CS 4803 / 7643: Deep Learning

Topics:

- Visualizing CNNs

Zsolt Kira
Georgia Tech

Plan for Today

- Visualizing CNNs
  - Visualizing filters
  - Last layer embeddings
  - Visualizing activations
  - Maximally activating patches
  - Occlusion maps
  - Salient or “important” pixels
    - Gradient-based visualizations
  - How to evaluate visualizations?
  - Creating “prototypical” images for a class
  - Creating adversarial images
  - Deep dream
  - Feature inversion
  - Style Transfer

What’s going on inside ConvNets?

Input Image:
3 x 224 x 224

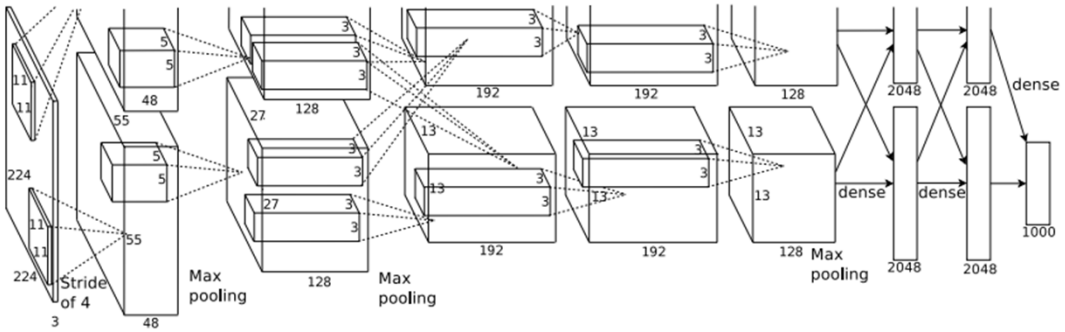

Class Scores:
1000 numbers

What are the intermediate features looking for?

Krizhevsky et al, “ImageNet Classification with Deep Convolutional Neural Networks”, NIPS 2012.
Figure reproduced with permission.

First Layer: Visualize Filters

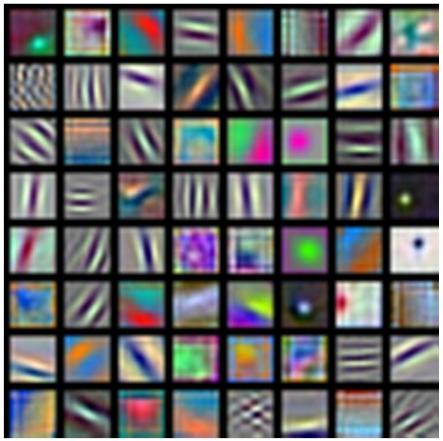

AlexNet:
64 x 3 x 11 x 11

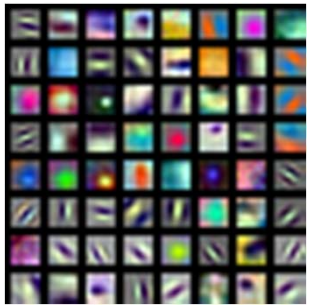

ResNet-18:
64 x 3 x 7 x 7

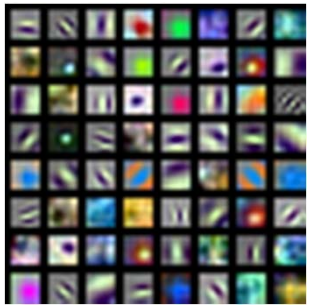

ResNet-101:
64 x 3 x 7 x 7

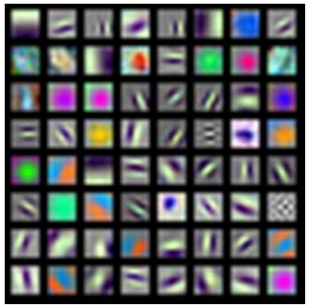

DenseNet-121:
64 x 3 x 7 x 7

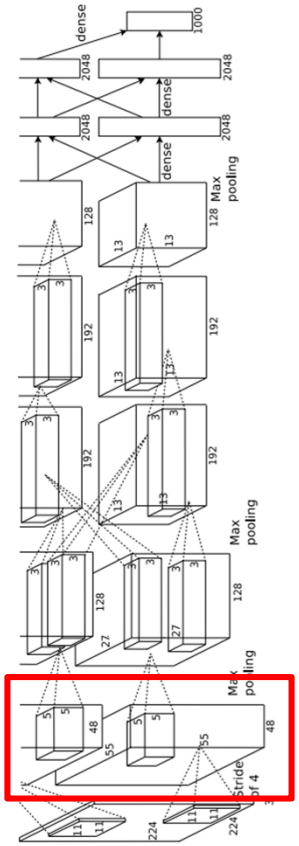

Krizhevsky, “One weird trick for parallelizing convolutional neural networks”, arXiv 2014
He et al, “Deep Residual Learning for Image Recognition”, CVPR 2016
Huang et al, “Densely Connected Convolutional Networks”, CVPR 2017

Visualize the filters/kernels (raw weights)

We can visualize filters at higher layers, but not that interesting

(these are taken from ConvNetJS CIFAR-10 demo)

layer 1 weights

16 x 3 x 7 x 7

layer 2 weights

20 x 16 x 7 x 7

layer 3 weights

20 x 20 x 7 x 7

Slide Credit: Fei-Fei Li, Justin Johnson, Serena Yeung, CS 231n

Last Layer

FC7 layer

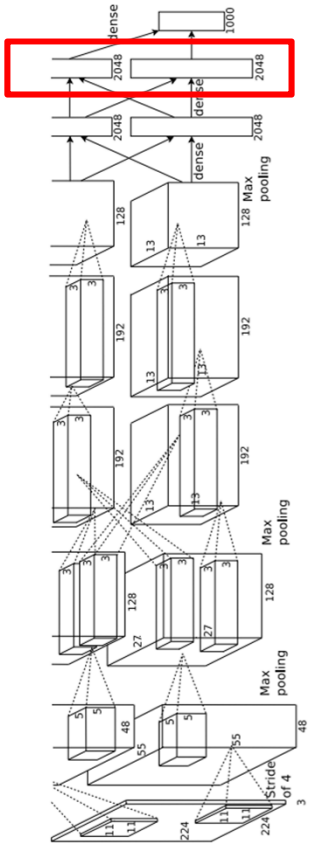

4096-dimensional feature vector for an image (layer immediately before the classifier)

Run the network on many images, collect the feature vectors

Last Layer: Nearest Neighbors

4096-dim vector

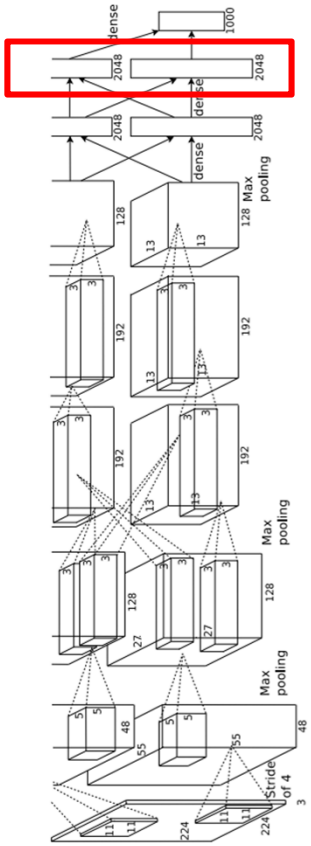

Test image

L2 Nearest neighbors in <u>feature</u> space

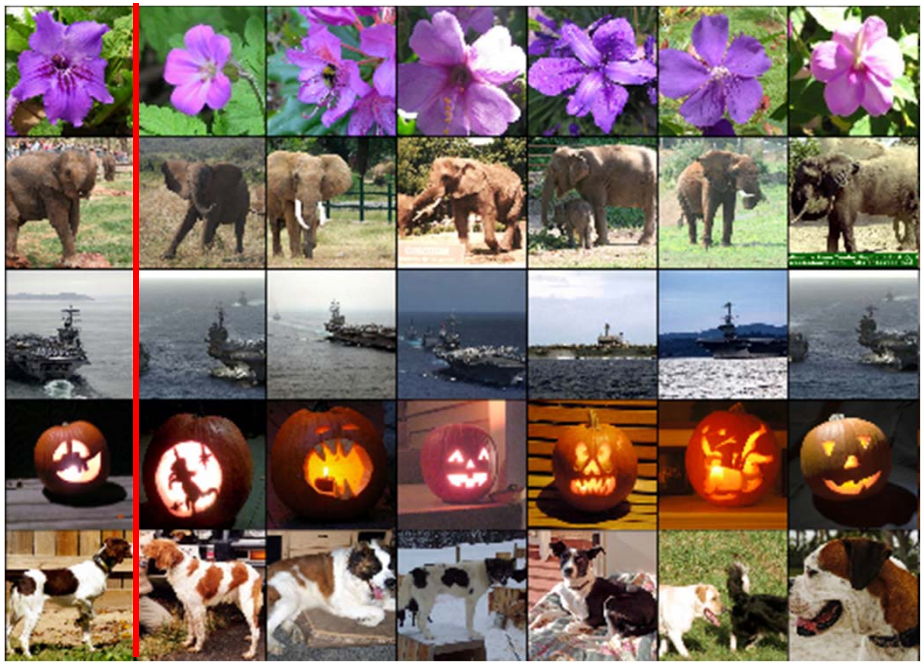

Recall: Nearest neighbors in <u>pixel</u> space

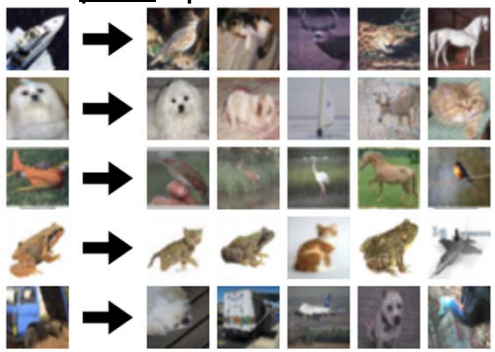

Krizhevsky et al, "ImageNet Classification with Deep Convolutional Neural Networks", NIPS 2012.
Figures reproduced with permission.

Last Layer: Dimensionality Reduction

Visualize the “space” of FC7 feature vectors by reducing dimensionality of vectors from 4096 to 2 dimensions

Simple algorithm: Principal Component Analysis (PCA)

More complex: **t-SNE**

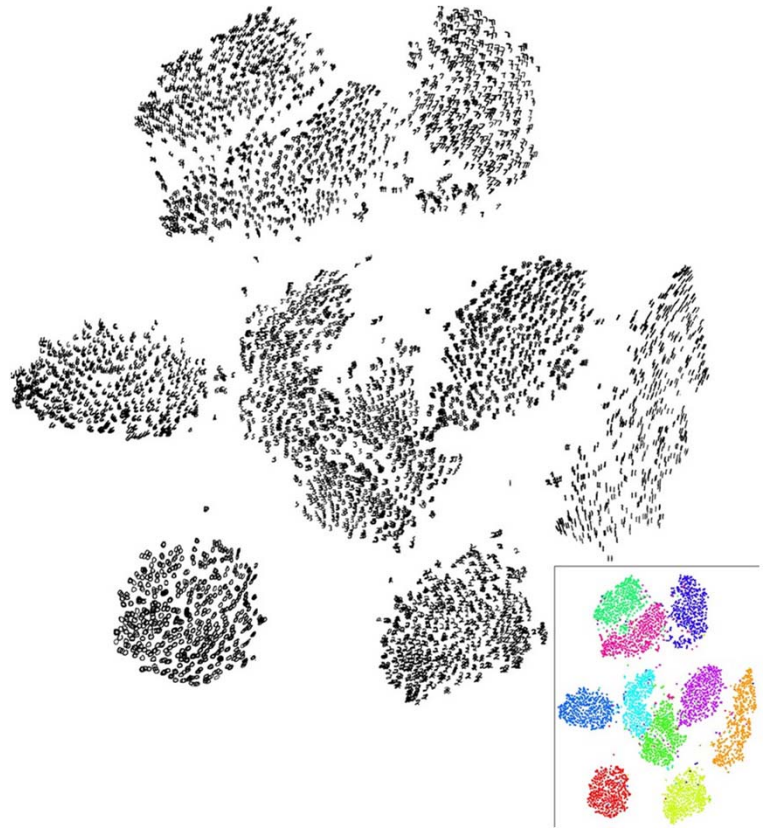

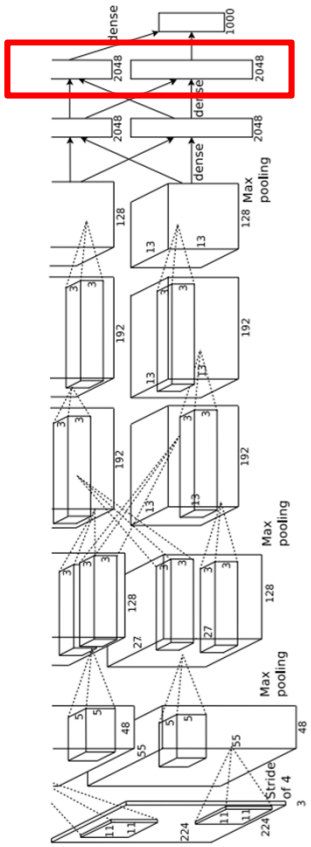

Van der Maaten and Hinton, “Visualizing Data using t-SNE”, JMLR 2008
Figure copyright Laurens van der Maaten and Geoff Hinton, 2008. Reproduced with permission.

Last Layer: Dimensionality Reduction

Van der Maaten and Hinton, “Visualizing Data using t-SNE”, JMLR 2008
Krizhevsky et al, “ImageNet Classification with Deep Convolutional Neural Networks”, NIPS 2012.
Figure reproduced with permission.

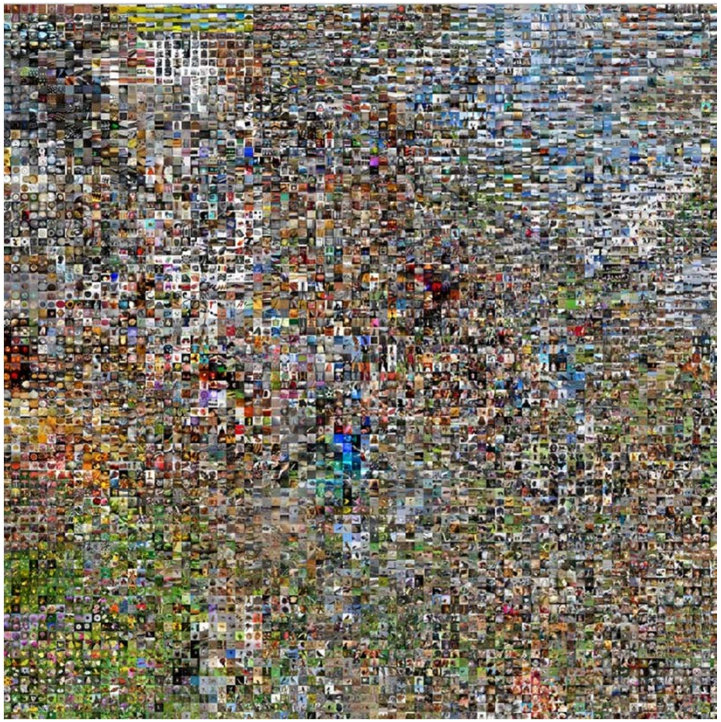

See high-resolution versions at
http://cs.stanford.edu/people/karpathy/cnnembed/

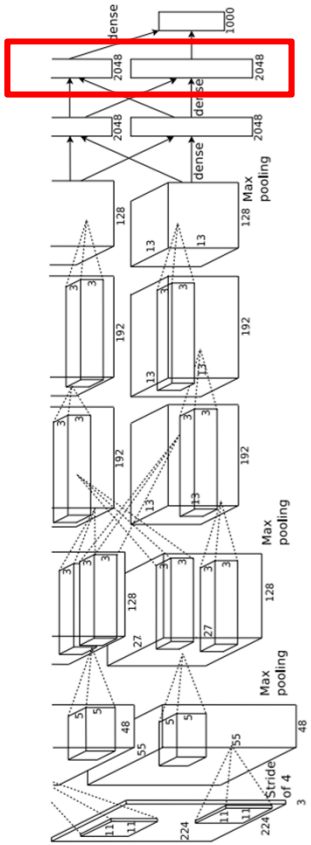

Visualizing Activations

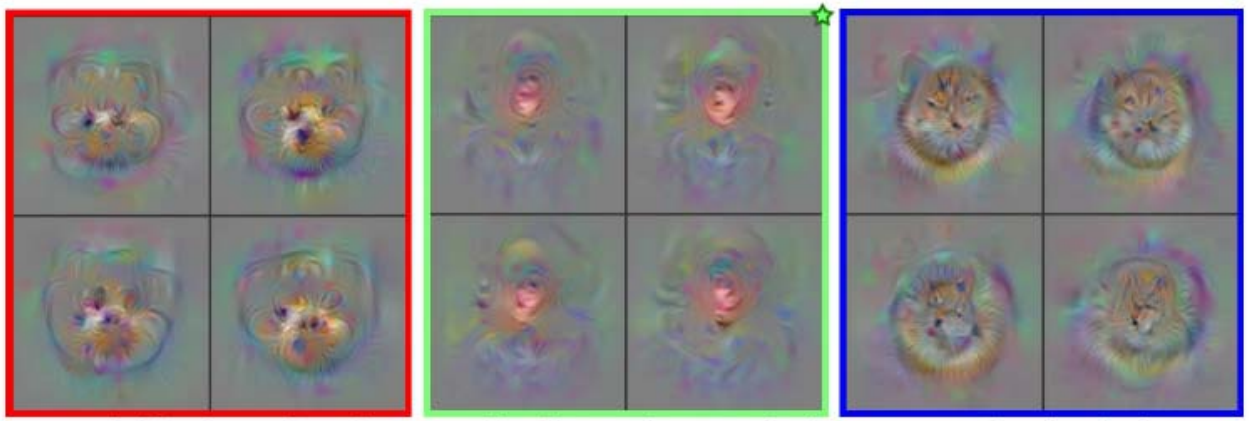

$conv5_2$ (dog face + flower) $conv5_{151}$ (human face + cat face) $conv5_{111}$ (cat face)

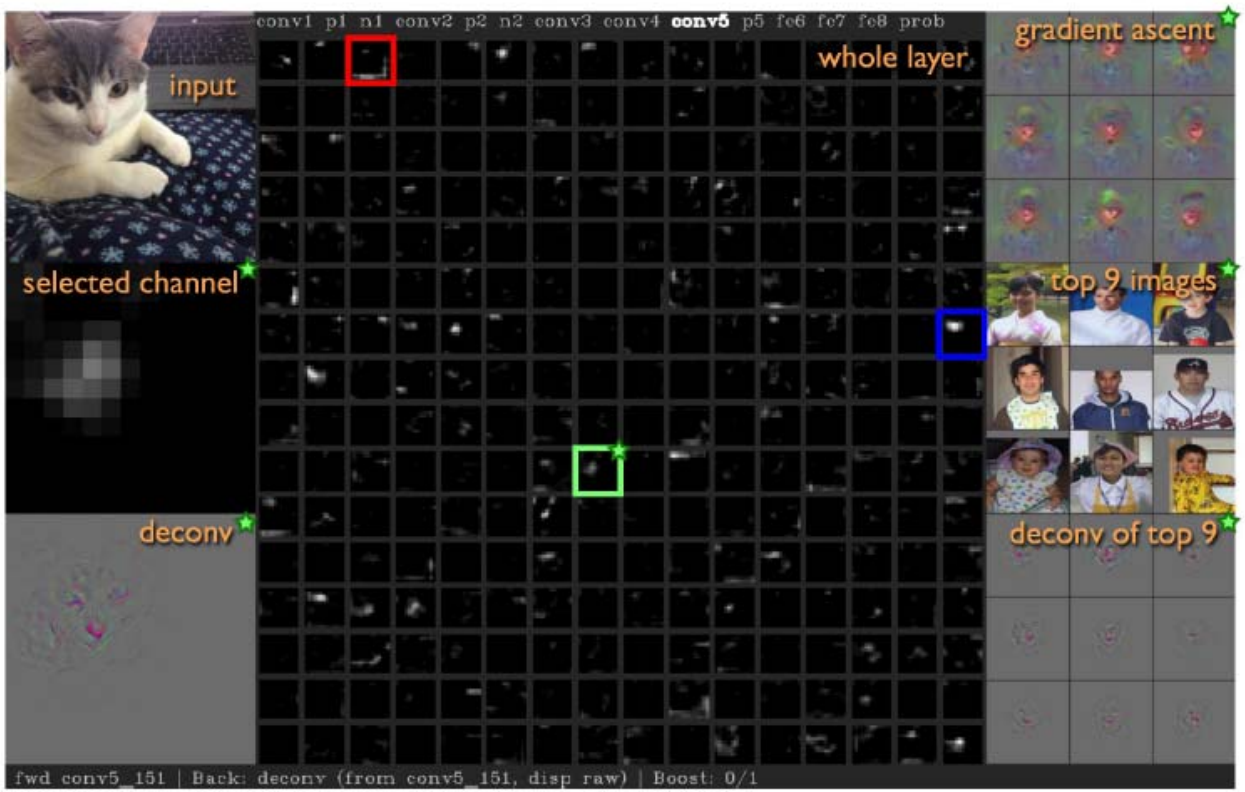

Yosinski et al, "Understanding Neural Networks Through Deep Visualization", ICML DL Workshop 2014.
Figure copyright Jason Yosinski, 2014. Reproduced with permission.

