Machine Learning Methods for Atmosphere, Ocean, and Climate Science

Lecture 3: Implementing CNNs in PyTorch

Mathematical modeling of Climate, Ocean, and Atmosphere processes
International Centre for Theoretical Sciences, TIFR, Bengaluru, India

Aman Gupta

Stanford University
DATAWAVE

Lecture 1

- Parametric estimation
- Introduction to deep neural networks
- The training algorithm

Lecture 2

- The PyTorch library
- Implementing artificial neural nets in PyTorch

PyTorch

Lecture 3

- Implementing Convolutional Neural Networks in PyTorch

Lecture 4

- Applications of ML in climate science

Convolutional Neural Networks (CNNs)

- Vanilla ANNs transformed the image to vectors. Not shape invariant.
- Makes it challenging, especially when identifying small-scale features in an image

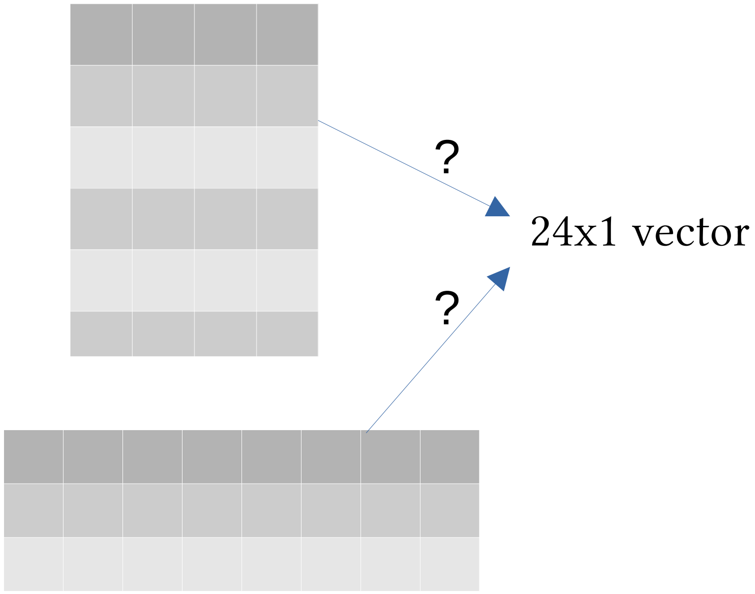

Convolutional Neural Networks (CNNs)

- Vanilla ANNs transformed the image to vectors. Not shape invariant.
- Makes it challenging, especially when identifying small-scale features in an image

CNNs are a special class of NNs, that allow working with images in a geometry-preserving way through the use of convolutions.

