



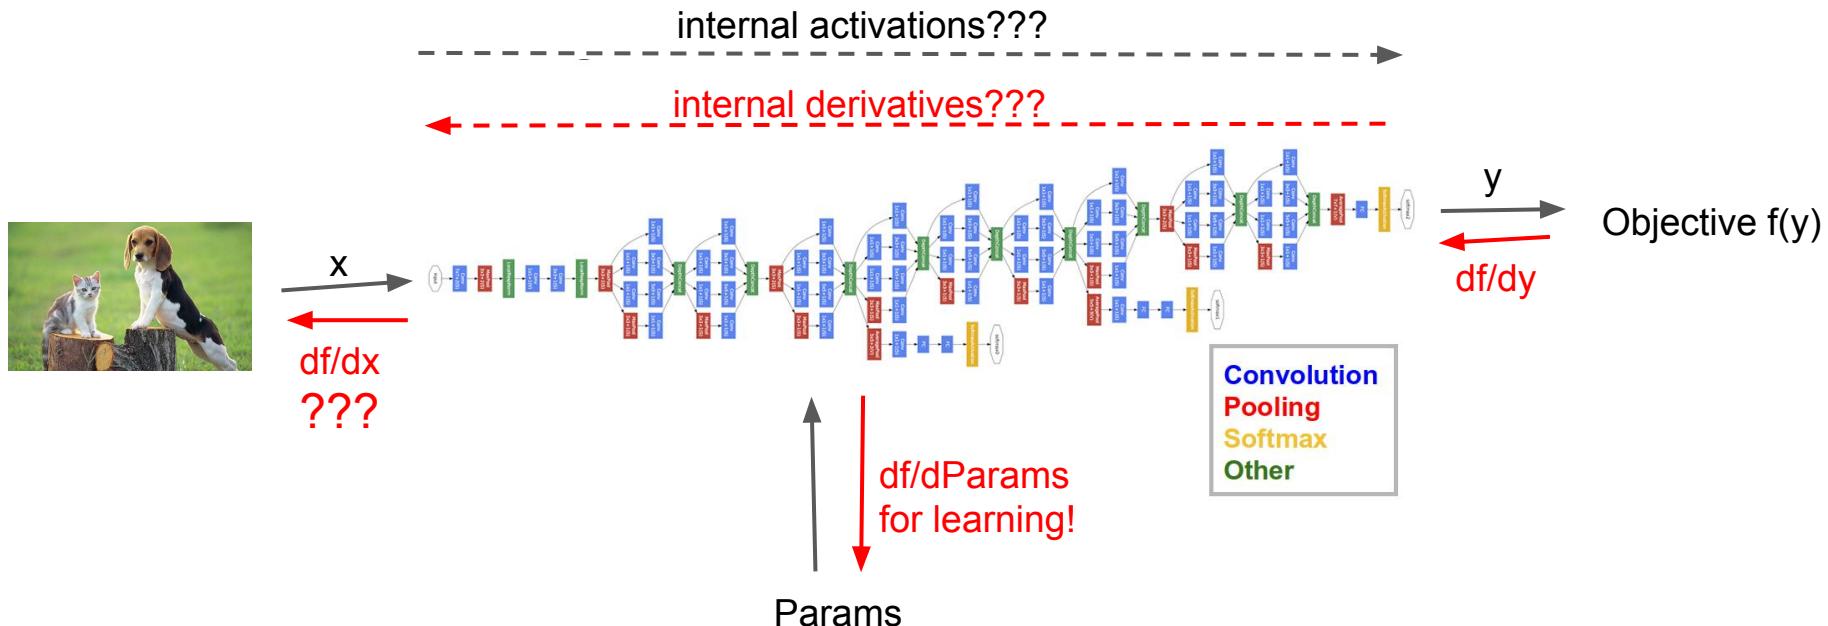
Beyond The DogSlug

Alexander Mordvintsev, Google Zurich

moralex@google.com

Neural Nets: transparent, but obscure box

What can we learn from ...



Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps

Karen Simonyan

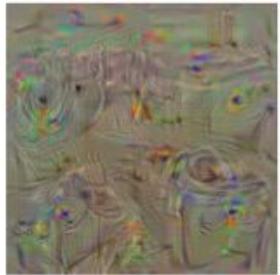
Andrea Vedaldi

Andrew Zisserman

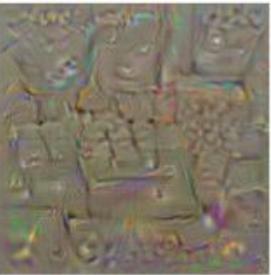
Visual Geometry Group, University of Oxford

{karen, vedaldi, az}@robots.ox.ac.uk

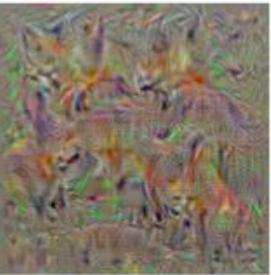
[link](#)



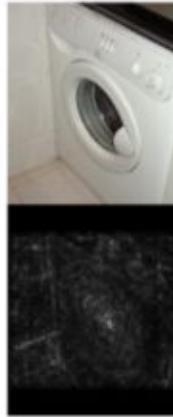
washing machine



computer keyboard



kit fox



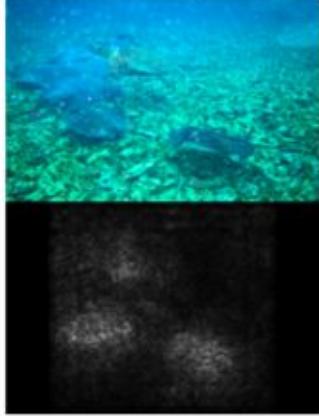
goose



ostrich



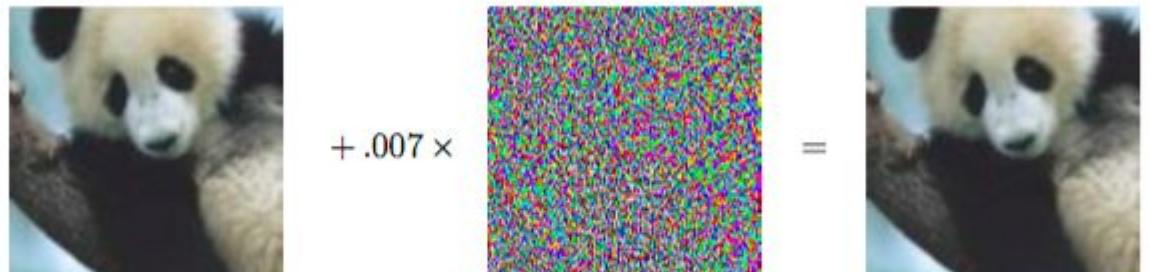
limousine



EXPLAINING AND HARNESSING ADVERSARIAL EXAMPLES

[link](#)

Ian J. Goodfellow, Jonathon Shlens & Christian Szegedy
Google Inc., Mountain View, CA
`{goodfellow,shlens,szegedy}@google.com`

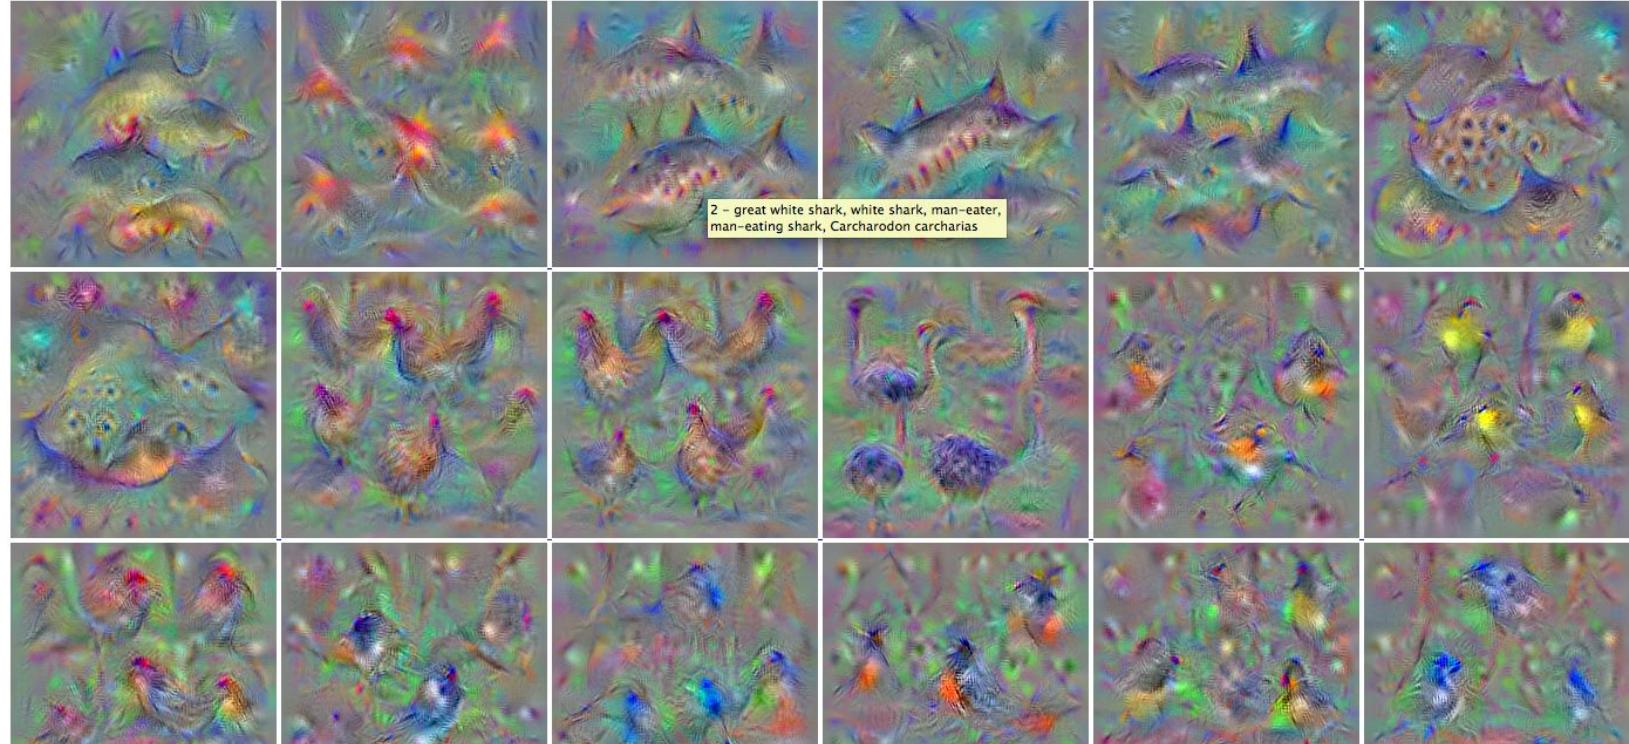
$$\begin{array}{ccc} \text{panda} & + .007 \times & \text{nematode} \\ x & & \text{sign}(\nabla_x J(\theta, x, y)) \\ & & \text{"nematode"} \\ & & 8.2\% \text{ confidence} \\ 57.7\% \text{ confidence} & & \end{array} = \begin{array}{c} \text{gibbon} \\ x + \epsilon \text{sign}(\nabla_x J(\theta, x, y)) \\ 99.3 \% \text{ confidence} \end{array}$$


1000 Classes

These visualizations made by maximizing the score, neural network assigns to a particular class, in a way, similar to [1].

