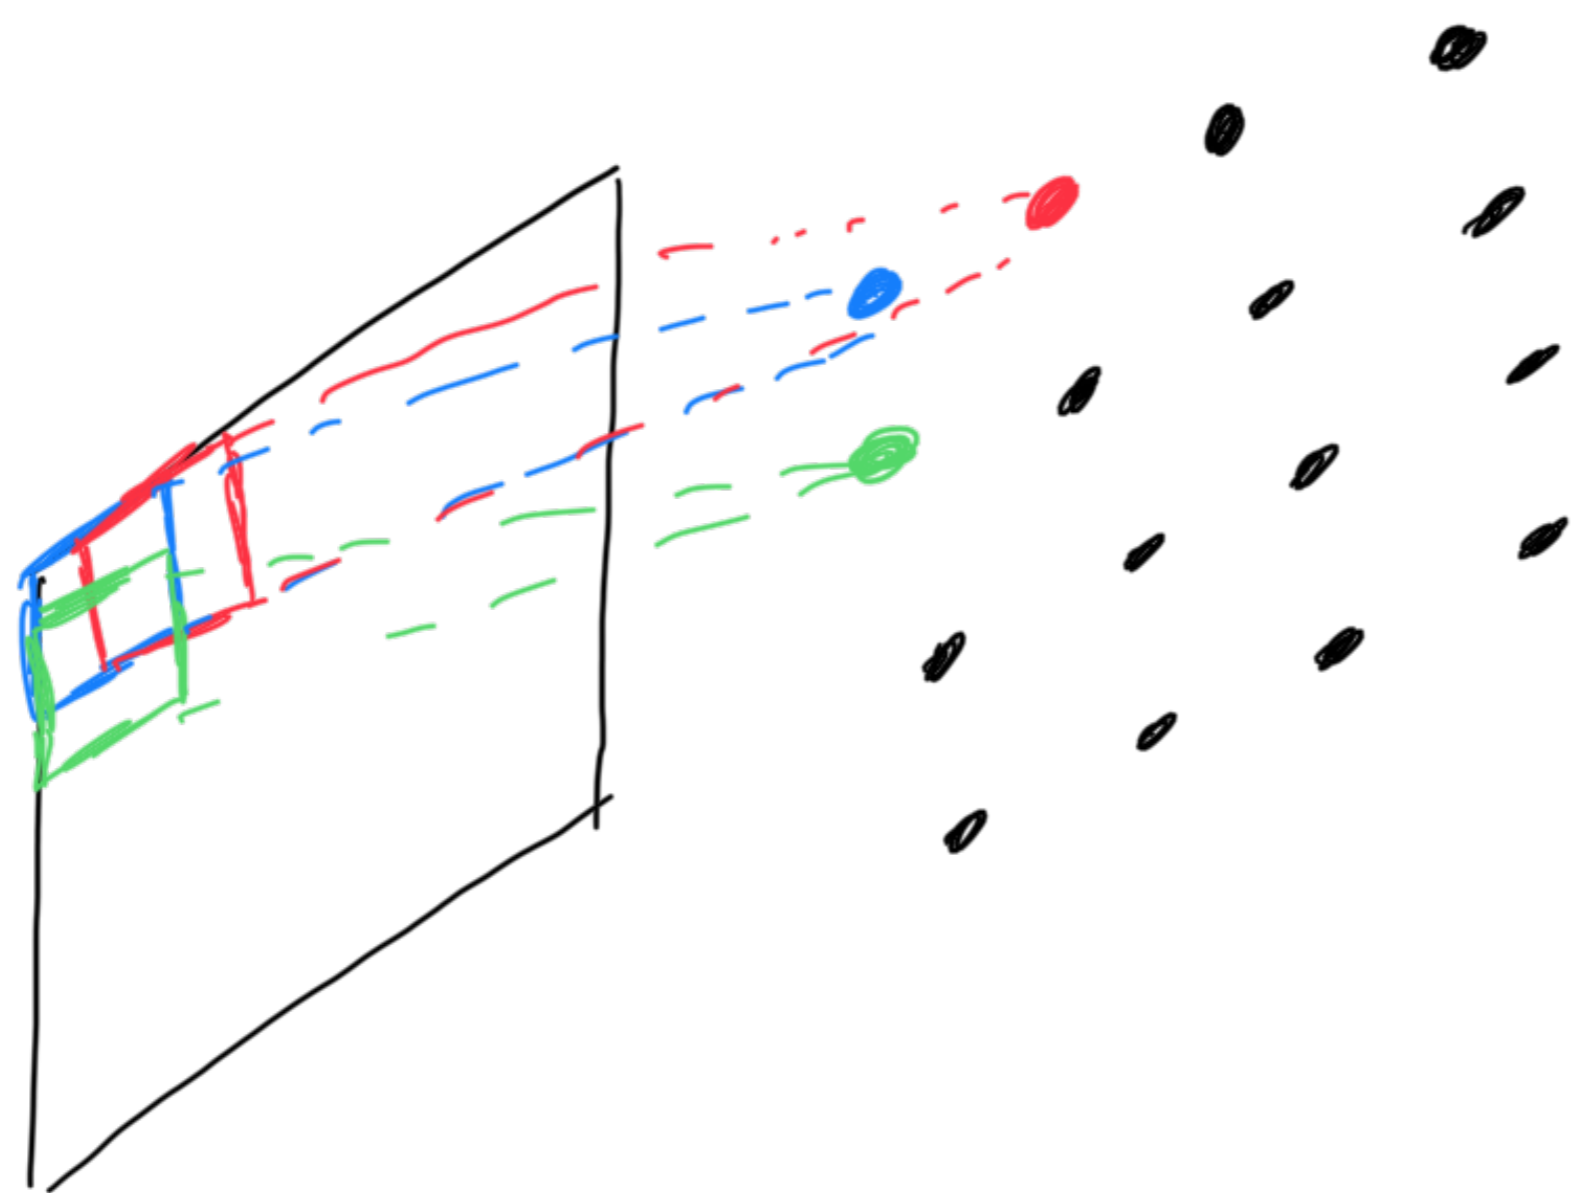


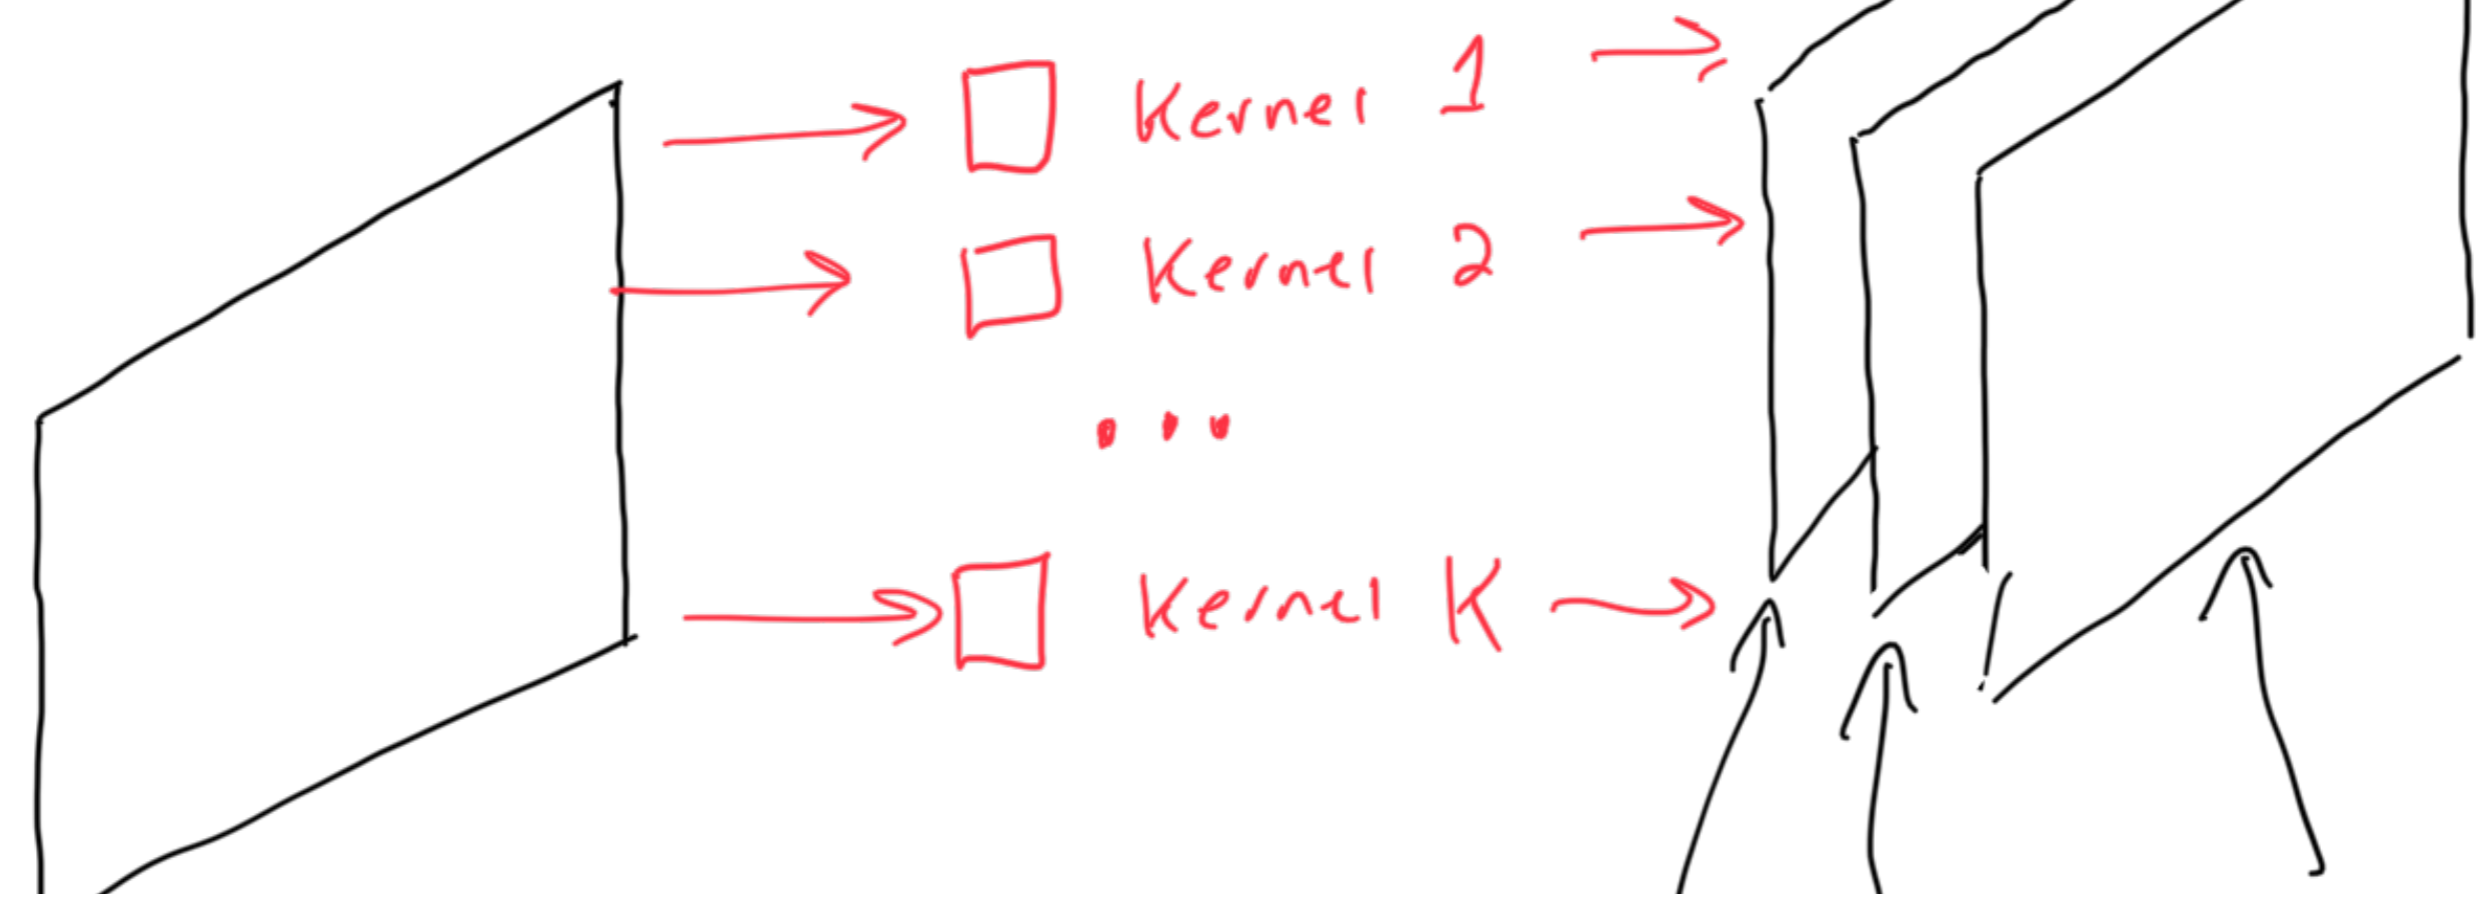
# Day 21: Convolutional Neural Networks



Key ideas of CNNs

- ① use several kernels
- ② learn the best kernels to use

Convolutional Layers



Feature map 1  
 Feature map 2  
 ...  
 Feature map k

Feature/activation map  $i$   
 is formed by convolving kernel  
 $i$  across the image

