We discuss the generation of $\phi\left(\operatorname{pad}\left(Z^{m, i}\right)\right)$

## im2col in Existing Packages I

- Due to the wide use of CNN, a subroutine for $\phi\left(\operatorname{pad}\left(Z^{m, i}\right)\right)$ has been available in some packages
- For example, MATLAB has a built-in function im2col that can generate $\phi\left(\operatorname{pad}\left(Z^{m, i}\right)\right)$ for

$$
s=1 \text { and } s=h \text { (width of filter) }
$$

- But this function cannot handle general s
- Can we do a reasonably efficient implementation by ourselves?


## im2col in Existing Packages II

- For an easy description we consider

$$
\operatorname{pad}\left(Z^{m, i}\right)=Z^{\mathrm{in}, i} \quad \rightarrow \quad Z^{\mathrm{out}, i}=\phi\left(Z^{\mathrm{in}, i}\right)
$$

## Linear Indices and an Example I

- For the matrix

$$
\begin{aligned}
& Z^{\text {in }, i} \\
& =\left[\begin{array}{ccccccc}
z_{1,1,1}^{i} & z_{2,1,1}^{i} & \ldots & z_{a^{\text {in }}, 1,1}^{i} & z_{1,2,1}^{i} & \ldots & z_{a^{\text {in }}, b^{\text {in }}, 1}^{i} \\
\vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \\
z_{1,1, j}^{i} & z_{2,1, j}^{i} & \ldots & z_{a^{\text {in }}, 1, j}^{i} & z_{1,2, j}^{i} & \ldots & z_{a^{\text {in }}, b^{\text {in }}, j}^{i} \\
\vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \\
z_{1,1, d^{\text {in }}}^{i} & z_{2,1, d^{\text {in }}}^{i} & \ldots & z_{a^{\text {in }}, 1, d^{\text {in }}}^{i} & z_{1,2, d^{\text {in }}}^{i} & \ldots & z_{a^{\text {in }}, b^{\text {in }}, d^{\text {in }}}^{i}
\end{array}\right]
\end{aligned}
$$

we count elements in a column-oriented way

## Linear Indices and an Example II

- This leads to the following column-oriented linear indices of $Z^{\text {in, } i}$ :

$$
\left[\begin{array}{cccc}
1 & d^{\text {in }}+1 & \ldots & \left(b^{\text {in }} a^{\text {in }}-1\right) d^{\text {in }}+1 \\
2 & d^{\text {in }}+2 & \ldots & \left(b^{\text {in }} a^{\text {in }}-1\right) d^{\text {in }}+2 \\
: & : & . & :
\end{array}\right] \in R^{d^{\text {in }} \times a^{\text {in }} b^{\text {in }}} .
$$

- We know every element in

$$
\phi\left(Z^{\text {in }, i}\right) \in R^{h h d^{\text {in }} \times a^{\text {out }} b^{\text {but }}}
$$

is extracted from $Z^{\text {in, }, i}$

## Linear Indices and an Example III

- Thus the task is to find the mapping between each element in $\phi\left(Z^{\text {in, }, i}\right)$ and a linear index of $Z^{\text {in, }, i}$.
- Consider an example with

$$
a^{\mathrm{in}}=3, b^{\mathrm{in}}=2, d^{\mathrm{in}}=1
$$

Because $d^{\text {in }}=1$, we omit the channel subscript.

- In addition, we omit the instance index $i$, so the image is

$$
\left[\begin{array}{ll}
z_{11} & z_{12}  \tag{2}\\
z_{21} & z_{22} \\
z_{31} & z_{32}
\end{array}\right] .
$$

## Linear Indices and an Example IV

- If

$$
h=2, s=1
$$

two sub-images are

$$
\left[\begin{array}{ll}
z_{11} & z_{12} \\
z_{21} & z_{22}
\end{array}\right] \text { and }\left[\begin{array}{ll}
z_{21} & z_{22} \\
z_{31} & z_{32}
\end{array}\right]
$$