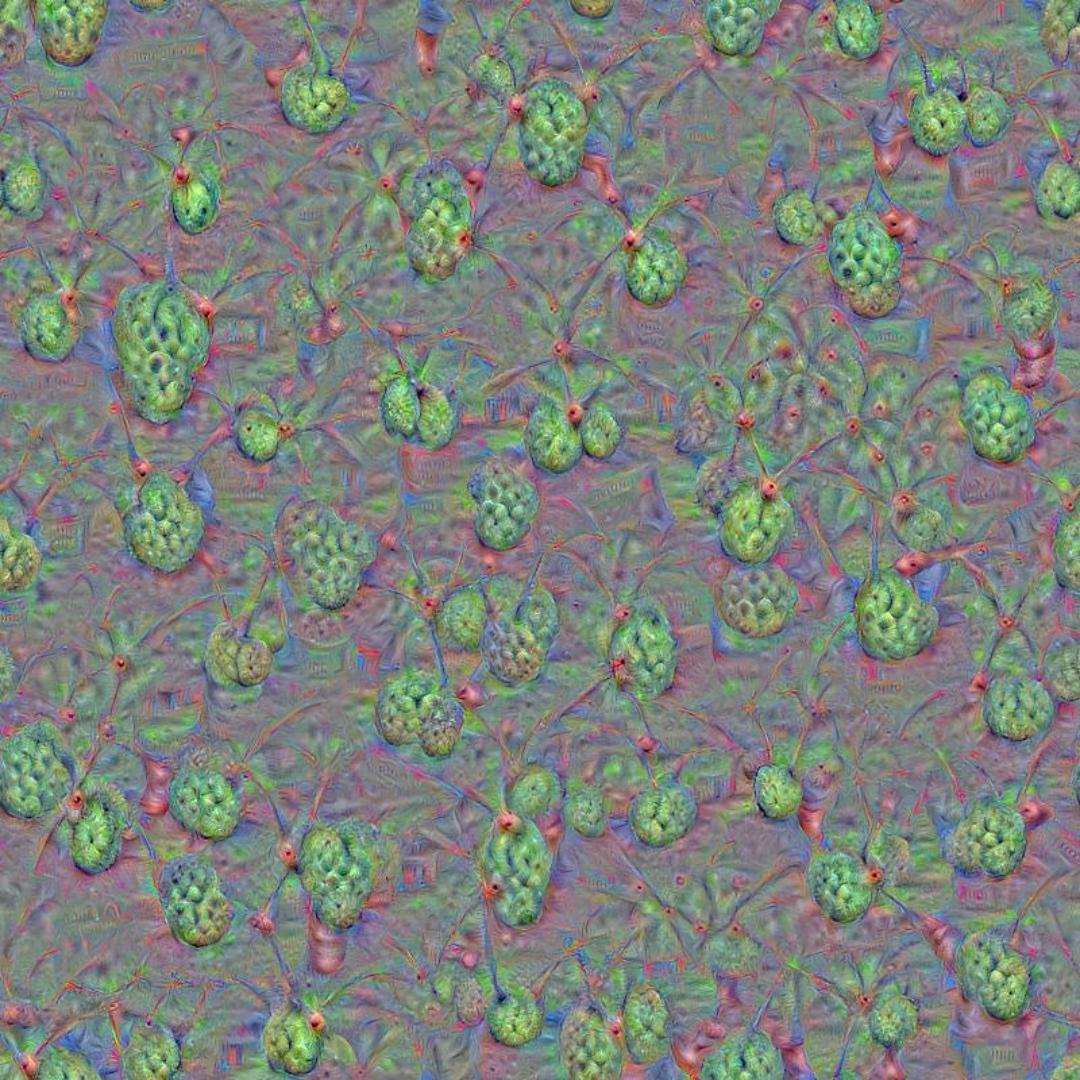
Hover over images to see labels, and click to navigate to ImageNet page for that class.

Contact moralex@ for details.

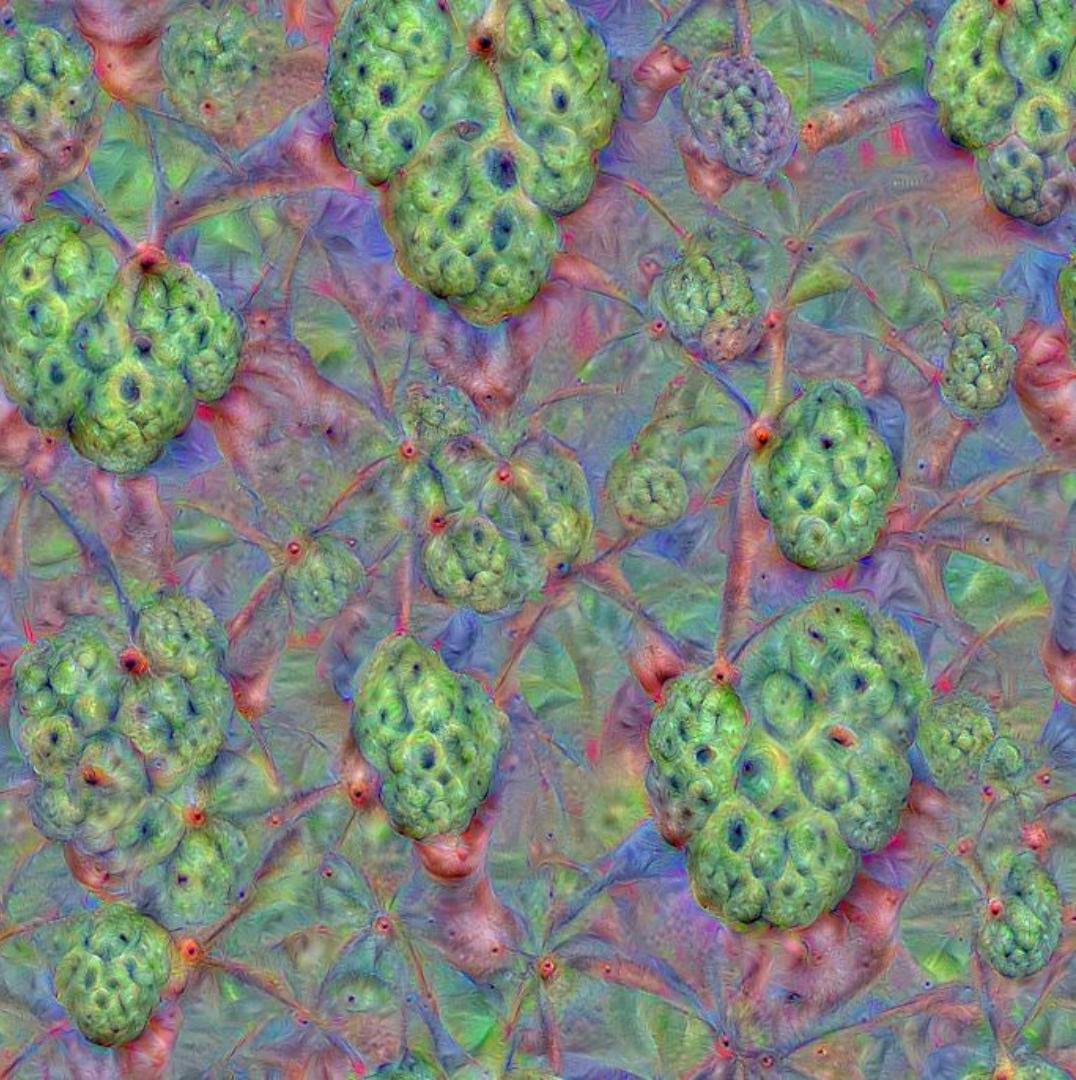
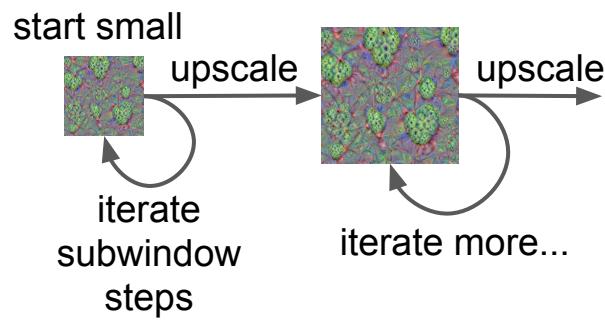


Bigger picture

- allocate big image
- apply **basic step** to random subwindows
- shows variation in the class
- somewhat lacks scale diversity...