Visualizing Activations

conv5 feature map is 128x13x13; visualize as 128 13x13 grayscale images

https://youtu.be/AgkfIQ4IGaM?t=92

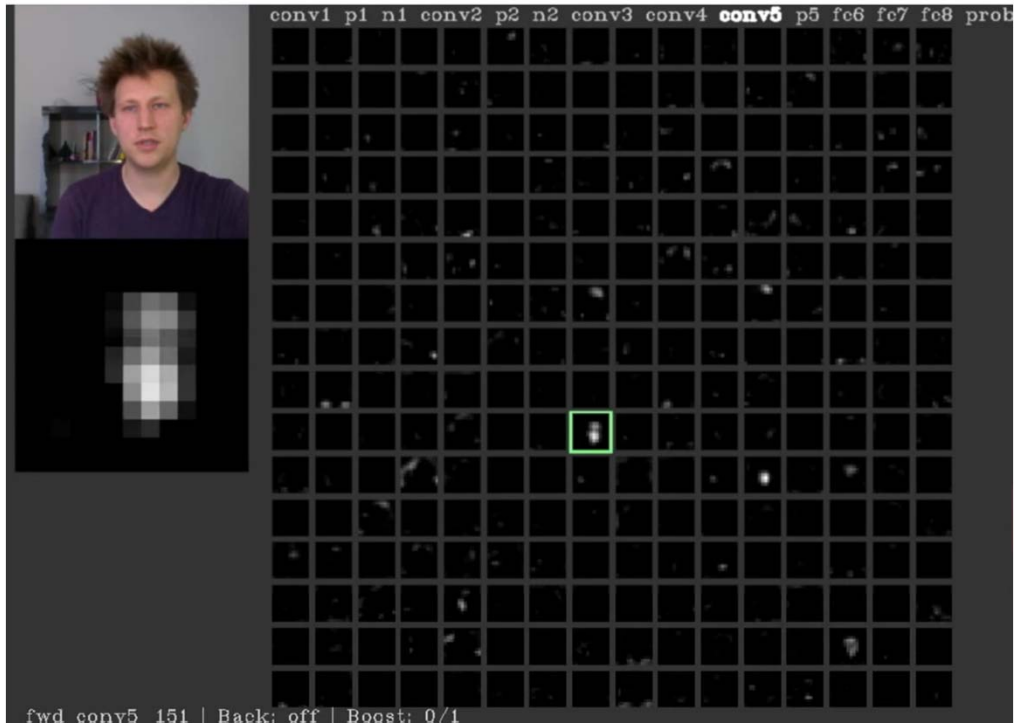

Yosinski et al, "Understanding Neural Networks Through Deep Visualization", ICML DL Workshop 2014.
Figure copyright Jason Yosinski, 2014. Reproduced with permission.

Maximally Activating Patches

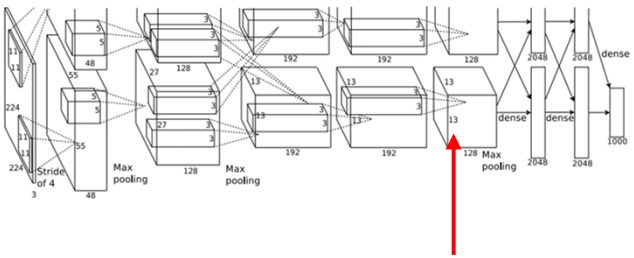

Pick a layer and a channel; e.g. conv5 is 128 x 13 x 13, pick channel 17/128

Run many images through the network, record values of chosen channel

Visualize image patches that correspond to maximal activations

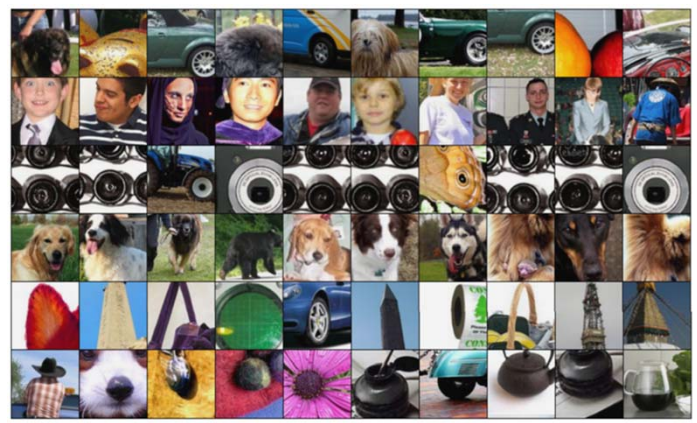

Springenberg et al, “Striving for Simplicity: The All Convolutional Net”, ICLR Workshop 2015
Figure copyright Jost Tobias Springenberg, Alexey Dosovitskiy, Thomas Brox, Martin Riedmiller, 2015; reproduced with permission.

Plan for Today

- Visualizing CNNs
  - Visualizing filters
  - Last layer embeddings
  - Visualizing activations
  - Maximally activating patches
  - **Occlusion maps**
  - Salient or “important” pixels
    - Gradient-based visualizations
  - How to evaluate visualizations?
  - Creating “prototypical” images for a class
  - Creating adversarial images
  - Deep dream
  - Feature inversion

Visual Explanations

Where does an intelligent system "look" to make its predictions?

Which pixels matter: Occlusion Maps

Mask part of the image before feeding to CNN,
check how much predicted probabilities change

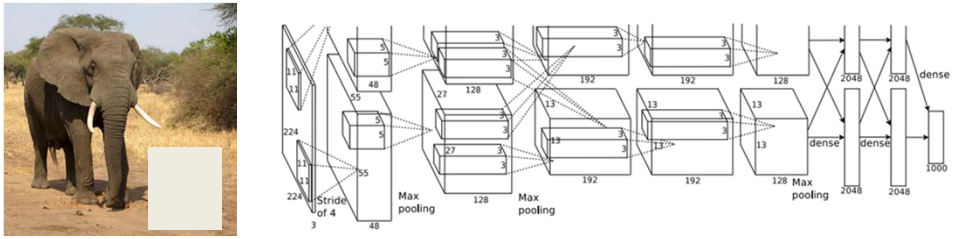

P(elephant) = 0.95

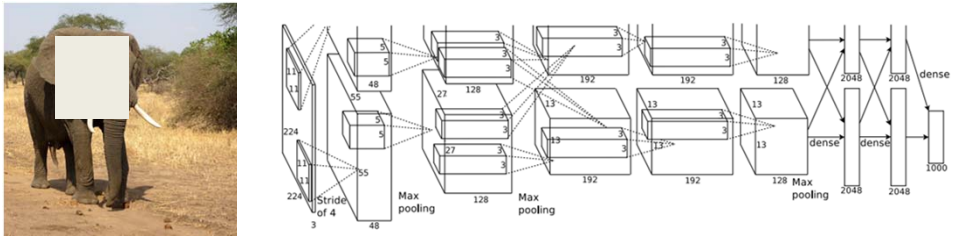

P(elephant) = 0.75

Zeiler and Fergus, “Visualizing and Understanding Convolutional Networks”, ECCV 2014

Which pixels matter: Occlusion Maps

Mask part of the image before feeding to CNN, check how much predicted probabilities change

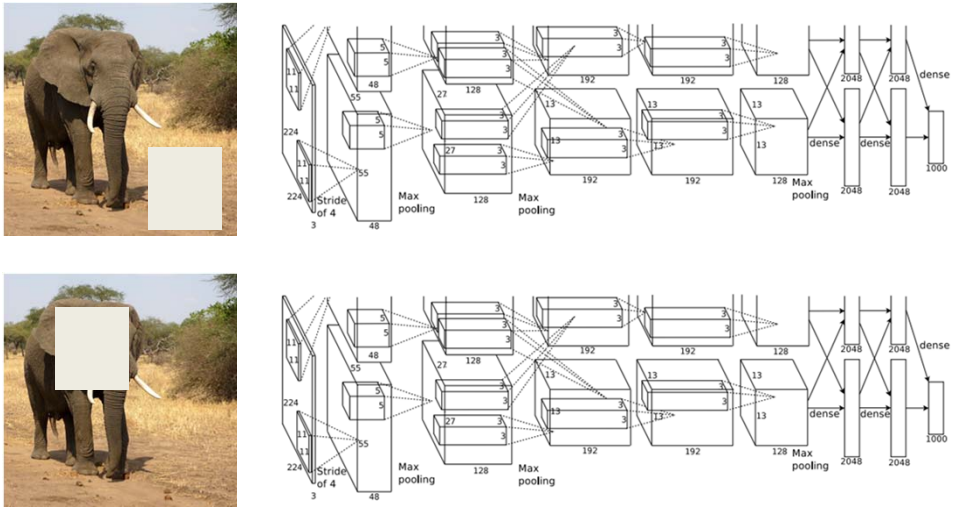

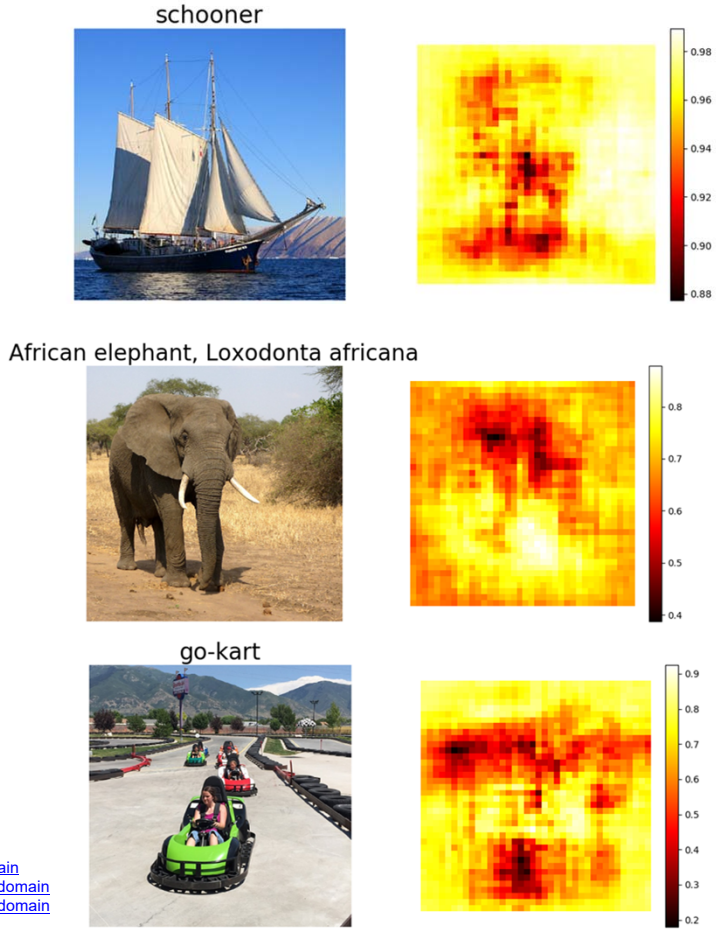

Zeiler and Fergus, "Visualizing and Understanding Convolutional Networks", ECCV 2014

What if our model was linear?

$$\langle \mathbf{w}_c \, , \, \mathbf{x} \rangle + b = S_c(\mathbf{x})$$

What if our model was linear?

$$\left\langle \begin{pmatrix} \mathbf{100} \\ 0.1 \\ -0.1 \\ \mathbf{510} \\ \mathbf{-200} \end{pmatrix}, \begin{pmatrix} 1 \\ 0.9 \\ -0.2 \\ 0.5 \\ -0.9 \end{pmatrix} \right\rangle + b = S_c(\mathbf{x})$$

But it's not ☹

$$\langle \mathbf{w}_c \, , \, \mathbf{x} \rangle + b = S_c(\mathbf{x})$$

Can we make it linear?

$$f(\mathbf{x}) = S_c(\mathbf{x})$$

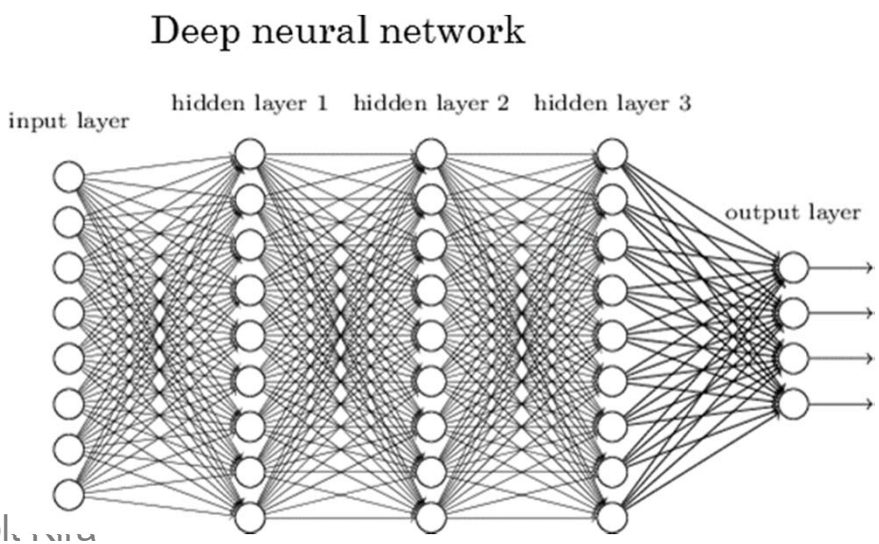

Taylor Series

$$f(x) \approx f(x_0) + f'(x_0)(x - x_0)$$

Feature Importance in Deep Models

$$\mathbf{w}_c = \left.\frac{\partial S_c}{\partial \mathbf{x}}\right|_{\mathbf{x}_0}$$

$$\langle \mathbf{w}_c \, , \, \mathbf{x} \rangle + b \approx S_c(\mathbf{x})$$

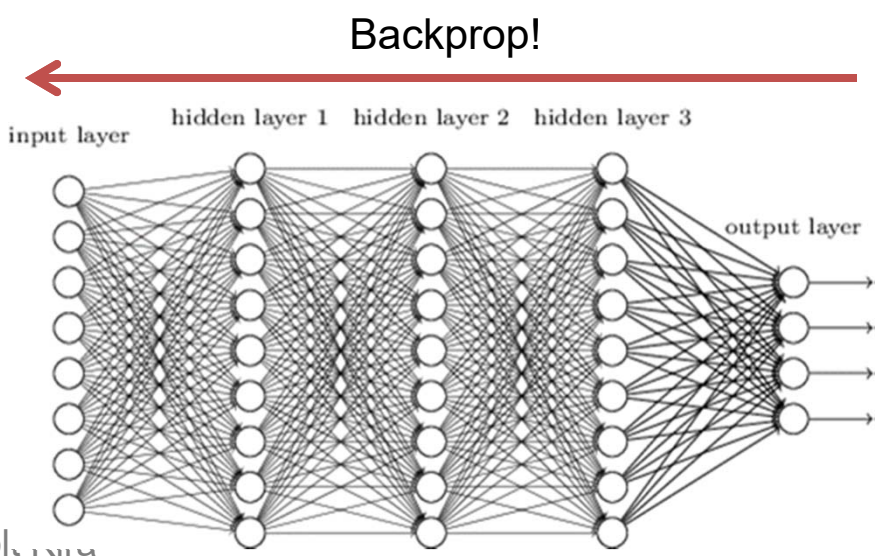

Which pixels matter: Saliency via Backprop

Forward pass: Compute probabilities

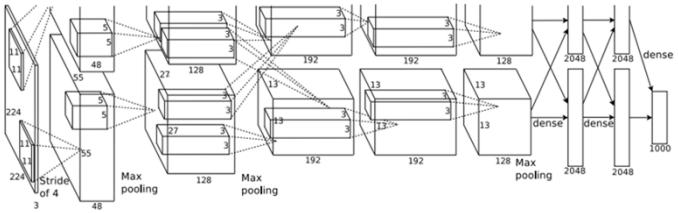

Dog

Simonyan, Vedaldi, and Zisserman, “Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps”, ICLR Workshop 2014.

Which pixels matter: Saliency via Backprop

Forward pass: Compute probabilities

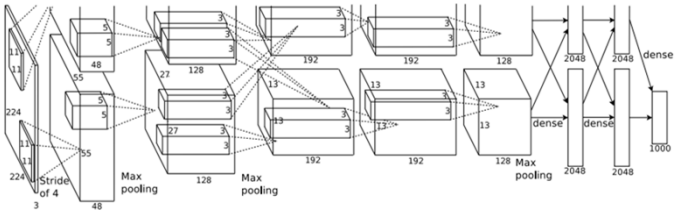

Dog

Compute gradient of (unnormalized) class score with respect to image pixels, take absolute value and max over RGB channels

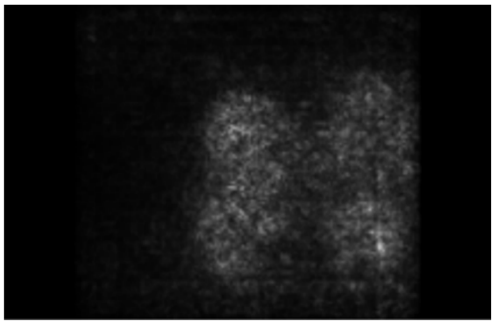

Simonyan, Vedaldi, and Zisserman, "Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps", ICLR Workshop 2014.

Why unnormalized class scores?