$$(f * g)(t) := \int_{-\infty}^{\infty} f(\tau) g(t - \tau) \, d\tau.$$

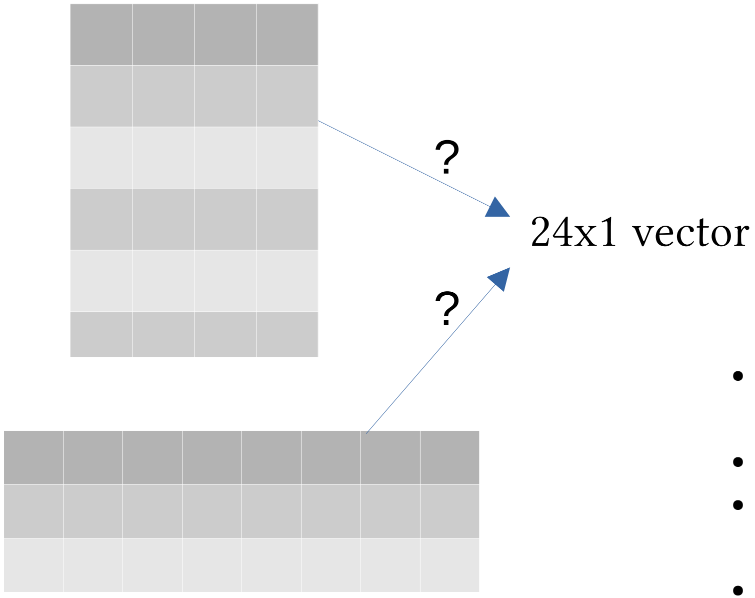

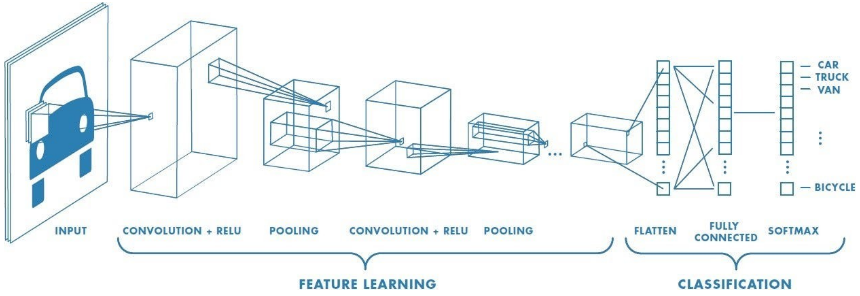

- Convolve the input image with a **kernel/filter by striding,** to create a **feature map**
- Can add padding to the image for shape preservation
- Downsample the image using **pooling** (helps with invariance): max pooling, min pooling, average pooling etc.
- Input downsamples maps into a **fully connected ANN**
- Train to get the optimal weights and optimal filters