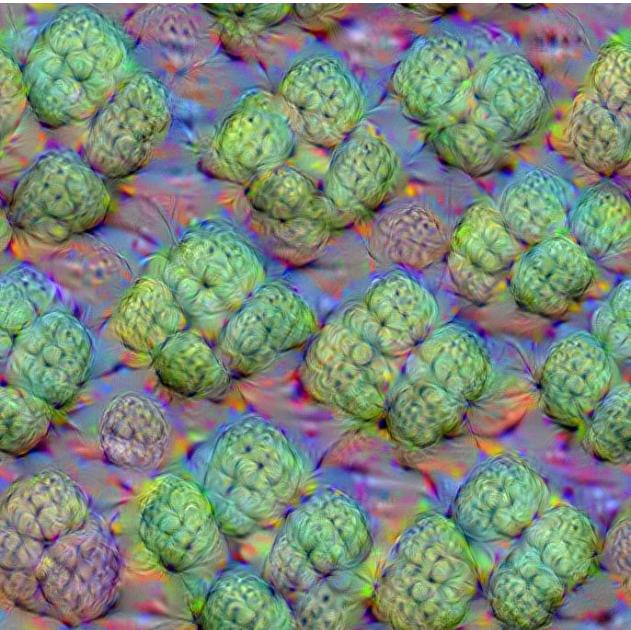


Multiscale

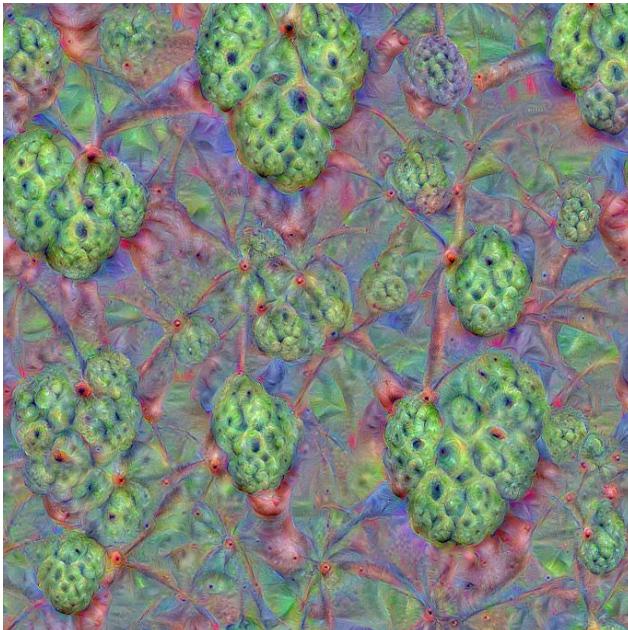


- Can also create “deep zoom” animations this way

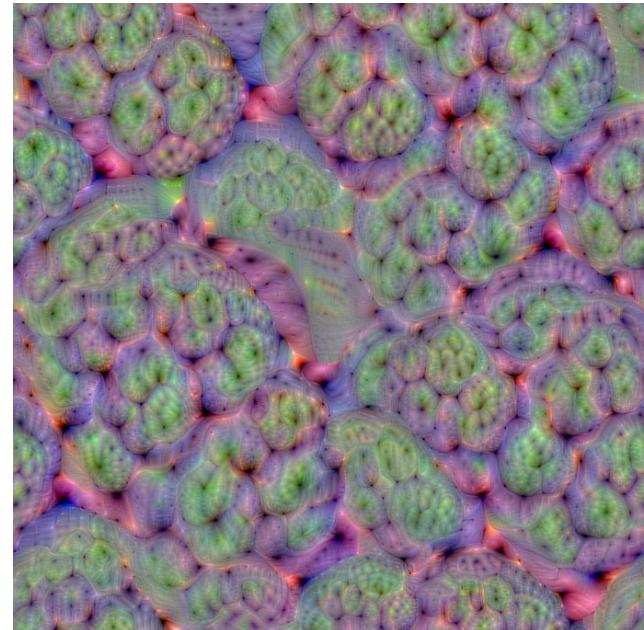
Same class, different models



CaffeNet

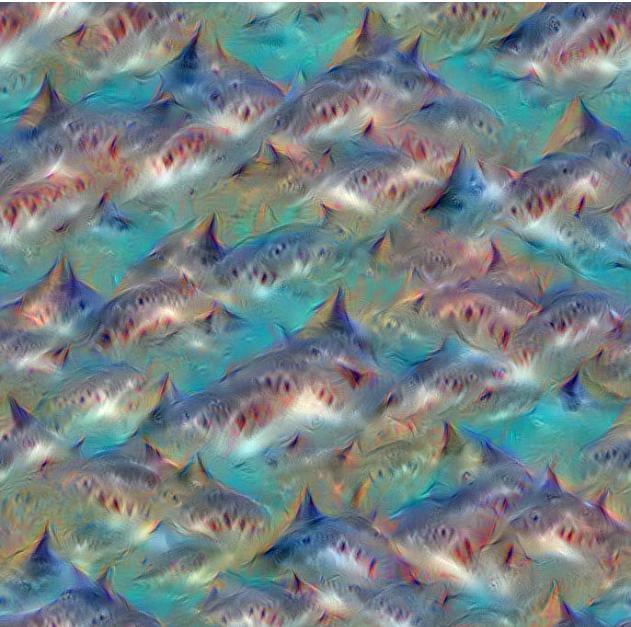


GoogleNet

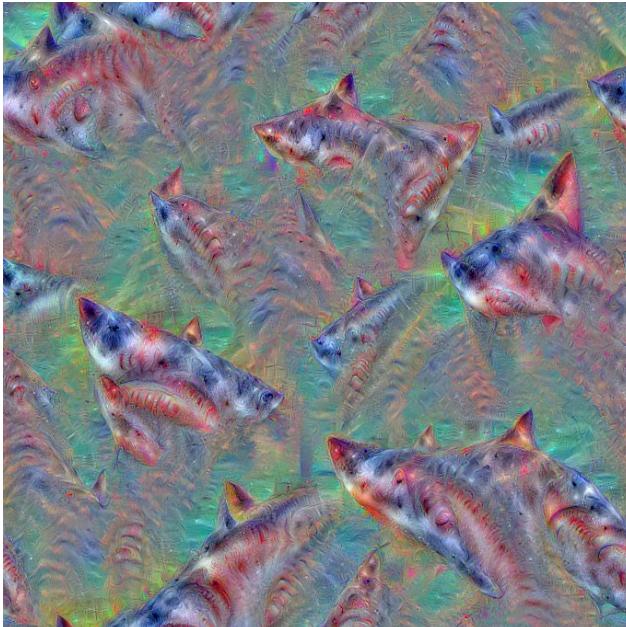


VGG19

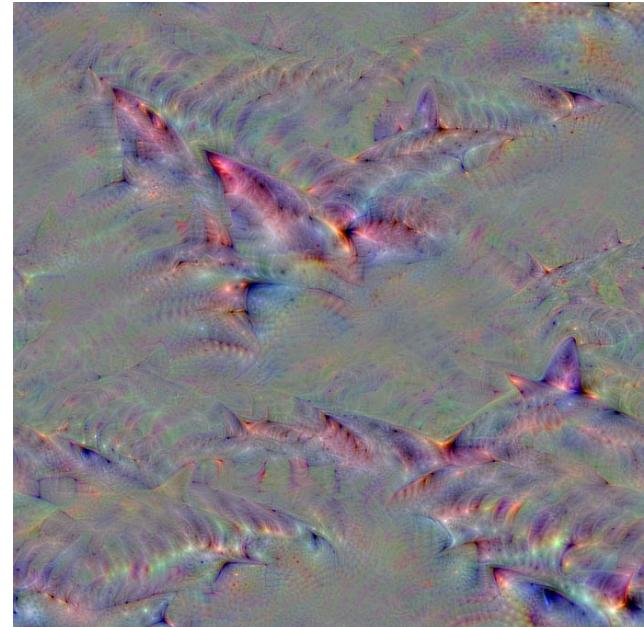
Same class, different models



CaffeNet



GoogleNet



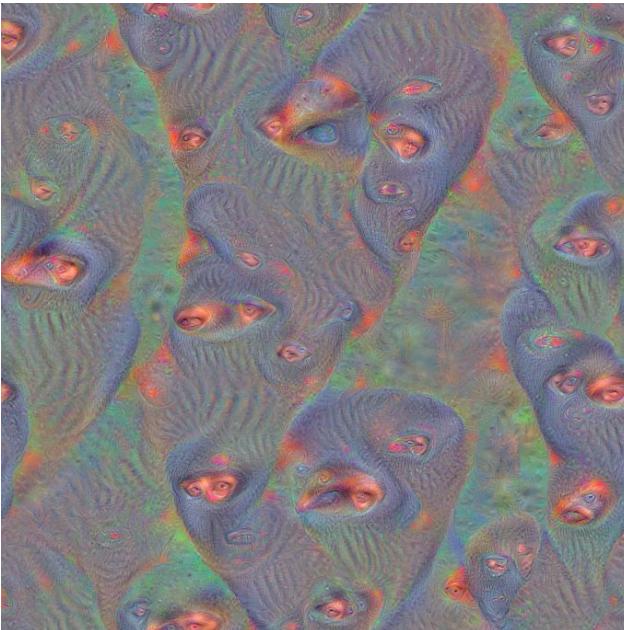
VGG19

Same class, different models

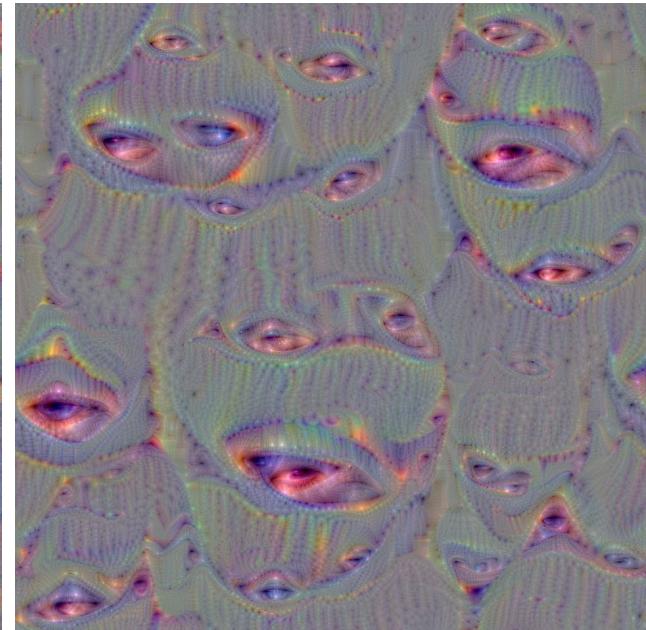
