# Administrative

- **Poster Session** on Wednesday!
  - http://cs231n.stanford.edu/postersession.html
  - Facebook group: https://www.facebook.com/events/1024309900974098/
- **If using AWS, use EBS**
  - Instructions on Piazza
- **A3 Grades** out soon

# Together, we've defined Score Functions...

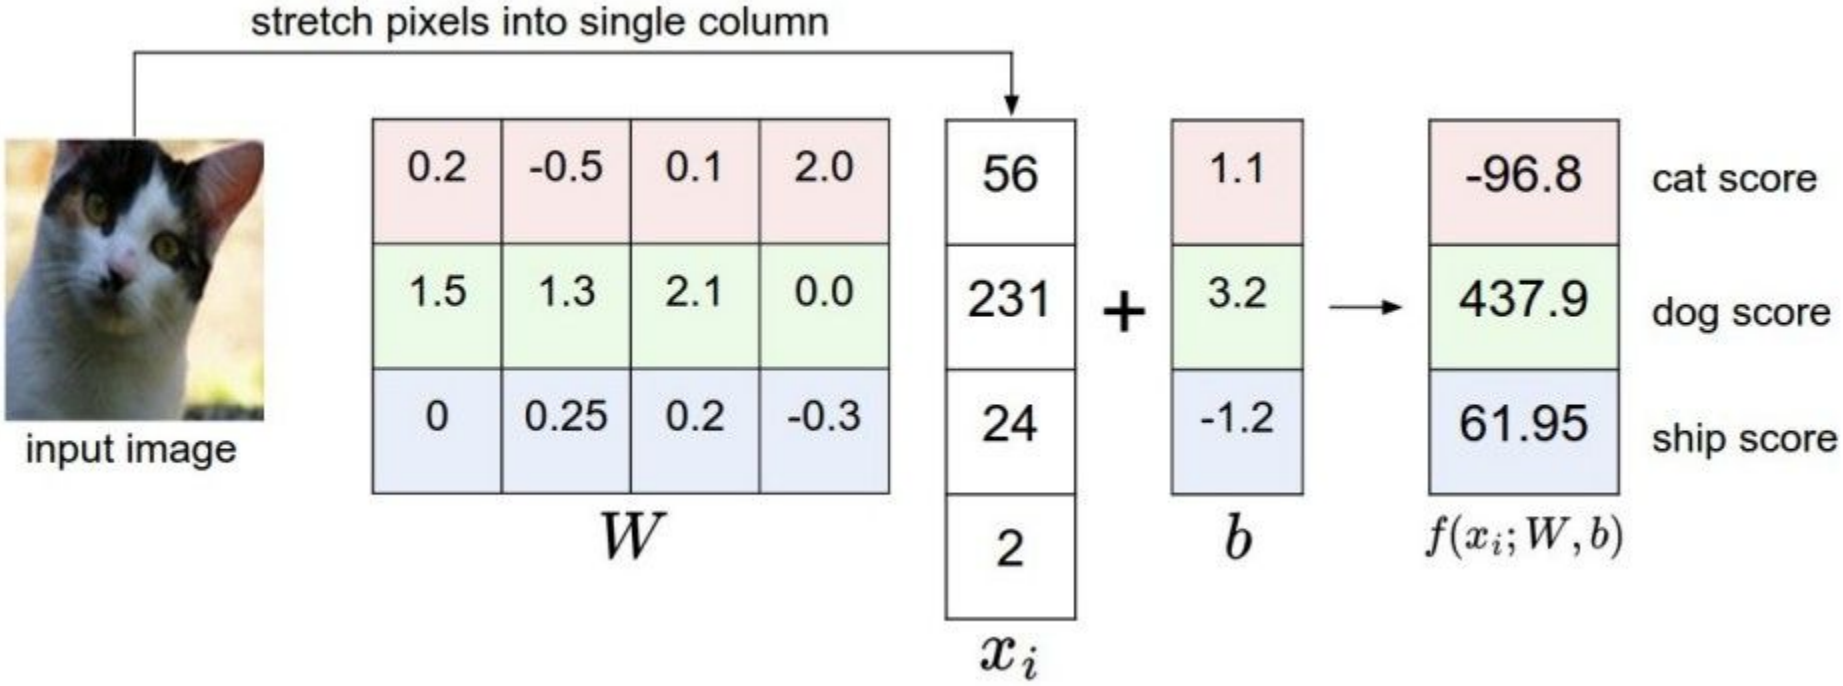

