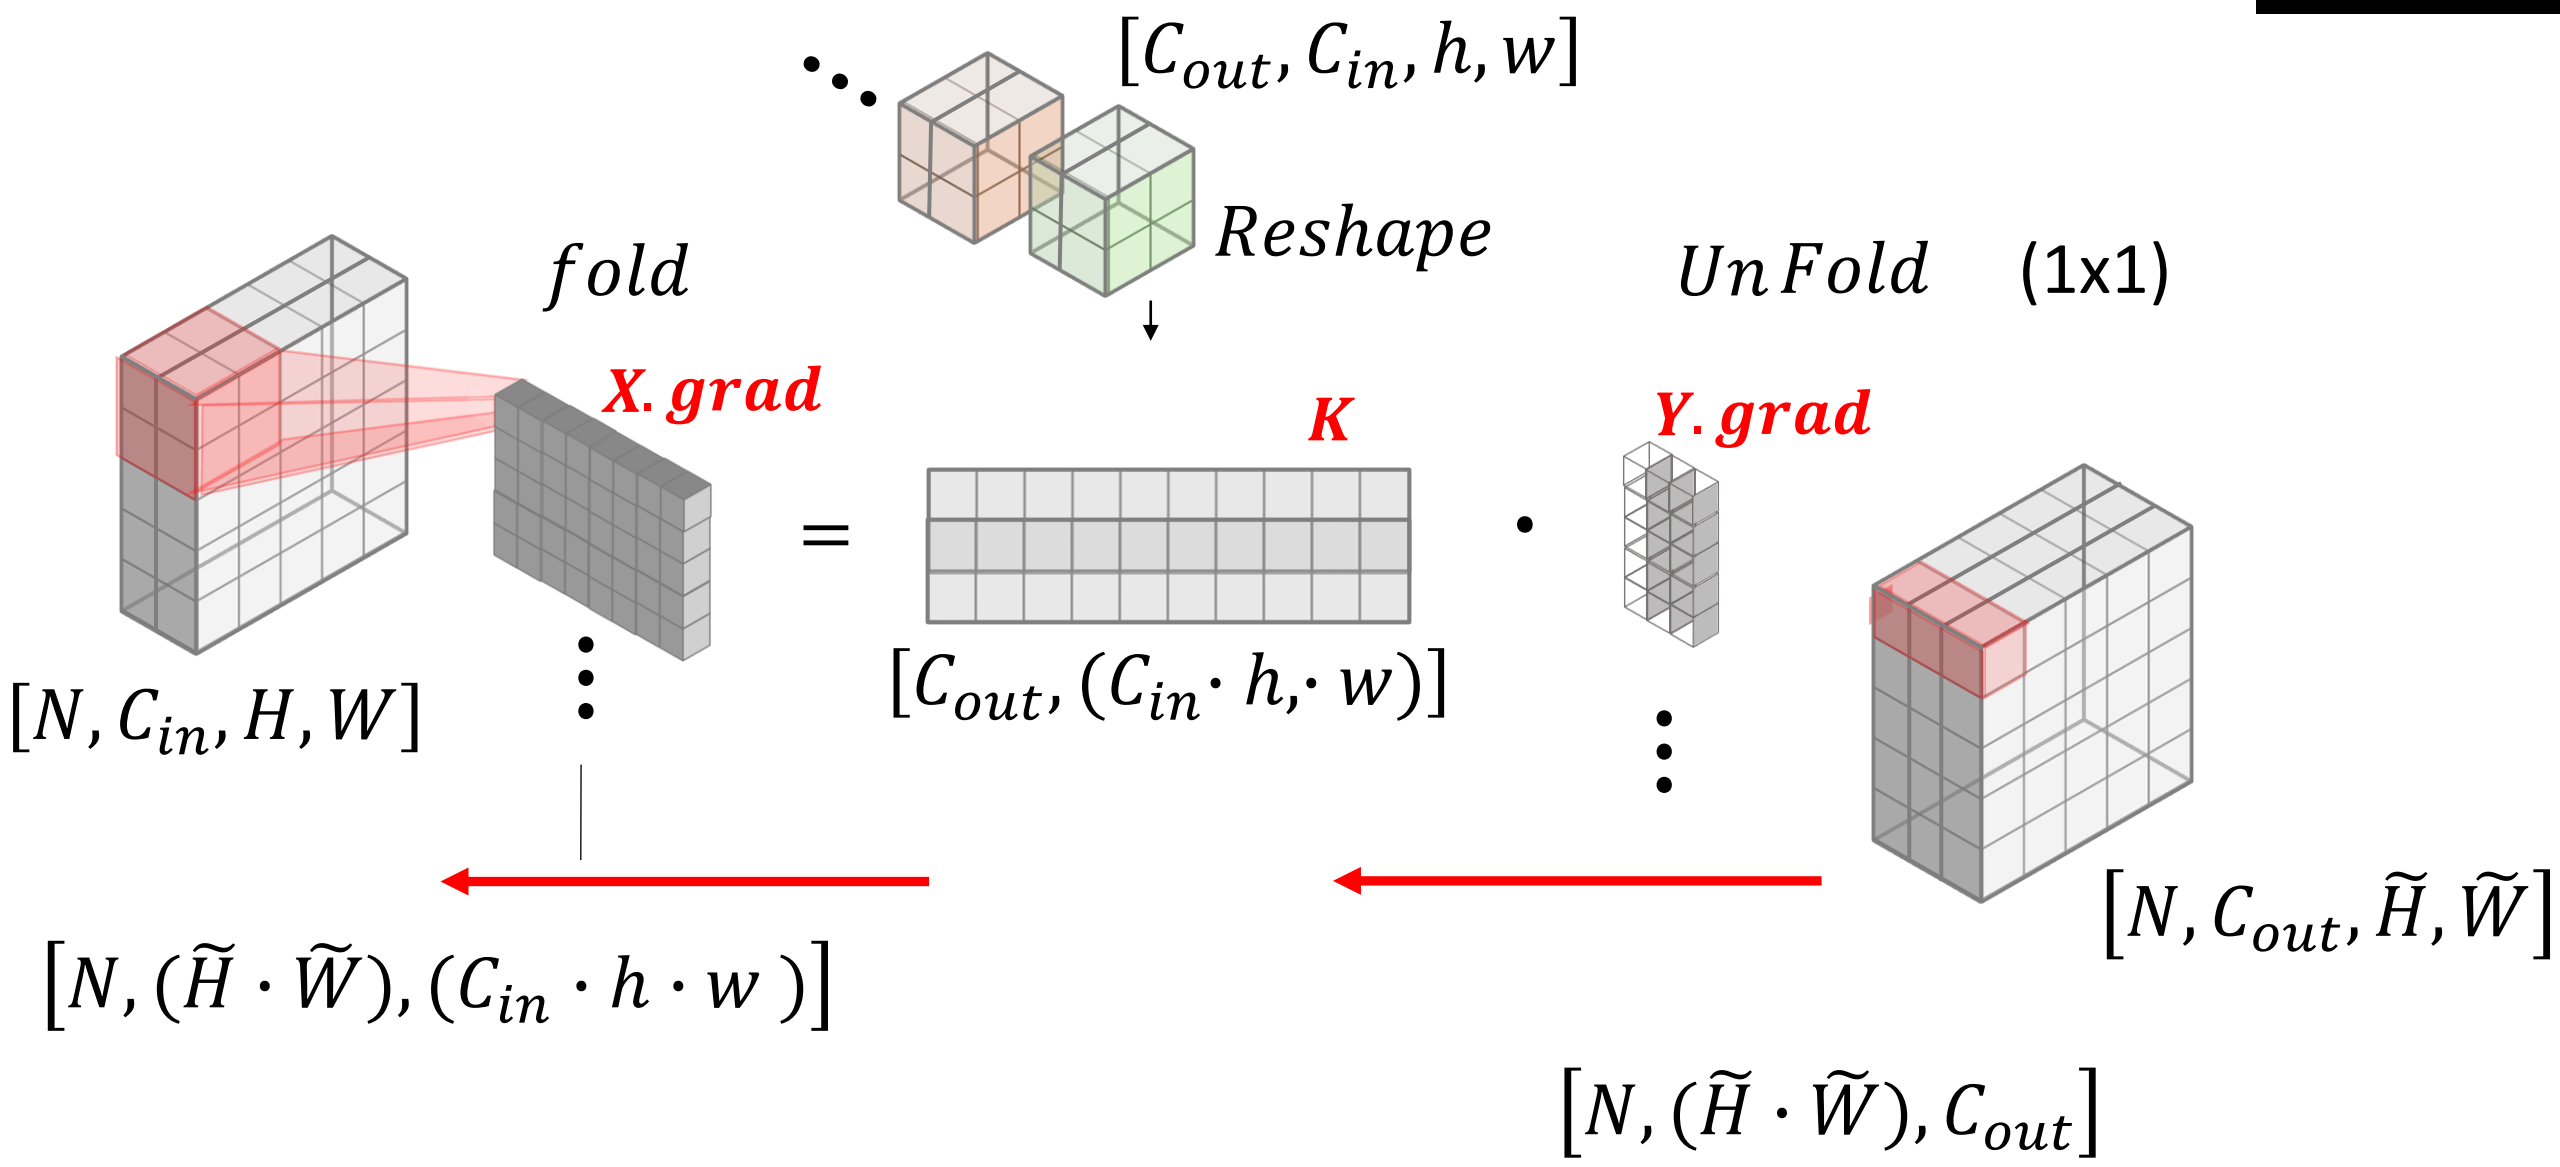
A note about the implementation of conv - BACKWARD