ski mask



CaffeNet



GoogleNet



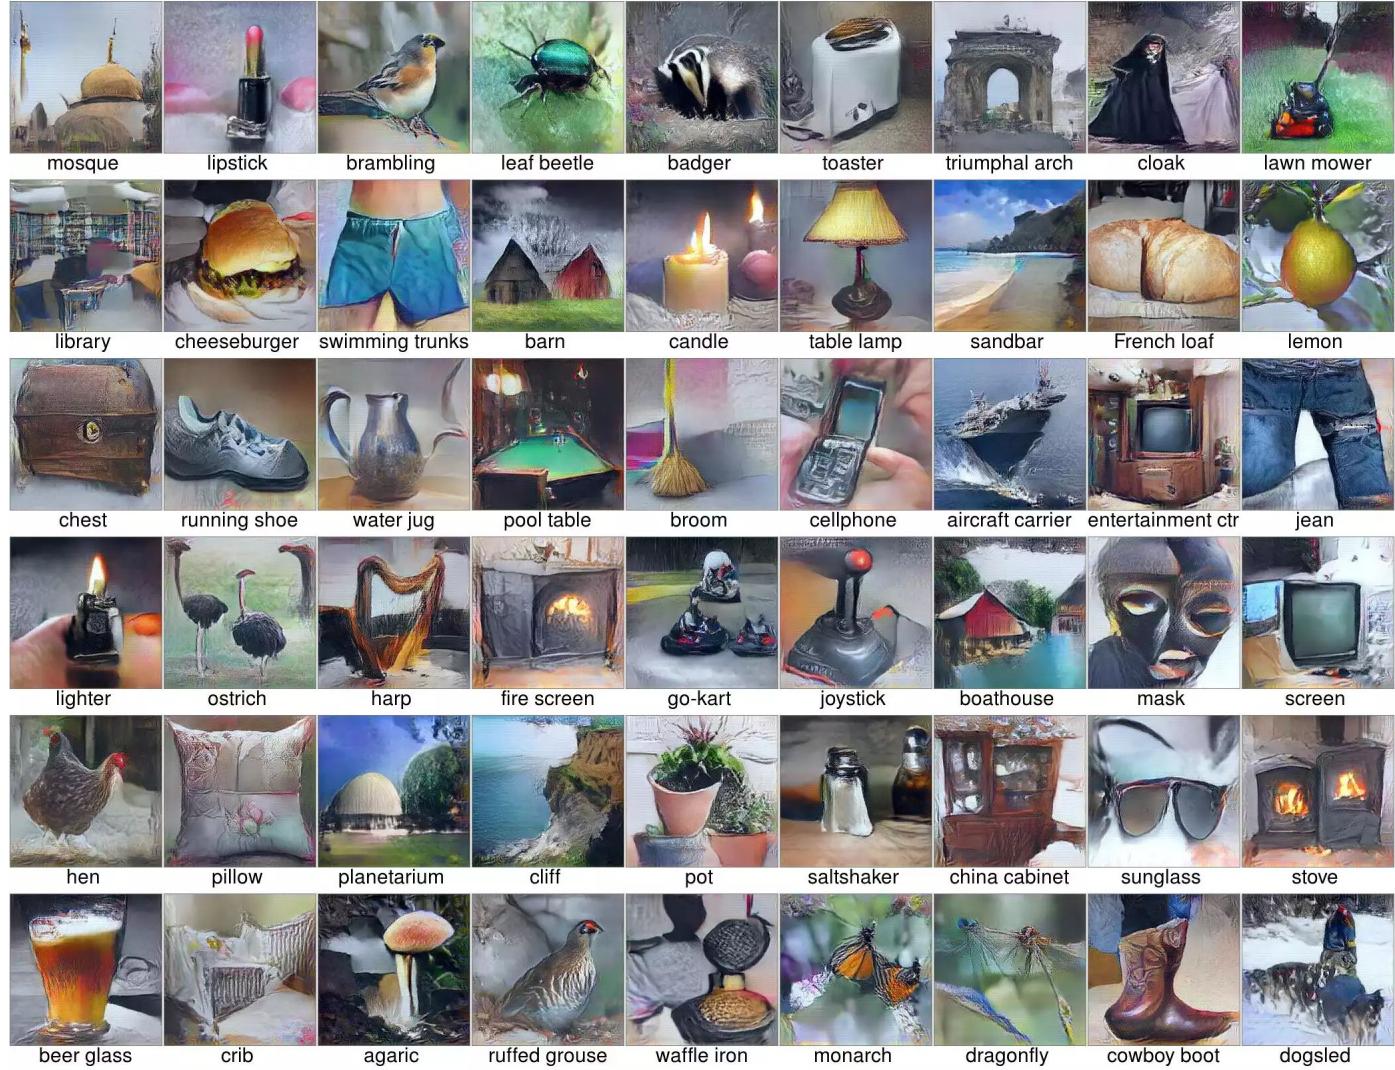
VGG19

Bilateral Filter Prior (by Mike Tyka, mtyka.github.io)

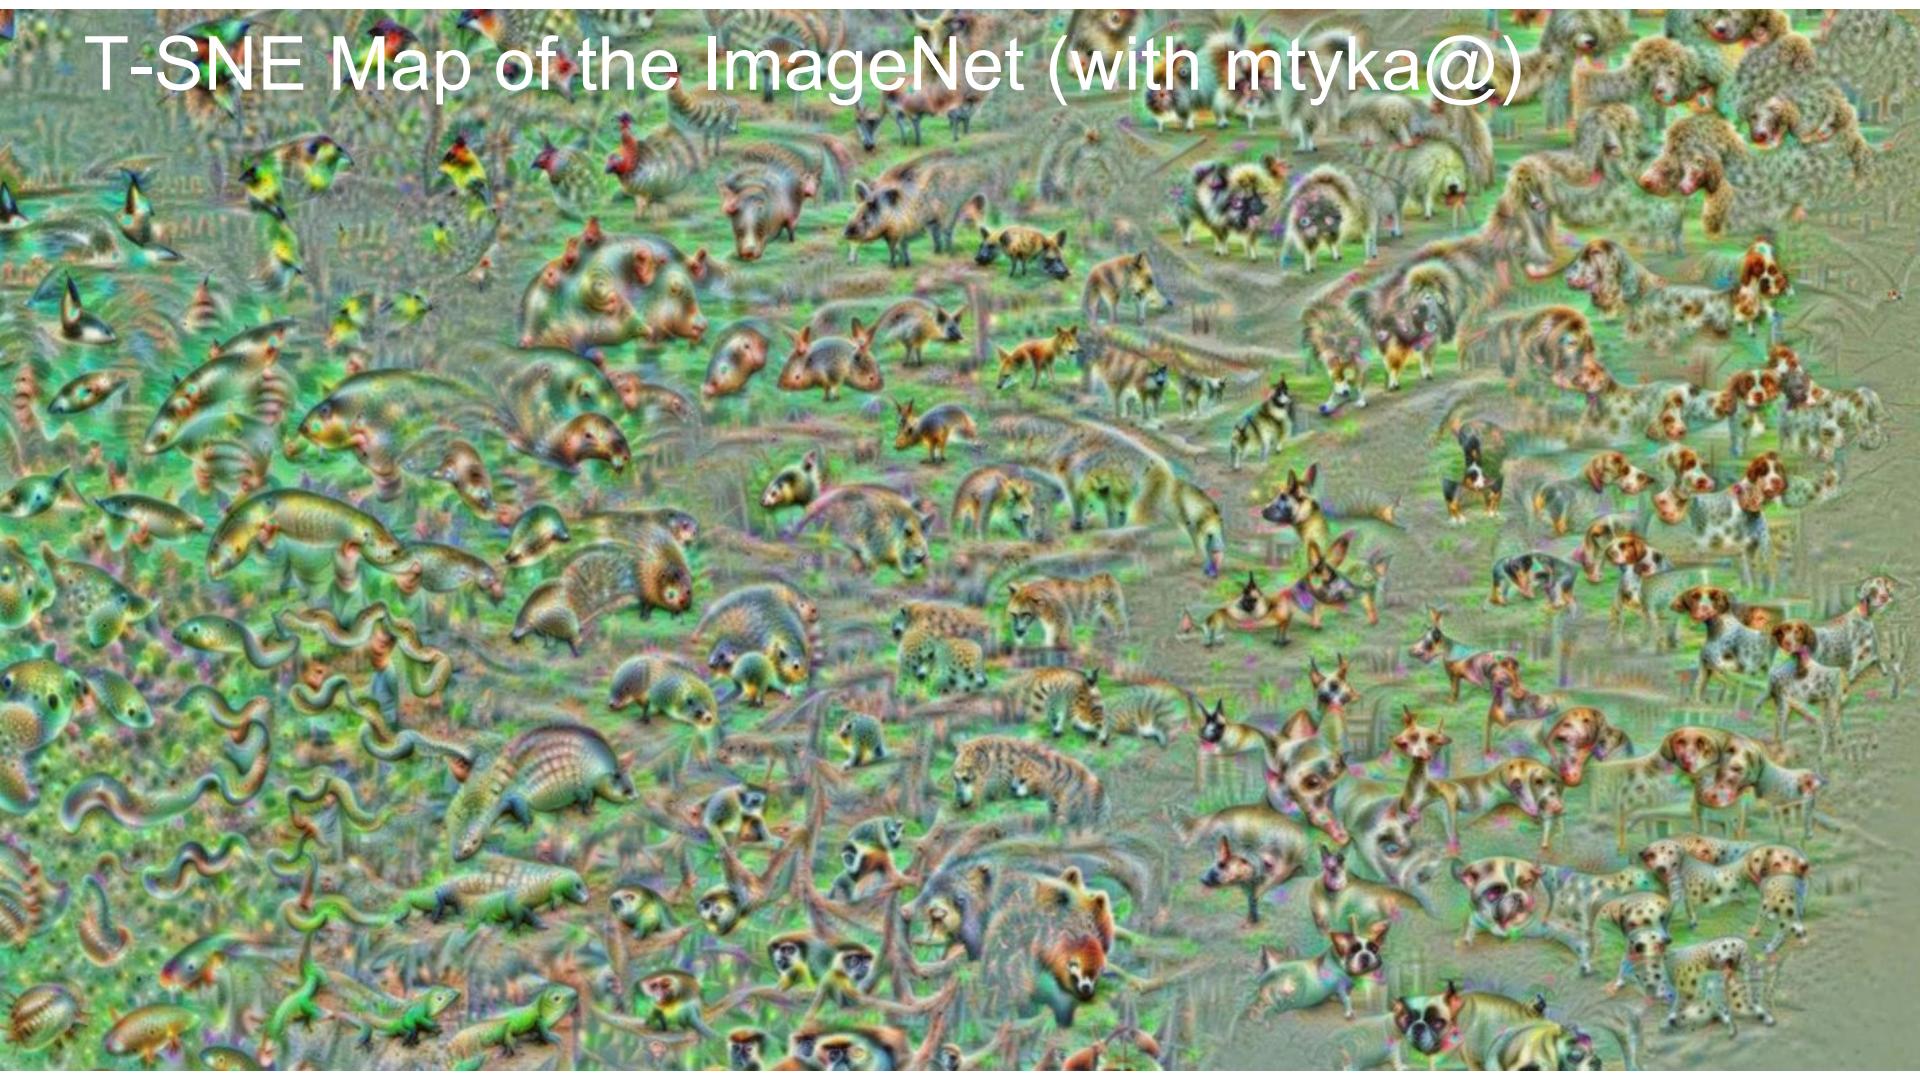


GAN prior

[Nguyen A, Dosovitskiy A, Yosinski J, Brox T, Clune J \(2016\).](#)
Synthesizing the preferred inputs for neurons in neural networks via deep generator networks.
Advances in Neural Information Processing Systems 29. ([pdf](#))

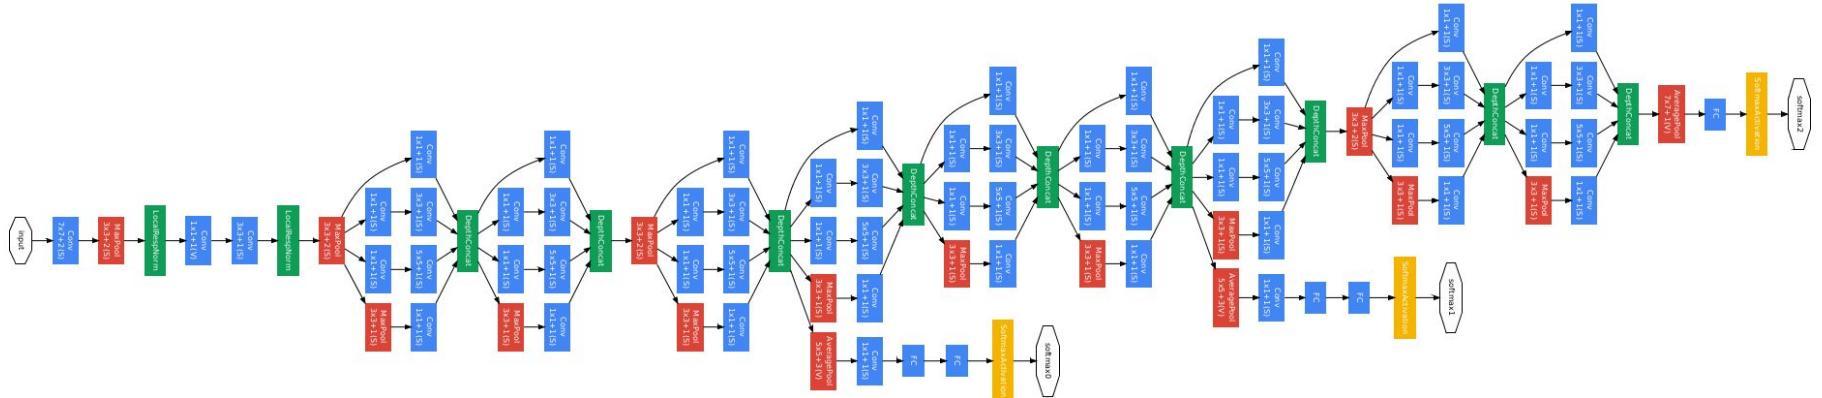


T-SNE Map of the ImageNet (with mtyka@)



What else can the network generate?

