


# Convolutions - II


# Recap

- Fully connected layer - too many parameters
- New building block:
  - Convolution operator
  - Locality
  - Parameter sharing

# Convolution Animation



# Padding + Strides Animation



# Above GIFs From

- [https://github.com/vdumoulin/conv\\_arithmetic](https://github.com/vdumoulin/conv_arithmetic)

# Code

2D Convolution layer:

- Keras:

```
Conv2D(  
    filter_size=24,  
    kernel_size=(2,2),  
    strides=(1,1),  
    padding='valid',  
    input_shape=(128, 128, 3),  
    channels_last=True  
)
```

# What Does This Mean

- 24 Filters
- Each filter is 2x2 - (patch size, FOV)
- Stride 1x1 - The adjacent patches are one over in both directions.
- Valid padding - don't add any padding (tensorflow bs)
- input\_shape - 128x128 color image (R, G ,B channels)
- Channels\_last - tensor shape - see explanation

# Parameters

- 24 Filters
- 2x2 filter
  - So patch has 4 pixels
  - Linear combination so **4 weights**
- **24 \* 4 = 96**



# Typically

- Conv2D is followed by:
  - Pooling layers
    - MaxPool
    - AvgPool
    - . . .

# Pooling

- Also Image -> Image
- Operates at patch level as well
- Patches **Don't Overlap**


# MaxPool

- Just take the max in a patch

# AvgPool

- Just take the average in a patch

# MaxPool



# Code

- keras:

```
MaxPooling2D(  
    pool_size=(2, 2),  
    strides=None,  
    padding='valid',  
    data_format=None  
)
```

# What Do You Get

A New **Shrunken(ish)** image

# Parameters

- None



# Hyperparameters

- Stride
- Patch size

# Now

- We discussed
  - Image -> Image transformations
  - Pooling, Conv2D

# How does it fit in


- Cross entropy not defined for image
- Regression not defined for image
- Etc. Etc.