A note about the implementation of conv - BACKWARD




A note about the implementation of conv - BACKWARD




A note about the implementation of conv - BACKWARD




A note about the implementation of conv - BACKWARD




A note about the implementation of conv - BACKWARD



Max Pooling

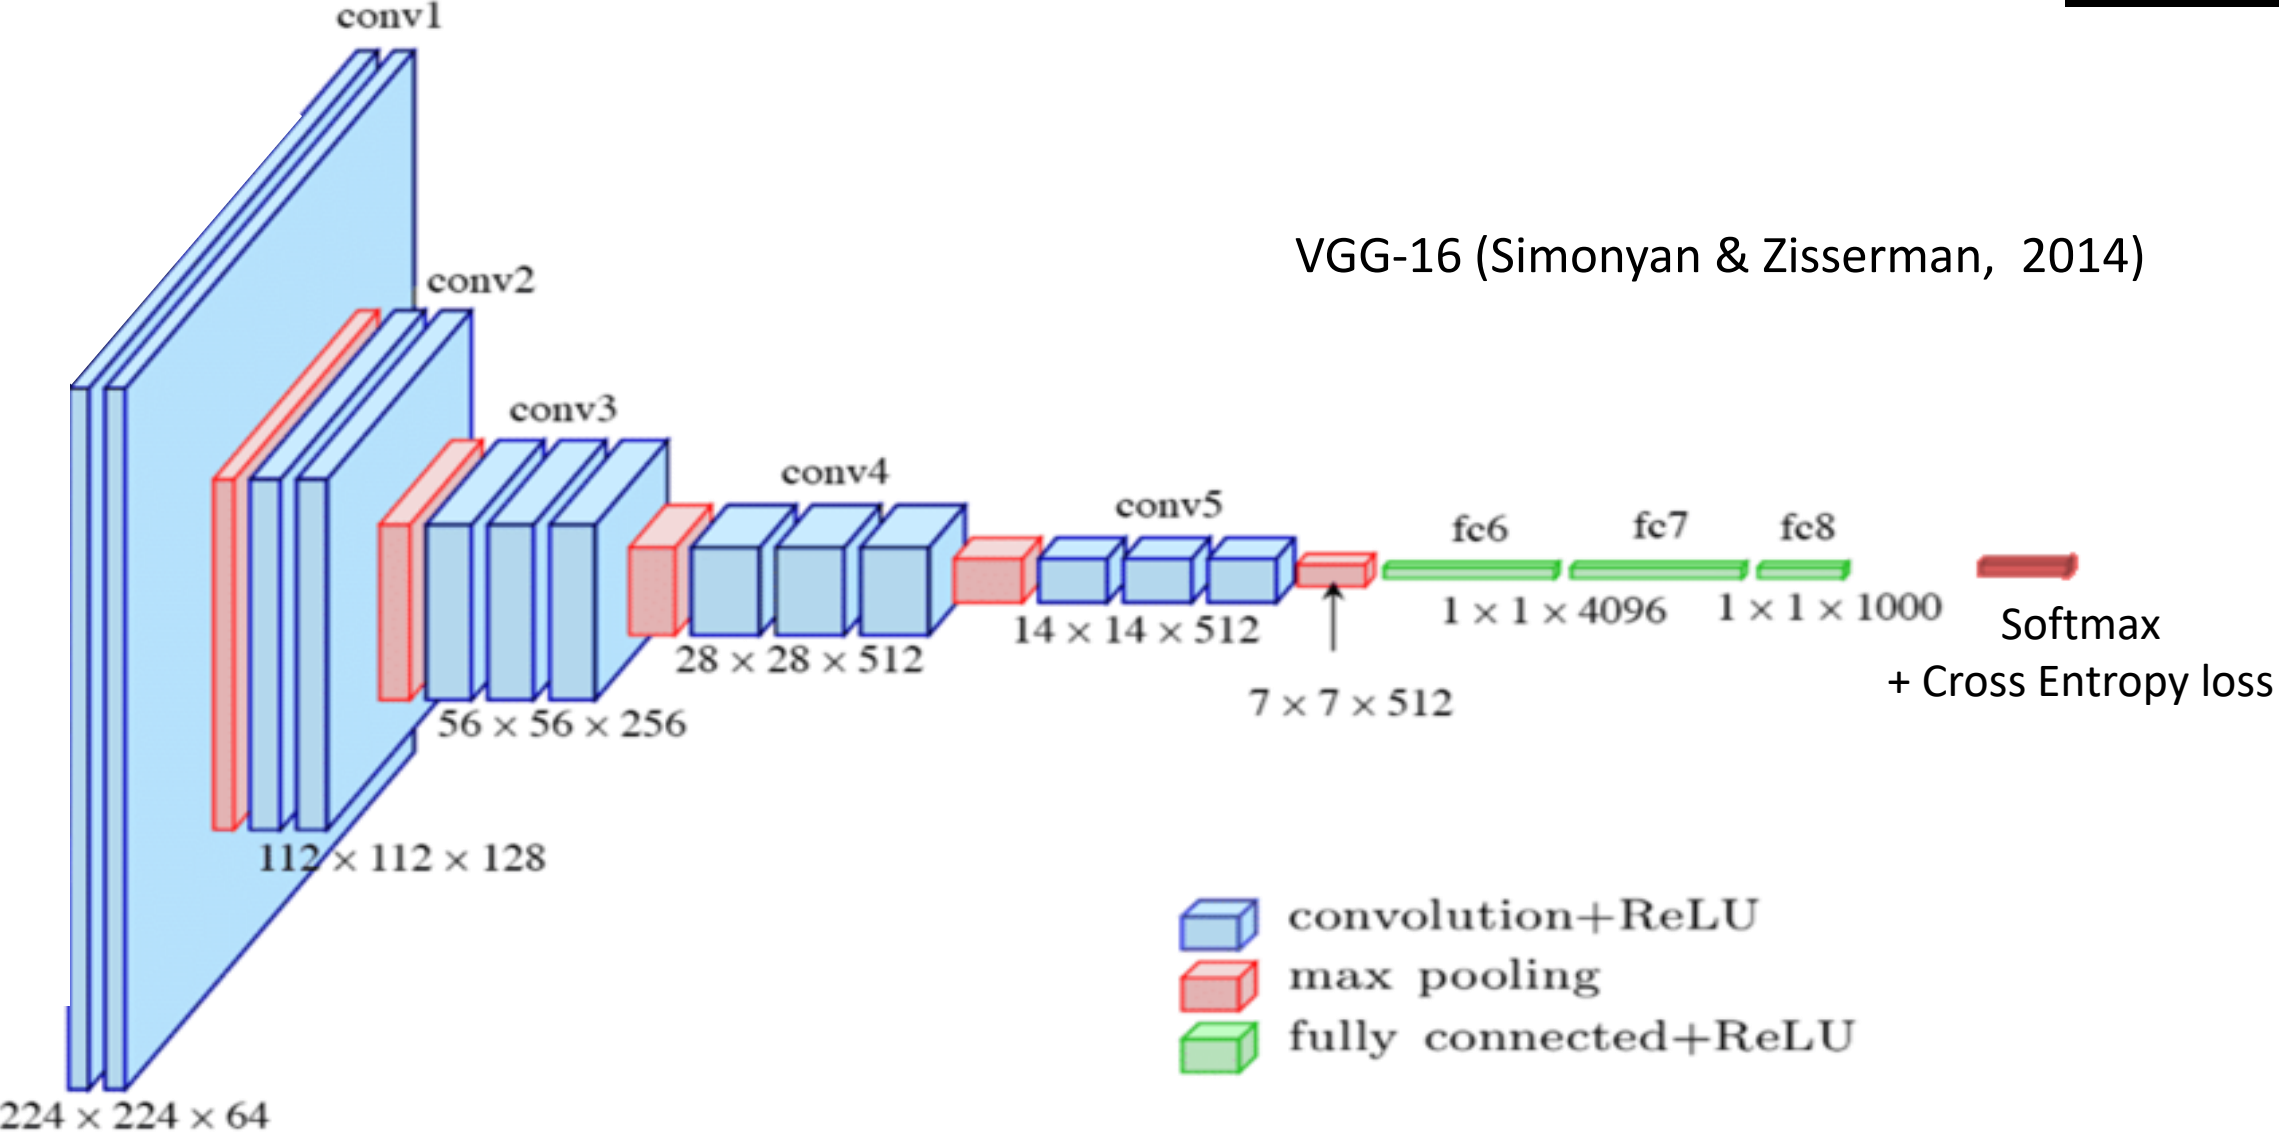


Max Pooling



- Usually stride=win-size, but not always.
- Each channel separately.

ConvNet Example




More special Convs!



More special Convs!




Depthwise Conv



More special Convs!

Group Conv (Krihzevsky 2012)




Depthwise Conv




More special Convs!

Group Conv (Krihzevsky 2012)

Depthwise Conv



Dilated Conv (Yu&Koltun 2016)





DL4CV@Weizmann

No Tutorial this week! (sigd)




No Tutorial this week! (sigd)

Next week's lecture:



Shai Bagon

Practical Training




No Tutorial this week! (sigd)

Next week's lecture:



Shai Bagon

Practical Training



Dror Moran

Next week's tutorial: CNN Architectures