Opensource Inception



59 conv layers
7548 feature channels total

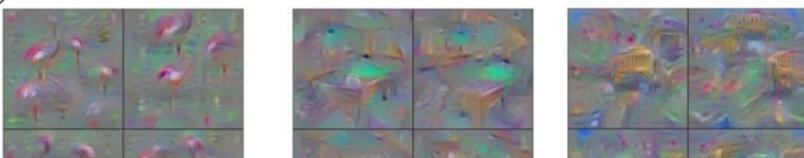
A lot of stuff to visualize!

Understanding Neural Networks Through Deep Visualization

Jason Yosinski, Jeff Clune, Anh Nguyen, Thomas Fuchs, and Hod Lipson

Quick links: [ICML DL Workshop paper](#) | [code](#) | [video](#)

Relevant



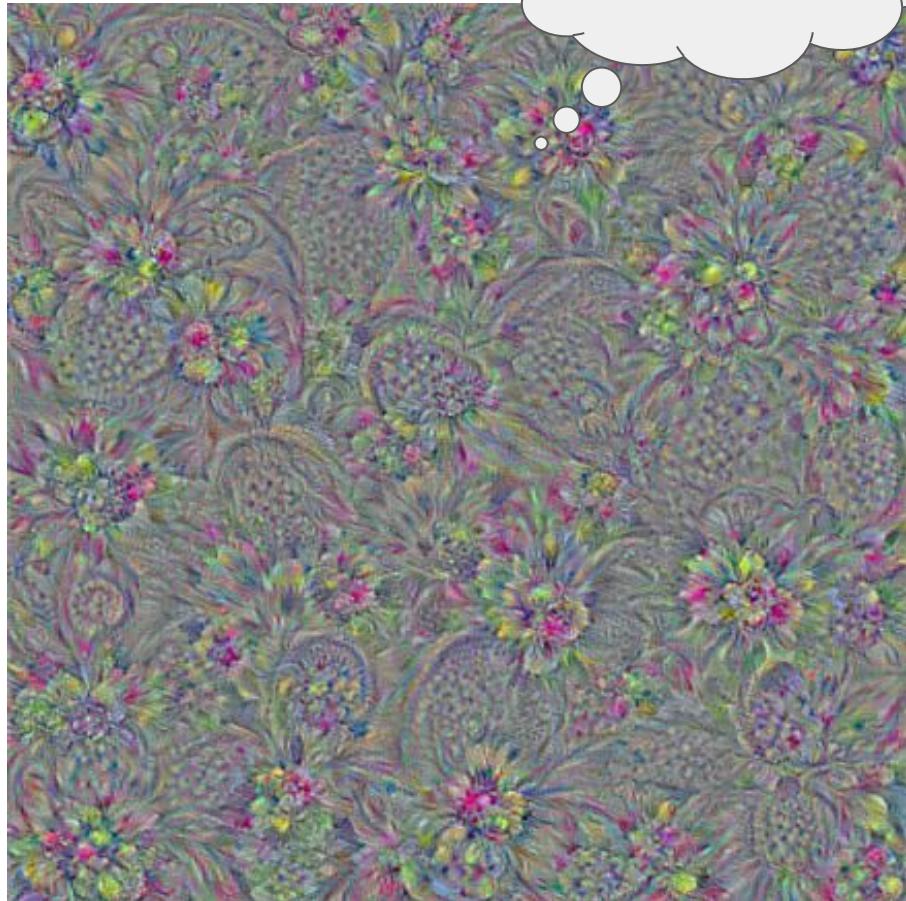
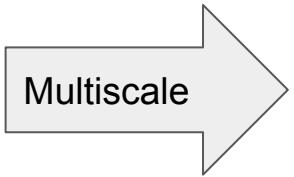
DeepDream: Super-resolution gone wrong



Image space gradient ascent

mixed4d_3x3_bottleneck_pre_relu,
channel 139

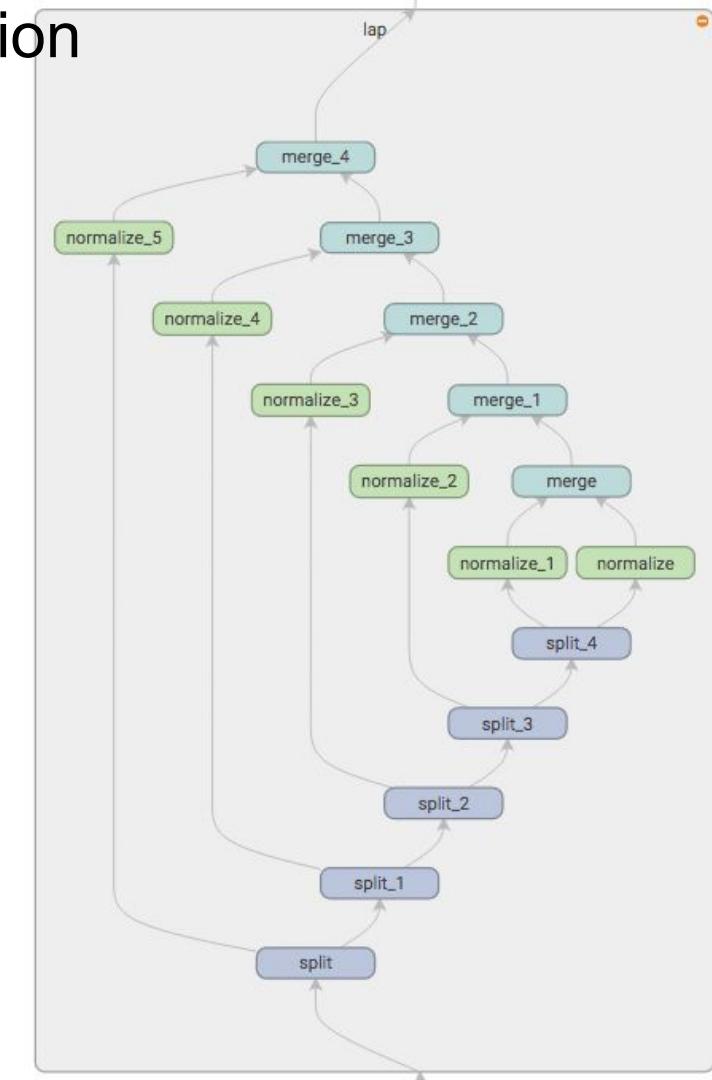
naive



Lacks low
frequencies...

Laplacian Pyramid Gradient Normalization

[TensorFlow code](#)