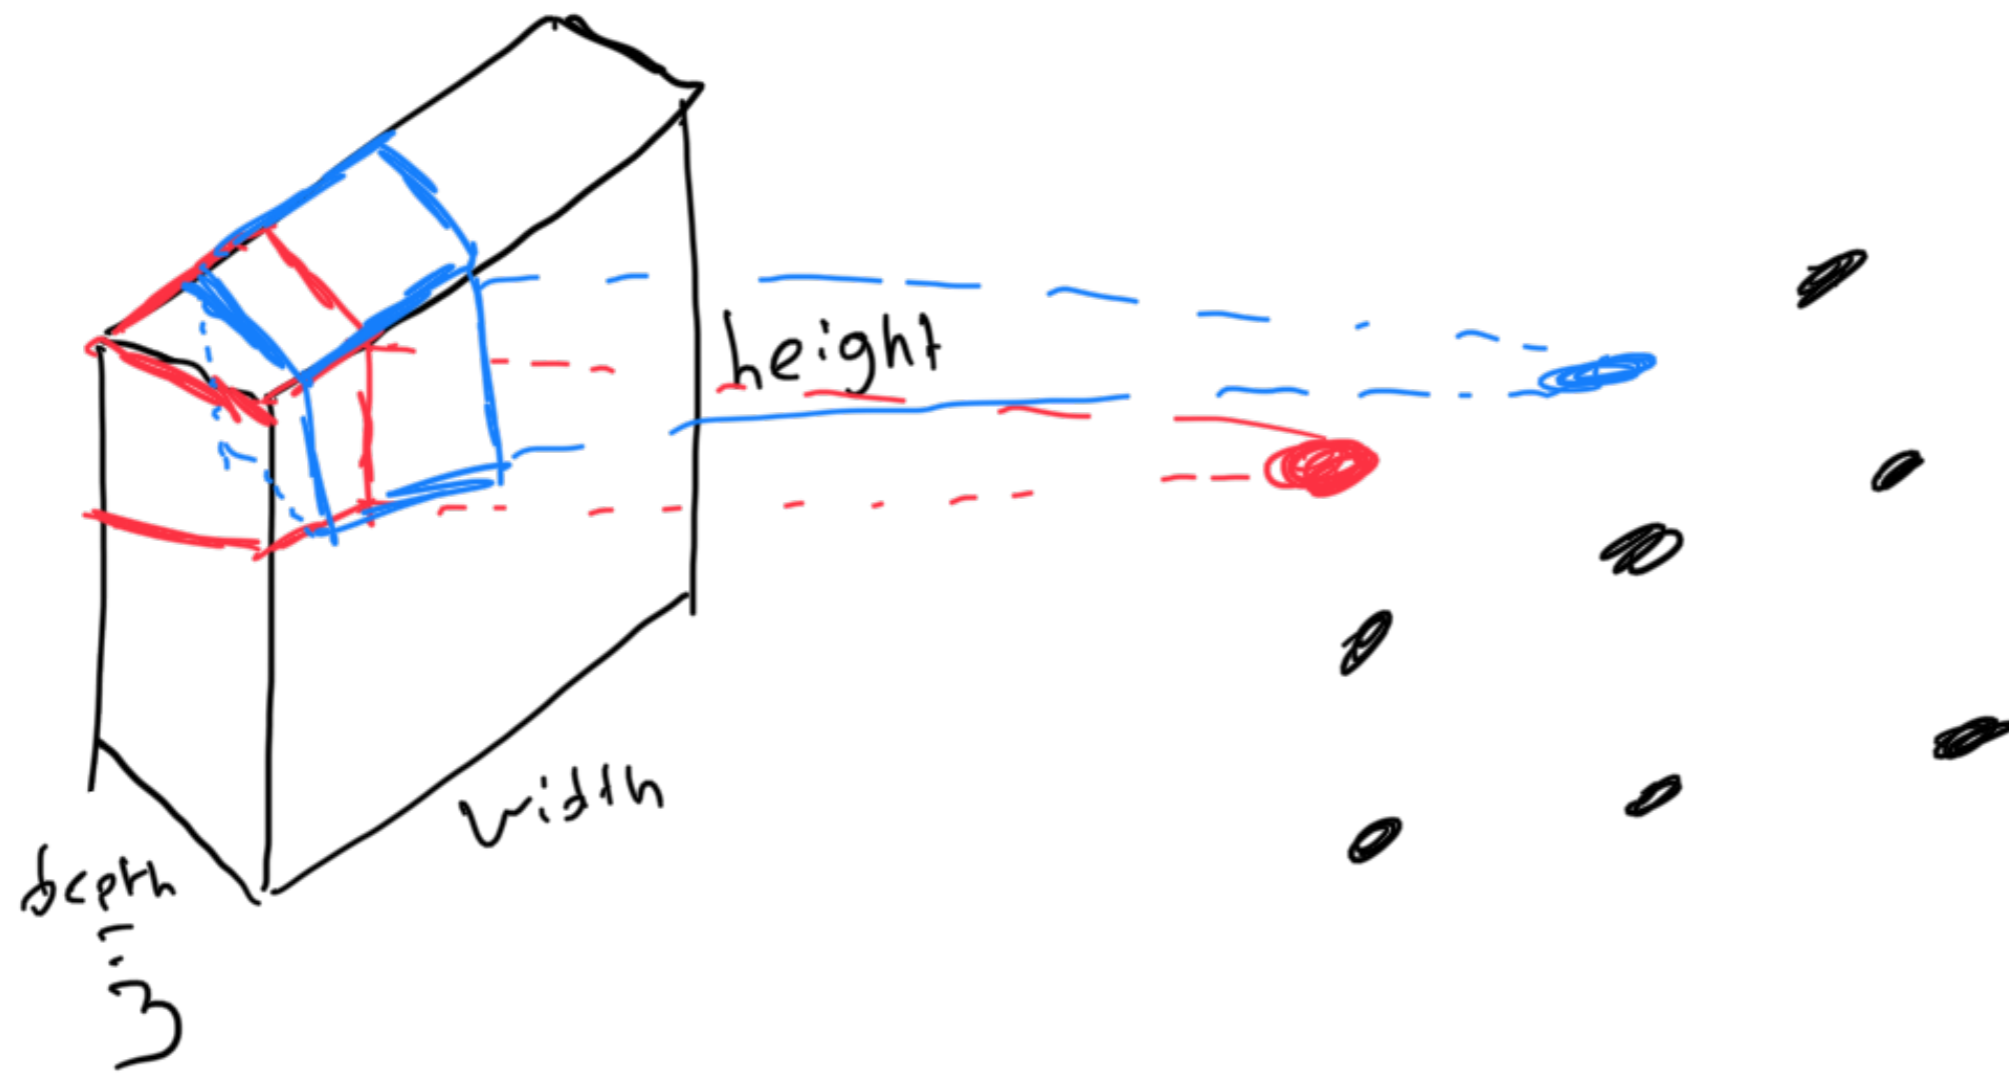
Recall: images are represented as RGB

$$\left[ \begin{array}{c} \begin{bmatrix} 255 & 255 & 255 \\ 0 & 255 & 0 \end{bmatrix} \\ \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \\ \begin{bmatrix} 255 & 255 & 255 \\ 0 & 255 & 0 \end{bmatrix} \end{array} \right]$$

R                      G                      B

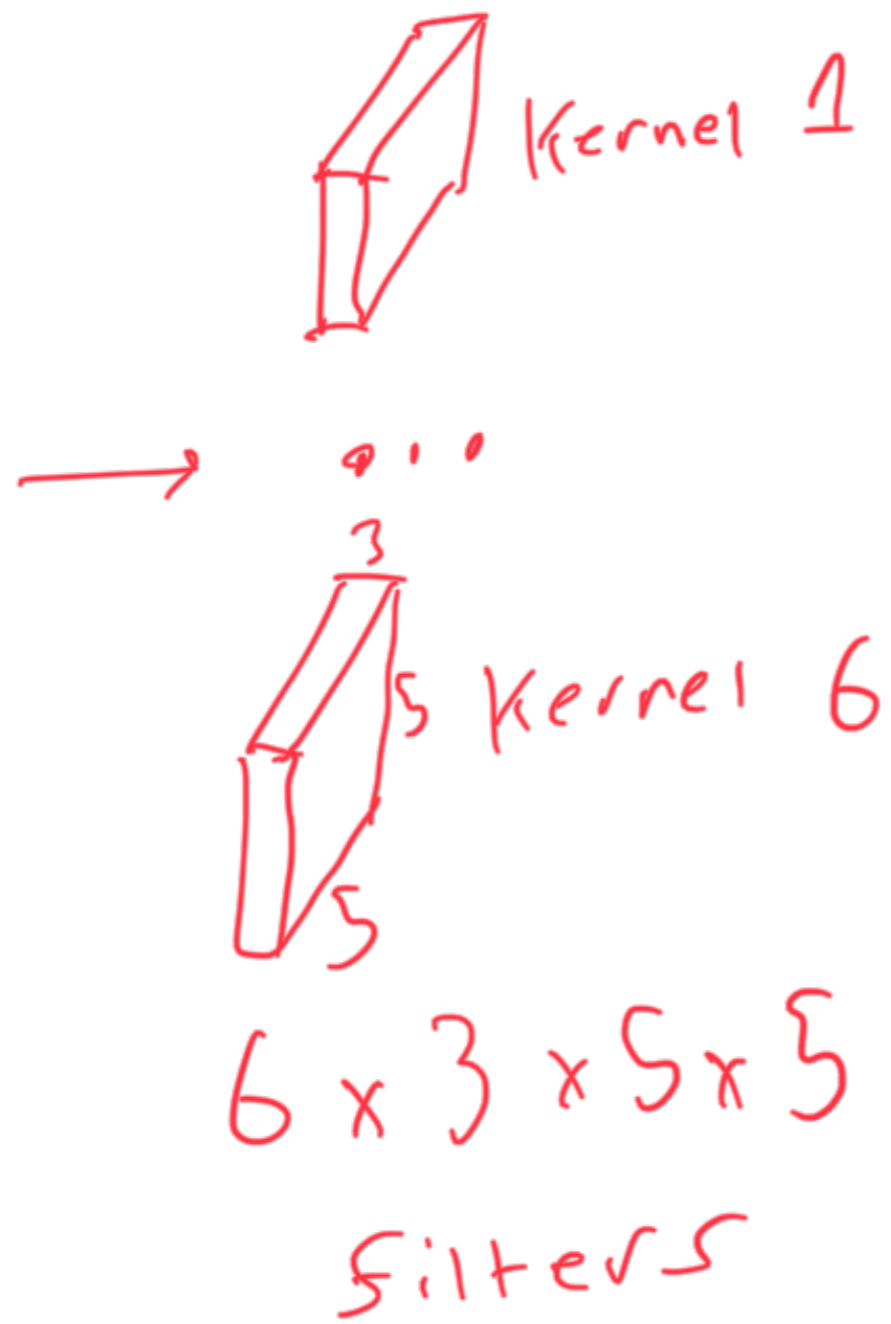
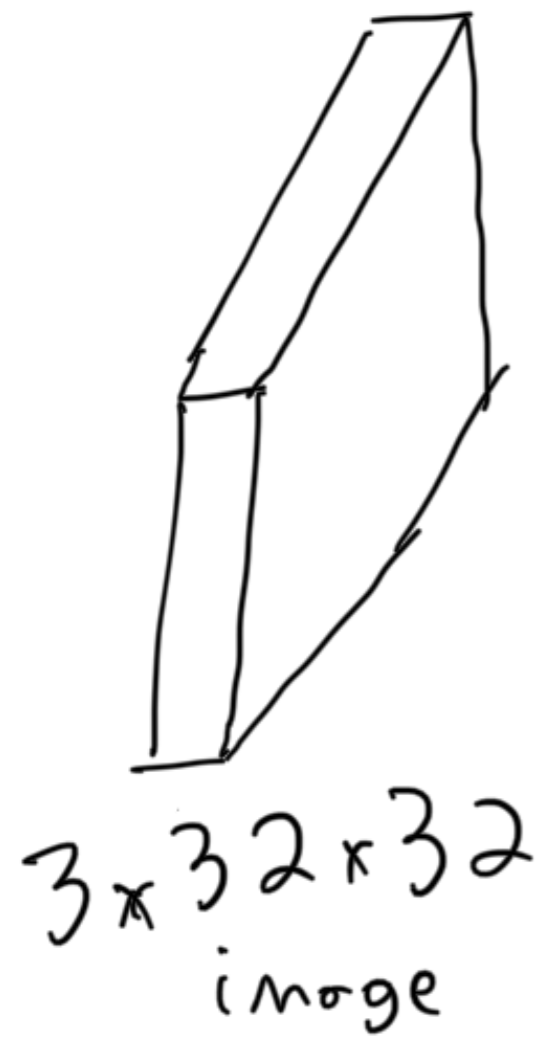


So for color images:



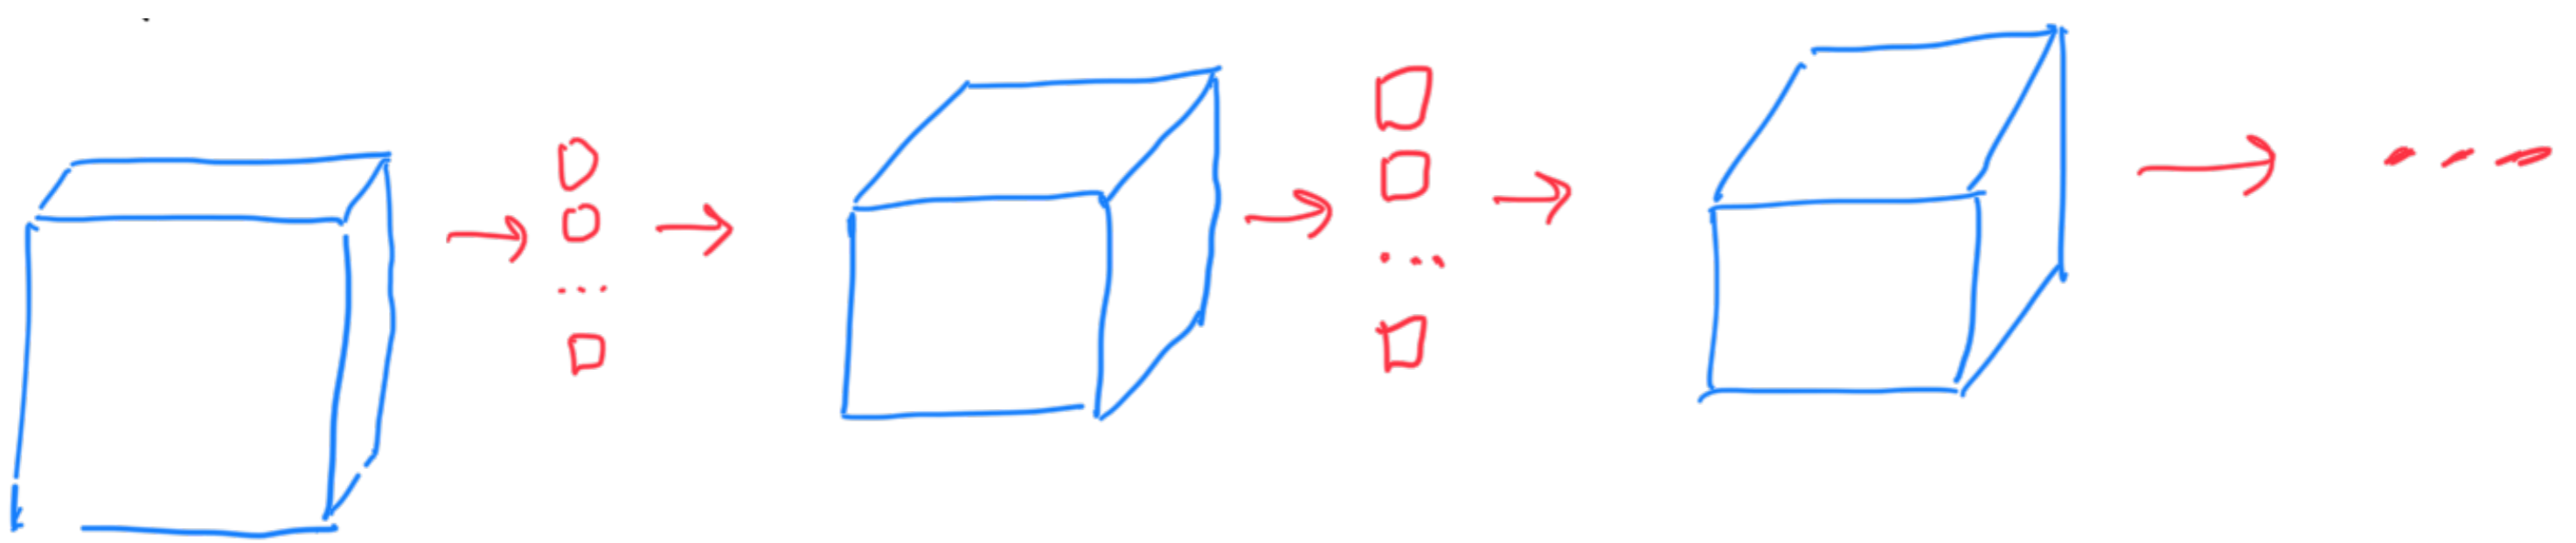
The weights/parameters are the  
kernels themselves

We stack the feature maps to get  
the input to the next layer.



("filter" = "kernel")  
 ("feature map" = "activation map")

A CNN starts with several convolutional layers;



In order to learn higher-level (<sup>more</sup> "abstract") representations, we can downsample the feature maps to extract the most salient information:

## Max Pooling Layers



7	3	0	1
-1	-2	4	0

Feature Map

Pooling

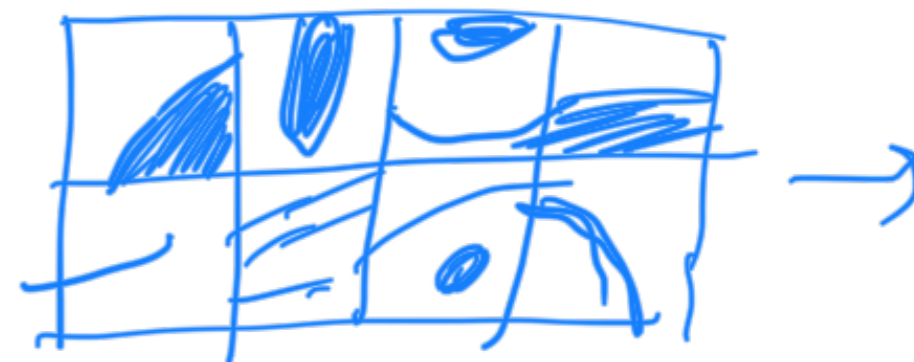
7	4
---	---

Take max element in  $N \times N (x N)$  window

Effect = each pixel in progressive layers represents larger portion of the original image



Earlier feature

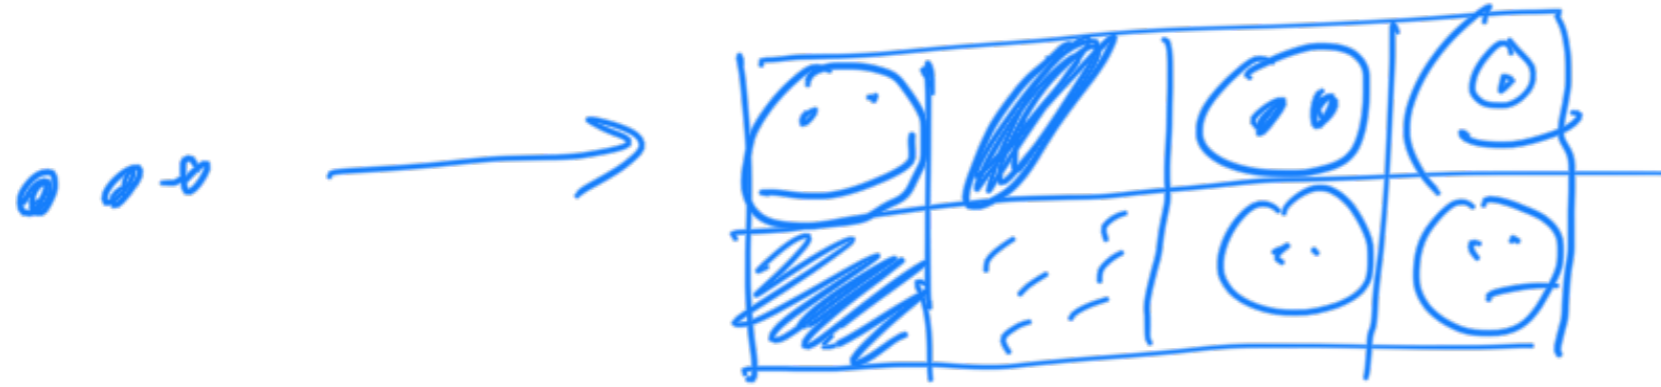


Mid-level

Input Image

maps:  
low-level  
features

features



High-level  
features

Putting it all together:

The "prototypical" CNN:





Image

Layer

Layer

Repeat multiple times

Repeat multiple times

automatic feature extraction

Prediction  
(MLP network from last class)