Saliency Maps

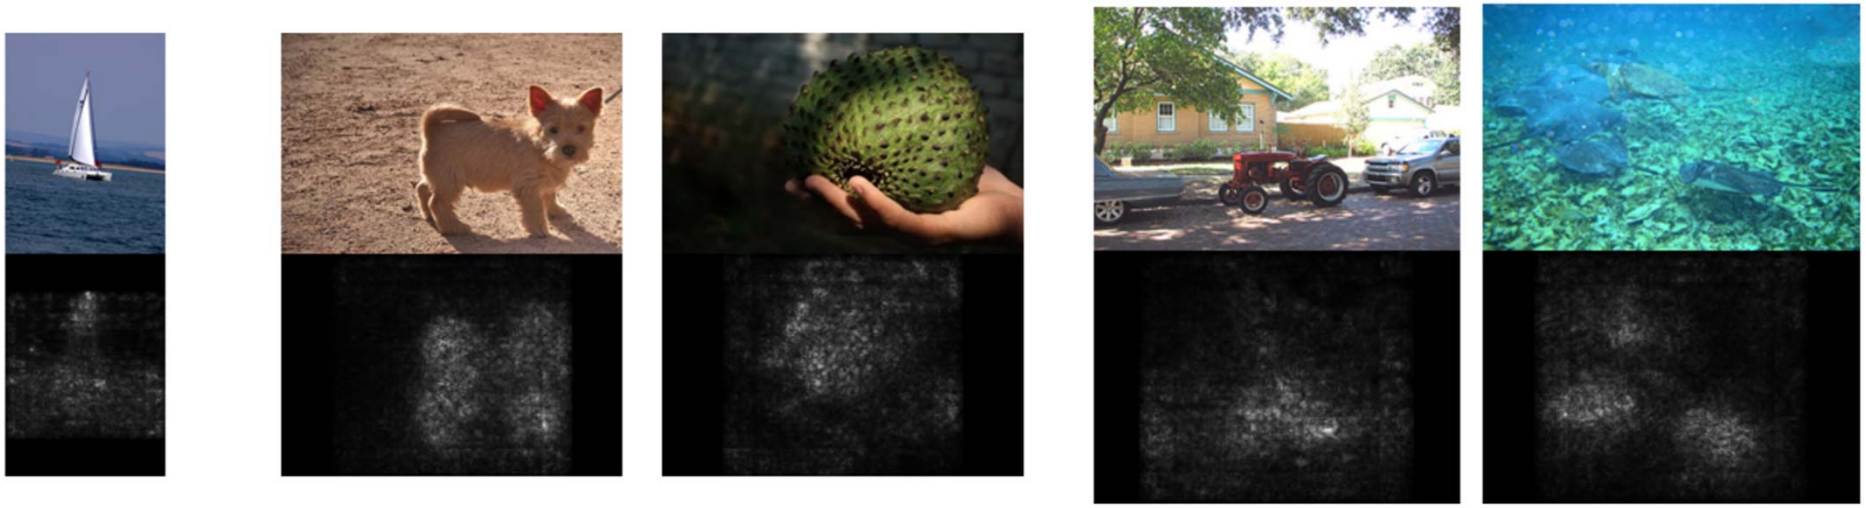

Simonyan, Vedaldi, and Zisserman, “Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps”, ICLR Workshop 2014.

Saliency Maps: Segmentation without supervision

Use GrabCut on saliency map

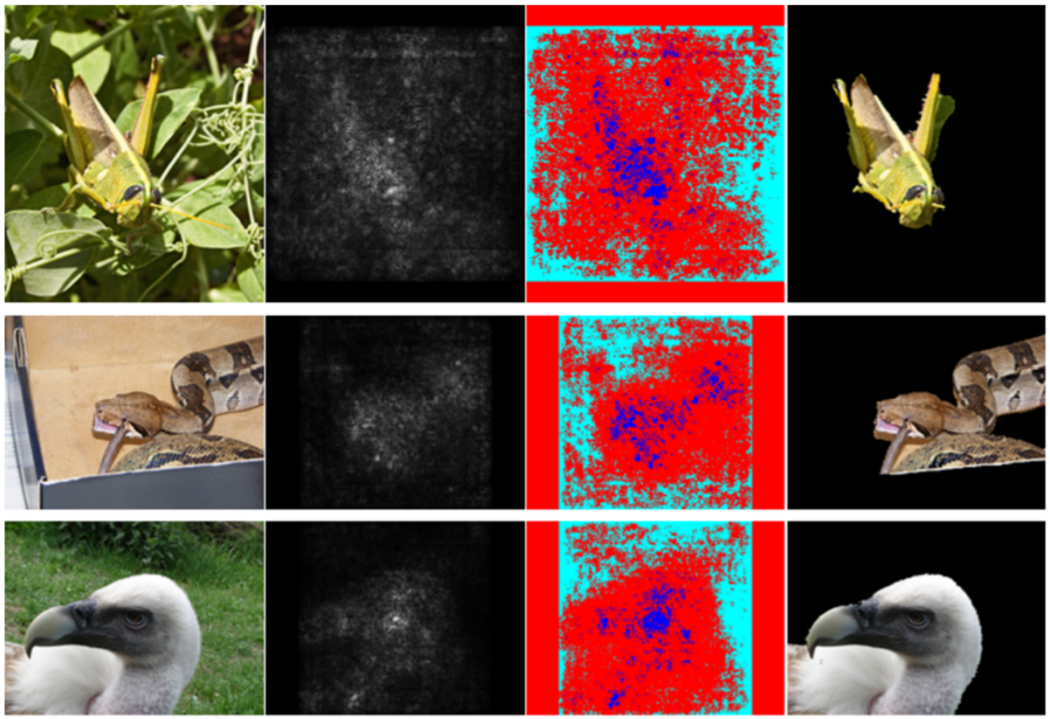

Simonyan, Vedaldi, and Zisserman, "Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps", ICLR Workshop 2014.
Figures copyright Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman, 2014; reproduced with permission.
Rother et al, "Grabcut: Interactive foreground extraction using iterated graph cuts", ACM TOG 2004

Aside: DeconvNet

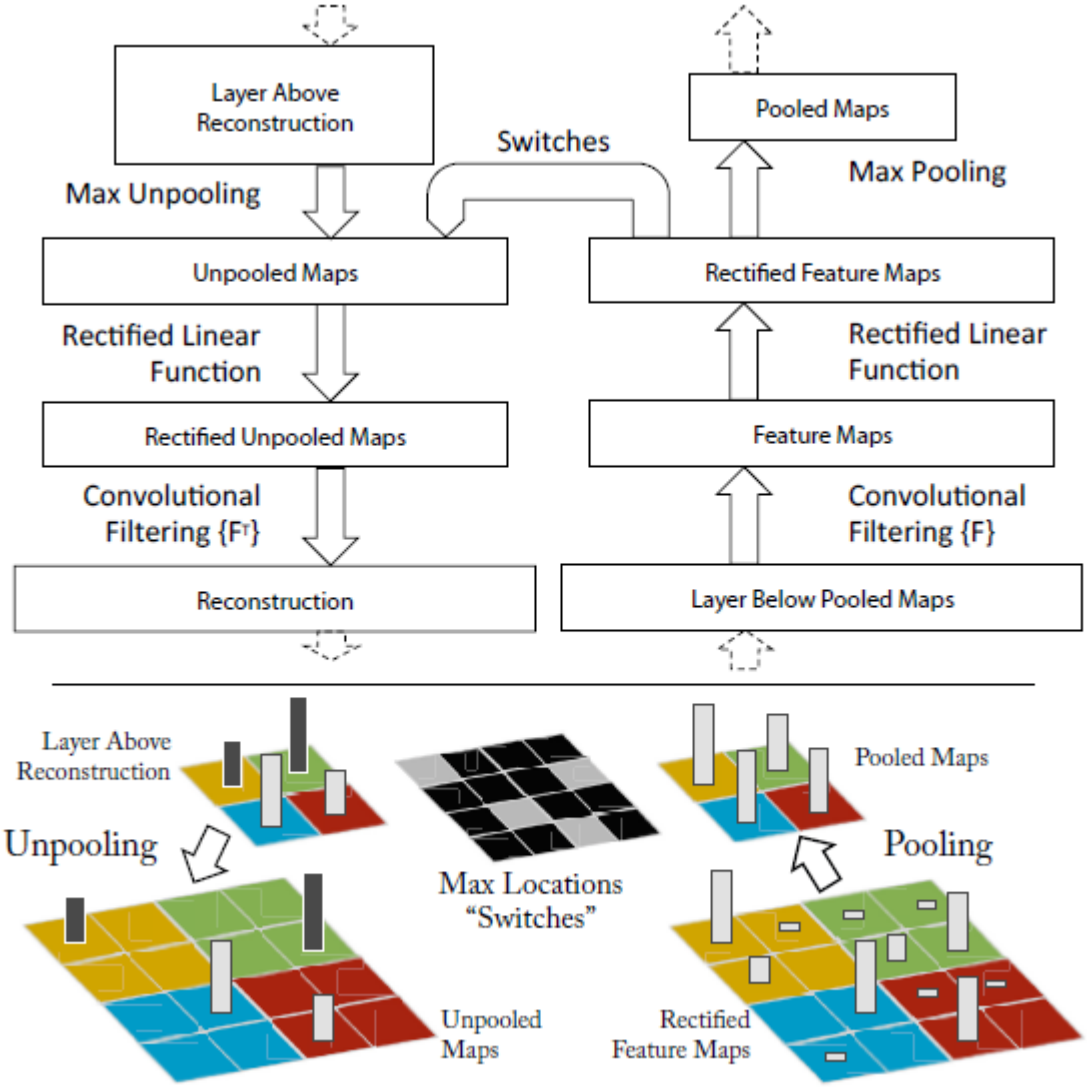

Zeiler and Fergus, "Visualizing and Understanding Convolutional Networks", ECCV 2014

Gradient-based Visualizations

Identical for all layers except ReLU

Remember ReLUs?

$$h^{l+1} = \mathrm{ReLU}(h^l) = \max\{0, h^l\}$$

$$\frac{\partial h^{l+1}}{\partial h^l} = \begin{cases} 0 & \text{if } h^l < 0 \\ 1 & \text{if } h^l > 0 \end{cases} = [\![h^l > 0]\!]$$

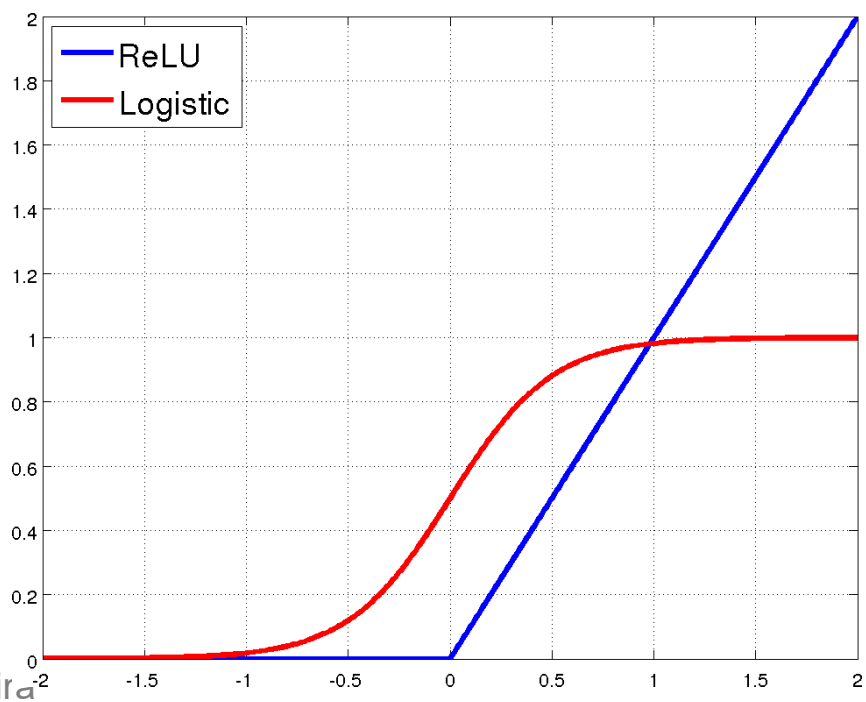

$$h^{l+1} = \max\{0, h^l\}$$

Forward pass

h^l

1	-1	5
2	-5	-7
-3	2	4

$\rightarrow$

1	0	5
2	0	0
0	2	4

h^{l+1}

$$\frac{\partial L}{\partial h^l} = [\![h^l > 0]\!] \frac{\partial L}{\partial h^{l+1}}$$

Backward pass: backpropagation

-2	0	-1
6	0	0
0	-1	3

$\leftarrow$

-2	3	-1
6	-3	1
2	-1	3

$\frac{\partial L}{\partial h^{l+1}}$

$$\frac{\partial L}{\partial h^l} = [\![h^{l+1} > 0]\!] \frac{\partial L}{\partial h^{l+1}}$$

Backward pass: "deconvnet"

0	3	0
6	0	1
2	0	3

$\leftarrow$

-2	3	-1
6	-3	1
2	-1	3

$\frac{\partial L}{\partial h^{l+1}}$

$$\frac{\partial L}{\partial h^l} = [\![(h^l > 0)\&\&(h^{l+1} > 0)]\!] \frac{\partial L}{\partial h^{l+1}}$$

Backward pass: *guided backpropagation*

0	0	0
6	0	0
0	0	3

$\leftarrow$

-2	3	-1
6	-3	1
2	-1	3

$\frac{\partial L}{\partial h^{l+1}}$

Backprop vs Deconv vs Guided BP

- Guided Backprop tends to be “cleanest”

Backprop

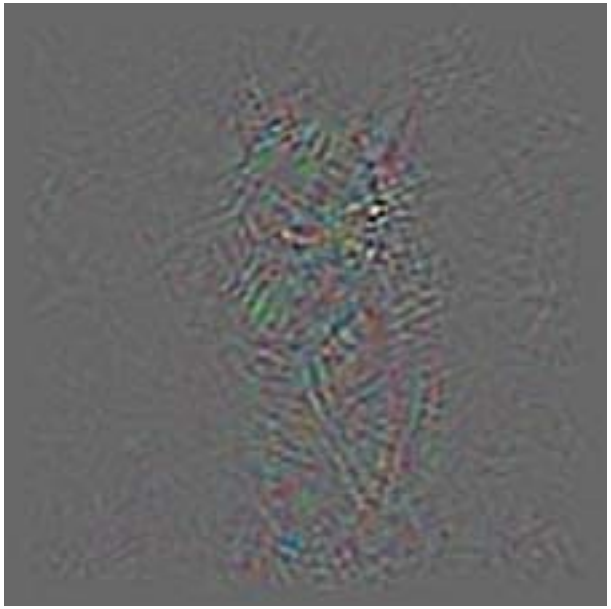

Deconv

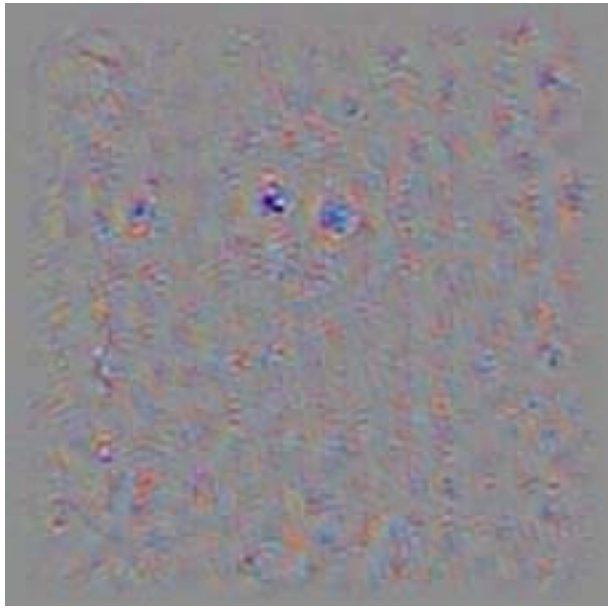

Guided Backprop

Intermediate features via (guided) backprop

Maximally activating patches
(Each row is a different neuron)

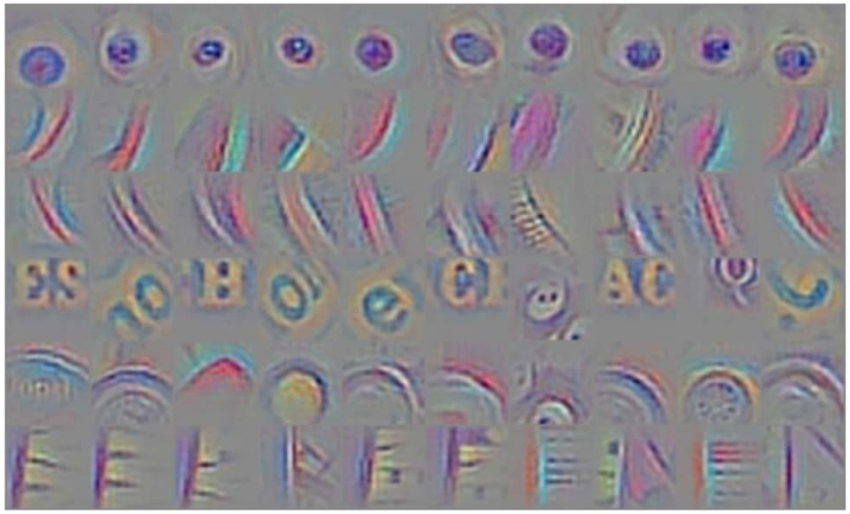

Guided Backprop

Zeiler and Fergus, “Visualizing and Understanding Convolutional Networks”, ECCV 2014
Springenberg et al, “Striving for Simplicity: The All Convolutional Net”, ICLR Workshop 2015

Intermediate features via (guided) backprop

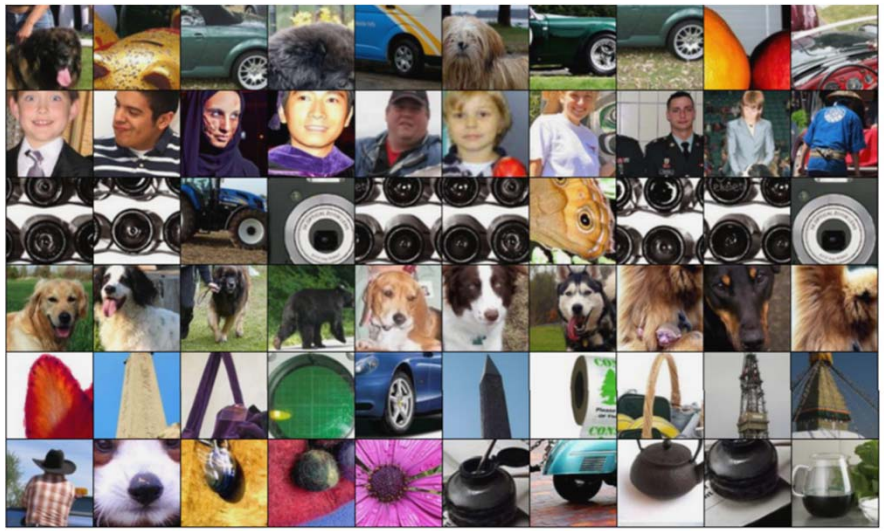

Maximally activating patches
(Each row is a different neuron)

Guided Backprop

Zeiler and Fergus, “Visualizing and Understanding Convolutional Networks”, ECCV 2014
Springenberg et al, “Striving for Simplicity: The All Convolutional Net”, ICLR Workshop 2015