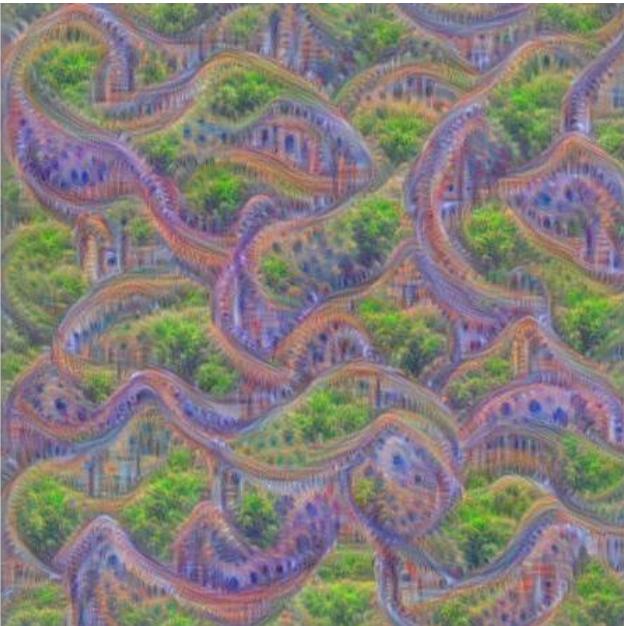


Feature combinations

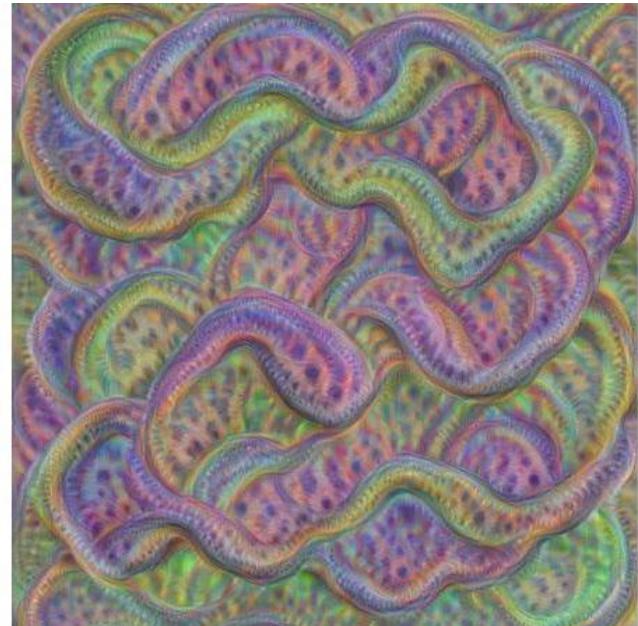
A



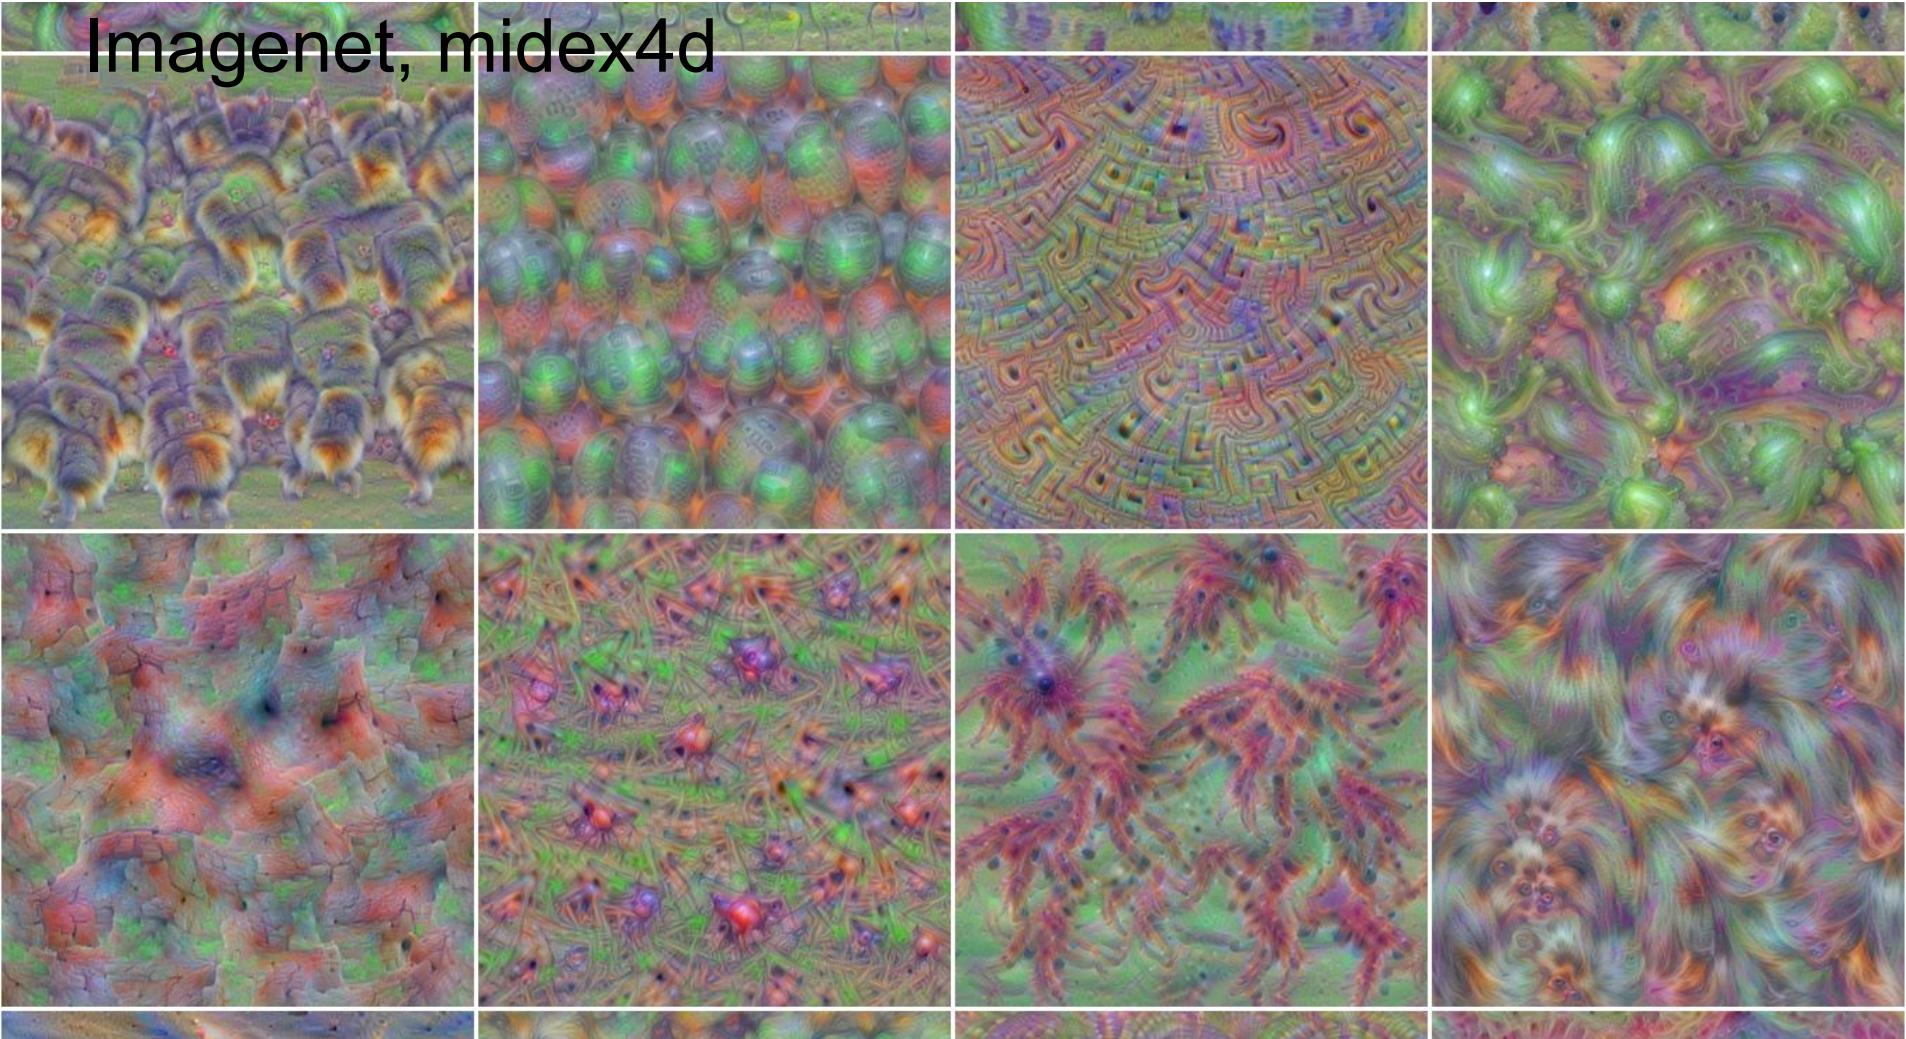
A+B



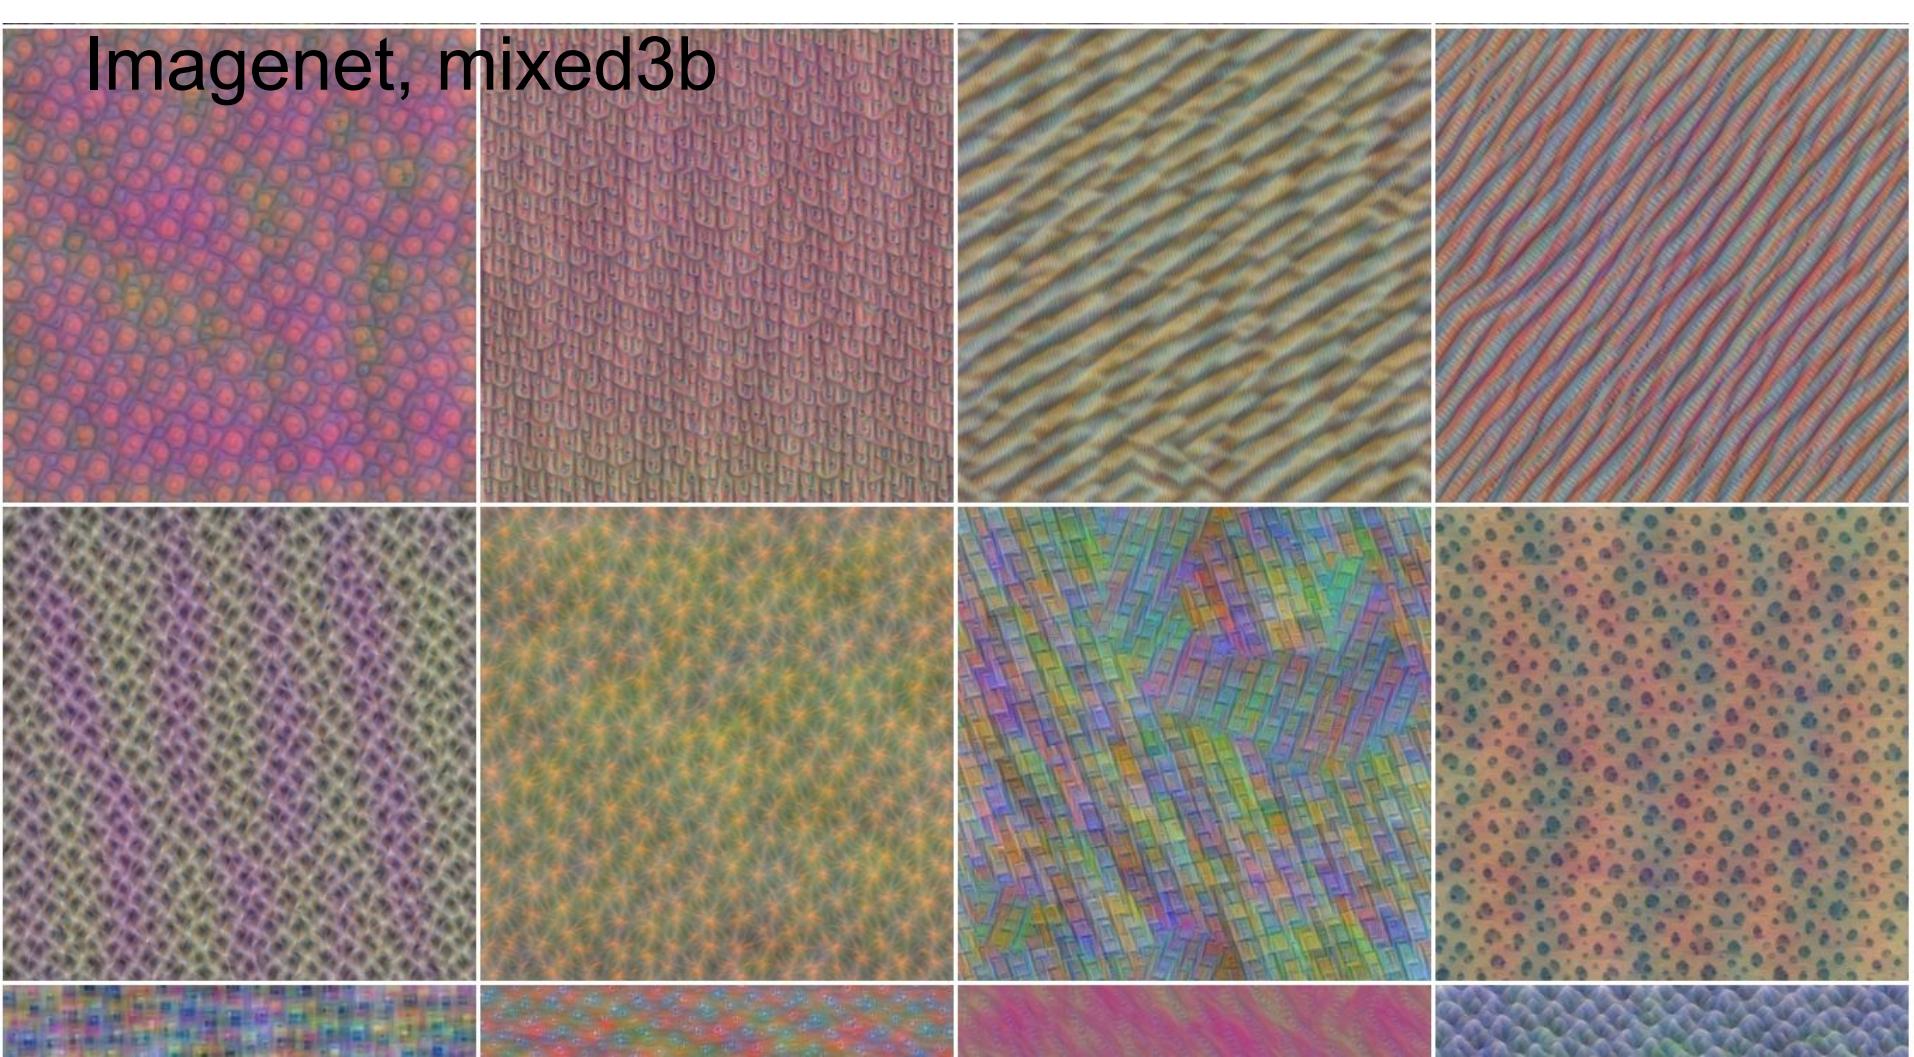
B



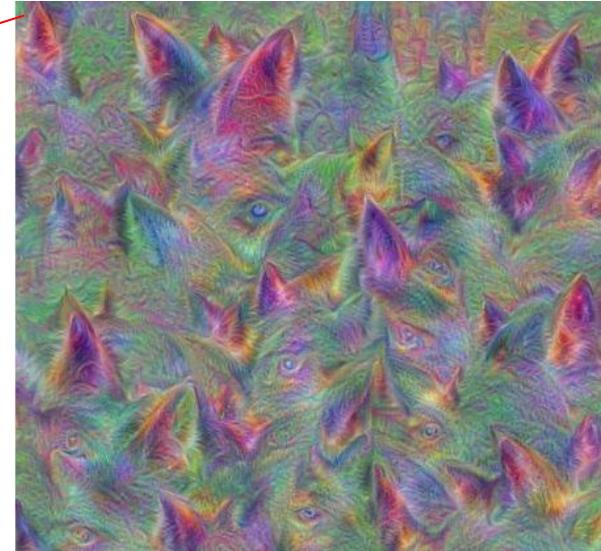
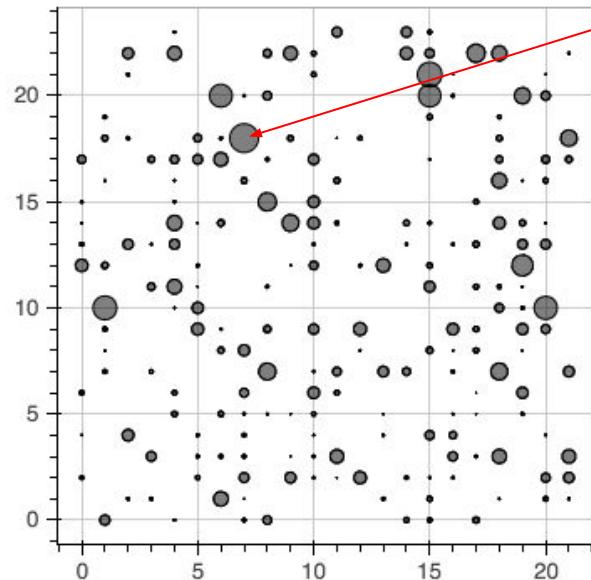
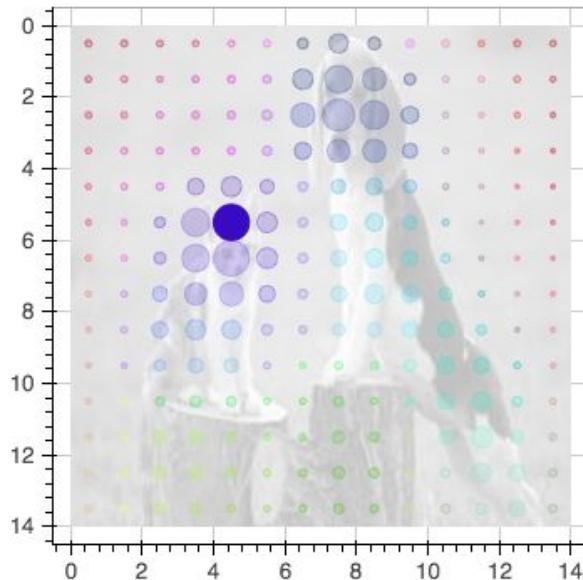
Imagenet, midex4d



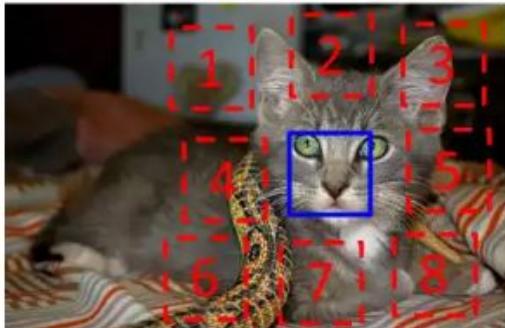
Imagenet, mixed3b



Visualizing internal activations