# And Loss Functions...

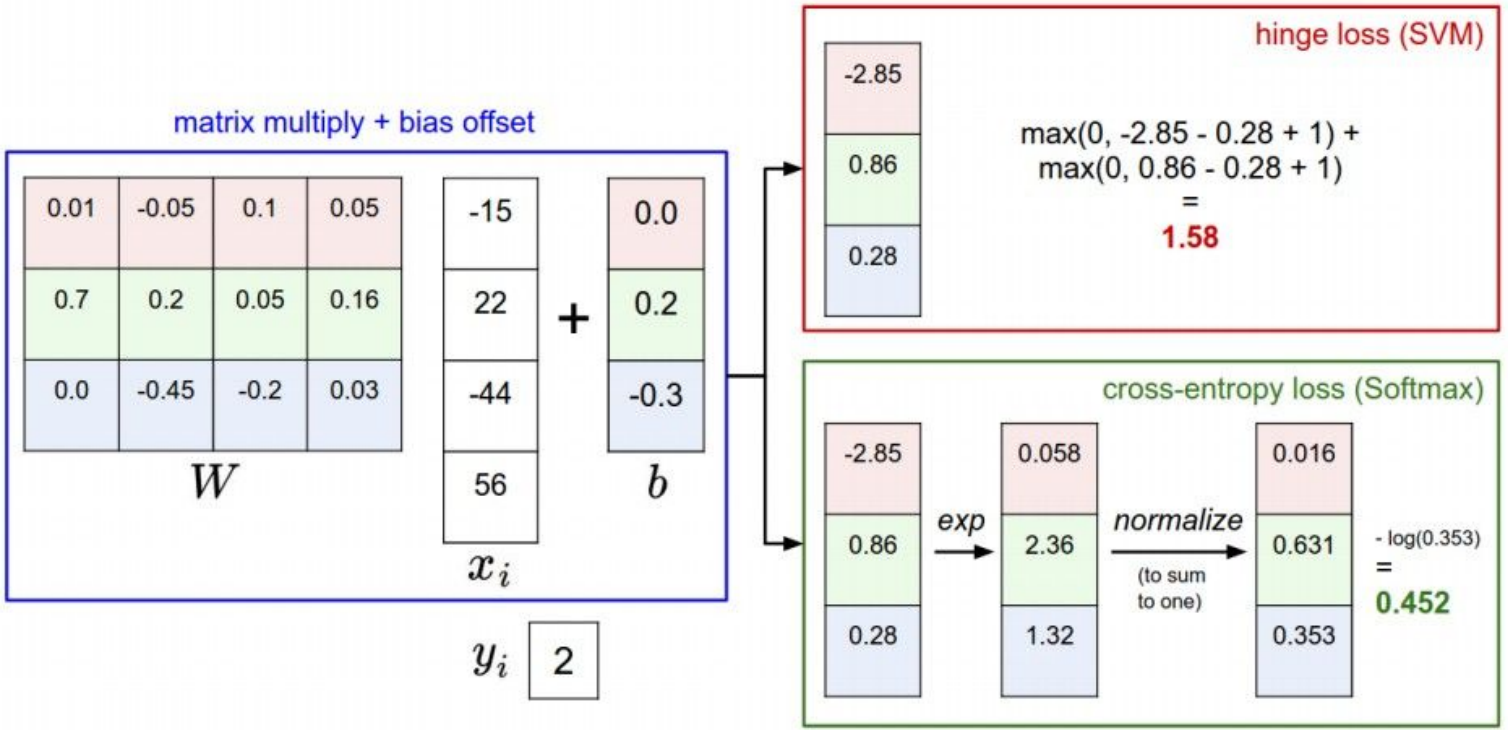

# We've learned how to optimize them...

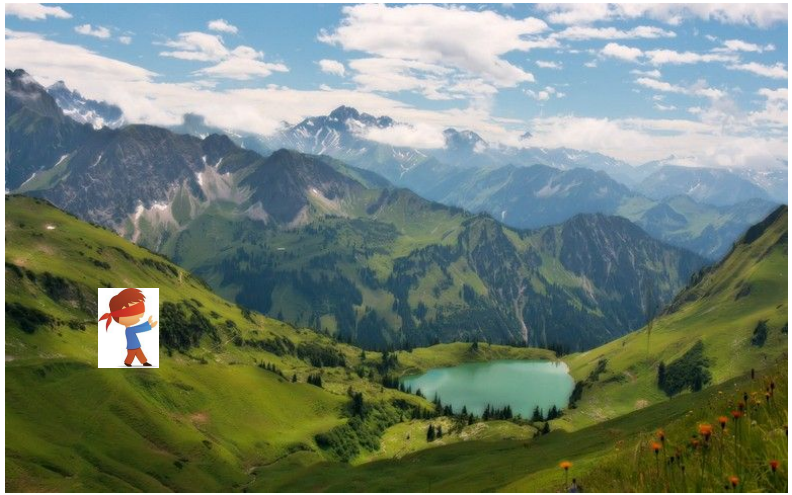

Chain rule:

$$\frac{\partial f}{\partial x} = \frac{\partial f}{\partial q}\frac{\partial q}{\partial x}$$