output from feature extraction is input into predictor

backpropagation

# Common Hyperparameters in CNNs :

## Dense Layers

- # of nodes
- Non-linear activation function

## Convolutional Layers

- # Kernels
- Kernel size
- padding



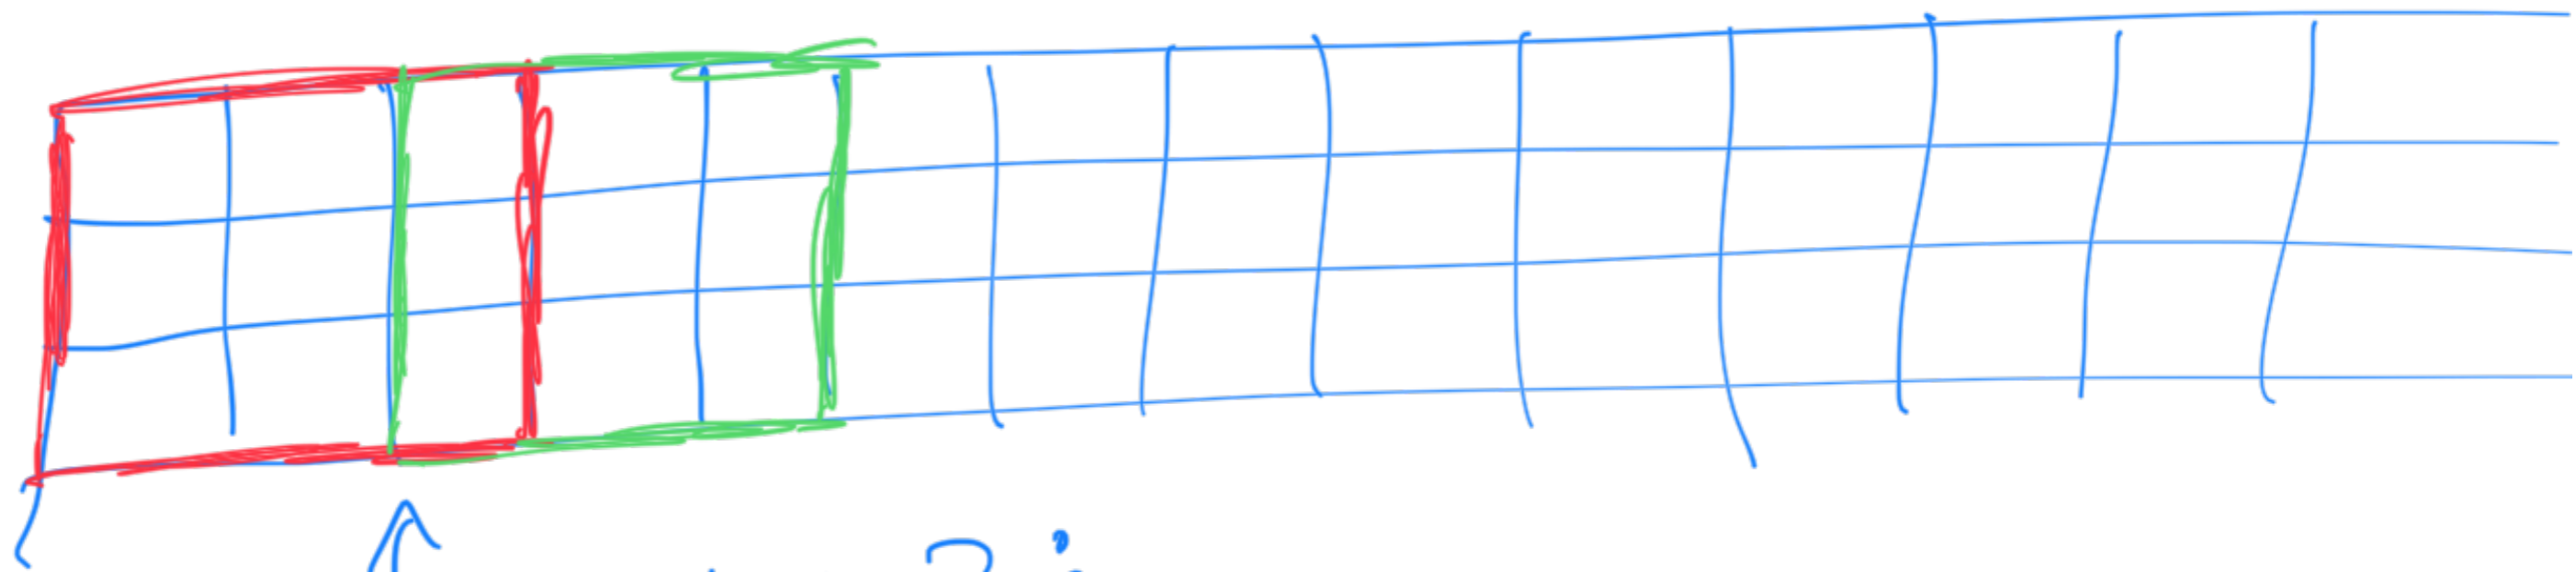
(can help ensure that output dimension =



dimension  
(input dimension)

• Stride

↳ how much to move the filter



Stride = 2

Step 1

Step 2

Pooling Layers

- window size
- padding
- stride

# Transfer Learning

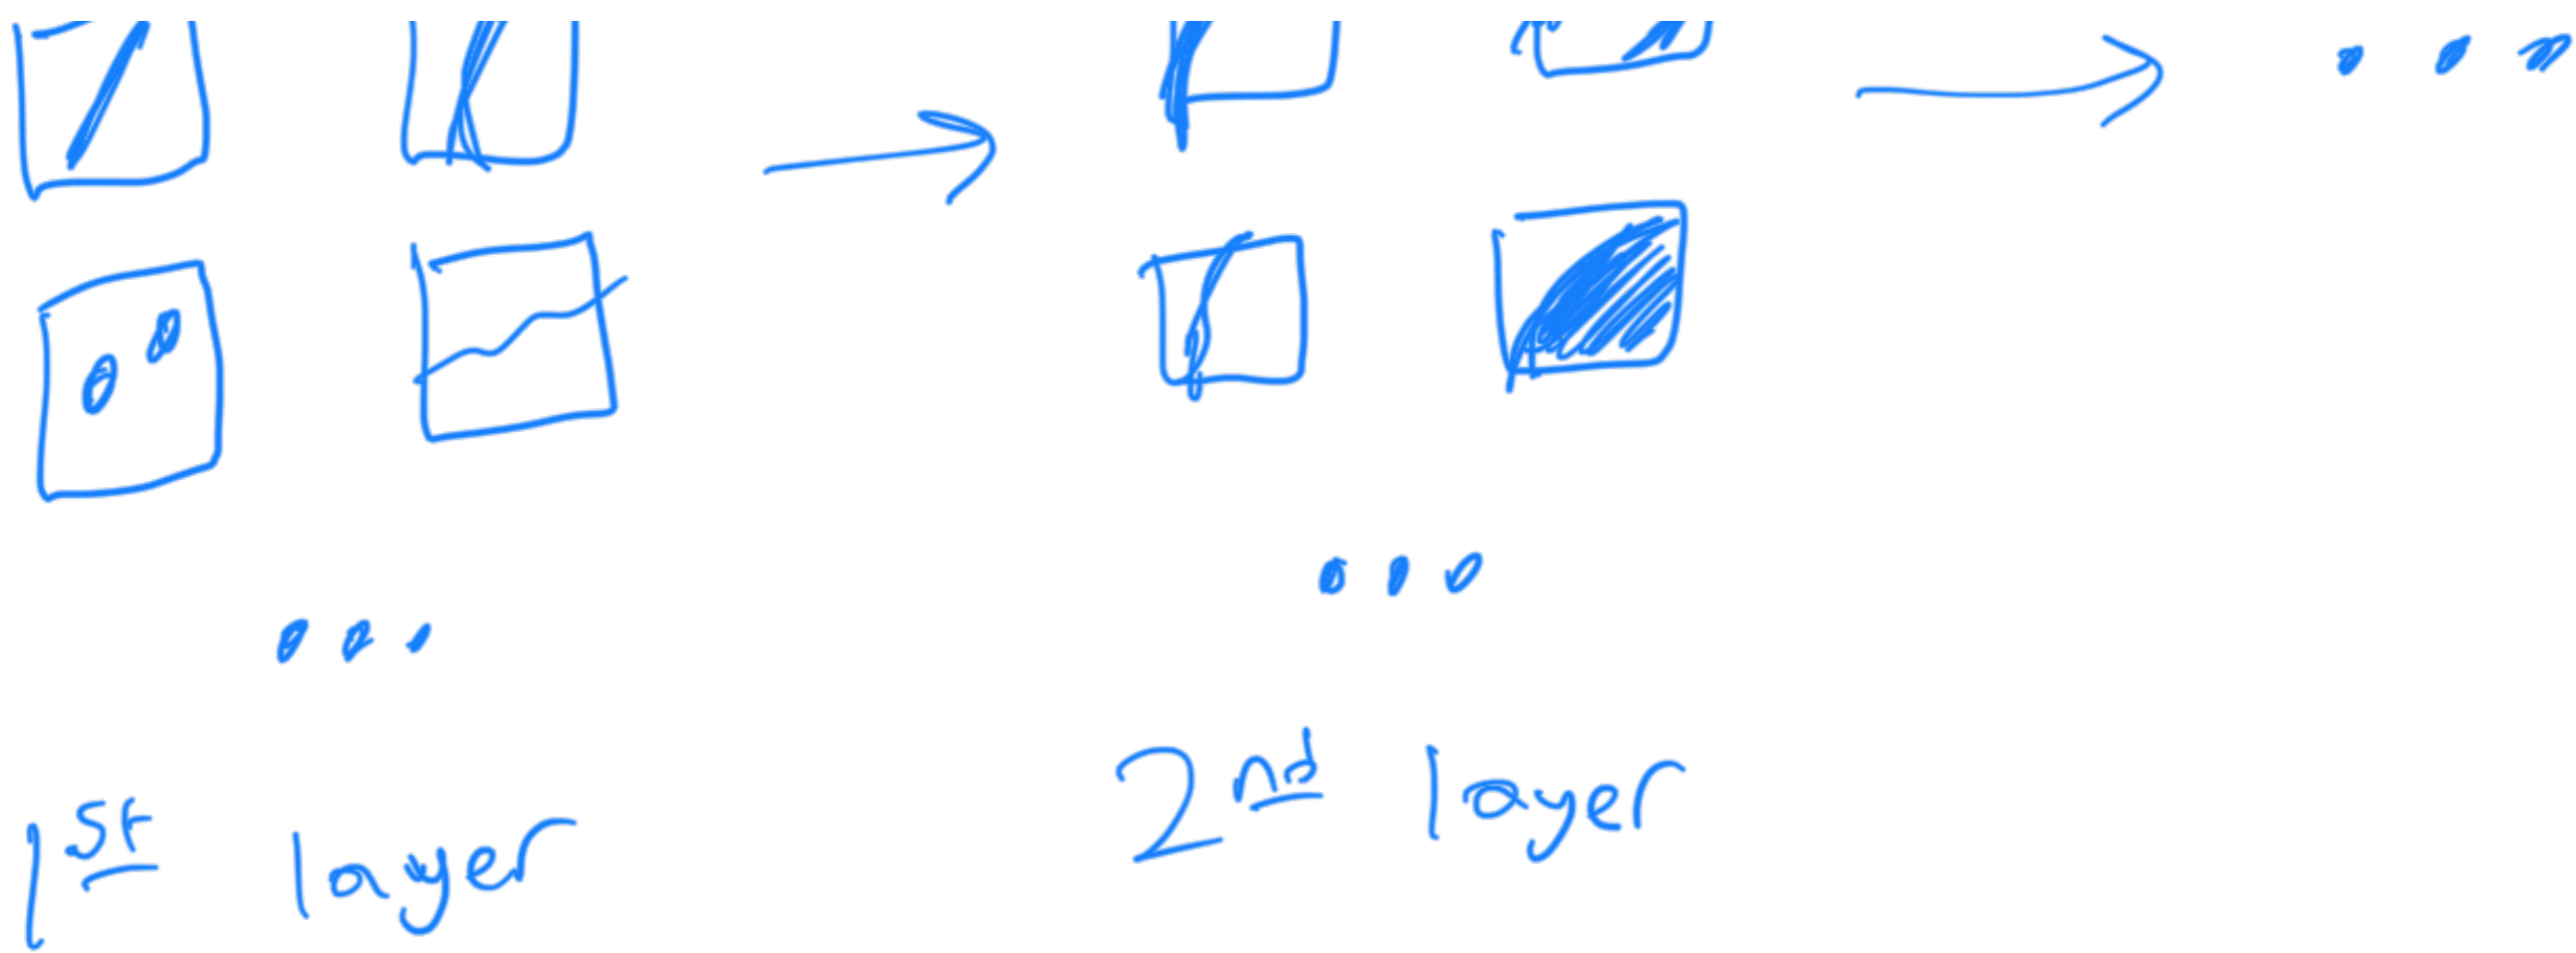
Many lower-level features can be re-used for a variety of purposes