UNSUPERVISED VISUAL REPRESENTATION LEARNING BY CONTEXT PREDICTION

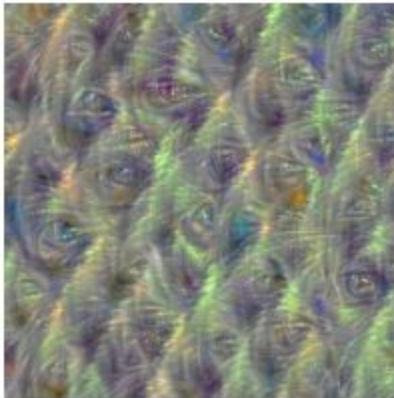
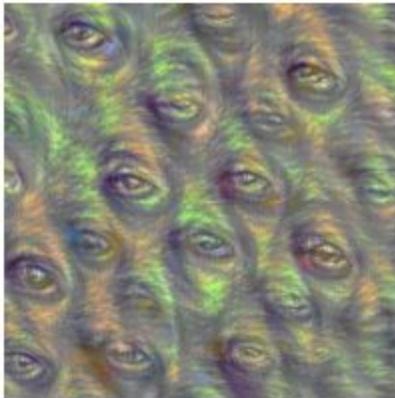


Presented at [ICCV 2015](#)

[Full gallery](#)

People

Carl Doersch
Abhinav Gupta
Alexei A. Efros



Colorful Image Colorization

Richard Zhang

Phillip Isola

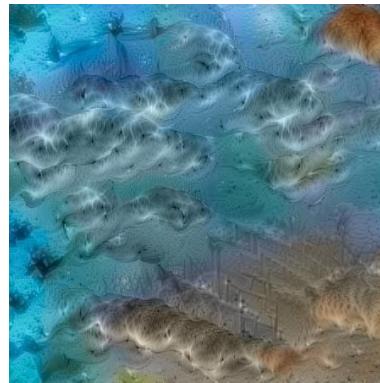
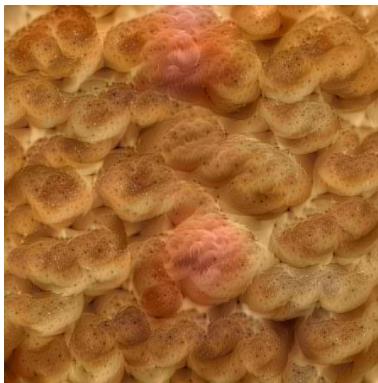
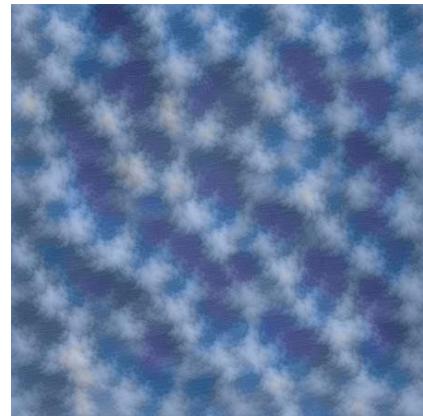
Alexei A. Efros

[Demo]

[GitHub]