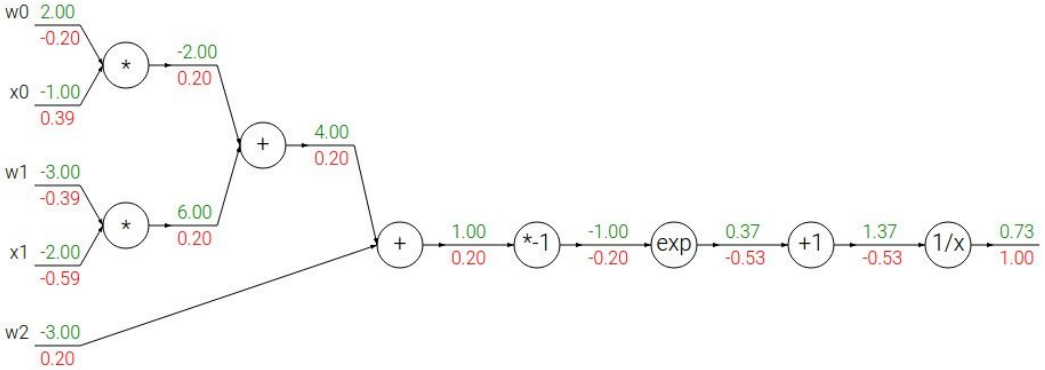

# We learned to express more powerful Score Functions...

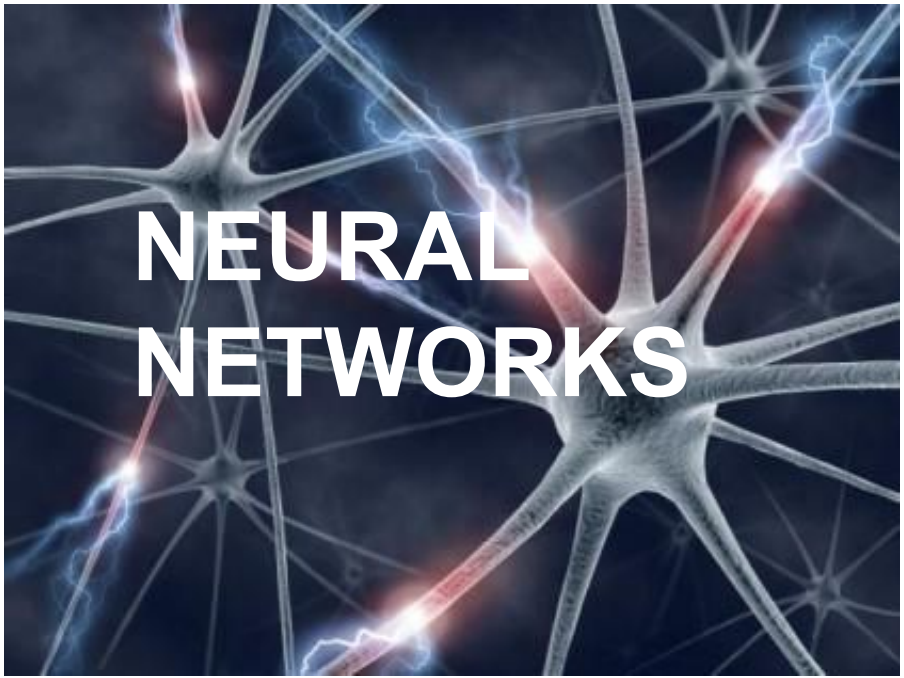


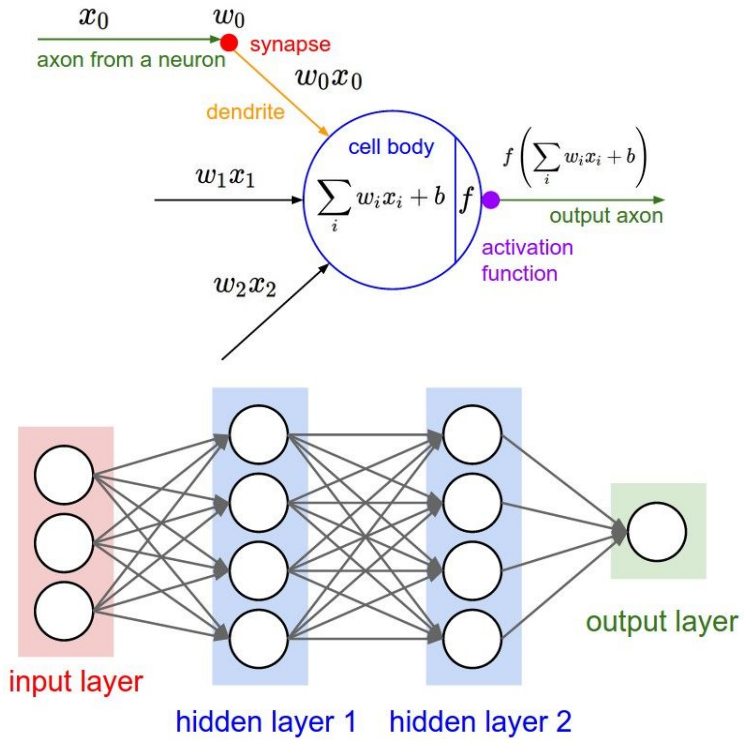

# For an extra wiggle...

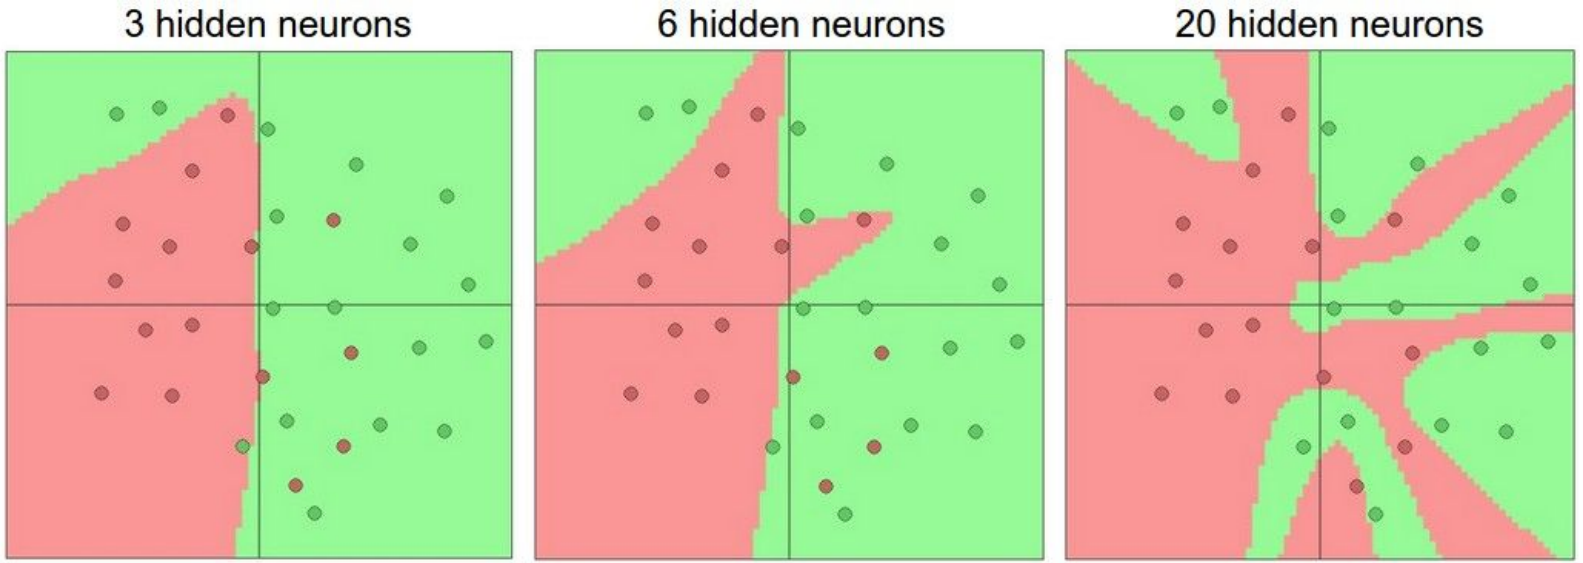

# Together we tamed the learning process...

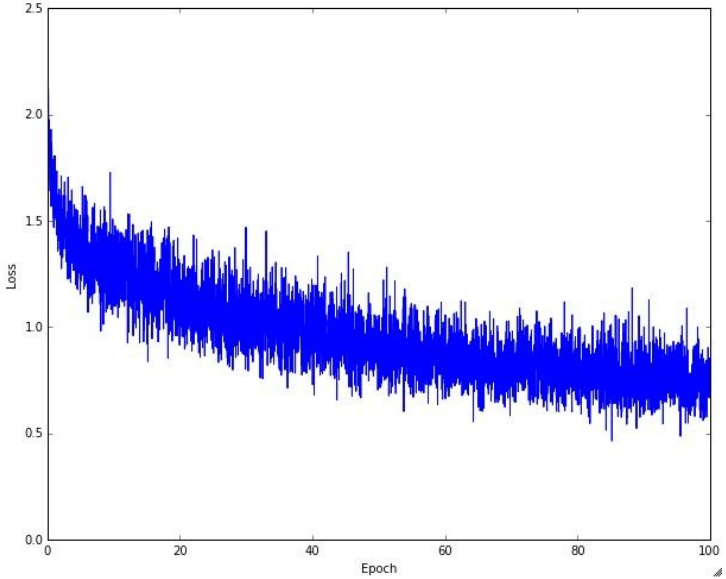


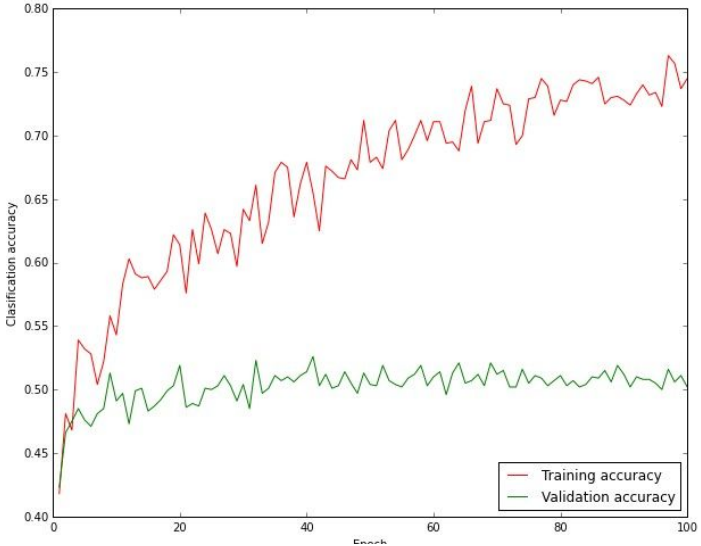

# Together we explored image-specific Neural Nets...

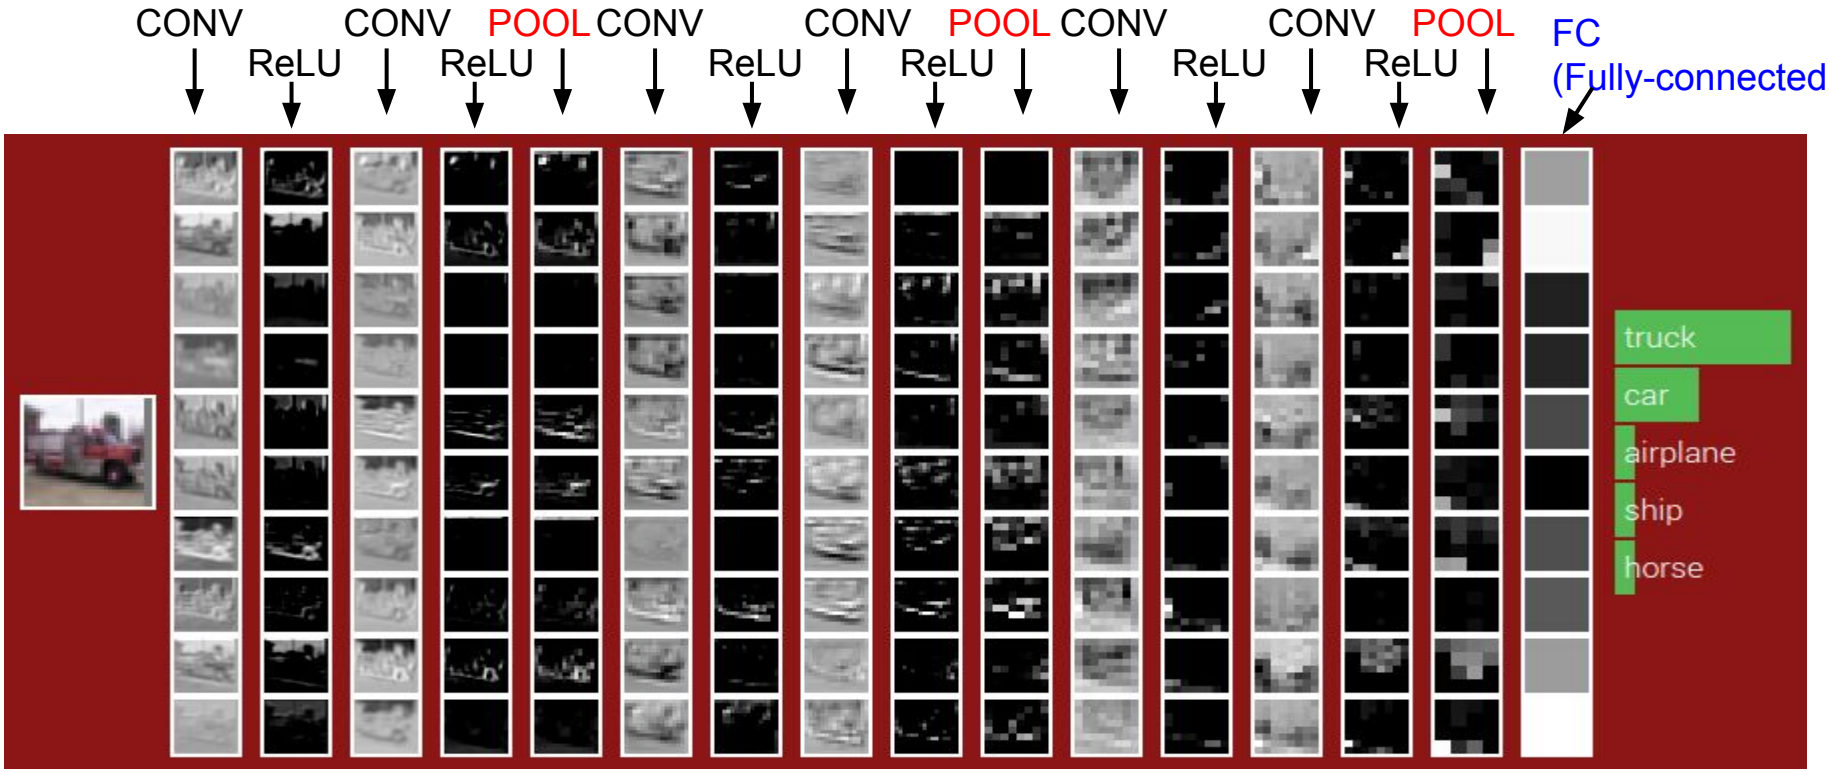

# We explored how they work...

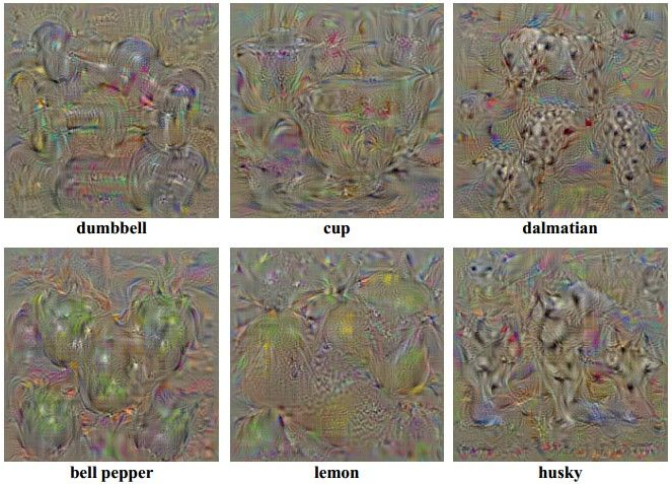


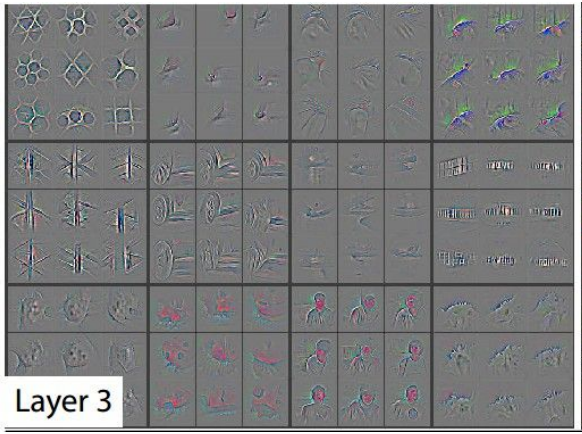


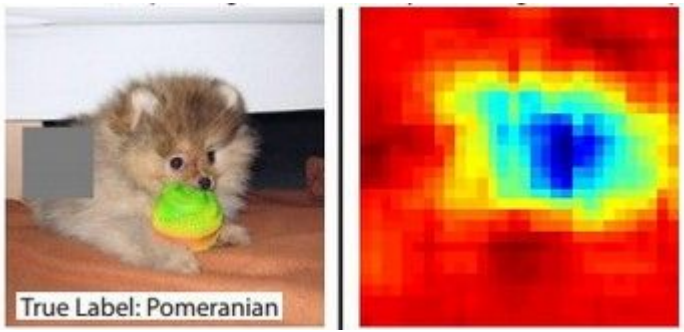