## Example:

Train a CNN to predict cat vs dog.

Feature maps:





The above features are useful to many prediction tasks besides cat vs. dog.

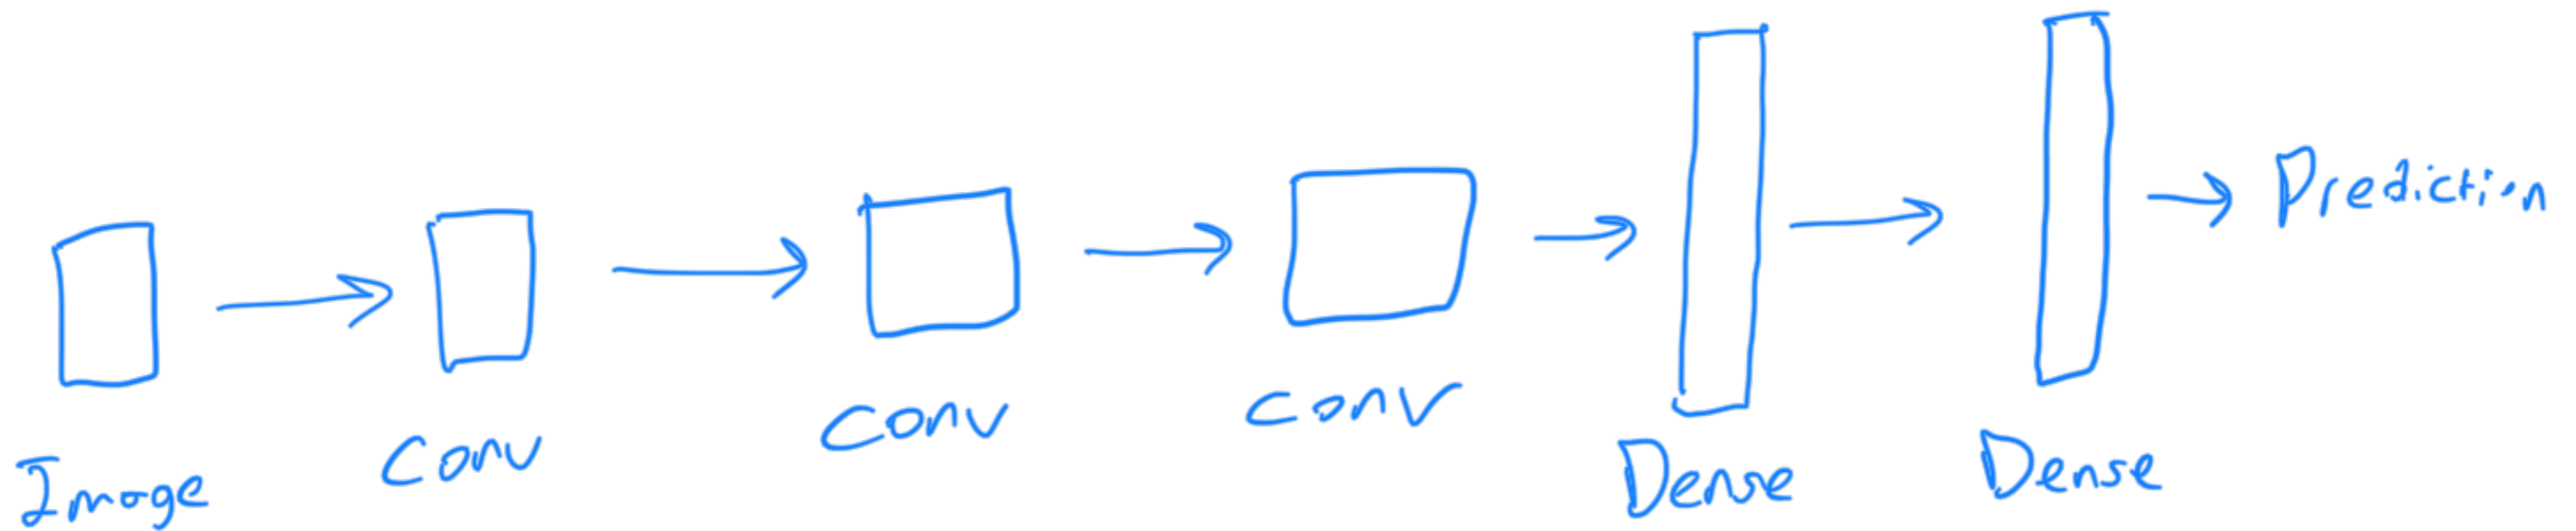
### Steps of transfer learning:

- ① "Pre-train" a CNN on a dataset
- ② "Freeze" the weights of the first  $N$

any dataset

# Convolutional layers

③ Train (aka transfer learn) the remaining weights on dataset of interest



Keep these weights fixed during gradient descent

only update these weights during GD

Most common pre-training dataset is

TensorFlow

# ImageNet

- > 14 million images
- > 20,000 classes
- each class is a common object (e.g, Chair, remote control, grasshopper, ...)