- By our earlier way of representing images,

$$
Z^{\mathrm{in}, i}=\left[\begin{array}{cccc}
z_{1,1,1}^{i} & z_{2,1,1}^{i} & \ldots & z_{a^{i n}, b^{\text {in }}, 1}^{i} \\
\vdots & \vdots & \ddots & \vdots \\
z_{1,1, d^{\text {in }}}^{i} & z_{2,1, \text { d }^{\text {in }}}^{i} & \ldots & z_{a^{\text {in }}, \text { bin }^{i n}, \text { din }^{i n}}^{i}
\end{array}\right]
$$

## Linear Indices and an Example V

we now have

$$
Z^{\text {in }}=\left[\begin{array}{llllll}
z_{11} & z_{21} & z_{31} & z_{12} & z_{22} & z_{32}
\end{array}\right]
$$

- The linear indices from (1) are

$$
\left[\begin{array}{llllll}
1 & 2 & 3 & 4 & 5 & 6
\end{array}\right] .
$$

## Linear Indices and an Example VI

- Recall that

$$
\begin{aligned}
& \phi\left(Z^{\text {in }, i}\right)= \\
& {\left[\begin{array}{cccc}
z_{1,1,1}^{i} & z_{1+s, 1,1}^{i} & & z_{1}^{i}+\left(a^{\text {out }}-1\right) s, 1+\left(b^{\text {out }}-1\right) s, 1 \\
z_{2,1,1}^{i} & z_{2+s, 1,1}^{i} & & z_{2}^{i}+\left(a^{\text {out }}-1\right) s, 1+\left(b^{\text {out }}-1\right) s, 1 \\
\vdots & \vdots & \cdots & \vdots \\
z_{h, h, 1}^{i} & z_{h+s, h, 1}^{i} & & z_{h+\left(a^{\text {out }}-1\right) s, h+\left(b^{\text {out }}-1\right) s, 1}^{i} \\
\vdots & \vdots & & \vdots \\
z_{h, h, d^{\text {in }}}^{i} & z_{h+s, h, d^{\text {in }}}^{i} & z_{h+\left(a^{\text {out }}-1\right) s, h+\left(b^{\text {out }}-1\right) s, d^{\text {in }}}^{i}
\end{array}\right]}
\end{aligned}
$$

## Linear Indices and an Example VII

- Therefore,

$$
\phi\left(Z^{\text {in }}\right)=\left[\begin{array}{ll}
z_{11} & z_{21}  \tag{3}\\
z_{21} & z_{31} \\
z_{12} & z_{22} \\
z_{22} & z_{32}
\end{array}\right]
$$

- Let's check linear indices using Matlab/Octave octave:8> reshape((1:6)', 3, 2) ans =

$$
\begin{array}{ll}
1 & 4 \\
2 & 5
\end{array}
$$

## Linear Indices and an Example VIII

## 3 <br> 6

This gives linear indices of the image in (2) octave:9> im2col(reshape((1:6)', 3, 2), [2,2], "sliding")
ans =

| 1 | 2 |
| :--- | :--- |
| 2 | 3 |
| 4 | 5 |
| 5 | 6 |

Here [2 2] is the filter size.

## Linear Indices and an Example IX

- Clearly,

$$
\begin{aligned}
& 1 \\
& 2 \\
& 4 \\
& 5
\end{aligned}
$$

correspond to linear indices of the first column in $\phi\left(Z^{\text {in }}\right)$; see (3)

## Linear Indices and an Example X

- To handle all instances together, we store

$$
Z^{\text {in }, 1}, \ldots, Z^{\text {in }, l}
$$

as

$$
\left[\begin{array}{lll}
\operatorname{vec}\left(Z^{\mathrm{in}, 1}\right) & \ldots & \operatorname{vec}\left(Z^{\mathrm{in}, l}\right) \tag{4}
\end{array}\right]
$$

- For our example,
$\operatorname{vec}\left(Z^{m, i}\right)=\left[\begin{array}{l}1 \\ 2 \\ 3 \\ 4 \\ 5\end{array}\right], \quad \operatorname{vec}\left(\phi\left(Z^{m, i}\right)\right)=\left[\begin{array}{l}1 \\ 2 \\ 4 \\ 5 \\ 2 \\ 3\end{array}\right]$