# And how they don’t… (but really they still do)

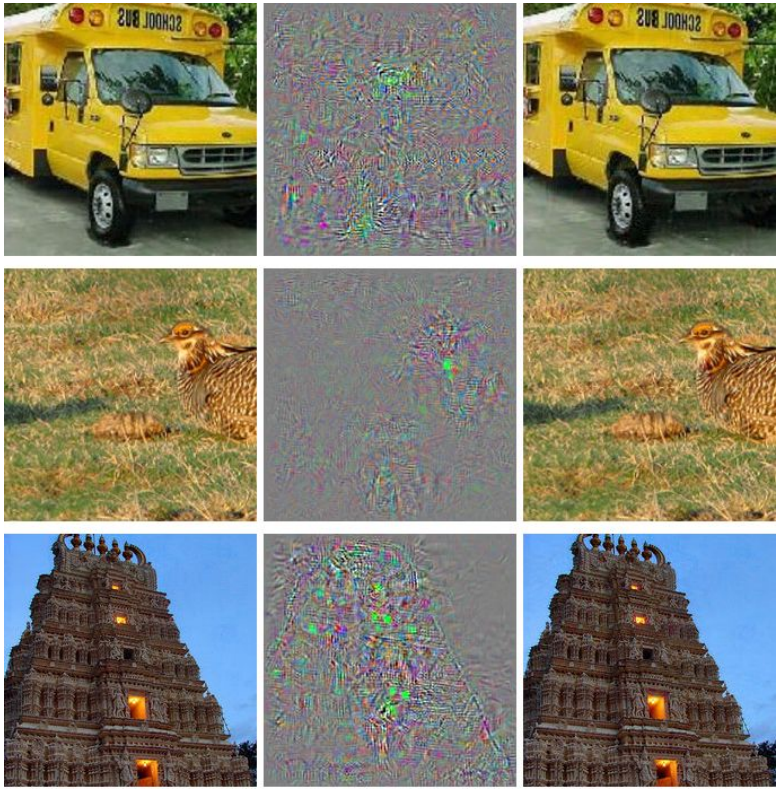

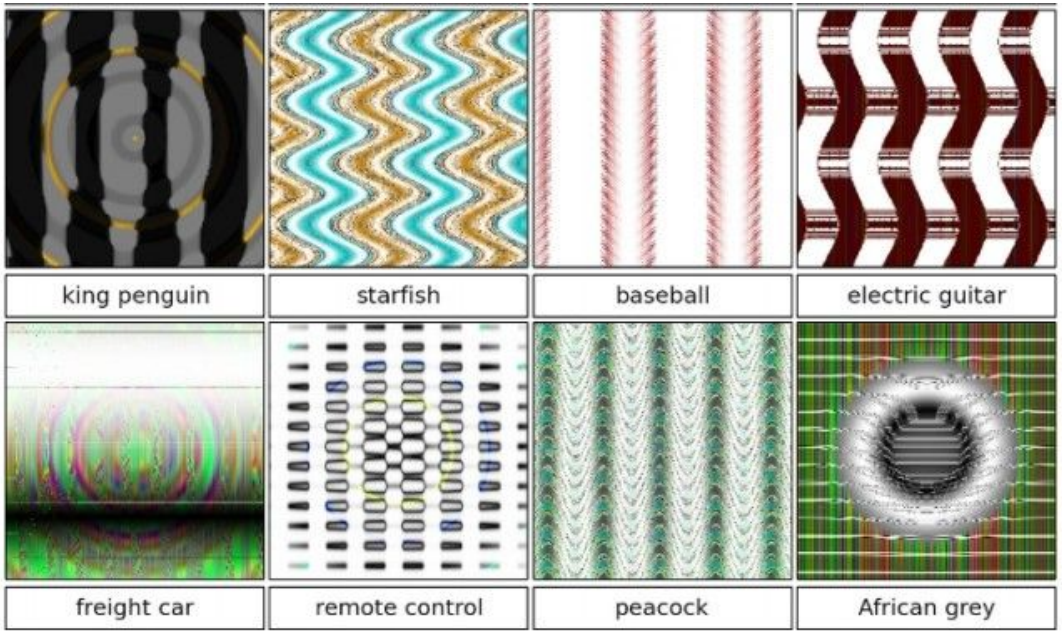

# And how to make art

# We looked at what makes ConvNets “tick”...

| ConvNet Configuration | | | | | |
|---|---|---|---|---|---|
| A | A-LRN | B | C | D | E |
| 11 weight layers | 11 weight layers | 13 weight layers | 16 weight layers | 16 weight layers | 19 weight layers |
| input (224 × 224 RGB image) | | | | | |
| conv3-64 | conv3-64 **LRN** | conv3-64 **conv3-64** | conv3-64 conv3-64 | conv3-64 conv3-64 | conv3-64 conv3-64 |
| maxpool | | | | | |
| conv3-128 | conv3-128 | conv3-128 **conv3-128** | conv3-128 conv3-128 | conv3-128 conv3-128 | conv3-128 conv3-128 |