Applying a filter
to Mario

1/9	1/9	1/9
1/9	1/9	1/9
1/9	1/9	1/9

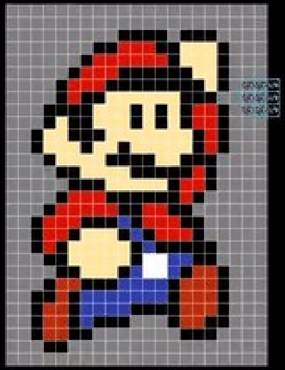

Image by @RetroArtist18

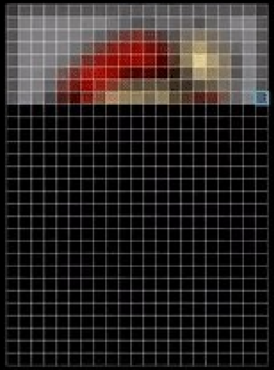

Creates a pixelated/blurred version of Mario

Applying a filter
to Mario

1/9	1/9	1/9
1/9	1/9	1/9
1/9	1/9	1/9

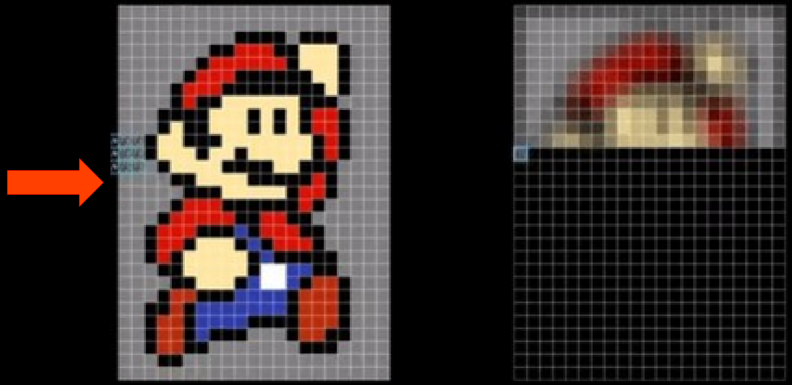

Creates a pixelated/blurred version of Mario

Applying a filter to Mario

1/9	1/9	1/9
1/9	1/9	1/9
1/9	1/9	1/9

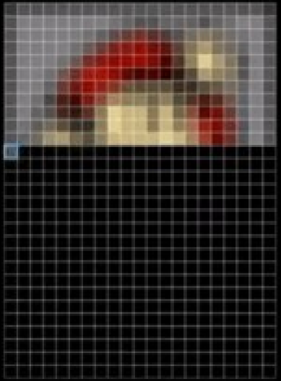

Creates a pixelated/blurred version of Mario

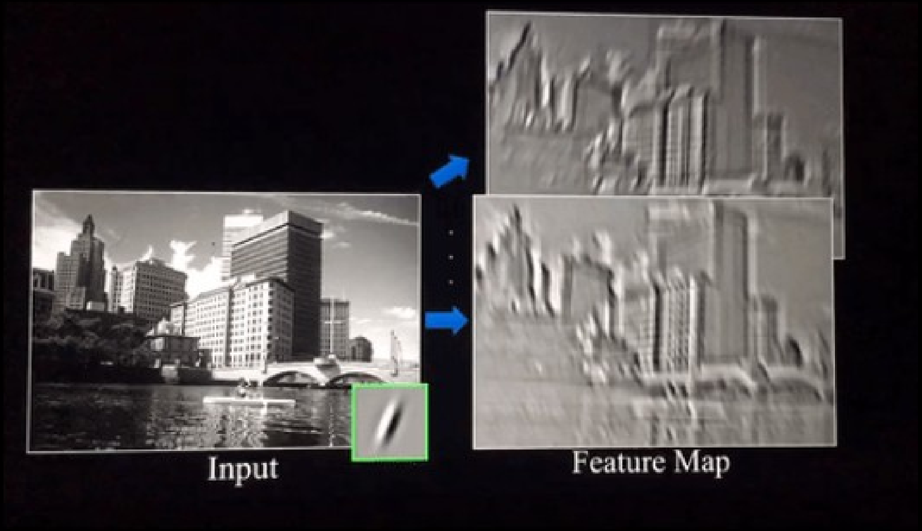

Different filters identify different features

Prewitt filter

-1/2	0	1/2
-1/2	0	1/2
-1/2	0	1/2

Sobel filter

-1/2	-1	-1/2
0	0	0
1/2	1	1/2

(Cat figures: www.paperspace.com)

Different filters identify different features

Applying a filter to Mario

1/9	1/9	1/9
1/9	1/9	1/9
1/9	1/9	1/9

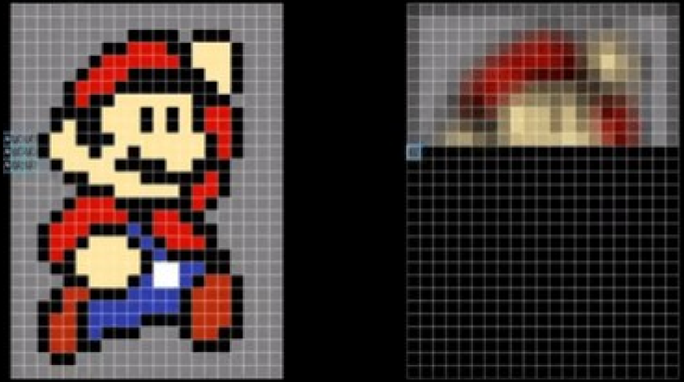

Creates a pixelated/blurred version of Mario

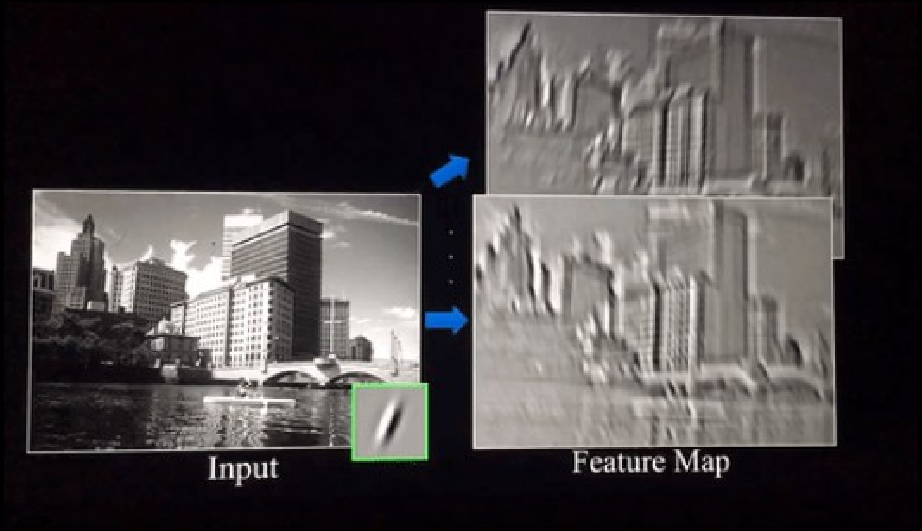

Laplacian filter

-1/9	-1/9	-1/9
-1/9	8/9	-1/9
-1/9	-1/9	-1/9

(Cat figures: www.paperspace.com)