Visualizing Activations

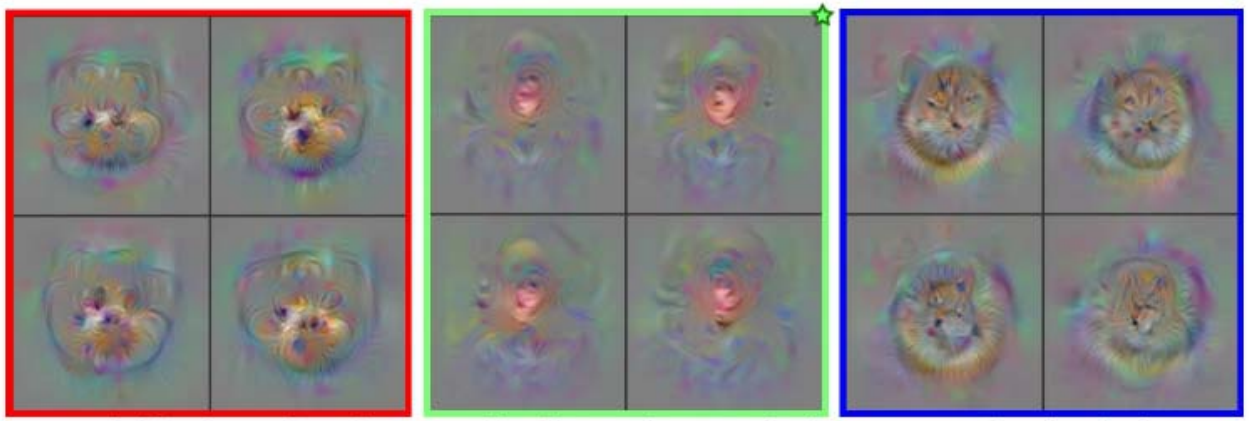

$conv5_2$ (dog face + flower) $conv5_{151}$ (human face + cat face) $conv5_{111}$ (cat face)

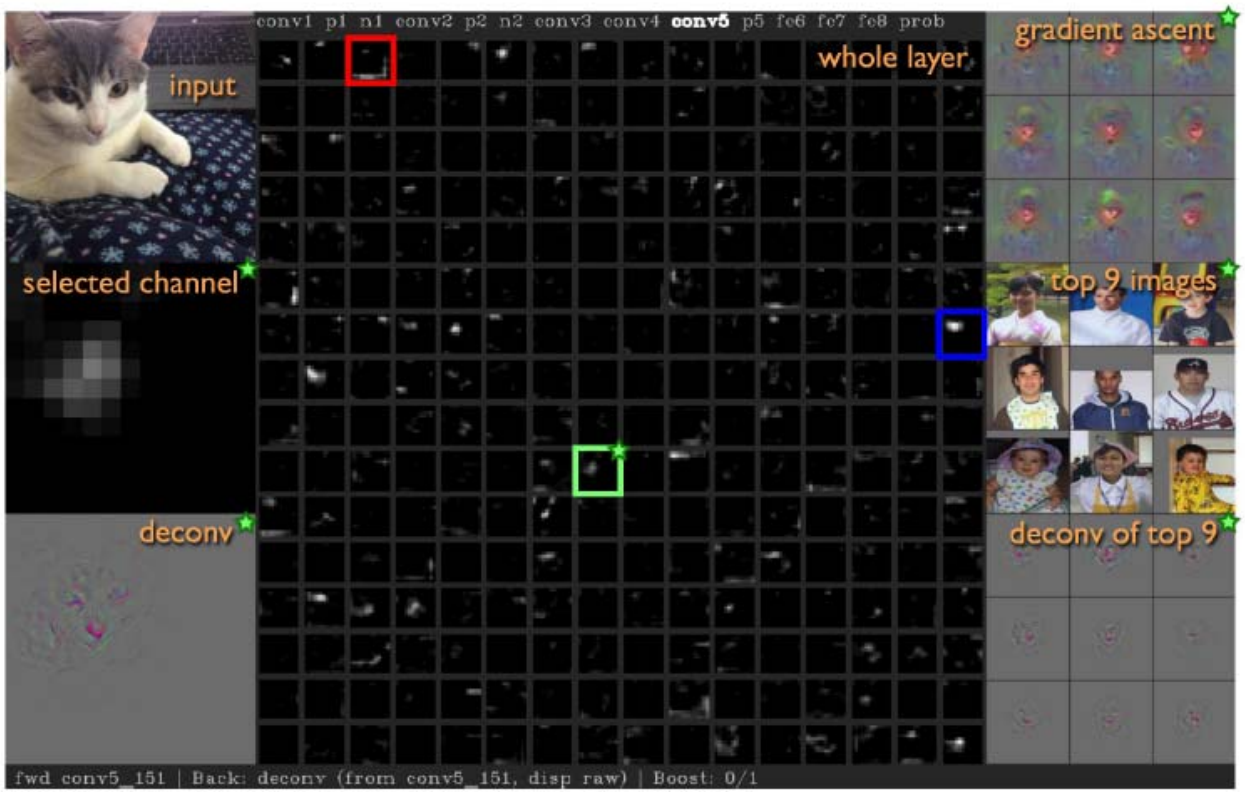

Yosinski et al, "Understanding Neural Networks Through Deep Visualization", ICML DL Workshop 2014.
Figure copyright Jason Yosinski, 2014. Reproduced with permission.

Problem with Guided Backup

- Not very “class-discriminative”

GB for “airliner”

GB for “bus”

Grad-CAM

Visual Explanations from Deep Networks via Gradient-based Localization

[ICCV '17]

Ramprasaath Selvaraju

Michael Cogswell

Abhishek Das

Ramakrishna Vedantam

Devi Parikh

Dhruv Batra

VirginiaTech

facebook research

Georgia Tech

Grad-CAM

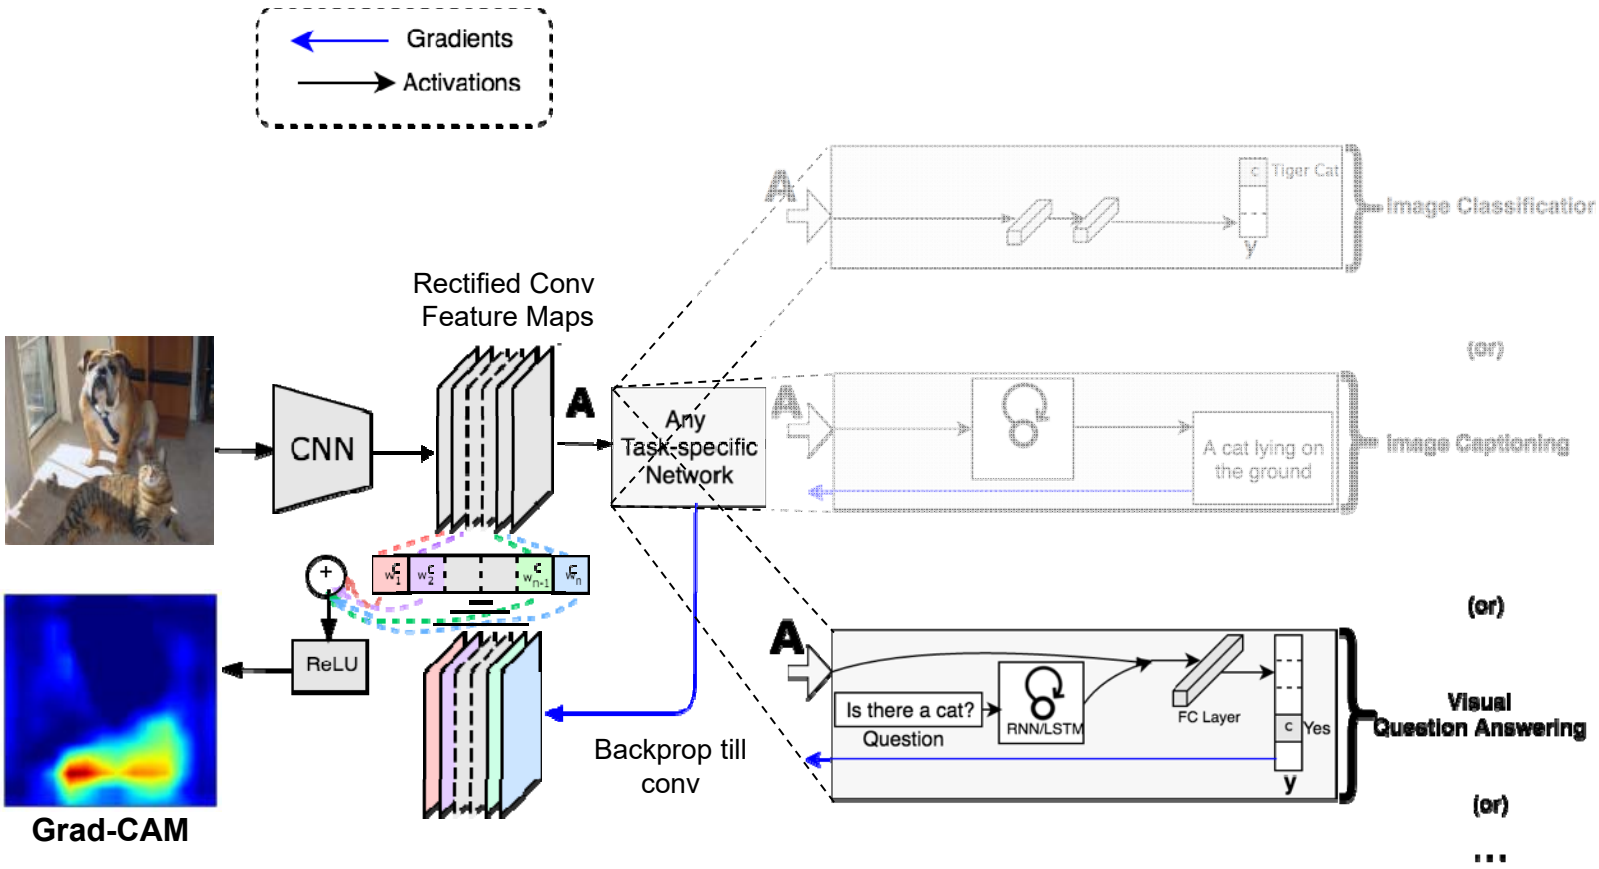
Gradients
Activations
Rectified Conv Feature Maps
CNN
A
Any Task-specific Network
ReLU
Grad-CAM
Backprop till conv
Tiger Cat
Image Classification
(or)
A cat lying on the ground
Image Captioning
(or)
Is there a cat?
Question
RNN/LSTM
FC Layer
Yes
Visual Question Answering
(or)
...

Guided Grad-CAM

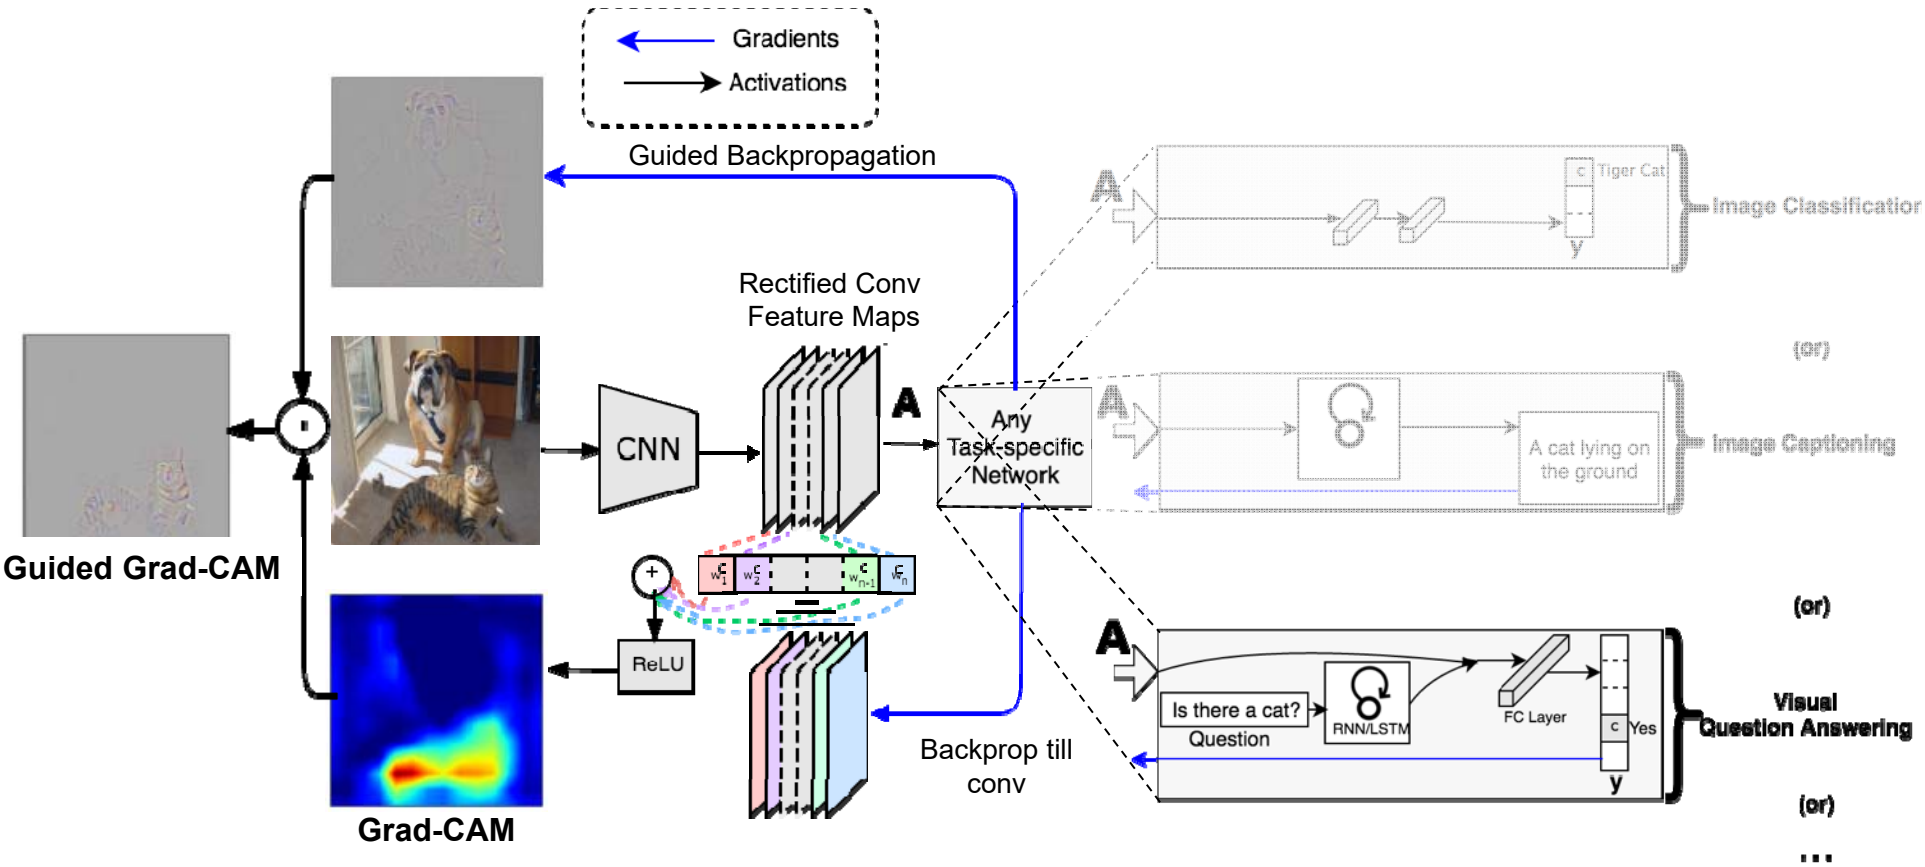

Analyzing Failure Modes with Grad-CAM

Predicted: *Car mirror*

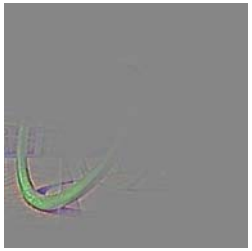

Predicted: *Vine snake*

Ground-truth: *Volcano*

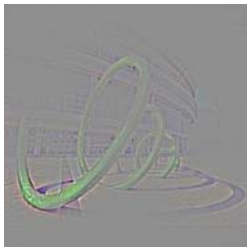

Ground-truth: *coil*

Reasonable predictions are made in many failure cases.