## Linear Indices and an Example XI

- Denote (4) as a MATLAB matrix
Z
- Then

$$
\left[\operatorname{vec}\left(\phi\left(Z^{m, 1}\right)\right) \ldots \operatorname{vec}\left(\phi\left(Z^{m, l}\right)\right)\right]
$$

is simply

$$
Z(P,:)
$$

in MATLAB, where we store the mapping by

$$
\mathrm{P}=\left[\begin{array}{llllllll}
1 & 2 & 4 & 5 & 2 & 3 & 5 & 6
\end{array}\right]^{T}
$$

## Linear Indices and an Example XII

- All instances are handled in one line
- Moreover, we hope that Matlab's implementation on this mapping operation is efficient
- But how to obtain P?
- Note that

$$
\left[\begin{array}{llllllll}
1 & 2 & 4 & 5 & 2 & 3 & 5 & 6
\end{array}\right]^{T}
$$

also corresponds to column indices of non-zero elements in $P_{\phi}^{m}$.

## Linear Indices and an Example XIII

$$
\left[\begin{array}{l}
z_{11} \\
z_{21} \\
z_{12} \\
z_{22} \\
z_{21} \\
z_{31} \\
z_{22} \\
z_{32}
\end{array}\right]=\left[\begin{array}{llllll}
1 & & & & & \\
& 1 & & & & \\
& & & 1 & & \\
& & & & 1 & \\
& 1 & & & & \\
& & 1 & & & \\
& & & & 1 & \\
& & & & & 1
\end{array}\right]\left[\begin{array}{l}
z_{11} \\
z_{21} \\
z_{31} \\
z_{12} \\
z_{22} \\
z_{32}
\end{array}\right]
$$

- Each row has a single non-zero
- The column index of the non-zero indicates the extraction of elements from $Z^{\text {in }}$ to $\phi\left(Z^{\text {in }}\right)$


## Finding the Mapping I

- We begin with checking how linear indices of $Z^{\text {in, } i}$ can be mapped to the first column of $\phi\left(Z^{\mathrm{in}, i}\right)$.
- For simplicity, we consider channel $j$ first
- From

$$
\begin{aligned}
& Z^{\text {in }, i}
\end{aligned}
$$

## Finding the Mapping II

for the first column of $\phi\left(Z^{\text {in, }, i}\right)$, for channel $j$, we select an $h \times h$ matrix from $Z^{\text {in, }, i}$.

## Finding the Mapping III

- The selected values and their linear indices are

| values | linear indices in $Z^{\text {in }, i}$ |
| :---: | :---: |
| $z_{1,1, j}$ | $j$ |
| $z_{2,1, j}$ | $d^{\text {in }}+j$ |
| $\vdots$ | $\vdots$ |
| $z_{h, 1, j}$ | $(h-1) d^{\text {in }}+j$ |
| $z_{1,2, j}$ | $a^{\text {in }} d^{\text {in }}+j$ |
| $\vdots$ | $\vdots$ |
| $z_{h, 2, j}$ | $\left((h-1)+a^{\text {in }}\right) d^{\text {in }}+j$ |
| $\vdots$ | $\vdots$ |
| $z_{h, h, j}$ | $\left((h-1)+(h-1) a^{\text {in }}\right) d^{\text {in }}+j$ |

## Finding the Mapping IV

- We take $d^{\text {in }}$ out and rewrite the earlier table as
$\left[\begin{array}{c}0+0 a^{\text {in }} \\ \vdots \\ (h-1)+0 a^{\text {in }} \\ 0+1 a^{\text {in }} \\ \vdots \\ (h-1)+1 a^{\text {in }} \\ \vdots \\ 0+(h-1) a^{\text {in }} \\ \vdots \\ (h-1)+(h-1) a^{\text {in }}\end{array}\right] d^{\text {in }}+j=\left[\begin{array}{c}0 \\ \vdots \\ h-1 \\ 0 \\ \vdots \\ h-1 \\ \vdots \\ 0 \\ \vdots \\ h-1\end{array}\right] d^{\text {in }}+j+\left[\begin{array}{c}0 \\ \vdots \\ 0 \\ 1 \\ \vdots \\ 1 \\ \vdots \\ h-1 \\ \vdots \\ h-1\end{array}\right] a^{\text {in }} d^{\text {in }}(6)$


## Finding the Mapping $V$

- We will see that the right-side of (6) will be used in the practical implementation
- Every linear index in (6) can be represented as

$$
\begin{equation*}
\left(p+q a^{\text {in }}\right) d^{\text {in }}+j, \tag{7}
\end{equation*}
$$

where

$$
p, q \in\{0, \ldots, h-1\}
$$

- Thus $(p+1, q+1)$ corresponds to the pixel position in the first sub-image (or say the convolutional filter)


## Finding the Mapping VI

- Next we consider other columns in $\phi\left(Z^{\text {in,i}}\right)$ by still fixing the channel to be $j$.