Different filters identify different features

Applying a filter to Mario

1/9	1/9	1/9
1/9	1/9	1/9
1/9	1/9	1/9

Creates a pixelated/blurred version of Mario

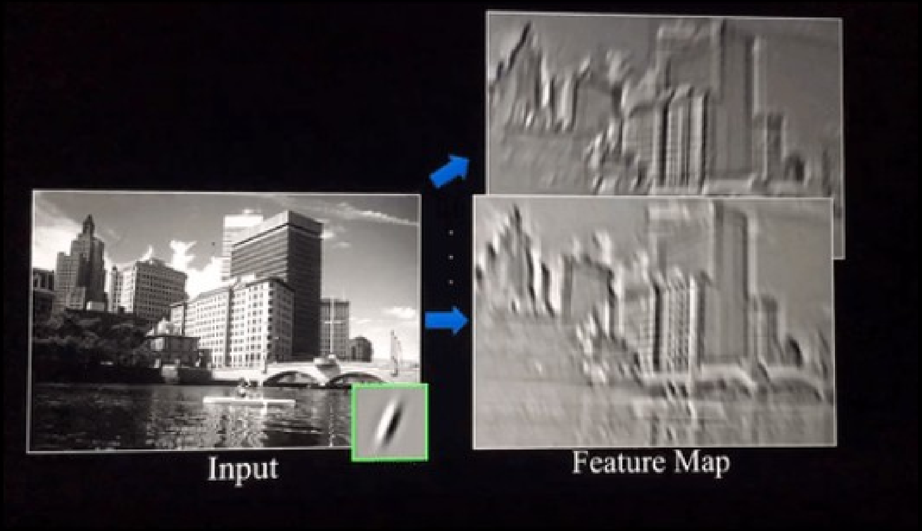

Prewitt filter

-1/2	0	1/2
-1/2	0	1/2
-1/2	0	1/2

Sobel filter

-1/2	-1	-1/2
0	0	0
1/2	1	1/2

Laplacian filter

-1/9	-1/9	-1/9
-1/9	8/9	-1/9
-1/9	-1/9	-1/9

1	0	0
0	1	0
0	0	1

?

Convolutional Neural Networks (CNNs) learn the optimal filters through training. They do not use these pre-defined filters

Feature Maps can be Complex

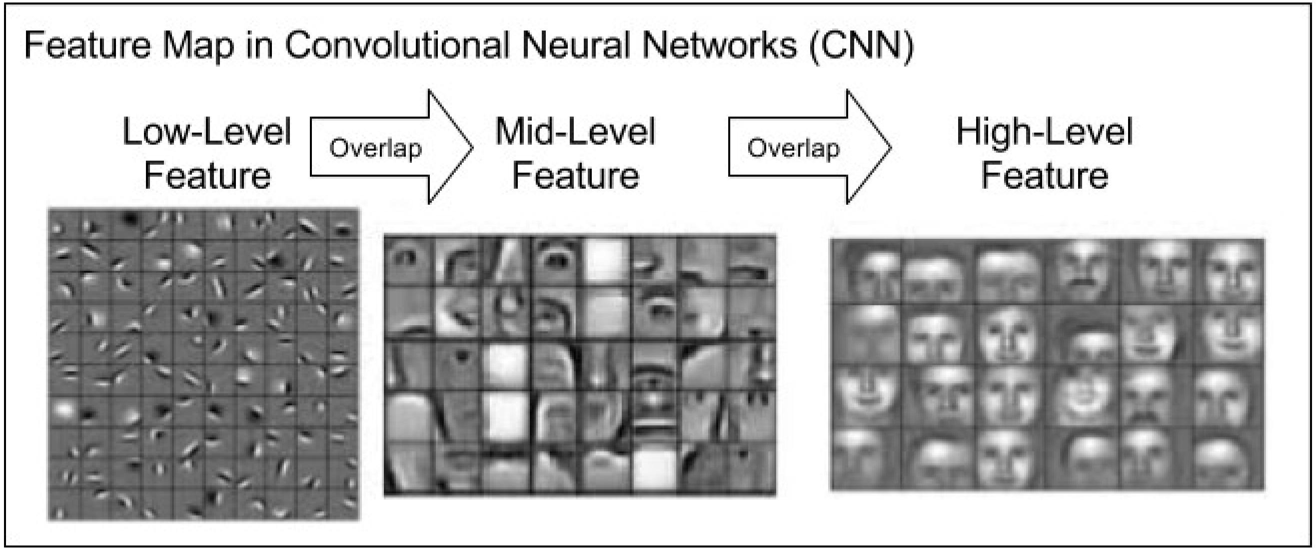

Feature maps for clouds (Phung and Rhee 2019)