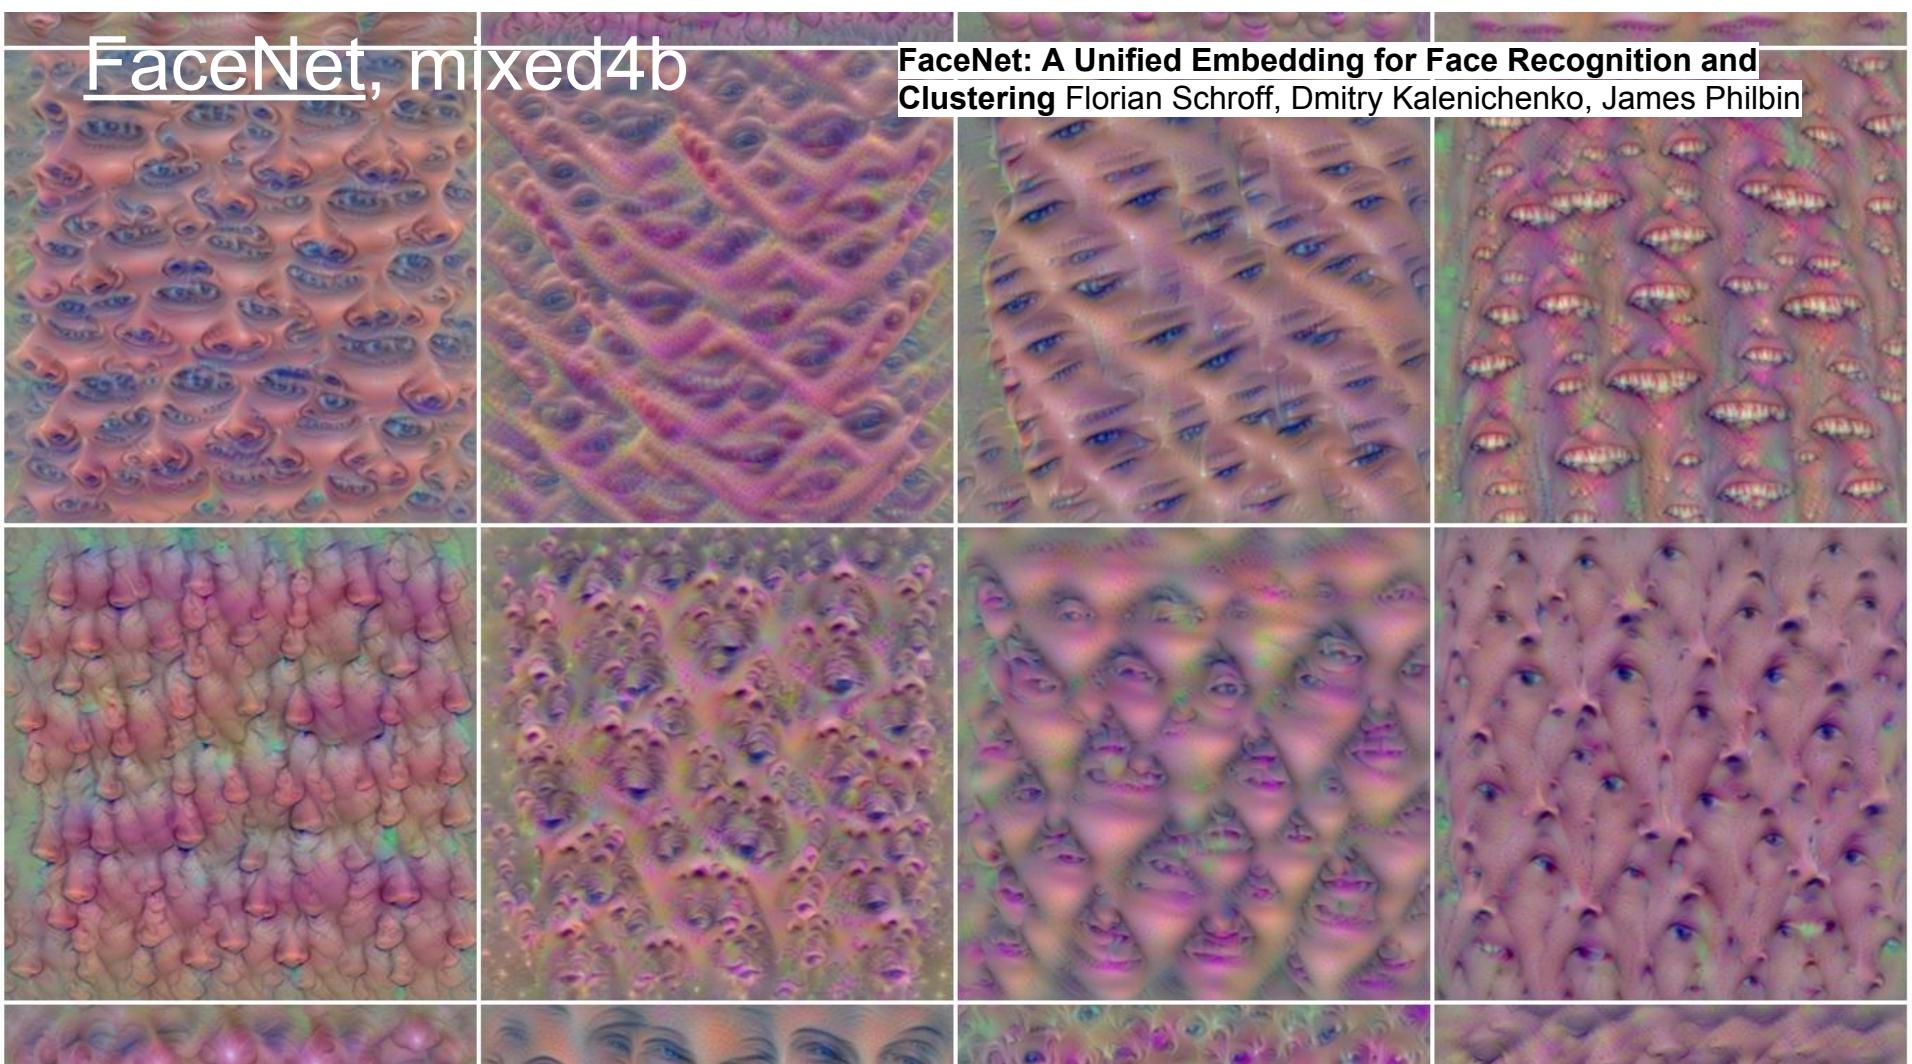
[Slides]

[Paper]



FaceNet, mixed4b

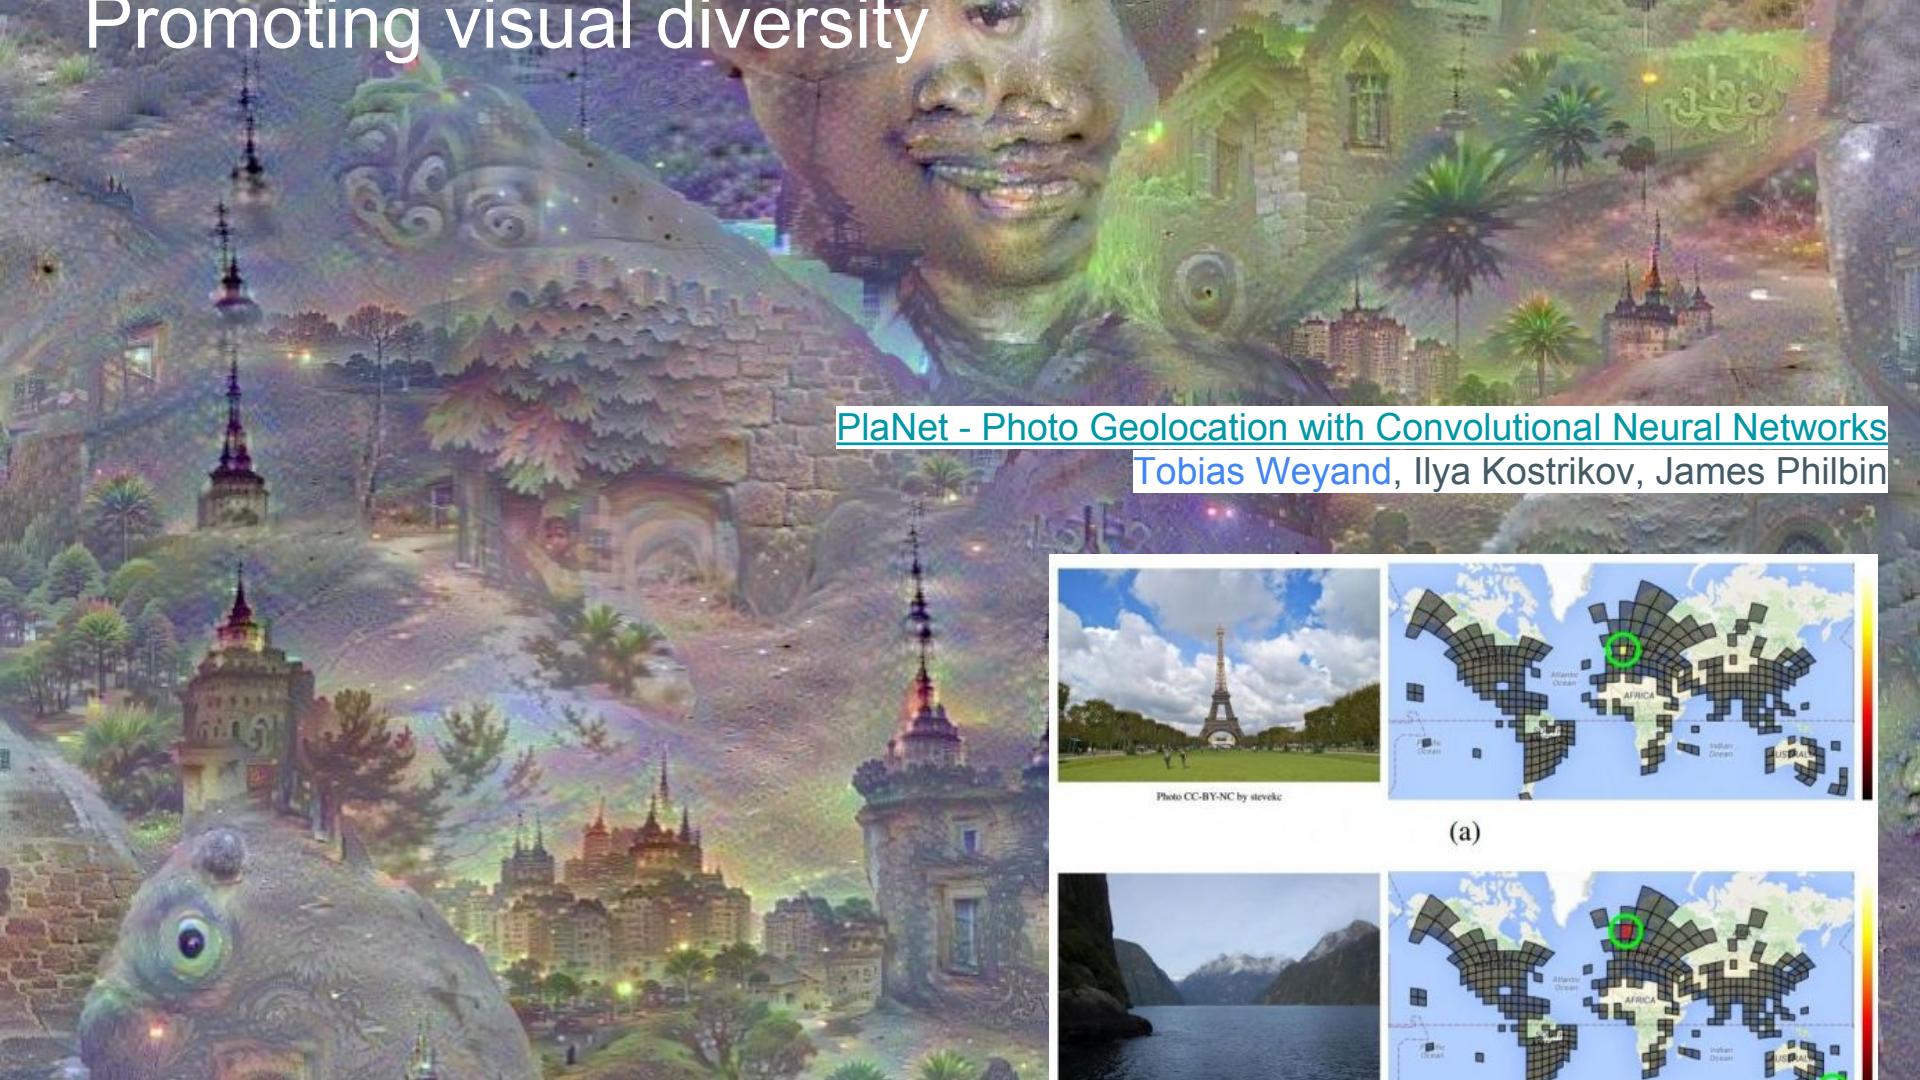
FaceNet: A Unified Embedding for Face Recognition and Clustering Florian Schroff, Dmitry Kalenichenko, James Philbin



Face emb-128 visualization (with sandler@)



Promoting visual diversity



[PlaNet - Photo Geolocation with Convolutional Neural Networks](#)
Tobias Weyand, Ilya Kostrikov, James Philbin



Photo CC-BY-NC by stevekc



(a)