## Finding the Mapping VII

- From

$$
\begin{aligned}
& \phi\left(Z^{\text {in }, i}\right)= \\
& {\left[\begin{array}{cccc}
\vdots & \vdots & & \vdots \\
z_{1,1, j}^{i} & z_{1+s, 1, j}^{i} & & z_{1+\left(a^{\text {out }}-1\right) s, 1+\left(b^{\text {out }}-1\right) s, j}^{i} \\
z_{2,1, j}^{i} & z_{2+s, 1, j}^{i} & & z_{2+\left(a^{\text {out }}-1\right) s, 1+\left(b^{\text {out }}-1\right) s, j}^{i} \\
\vdots & \vdots & \cdots & \vdots \\
z_{h, h, j}^{i} & z_{h+s, h, j}^{i} & & z_{h+\left(a^{\text {out }}-1\right) s, h+\left(b^{\text {out }}-1\right) s, j}^{i} \\
\vdots & \vdots & & \vdots
\end{array}\right]}
\end{aligned}
$$

## Finding the Mapping VIII

each column contains the following elements from the $j$ th channel of $Z^{\text {in, }, i}$.

$$
\begin{array}{ll}
z_{1+p+a s, 1+q+b s, j}, & a=0,1, \ldots, a^{\text {out }}-1 \\
& b=0,1, \ldots, b^{\text {out }}-1 \\
& p, q \in\{0, \ldots, h-1\}
\end{array}
$$

Clearly, when

$$
p=q=0
$$

we have that

$$
(1+a s, 1+b s)
$$

## Finding the Mapping IX

is the top-left position of a sub-image in the channel $j$ of $Z^{\text {in, } i}$.

- This is reasonable as we now use stride $s$ to generate $a^{\text {out }} b^{\text {out }}$ sub-images
- The linear index of each element in (8) is

$$
\begin{align*}
& \underbrace{\left((1+p+a s-1)+(1+q+b s-1) a^{\text {in }}\right)}_{\text {column index in } Z^{\mathrm{i}, i},-1} d^{\text {in }}+j \\
= & \left((p+a s)+(q+b s) a^{\text {in }}\right) d^{\text {in }}+j \\
= & \left(a+b a^{\text {in }}\right) s d^{\text {in }}+\underbrace{\left(p+q a^{\text {in }}\right) d^{\text {in }}+j}_{\text {see }(7)} .
\end{align*}
$$

## Finding the Mapping $X$

- Now we have known for each element of $\phi\left(Z^{\text {in, }, i}\right)$ what the corresponding linear index in $Z^{\text {in, } i}$ is.
- Next we discuss the implementation details


## A MATLAB Implementation I

- First, we compute elements in (6) with $j=1$ by MATLAB's '+' operator
- This operator has implicit expansion behavior to compute the outer sum of two arrays.
- From the second part of (6), we calculate the outer sum of

$$
\left[\begin{array}{c}
0 \\
\vdots \\
h-1
\end{array}\right] d^{\text {in }}+1
$$

and

$$
\left[\begin{array}{lll}
0 & \ldots & h-1
\end{array}\right] a^{\text {in }} d^{\text {in }}
$$

## A MATLAB Implementation II

- The result is the following matrix

$$
\left[\begin{array}{cccc}
1 & a^{\text {in }} d^{\text {in }}+1 & \cdots & (h-1) a^{\text {in }} d^{\text {in }}+1 \\
d^{\text {in }}+1 & \left(1+a^{\text {in }}\right) d^{\text {in }}+1 & \cdots & \left(1+(h-1) a^{\text {in }}\right) d^{\text {in }}+1 \\
\vdots & \vdots & \cdots & \vdots \\
(h-1) d^{\text {in }}+1 & \left((h-1)+a^{\text {in }}\right) d^{\text {in }}+1 & \cdots & \left((h-1)+(h-1) a^{\text {in }}\right) d^{\text {in }}+1
\end{array}\right]
$$