| maxpool | | | | | |
| conv3-256 conv3-256 | conv3-256 conv3-256 | conv3-256 conv3-256 | conv3-256 conv3-256 **conv1-256** | conv3-256 conv3-256 **conv3-256** | conv3-256 conv3-256 conv3-256 **conv3-256** |
| maxpool | | | | | |
| conv3-512 conv3-512 | conv3-512 conv3-512 | conv3-512 conv3-512 | conv3-512 conv3-512 **conv1-512** | conv3-512 conv3-512 **conv3-512** | conv3-512 conv3-512 conv3-512 **conv3-512** |
| maxpool | | | | | |
| conv3-512 conv3-512 | conv3-512 conv3-512 | conv3-512 conv3-512 | conv3-512 conv3-512 **conv1-512** | conv3-512 conv3-512 **conv3-512** | conv3-512 conv3-512 conv3-512 **conv3-512** |
| maxpool | | | | | |
| FC-4096 | | | | | |
| FC-4096 | | | | | |
| FC-1000 | | | | | |
| soft-max | | | | | |

Table 2: **Number of parameters** (in millions).

| Network | A,A-LRN | B | C | D | E |
|---|---|---|---|---|---|
| Number of parameters | 133 | 133 | 134 | 138 | 144 |

| Error % | Train Top-1 | Val Top-1 | Val Top-5 |
|---|---|---|---|
| Our replication of (Krizhevsky et al., 2012), 1 convnet | 35.1 | 40.5 | 18.1 |
| Removed layers 3,4 | 41.8 | 45.4 | 22.1 |
| Removed layer 7 | 27.4 | 40.0 | 18.4 |
| Removed layers 6,7 | 27.4 | 44.8 | 22.4 |
| Removed layer 3,4,6,7 | 71.1 | 71.3 | 50.1 |
| Adjust layers 6,7: 2048 units | 40.3 | 41.7 | 18.8 |
| Adjust layers 6,7: 8192 units | 26.8 | 40.0 | 18.1 |

| | | | |
|---|---|---|---|
| Our Model (as per Fig. 3) | 33.1 | 38.4 | 16.5 |
| Adjust layers 6,7: 2048 units | 38.2 | 40.2 | 17.6 |
| Adjust layers 6,7: 8192 units | 22.0 | 38.8 | 17.0 |
| Adjust layers 3,4,5: 512,1024,512 maps | 18.8 | **37.5** | **16.0** |
| Adjust layers 6,7: 8192 units and Layers 3,4,5: 512,1024,512 maps | **10.0** | 38.3 | 16.9 |

| | | | | | | |
|---|---|---|---|---|---|---|
| (p) CNN M | – | (C) | f | s | 4K | **79.89** |
| (x) CNN M 2048 | – | (C) | f | s | 2K | **80.10** |
| (y) CNN M 1024 | – | (C) | f | s | 1K | **79.91** |
| (z) CNN M 128 | – | (C) | f | s | 128 | **78.60** |