Many CNN architectures pre-trained on ImageNet, are readily available

Starting with ImageNet weights is common practice

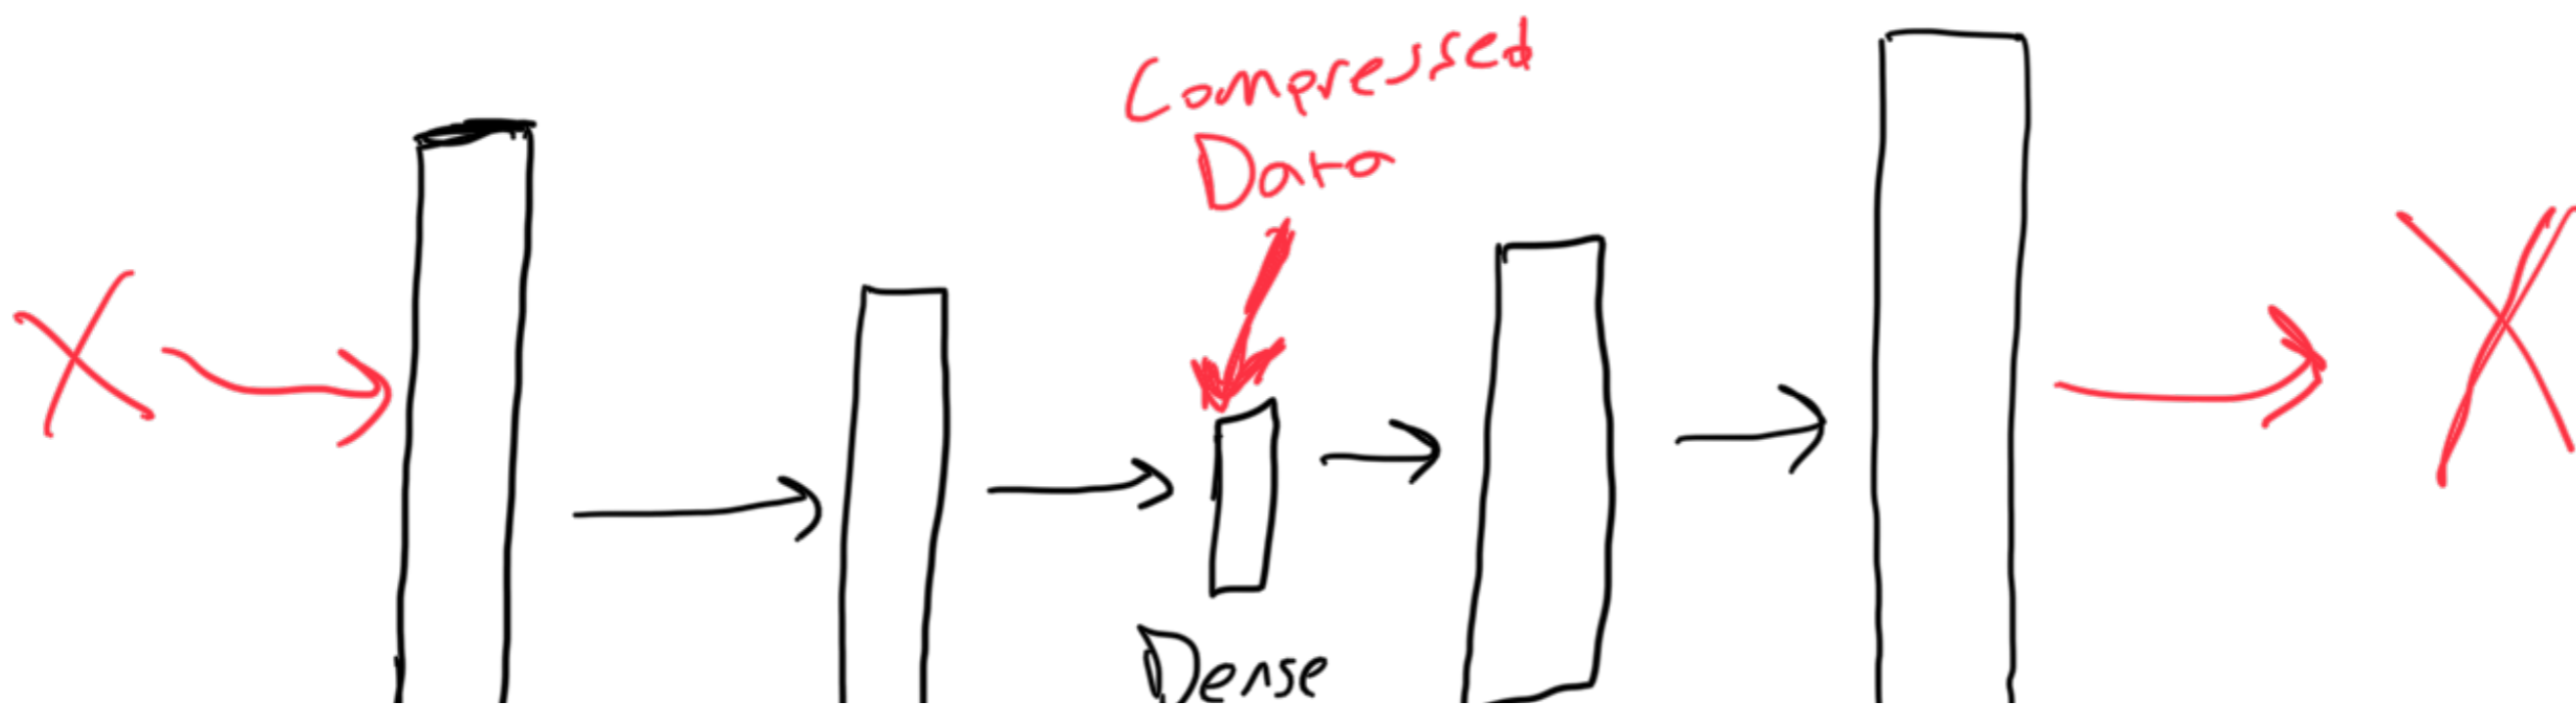


DL enables way more than  
just supervised learning

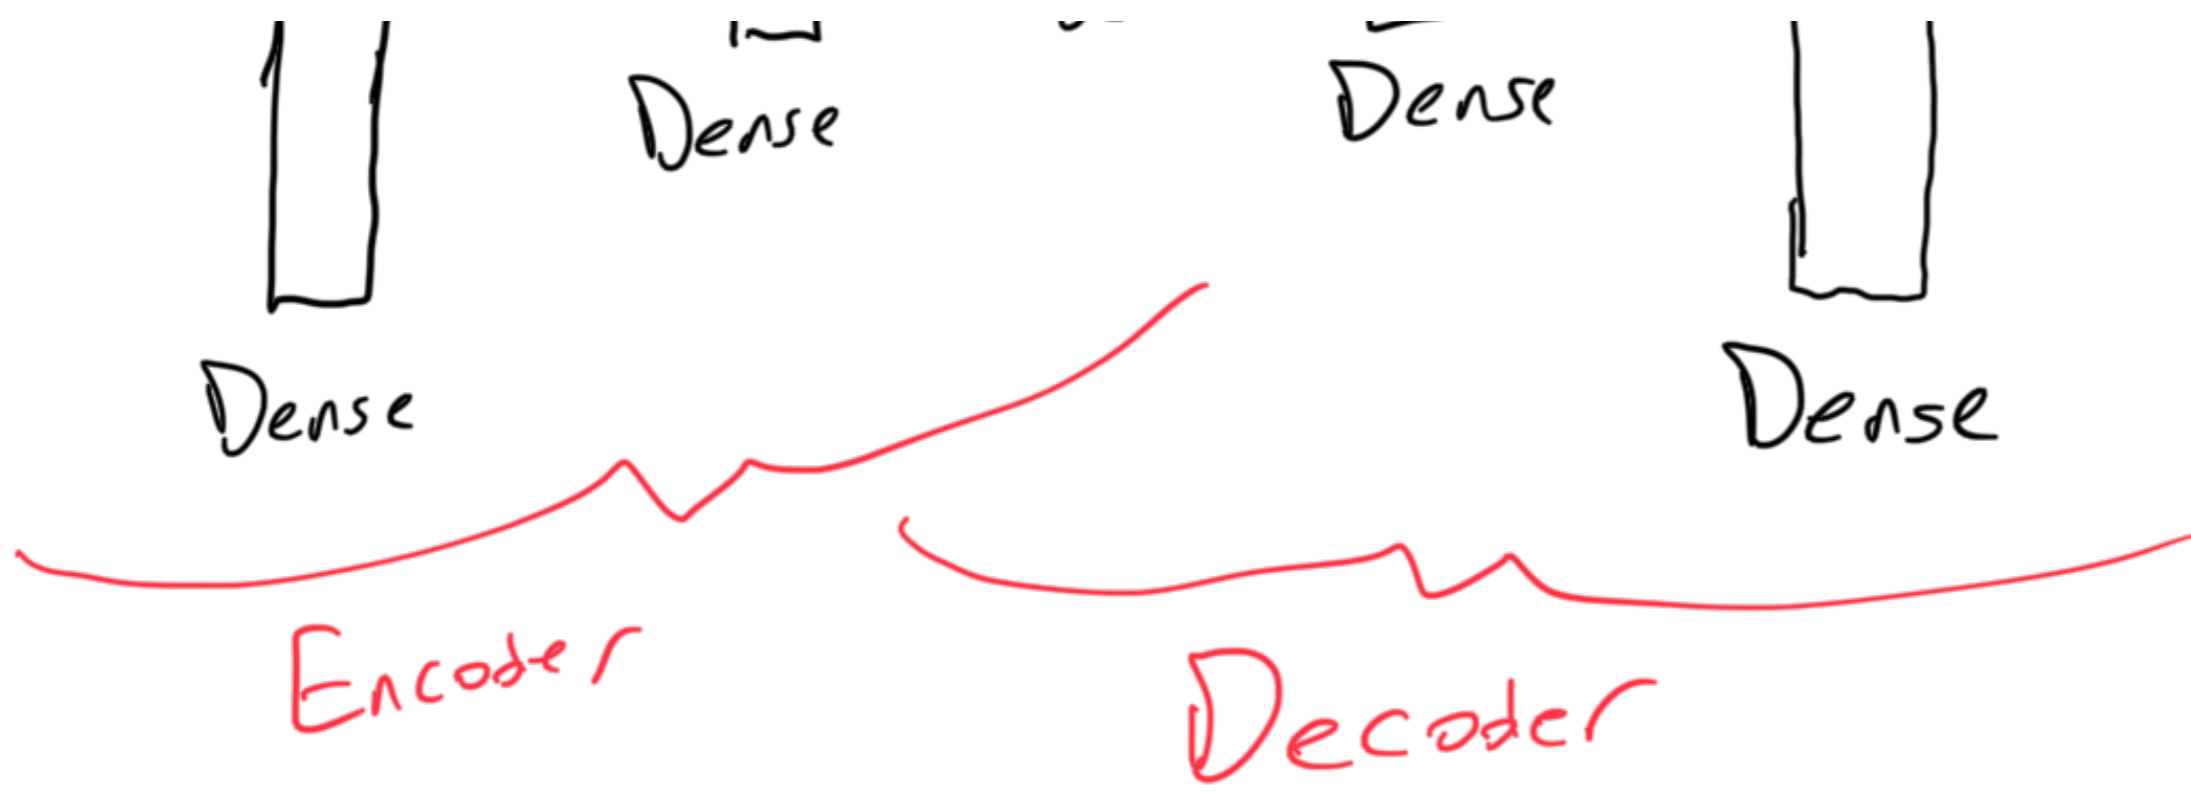
1st class example of many!