- If columns are concatenated, we get (6) with $j=1$
- To get (7) for all channels $j=1, \ldots, d^{\text {in }}$, we have

$$
\left[\begin{array}{c}
\operatorname{vec}((10))+0 \\
\vdots \\
\operatorname{vec}((10))+d^{\text {in }}-1
\end{array}\right]
$$

## A MATLAB Implementation III

- This can be computed by the following outer sum:

$$
\operatorname{vec}((10))+\left[\begin{array}{llll}
0 & 1 & \ldots & d^{\text {in }}-1 \tag{11}
\end{array}\right]
$$

Note that in (11) we have

$$
0, \ldots, d^{\text {in }}-1
$$

instead of

$$
1, \ldots, d^{\text {in }}
$$

because in (10) we have already done " +1 " for $j=1$

## A MATLAB Implementation IV

- Now we have linear indices of the first column of $\phi\left(Z^{\mathrm{in}, i}\right)$
- Next, we obtain other columns in $\phi\left(Z^{\text {in }, i}\right)$
- In the linear indices in (9), the second term corresponds to indices of the first column, while the first term is the following column offset

$$
\begin{aligned}
\left(a+b a^{\text {in }}\right) s d^{\text {in }}, \forall a & =0,1, \ldots, a^{\text {out }}-1 \\
b & =0,1, \ldots, b^{\text {out }}-1 .
\end{aligned}
$$

## A MATLAB Implementation $V$

- This column offset is the outer sum of the following two arrays.

$$
\left[\begin{array}{c}
0  \tag{12}\\
\vdots \\
a^{\text {out }}-1
\end{array}\right] \times s d^{\text {in }} \quad \text { and } \quad\left[\begin{array}{lll}
0 & \ldots & b^{\text {out }}-1
\end{array}\right] \times a^{\text {in }} s d^{\text {in }}
$$

- Finally, we compute the outer sum of
- the column offset and
- the linear indices in the first column of $\phi\left(Z^{\mathrm{in}, i}\right)$


## A MATLAB Implementation VI

in the following operation

$$
\begin{equation*}
\operatorname{vec}((12))^{T}+\operatorname{vec}((11)) \tag{13}
\end{equation*}
$$

where

$$
\operatorname{vec}((12)) \in R^{\text {aut }^{\text {out }} b^{\text {out }} \times 1} \text { and } \operatorname{vec}((11)) \in R^{h h d^{\text {in }} \times 1}
$$

- In the end we store

$$
\operatorname{vec}((13)) \in R^{\text {hhd } d^{\text {in }} a^{\text {out }} b^{\text {out }} \times 1}
$$

It is a vector collecting

## A MATLAB Implementation VII

column index of the non-zero in each row of $P_{\phi}^{m}$

- Note that each row in the $0 / 1$ matrix $P_{\phi}^{m}$ contains exactly only one non-zero element.
- See the example in (5)
- The obtained linear indices are independent of the values of $Z^{\text {in, } i}$.
- Thus the above procedure only needs to be run once in the beginning.


## A Simple MATLAB Code I

function idx = find_index_phiZ(a,b,d,h,s)
first_channel_idx = ([0:h-1]'*d+1) + [0:h-1]*a*d;
first_col_idx = first_channel_idx(:) + [0:d-1];
a_out = floor ((a -h$) / \mathrm{s})+1$;
b_out = floor ( $(\mathrm{b}-\mathrm{h}) / \mathrm{s})+1$;
column_offset $=\left(\left[0: a \_o u t-1\right]\right.$ ' +
[0:b_out-1]*a) *s*d;
idx = column_offset(:)' + first_col_idx(:); idx = idx(:);

## Discussion

- The code is simple and short
- We assume that Matlab operations used here are efficient and so is our resulting code
- But is that really the case?
- We will do experiments to check this
- Some works have tried to do similar things (e.g., https://github.com/wiseodd/hipsternet), though we don't see complete documents and evaluation