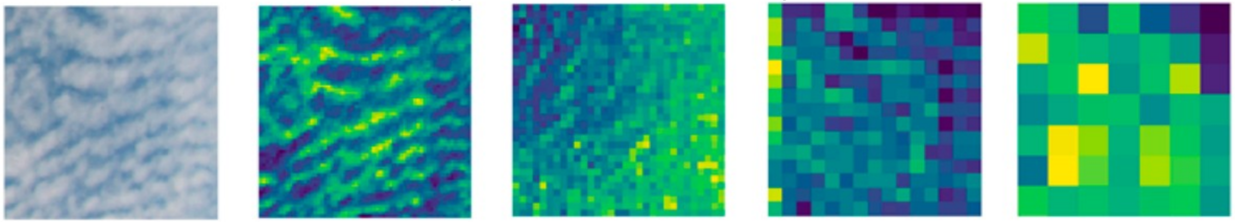

Let's Code our First Convolutional Neural Network!

Jupyter Notebook URL: tiny.cc/coaps_lec3

tiny.cc/coaps_html

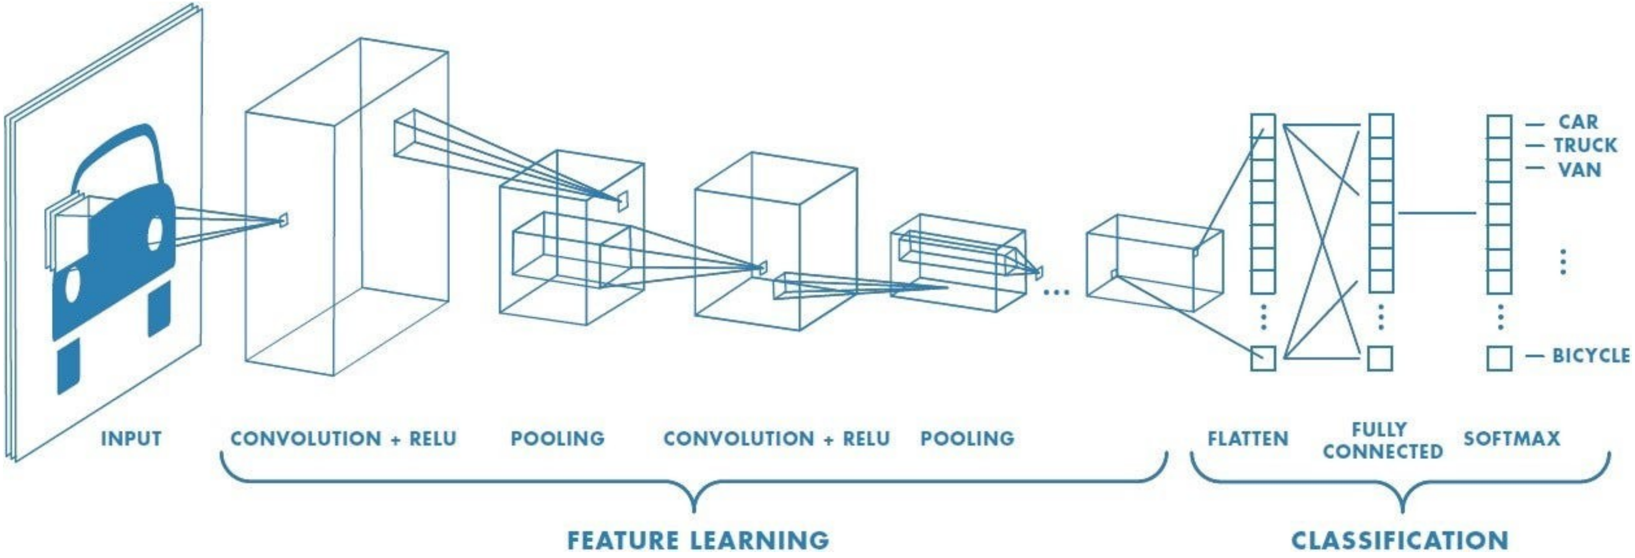

Let's Code our First Convolutional Neural Network!