Dimensionality Reduction with  
Neural Networks: Autoencoders

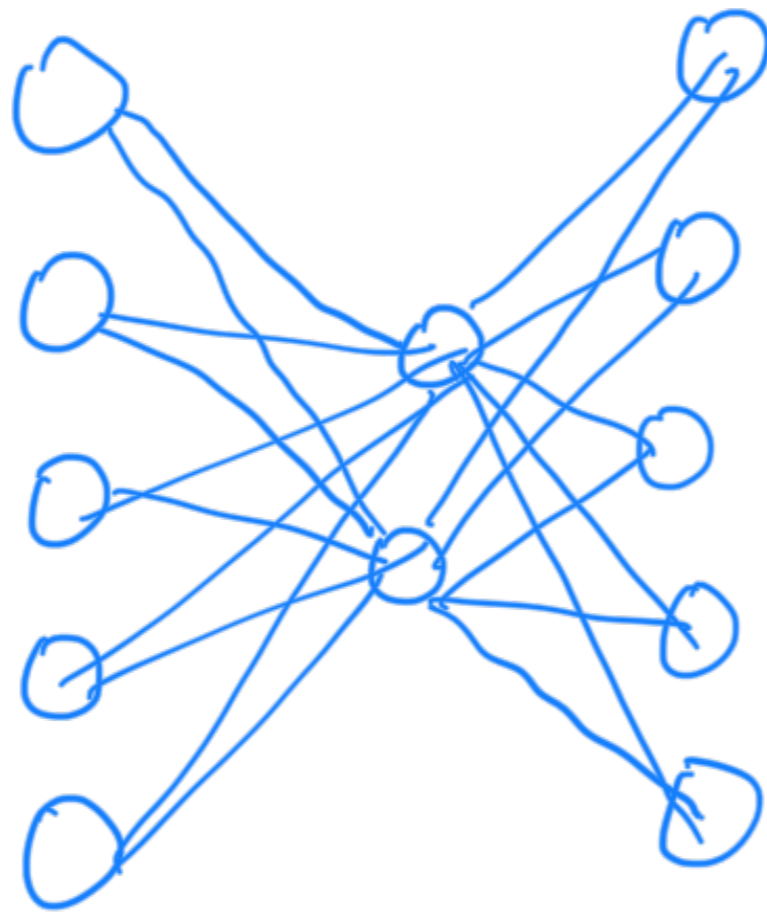
---







Example:



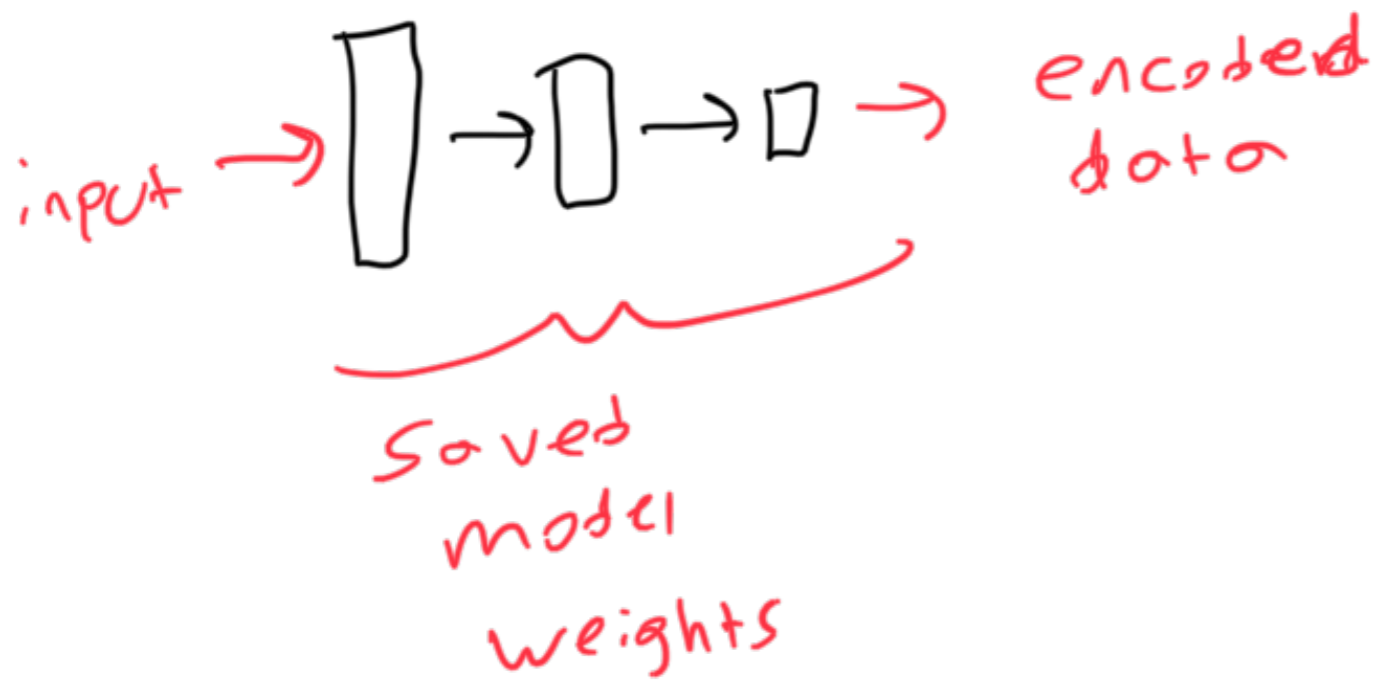
Distance Loss: how similar is

# Reconstruction

$X_{\text{train}}$  to  $X_{\text{train}}$ ?

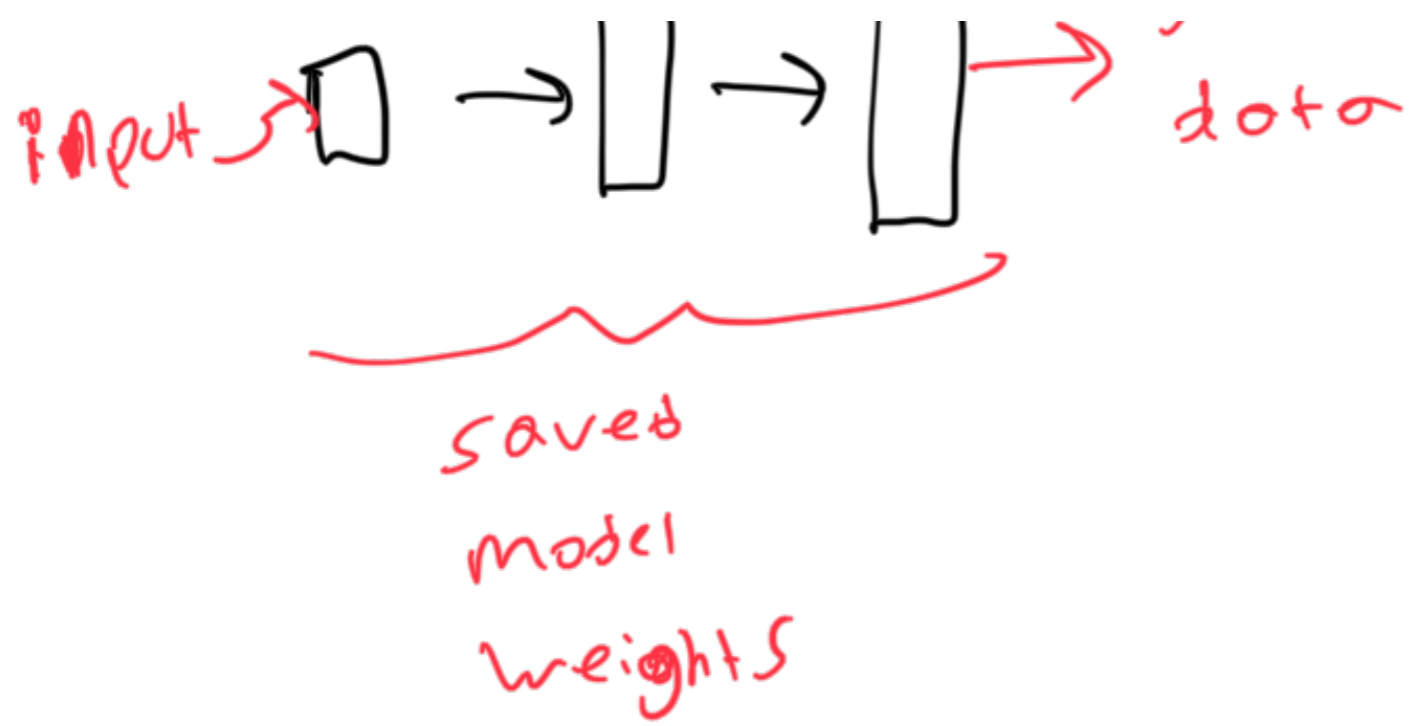
Cross Entropy ( $X_{\text{train}}, X_{\text{train}}$ )

## To encode data:



## To decode data:





Convolutional Autoencoders  
w/ image data:

work well