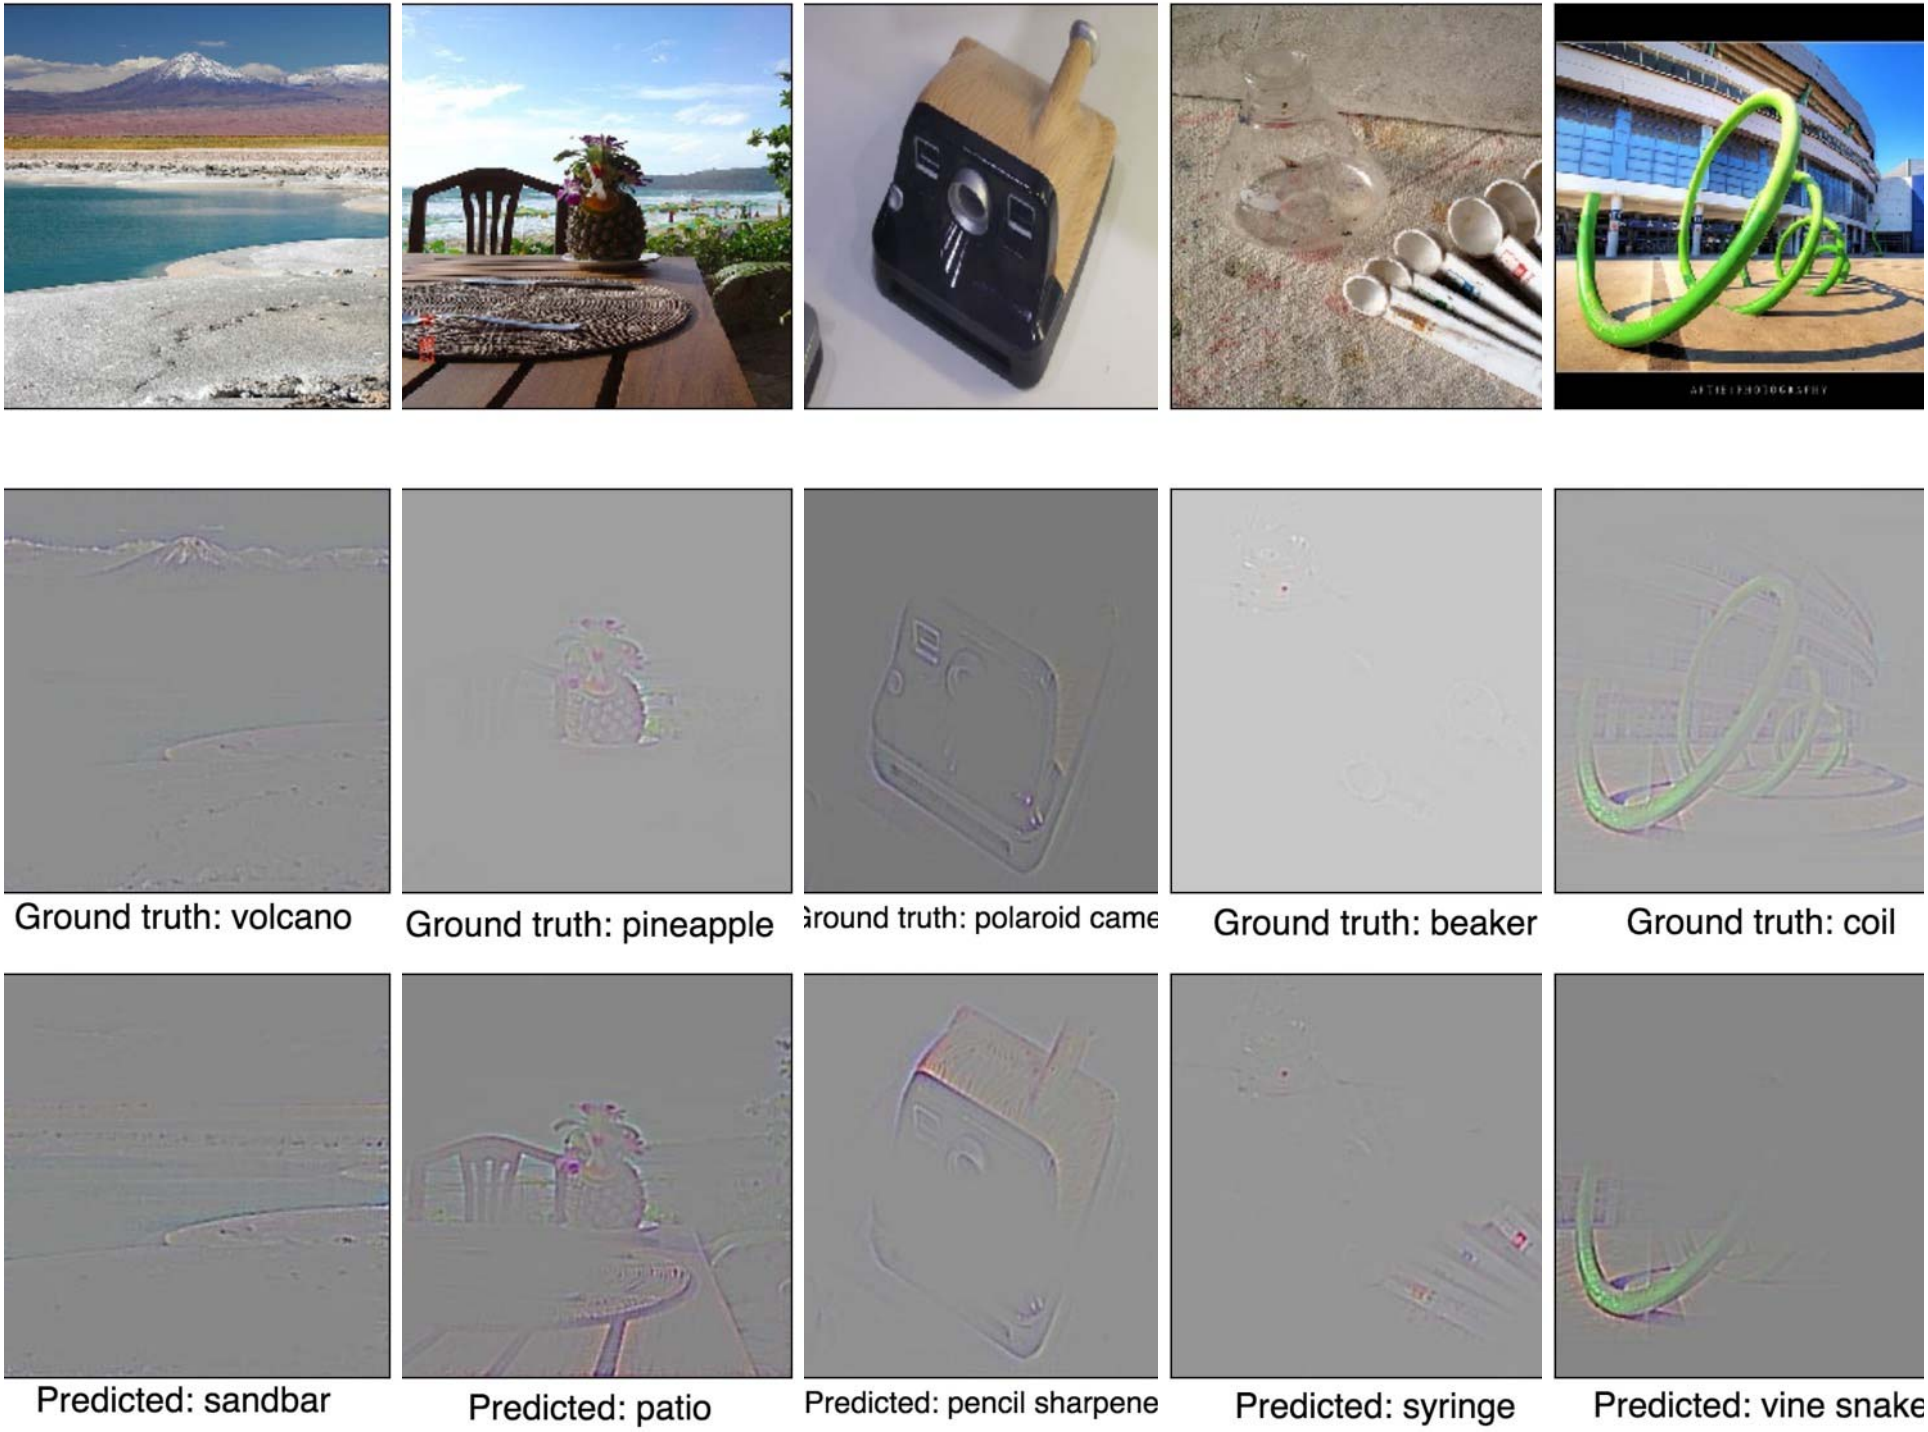
Ground truth: volcano
Ground truth: pineapple
Ground truth: polaroid came
Ground truth: beaker
Ground truth: coil
Predicted: sandbar
Predicted: patio
Predicted: pencil sharpene
Predicted: syringe
Predicted: vine snake

Grad-CAM Visual Explanations for Captioning

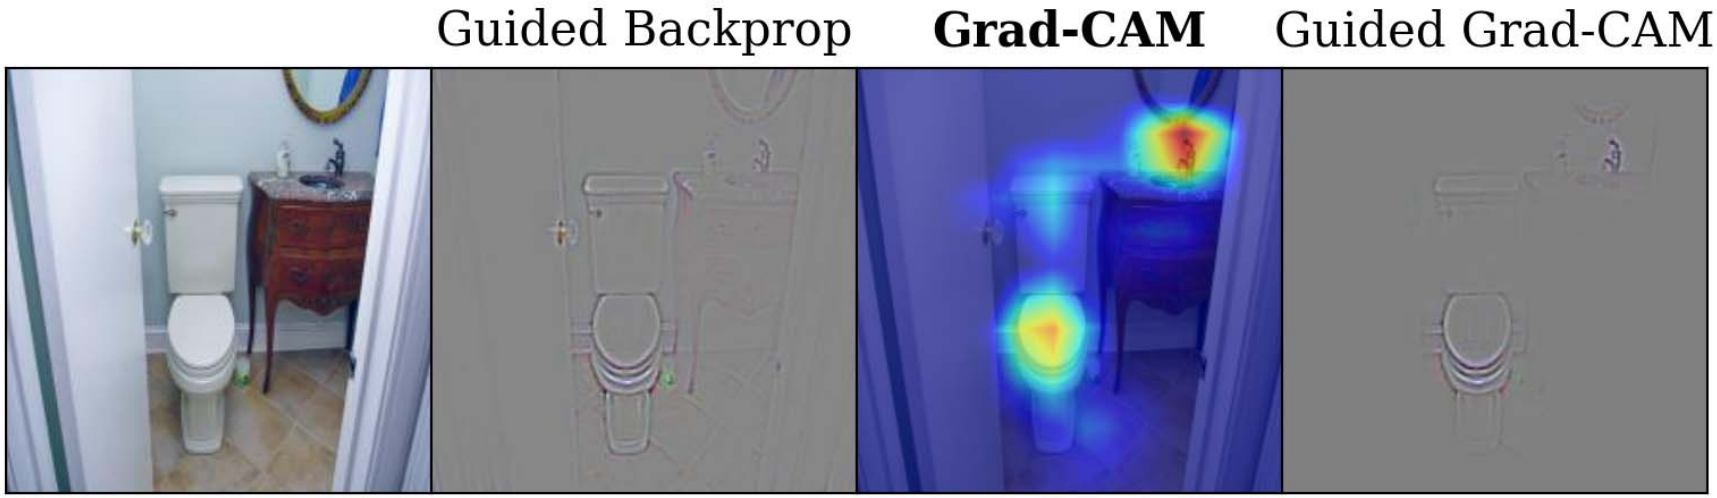

A bathroom with a toilet and a sink

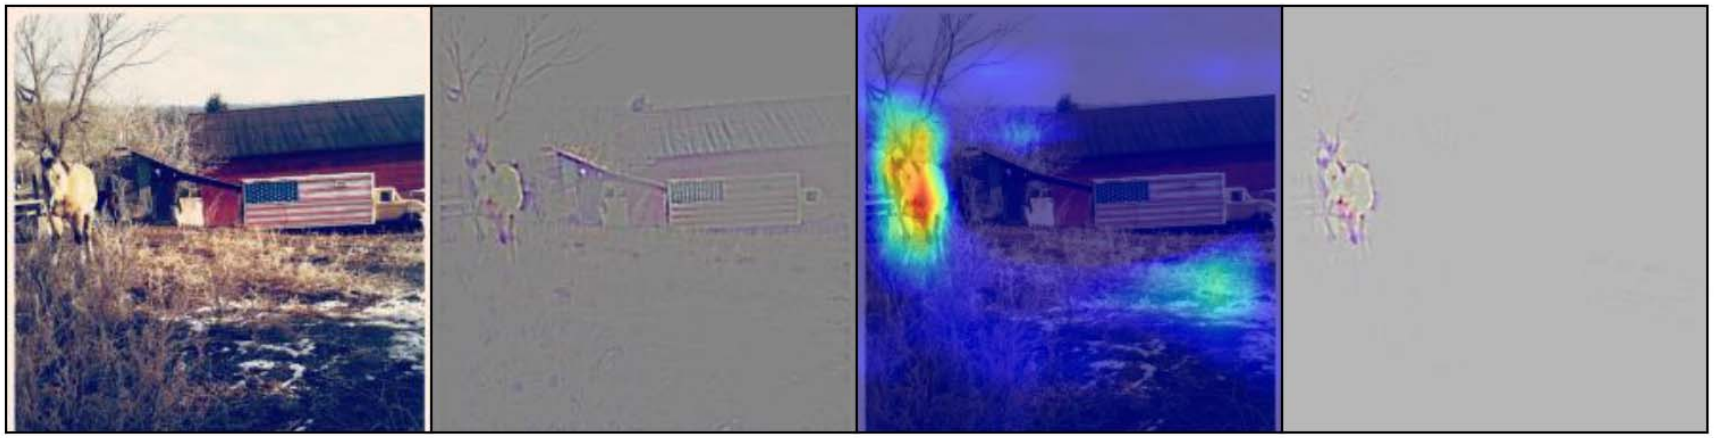

A horse is standing in a field with a fence in the background

Slide Credit: Ram Selvaraju

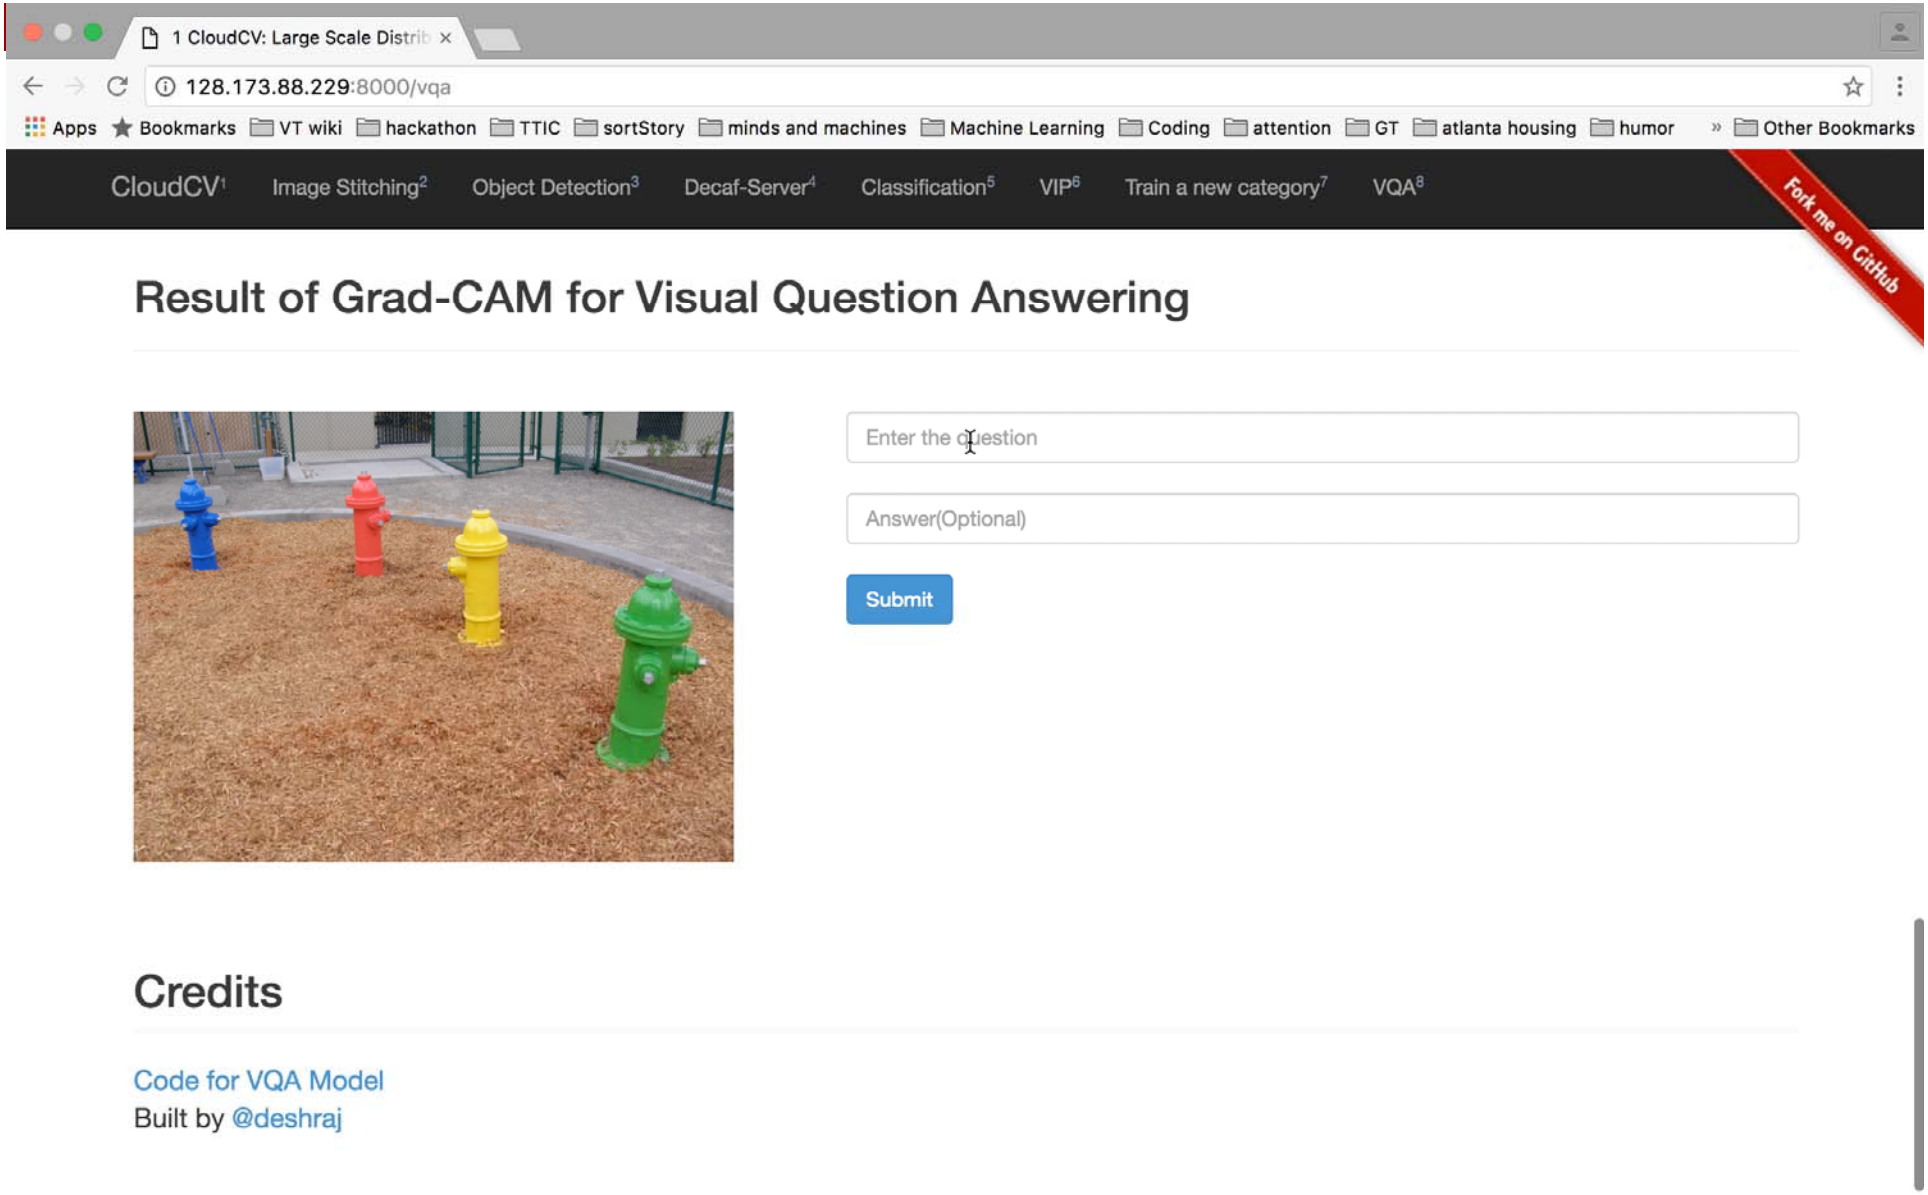
1 CloudCV: Large Scale Distrib
128.173.88.229:8000/vqa
Apps Bookmarks VT wiki hackathon TTIC sortStory minds and machines Machine Learning Coding attention GT atlanta housing humor Other Bookmarks
CloudCV1 Image Stitching2 Object Detection3 Decaf-Server4 Classification5 VIP6 Train a new category7 VQA8
Fork me on GitHub
Result of Grad-CAM for Visual Question Answering
Enter the question
Answer(Optional)
Submit
Credits
Code for VQA Model
Built by @deshraj

Plan for Today

- Visualizing CNNs
  - Visualizing filters
  - Last layer embeddings
  - Visualizing activations
  - Maximally activating patches
  - Occlusion maps
  - Salient or “important” pixels
    - Gradient-based visualizations
  - How to evaluate visualization?
  - Creating “prototypical” images for a class
  - Creating adversarial images
  - Deep dream
  - Feature inversion

How we evaluate explanations?

- Class-discriminative?
  - Show what they say they found?
- Building Trust with a User?
  - Help users?
- Human-like?
  - Do machines look where humans look?

Is Grad-CAM more class discriminative?

- Can people tell which class is being visualized?
  - Images from Pascal VOC'07 with exactly 2 categories.