Jupyter Notebook URL: tiny.cc/coaps_lec3
tiny.cc/coaps_html

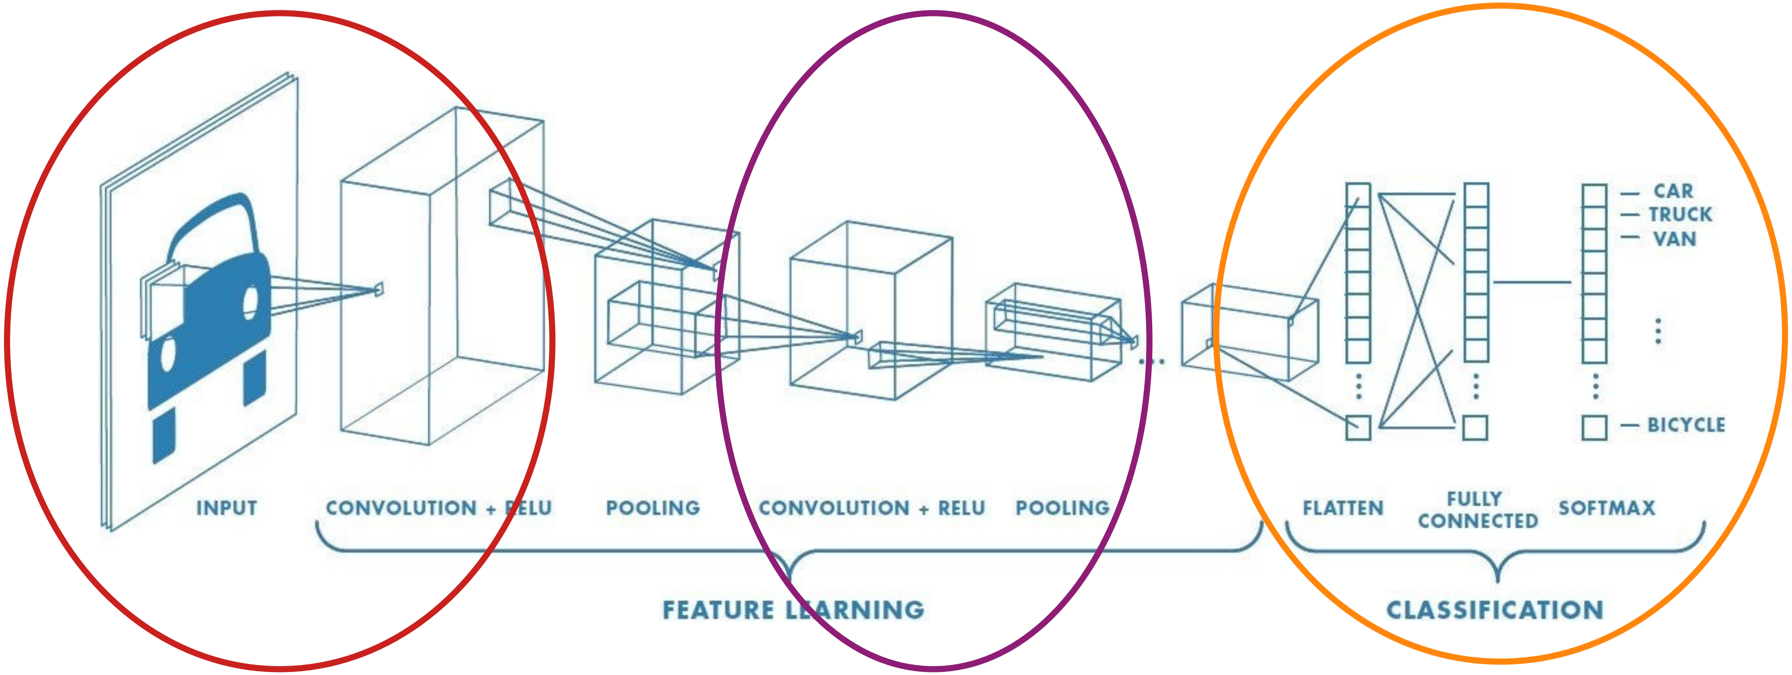

We know this part already!