# Deep Convnet

- Series of image transformations
- At the end you get an image

# Just

- Flatten the resulting image
- Use normal fully connected layers from there on



# Code Example

- MNIST

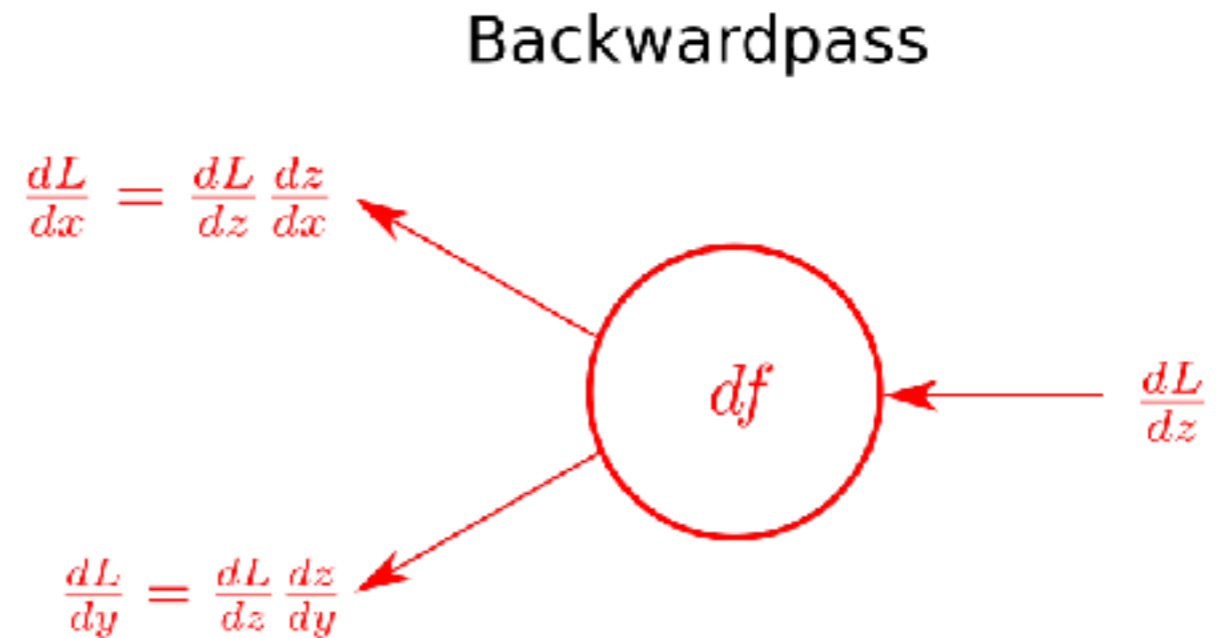

# Code Example

- MNIST Tensorflow
- [https://github.com/tensorflow/tensorflow/blob/master/tensorflow/examples/tutorials/mnist/mnist\\_deep.py](https://github.com/tensorflow/tensorflow/blob/master/tensorflow/examples/tutorials/mnist/mnist_deep.py)



# Training Details

# The Backprop Picture



Here

$$\begin{bmatrix} x_1^1 & x_1^2 & x_1^3 \\ x_2^1 & x_2^2 & x_2^3 \\ x_3^1 & x_3^2 & x_3^3 \end{bmatrix} \begin{bmatrix} w_1^1 & w_1^2 \\ w_2^1 & w_2^2 \end{bmatrix}$$

$$\begin{bmatrix} h_1^1 & h_1^2 \\ h_2^1 & h_2^2 \end{bmatrix}$$

# Backward Pass

$$\begin{bmatrix} \partial L / \partial h_1^1 & \partial L / \partial h_1^2 \\ \partial L / \partial h_2^1 & \partial L / \partial h_2^2 \end{bmatrix}$$

# Observe

$$h_1^1 = w_1^1 x_1^1 + w_1^2 x_1^2 + w_2^1 x_2^1 + w_2^2 x_2^2$$

We Are Given

$$\partial L / \partial h_1^1$$

We Want

$$\partial L / \partial w_i^j$$

$$\partial L / \partial x_i^j$$

**i.e. the partial derivative w.r.t the filter and the input**

# Consider

$$\partial L / \partial w_1^1 = \partial L / \partial h_1^1 \cdot X_1^1 + \partial L / \partial h_1^2 \cdot X_1^2 + \partial L / \partial h_2^1 \cdot X_2^1 + \partial L / \partial h_2^2 \cdot X_2^2$$



# Gist

- Backward pass is also a convolution

# Initialization

- Guess how?

# Gist


- Error gradient is also a convolution

# Arsenal

- A few loss functions
- Fully connected layers
- Convolutional layers


Examples

# LeNet



# AlexNet

- Imagenet LSVRC-2010 Challenge
- 1.2 Million images
- 1000 classes
- The first famous convnet





# VGG



# VGG

- All kernels are 3x3
- All pooling layers are 2x2

# Thought Experiments

# Flickr Example

- Images
- Attributes per image
- Figure out tags

# Siamese Net Example

- Face pair
- Same person?

# Object Detection Example

- Simple technique

# Convolution Wins

- Big wins recently in
  - Machine translation
  - Audio synthesis
  - Video
  - etc.

Debugging



# Recap

- Fully Connected Layer
- ReLUs fold the space

# What Are The Convolutions Doing?

- Typical approaches are empirical
- Plot filter responses
- Visualize embeddings (representations)
- Visualize weights

# Visualizing Responses

## Deep Visualization Toolbox

[yosinski.com/deepvis](http://yosinski.com/deepvis)

#deepvis



Jason Yosinski



Jeff Clune



Anh Nguyen



Thomas Fuchs



Hod Lipson



<https://www.youtube.com/watch?v=AgkfIQ4IGaM>

# Intuition

- Pick some layer you want to “enhance”
- Backprop all the way to the image
- Result = trippy images

# MNIST TSNE

- Demo

# Deepdream

- <https://github.com/google/deepdream/blob/master/dream.ipynb>

# Deepdream

- Enhance responses
- how?

# Acknowledgements

- Ruslan Salakhutdinov - [http://www.cs.cmu.edu/~rsalakhu/10707/Lectures/Lecture\\_Conv1.pdf](http://www.cs.cmu.edu/~rsalakhu/10707/Lectures/Lecture_Conv1.pdf)
- Some figures from CS231N @ Stanford
- Good vibes from the Keras and TF teams
- Vincent Dumoulin - [https://github.com/vdumoulin/conv\\_arithmetic](https://github.com/vdumoulin/conv_arithmetic)