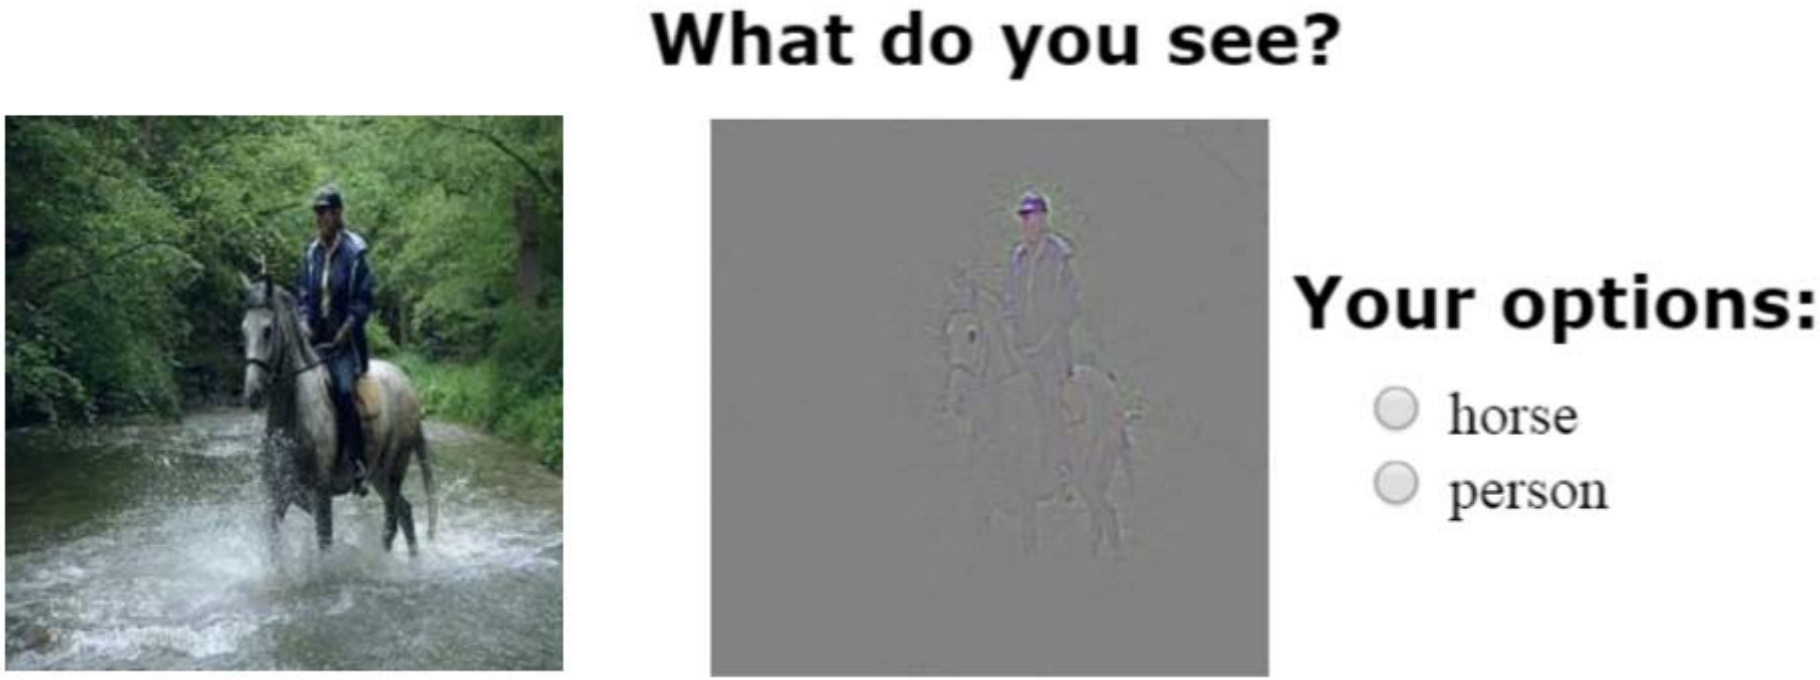

- Intuition: A good explanation produces discriminative visualizations for the class of interest.

Slide Credit: Ram Selvaraju

Is Grad-CAM more class discriminative?

- Human accuracy for 2-class classification

Method	Human Classification Accuracy
Guided Backpropagation	44.44
Guided Grad-CAM	61.23

+17%

Grad-CAM makes existing visualizations class discriminative.

Slide Credit: Ram Selvaraju

Help establish trust with a user?

- Given explanations from 2 models,
  - VGG16 and AlexNet

  which one is more trustworthy?

- Pick images where both models = correct prediction
- Show these to AMT workers and evaluate

Slide Credit: Ram Selvaraju

Help establish trust in a user?

Both robots predicted: horse

Robot A based it's decision on **Robot B** based it's decision on

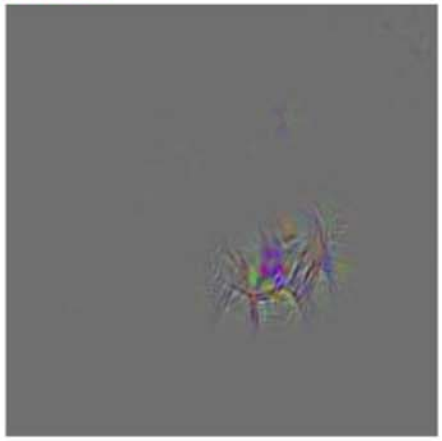

Which robot is more reasonable?

1. Robot A seems clearly more reasonable than robot B
2. Robot A seems slightly more reasonable than robot B
3. Both robots seem equally reasonable
4. Robot B seems slightly more reasonable than robot A
5. Robot B seems clearly more reasonable than robot A

Method	Relative Reliability
Guided Backpropagation	+1.00
Guided Grad-CAM	+1.27

Users place higher trust in a model that generalizes better.

Slide Credit: R

Where do humans choose to look to answer visual questions?

VQA-HAT (Human ATtention)

Question: How many players are visible in the image?

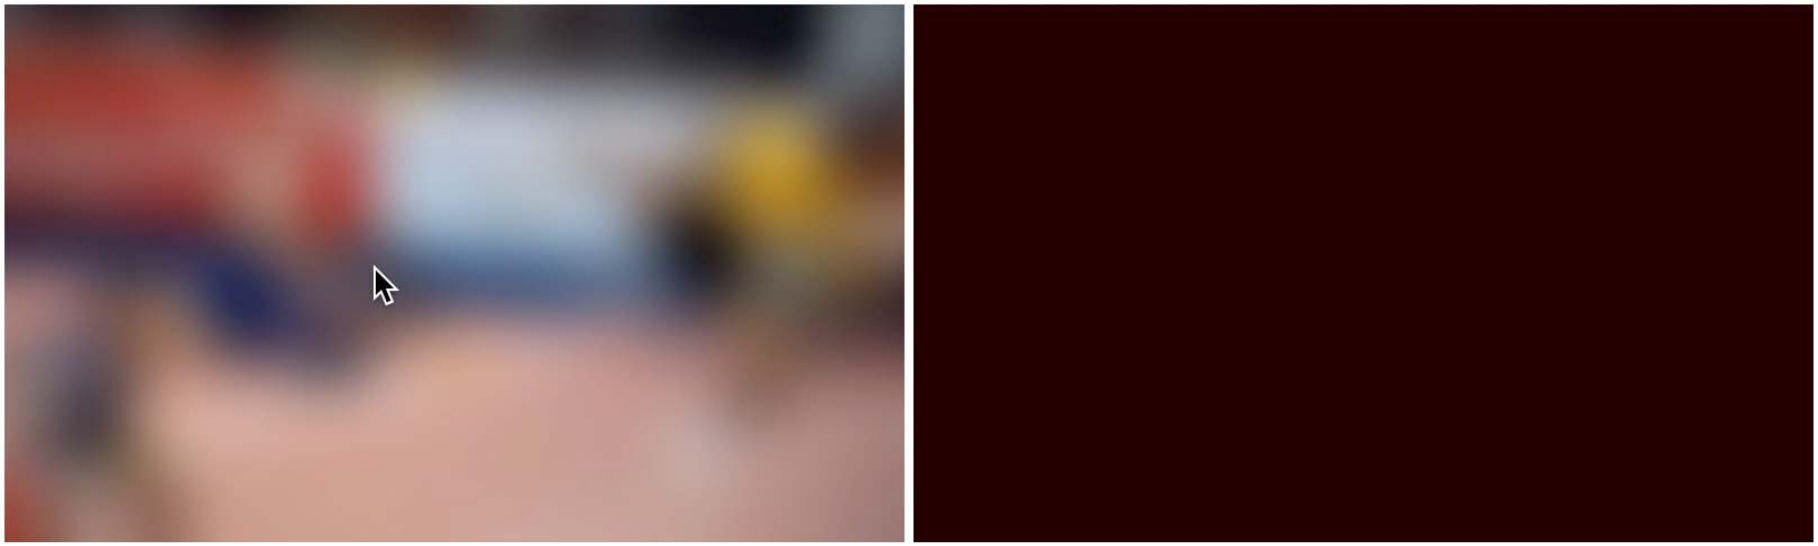

BLUR IMAGE

Answer:

3

VQA-HAT (Human ATtention)

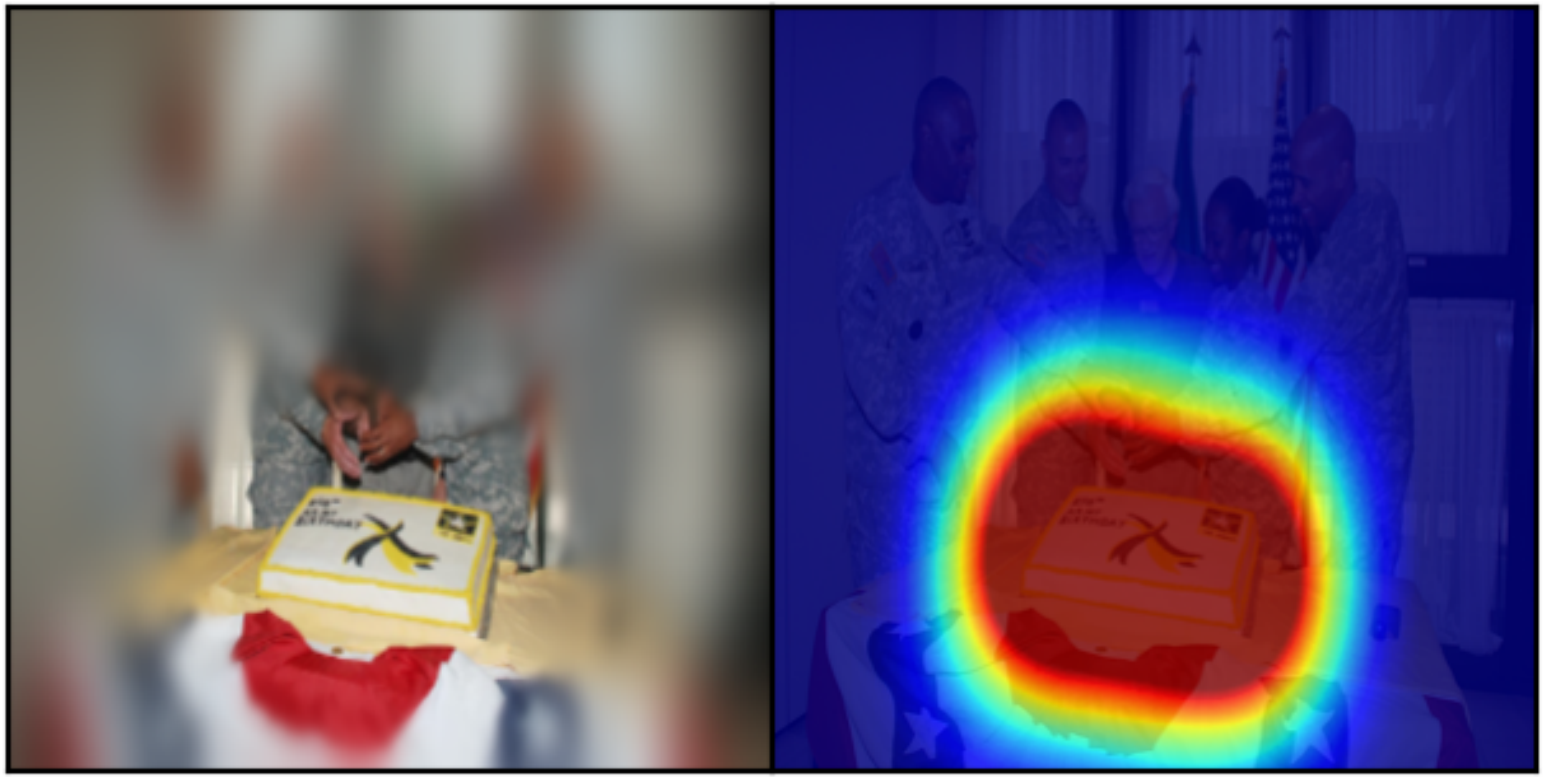

What food is on the table? Cake

Slide Credit: Abhishek Das

VQA-HAT (Human ATtention)

What animal is she riding? Horse

Slide Credit: Abhishek Das

VQA-HAT (Human ATtention)

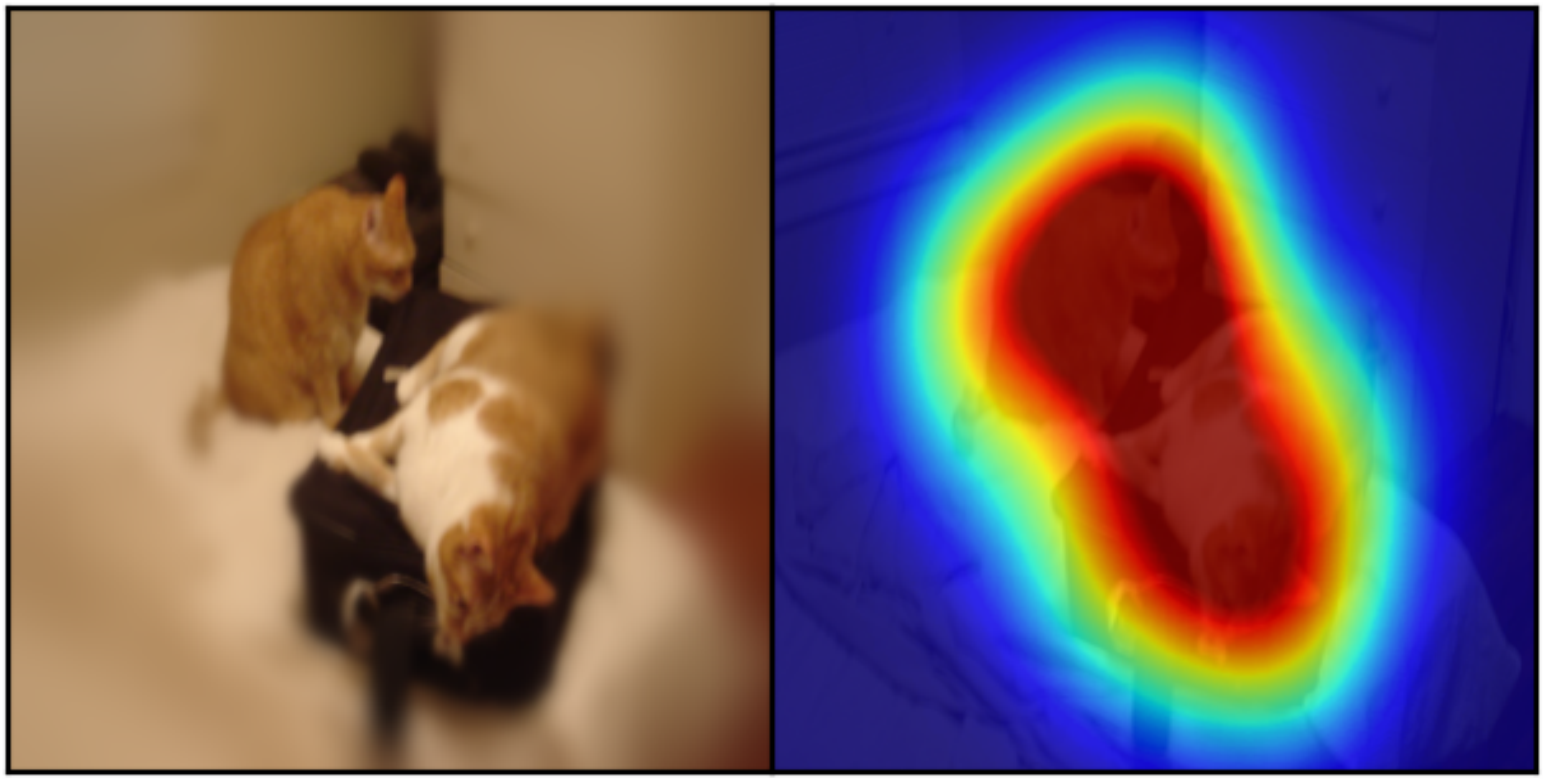

What number of cats are laying on the bed? 2

Slide Credit: Abhishek Das

Are Grad-CAM explanations human-like?

- Correlation with human attention maps
  [Das & Agarwal et al. EMNLP'16]

What are they doing?

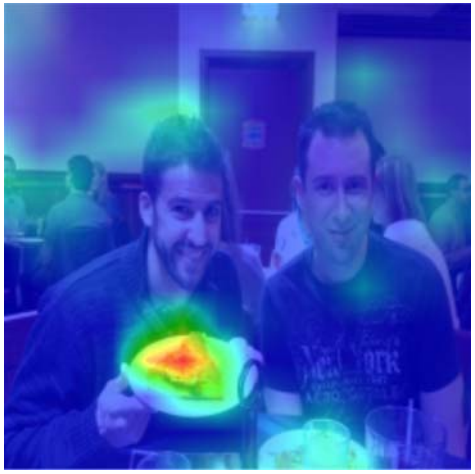

Grad-CAM for '*eating*'

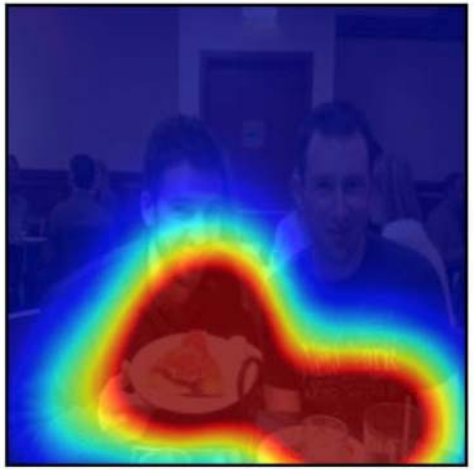

Human ATtention map (HAT) for '*eating*'

Method	Rank Correlation w/ HAT
Guided Backpropagation	0.122
Guided Grad-CAM	**0.136**

Current models look at regions more similar to humans than baselines

Plan for Today

- Visualizing CNNs
  - Visualizing filters
  - Last layer embeddings
  - Visualizing activations
  - Maximally activating patches
  - Occlusion maps
  - Salient or “important” pixels
    - Gradient-based visualizations
  - How to evaluate visualizations?
  - Creating “prototypical” images for a class
  - Creating adversarial images
  - Deep dream
  - Feature inversion

Visualizing CNN features: Gradient Ascent on Pixels

(Guided) backprop:
Find the part of an image that a neuron responds to

Gradient ascent on pixels:
Generate a synthetic image that maximally activates a neuron

$$I^* = \arg\max_I \boxed{f(I)} + \boxed{R(I)}$$

Neuron value

Natural image regularizer

Slide Credit: Fei-Fei Li, Justin Johnson, Serena Yeung, CS 231n

Visualizing CNN features: Gradient Ascent on Pixels

$$\arg\max_I \boxed{S_c(I)} - \lambda \|I\|_2^2$$

score for class c (before Softmax)

1. Initialize image to zeros

zero image

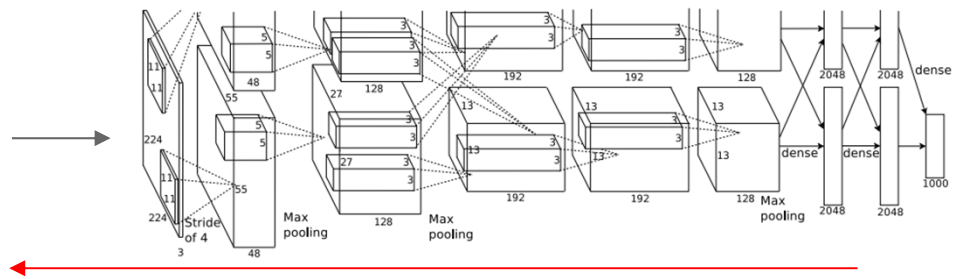

Repeat:

2. Forward image to compute current scores
3. Backprop to get gradient of neuron value with respect to image pixels
4. Make a small update to the image

Visualizing CNN features: Gradient Ascent on Pixels

$$\arg\max_I S_c(I) - \lambda \|I\|_2^2$$

Simple regularizer: Penalize L2 norm of generated image

Simonyan, Vedaldi, and Zisserman, "Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps", ICLR Workshop 2014.