# We learned tips/tricks for making ConvNets work well in practice

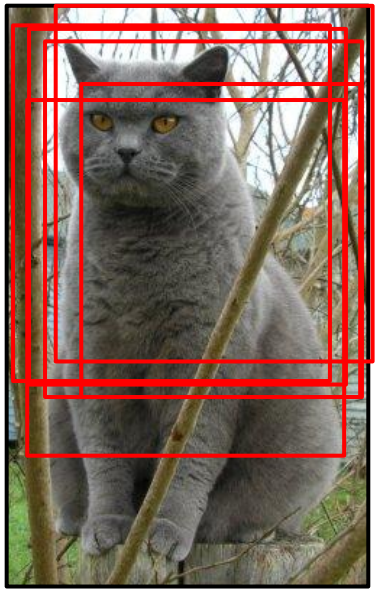

# And explored their practical bottlenecks...

Moving parts lol

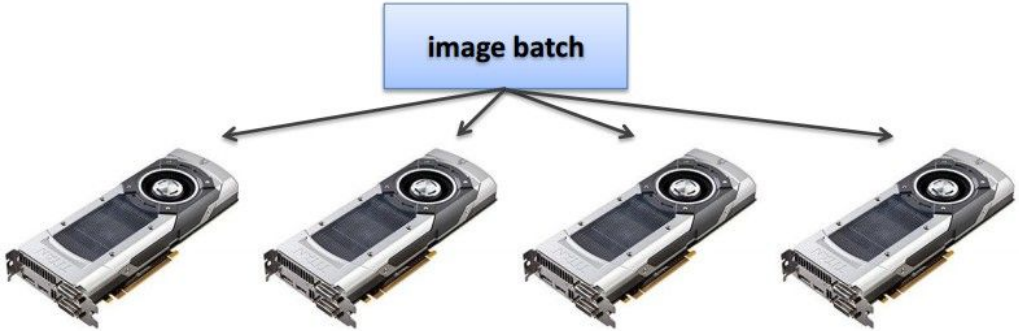



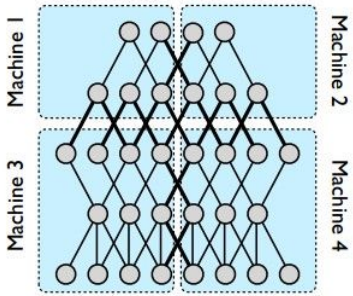


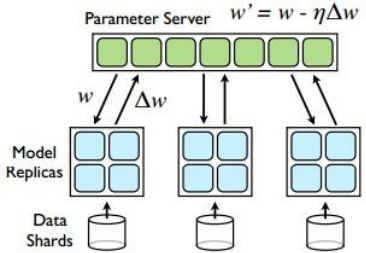

# And we bravely ventured beyond Image Classification...

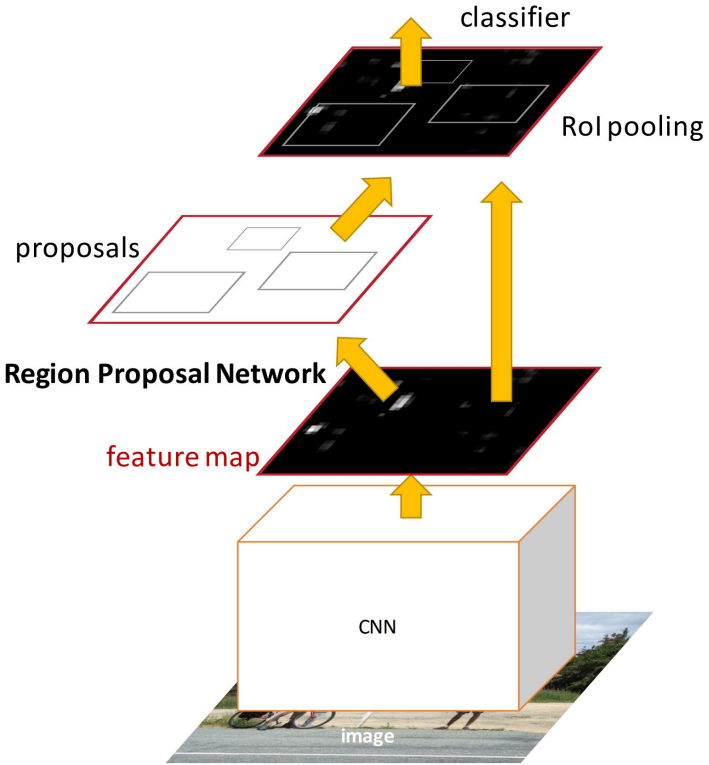


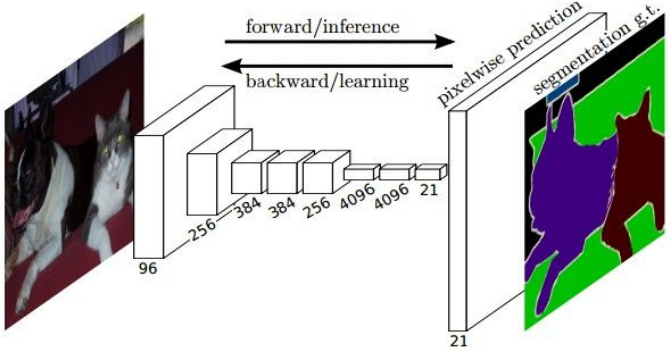


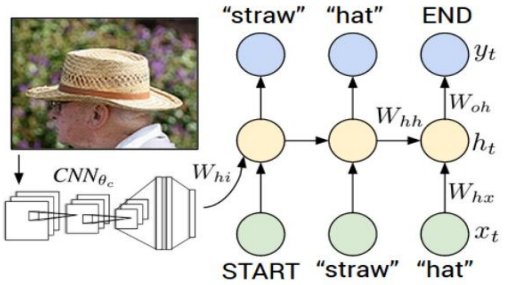


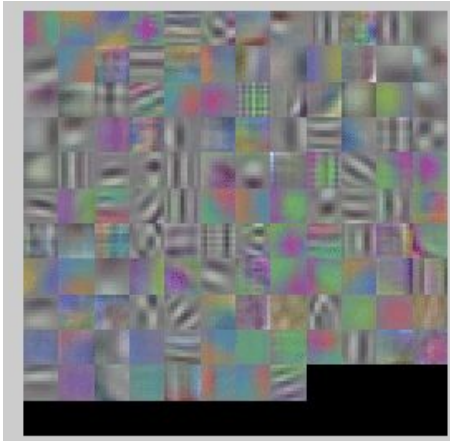

# And developed an understanding of cutting-edge research

We saw **2015 & 2016** citations...
e.g.

**Deep Residual Learning for Image Recognition**
*[He et al., CVPR 2016] (MSR)*

**Inception-v4, Inception-ResNet and the Impact of Residual Connections on Learning**
*[Szegedy et al., 2016] (Google)*

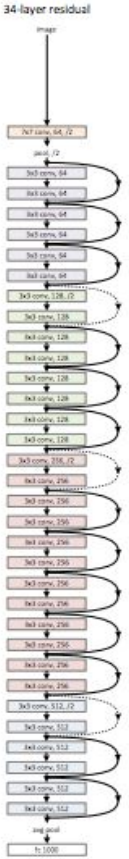

# You are now ready.

# You are now ready.

# You are now ready.

END

## Course Instructors


Fei-Fei Li


Andrej Karpathy


Justin Johnson

## Teaching Assistants


Serena Yeung


Subhasis Das


Song Han


Albert Haque


Bharath Ramsundar


Hieu Pham


Irwan Bello


Namrata Anand


Lane McIntosh


Catherine Dong


Kyle Griswold