Visualizing CNN features: Gradient Ascent on Pixels

$$\arg\max_I S_c(I) - \lambda \|I\|_2^2$$

Simple regularizer: Penalize L2 norm of generated image

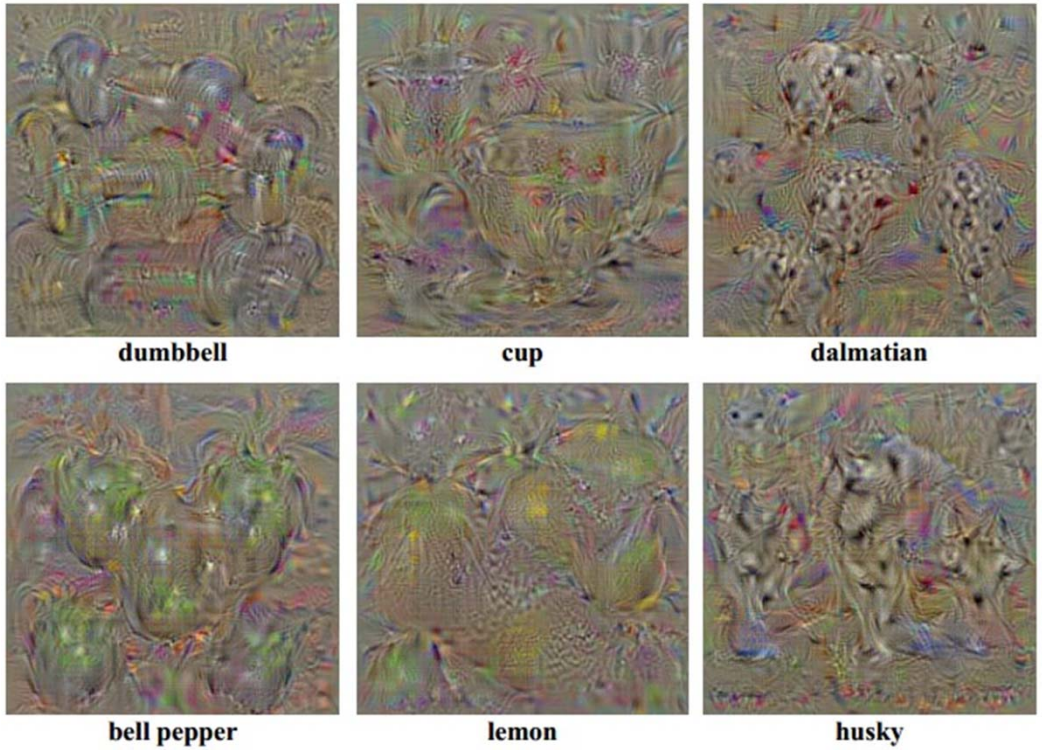

Simonyan, Vedaldi, and Zisserman, "Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps", ICLR Workshop 2014.

Visualizing CNN features: Gradient Ascent on Pixels

$$\arg\max_I S_c(I) - \lambda \|I\|_2^2$$

Simple regularizer: Penalize L2 norm of generated image

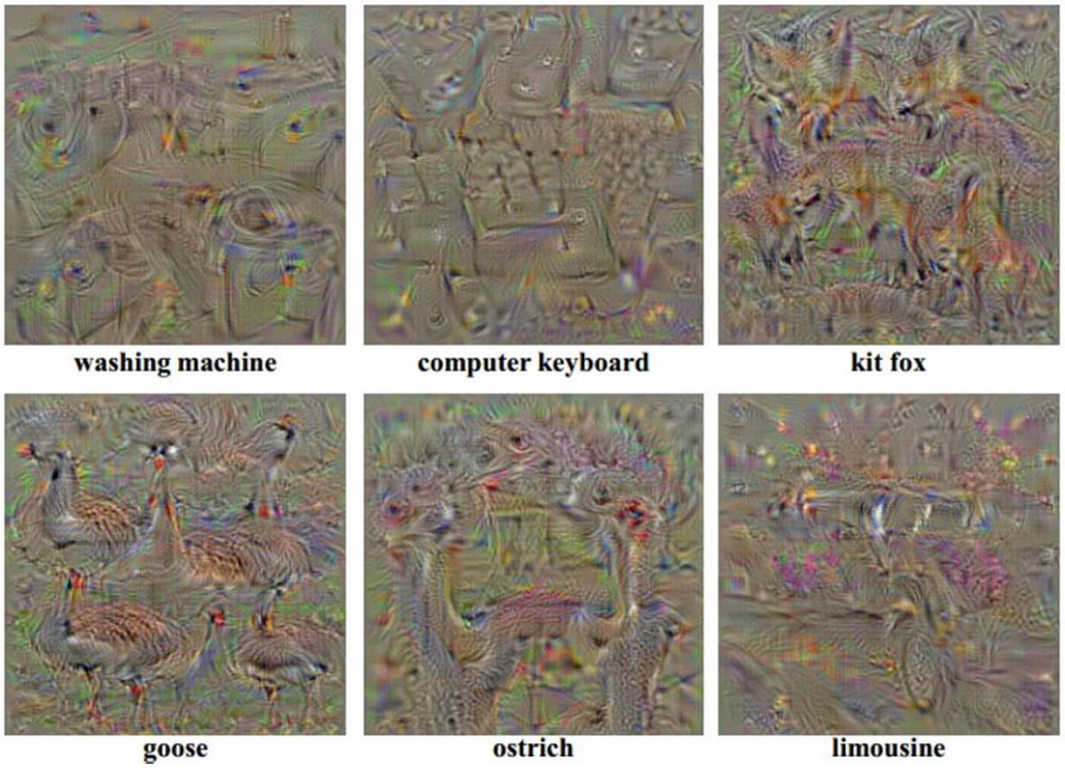

Yosinski et al, "Understanding Neural Networks Through Deep Visualization", ICML DL Workshop 2014. Figure copyright Jason Yosinski, Jeff Clune, Anh Nguyen, Thomas Fuchs, and Hod Lipson, 2014. Reproduced with permission.

Fooling Images / Adversarial Examples

(1)Start from an arbitrary image
(2)Pick an arbitrary class
(3)Modify the image to maximize the class
(4)Repeat until network is fooled

Fooling Images / Adversarial Examples

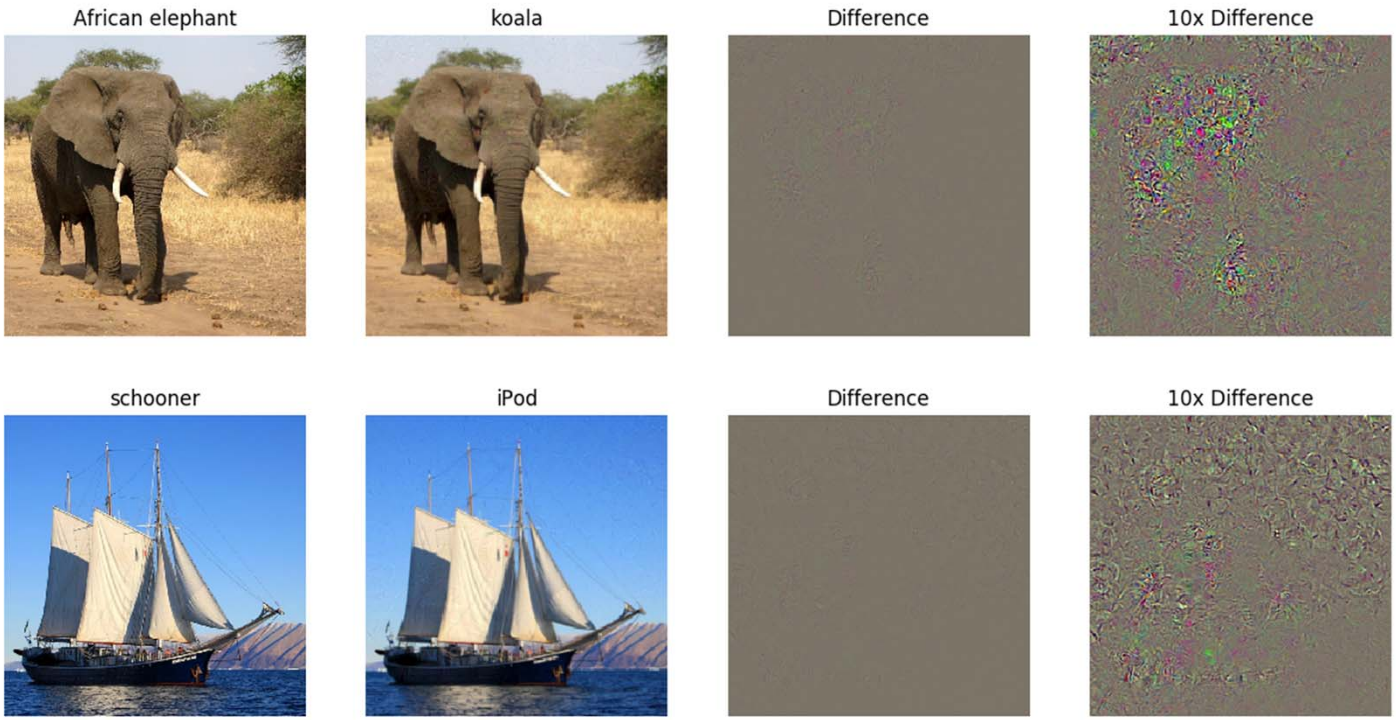

Slide Credit: Fei-Fei Li, Justin Johnson, Serena Yeung, CS 231n

Plan for Today

- Visualizing CNNs
  - Visualizing filters
  - Last layer embeddings
  - Visualizing activations
  - Maximally activating patches
  - Occlusion maps
  - Salient or “important” pixels
    - Gradient-based visualizations
  - How to evaluate visualizations?
  - Creating “prototypical” images for a class
  - Creating adversarial images
  - Deep dream
  - Feature inversion

DeepDream: Amplify existing features

Rather than synthesizing an image to maximize a specific neuron, instead try to **amplify** the neuron activations at some layer in the network

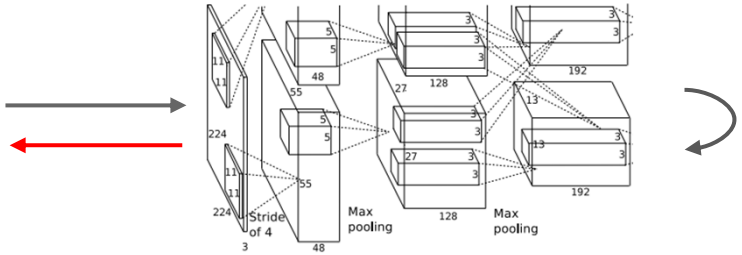

Choose an image and a layer in a CNN; repeat:

1. Forward: compute activations at chosen layer
2. Set gradient of chosen layer *equal to its activation*
3. Backward: Compute gradient on image
4. Update image

DeepDream: Amplify existing features

Rather than synthesizing an image to maximize a specific neuron, instead try to **amplify** the neuron activations at some layer in the network

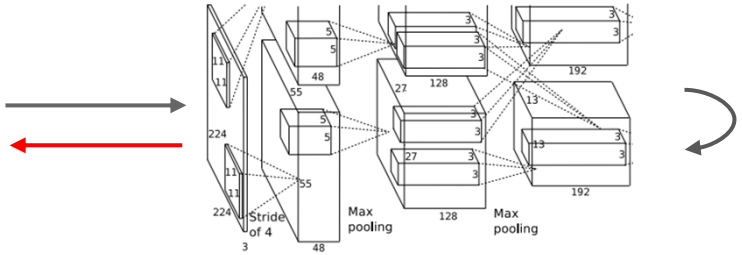

Choose an image and a layer in a CNN; repeat:

1. Forward: compute activations at chosen layer
2. Set gradient of chosen layer *equal to its activation*
3. Backward: Compute gradient on image
4. Update image

Equivalent to:

$$I^* = \arg\max_I \sum_i f_i(I)^2$$

Mordvintsev, Olah, and Tyka, "Inceptionism: Going Deeper into Neural Networks", Google Research Blog. Images are licensed under CC-BY 4.0

Slide Credit: Fei-Fei Li, Justin Johnson, Serena Yeung, CS 231n

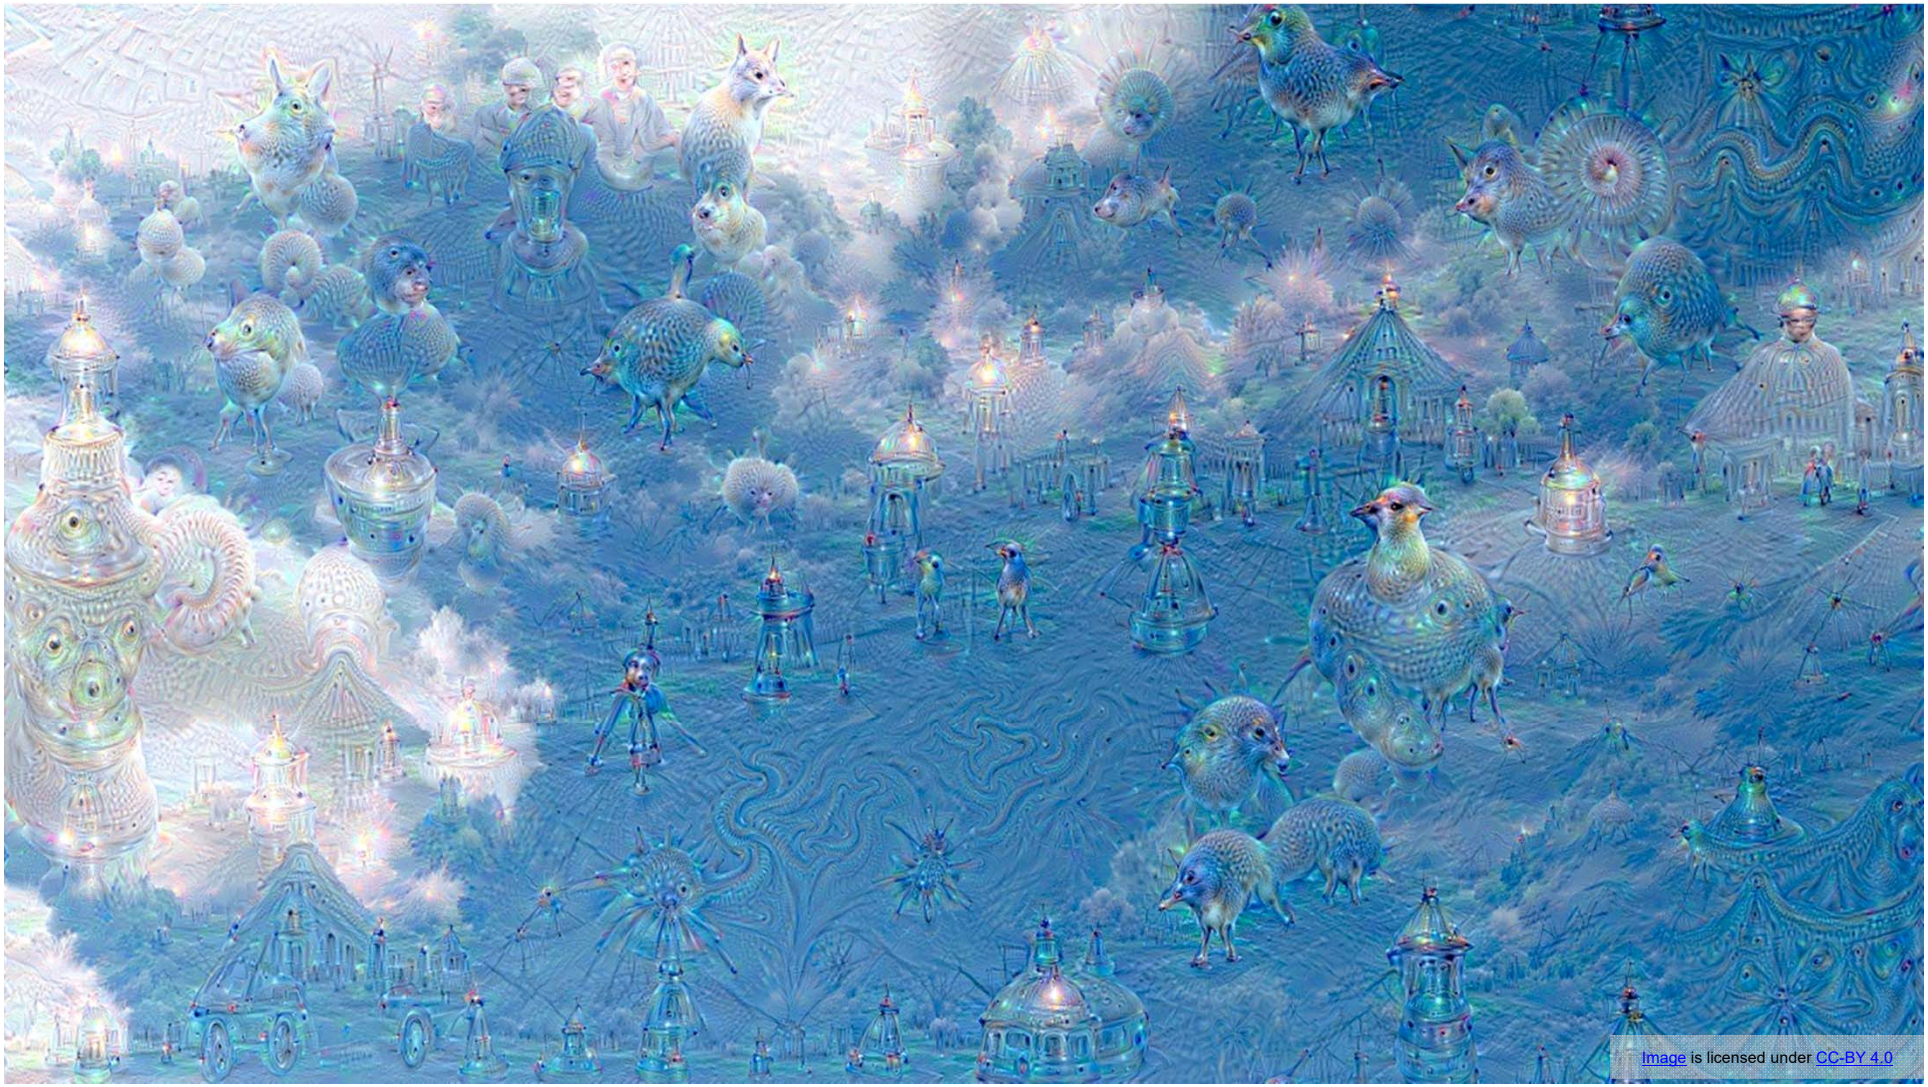

Slide Credit: Fei-Fei Li, Justin Johnson, Serena Yeung, CS 231n

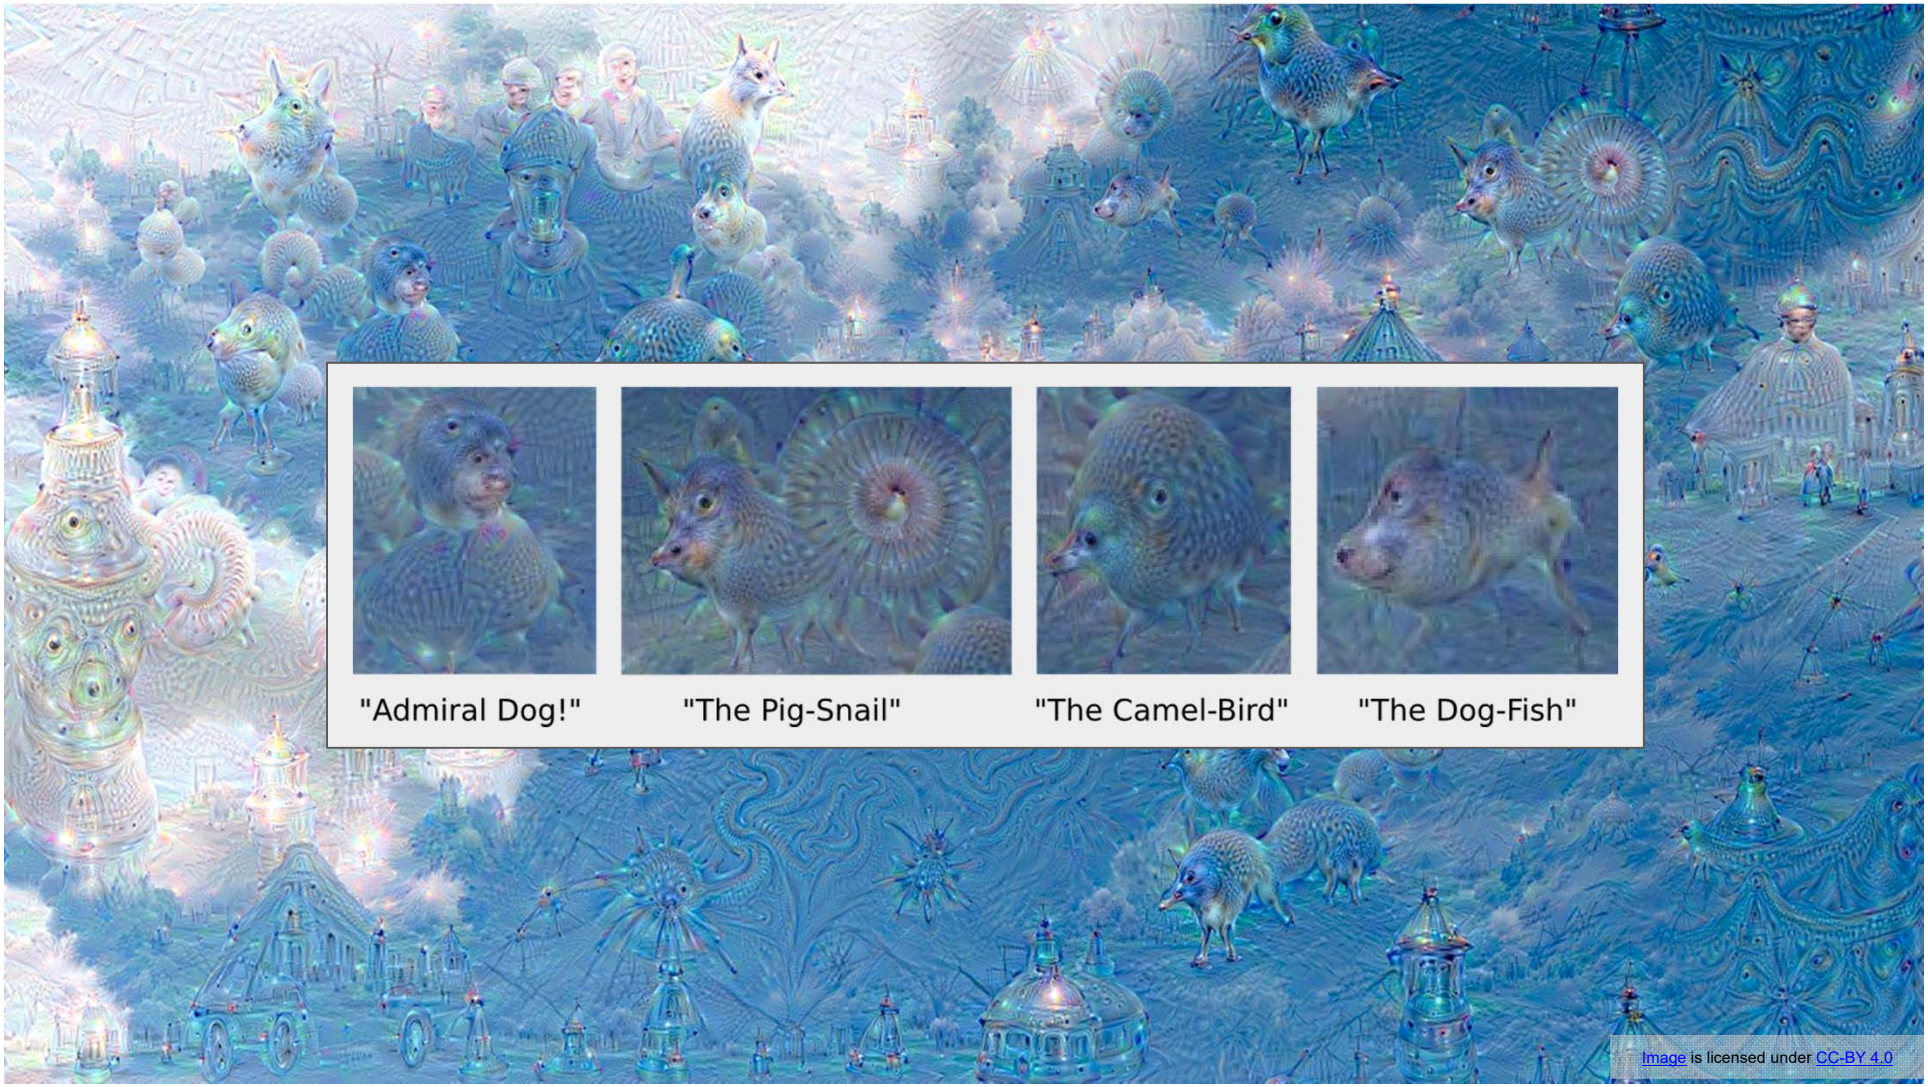

Slide Credit: Fei-Fei Li, Justin Johnson, Serena Yeung, CS 231n

Feature Inversion

Given a CNN feature vector for an image, find a new image that:

- Matches the given feature vector
- “looks natural” (image prior regularization)

$$\mathbf{x}^* = \underset{\mathbf{x} \in \mathbb{R}^{H \times W \times C}}{\operatorname{argmin}} \ \ell(\Phi(\mathbf{x}), \Phi_0) + \lambda \mathcal{R}(\mathbf{x})$$

Given feature vector

Features of new image

$$\ell(\Phi(\mathbf{x}), \Phi_0) = \|\Phi(\mathbf{x}) - \Phi_0\|^2$$

$$\mathcal{R}_{V^\beta}(\mathbf{x}) = \sum_{i,j} \left((x_{i,j+1} - x_{ij})^2 + (x_{i+1,j} - x_{ij})^2 \right)^{\frac{\beta}{2}}$$

Total Variation regularizer
(encourages spatial smoothness)

Mahendran and Vedaldi, “Understanding Deep Image Representations by Inverting Them”, CVPR 2015

Feature Inversion

Reconstructing from different layers of VGG-16

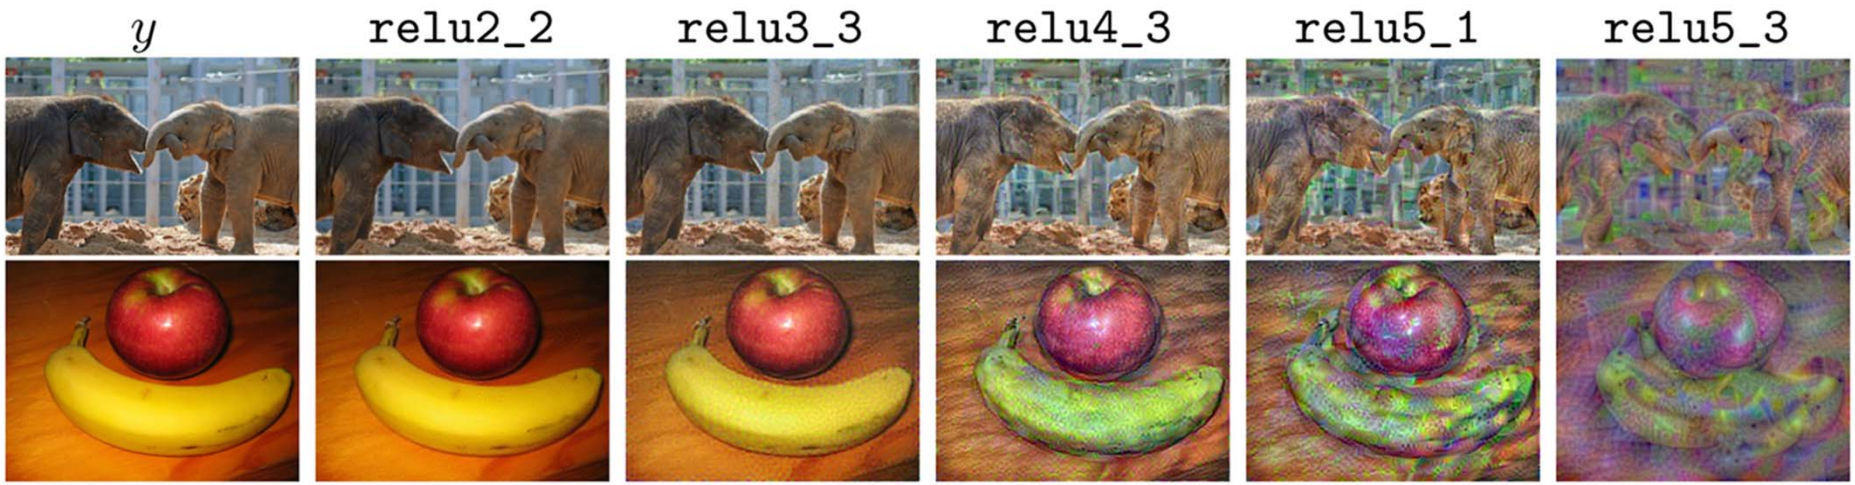

Mahendran and Vedaldi, “Understanding Deep Image Representations by Inverting Them”, CVPR 2015
Figure from Johnson, Alahi, and Fei-Fei, “Perceptual Losses for Real-Time Style Transfer and Super-Resolution”, ECCV 2016. Copyright Springer, 2016.
Reproduced for educational purposes.

Side-effect - style transfer

- *Content representation:* feature map at each layer
- *Style representation:* Covariance matrix at each layer
  - Spatially invariant
  - Average second-order statistics

- Idea: Optimize x to match c $G^l_{ij} = \sum_k F^l_{ik} F^l_{jk}.$ nage and style of another

Gatys, Leon A., Alexander S. Ecker, and Matthias Bethge. "A neural algorithm of artistic style." *arXiv preprint arXiv:1508.06576* (2015).

Neural Style

Step 1: Extract **content targets** (ConvNet activations of all layers for the given content image)

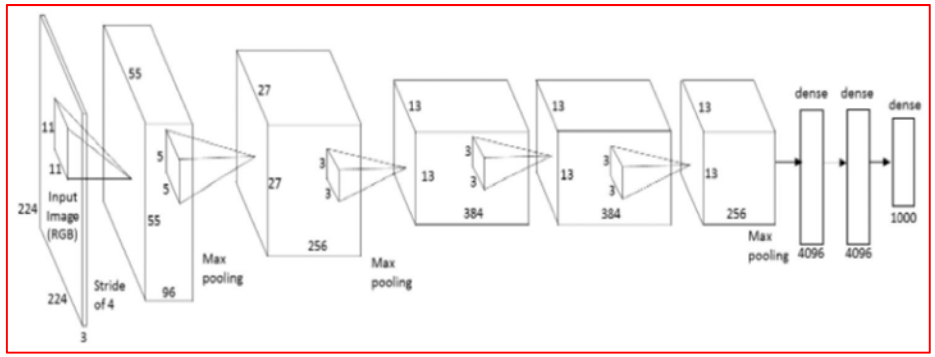

content activations

e.g.
at CONV5_1 layer we would have a [14x14x512] array of target activations

Neural Style

Step 2: Extract **style targets** (Gram matrices of ConvNet activations of all layers for the given style image)

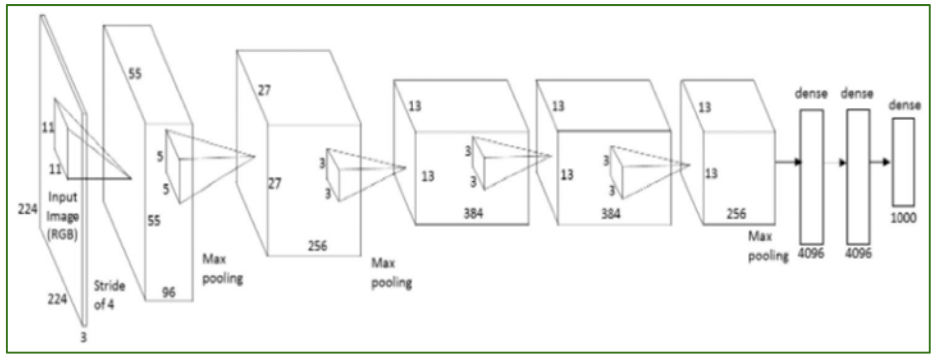

style gram matrices

$$G = V^{\mathrm{T}} V$$

e.g.
at CONV1 layer (with [224x224x64] activations) would give a [64x64] Gram matrix of all pairwise activation covariances (summed across spatial locations)

Neural Style

Step 3: Optimize over image to have:

- The **content** of the content image (activations match content)
- The **style** of the style image (Gram matrices of activations match style)

$$\mathcal{L}_{total}(\vec{p}, \vec{a}, \vec{x}) = \alpha \mathcal{L}_{content}(\vec{p}, \vec{x}) + \beta \mathcal{L}_{style}(\vec{a}, \vec{x})$$

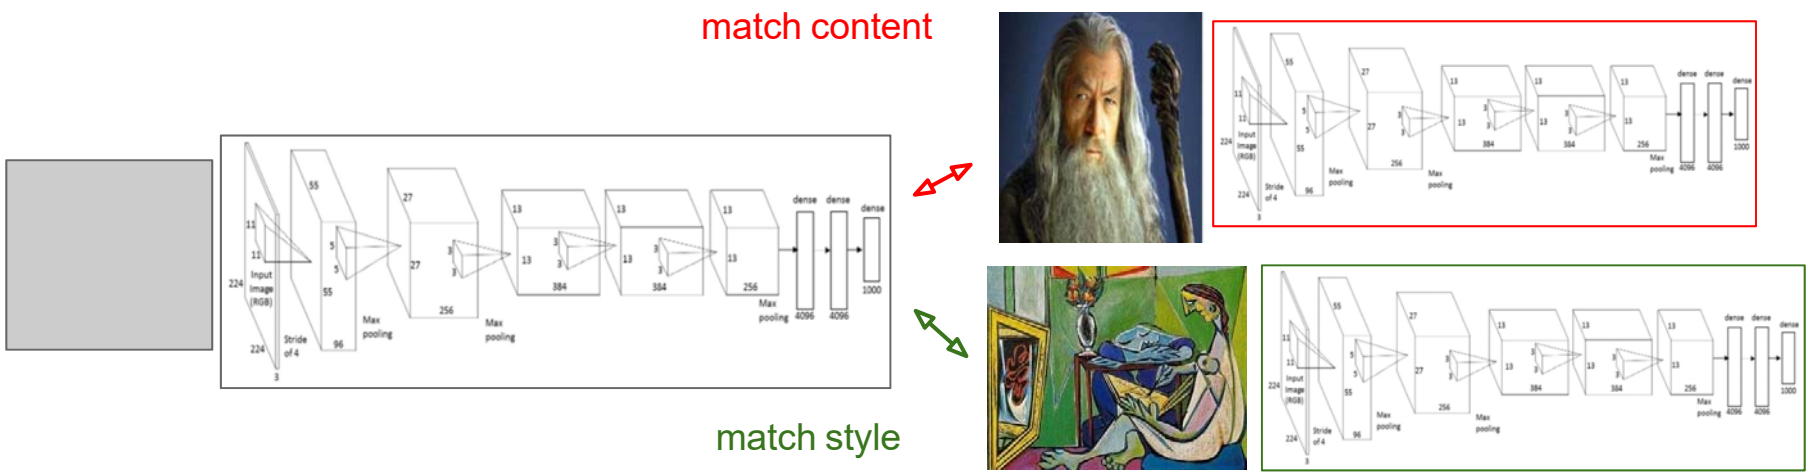

Adapted from Andrej Karpathy

Style transfer

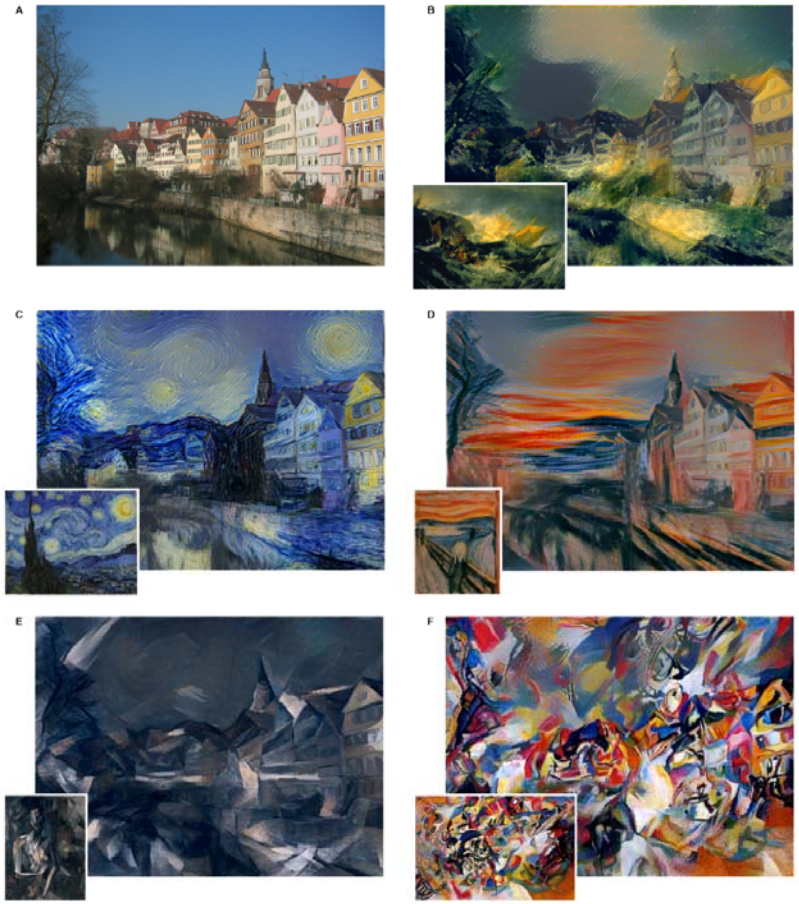

Plan for Today

- Visualizing CNNs
  - Visualizing filters
  - Last layer embeddings
  - Visualizing activations
  - Maximally activating patches
  - Occlusion maps
  - Salient or “important” pixels
    - Gradient-based visualizations
  - How to evaluate visualizations?
  - Creating “prototypical” images for a class
  - Creating adversarial images
  - Deep dream
  - Feature inversion
  - Style Transfer