

Lecture 15:

Adversarial examples, texture synthesis, and style transfer

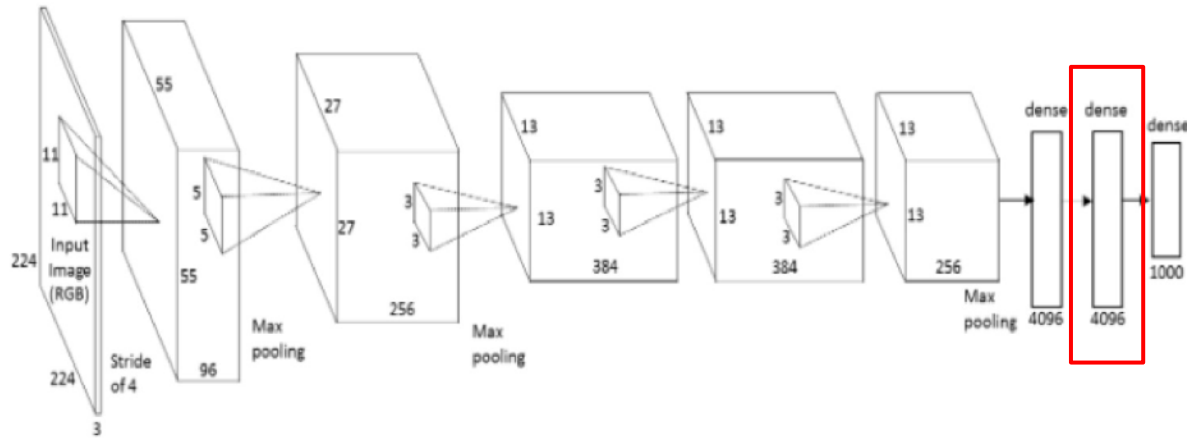
Agenda

- Recap
- Adversarial examples
- Texture synthesis and style transfer
- Bonus

Last lecture: Understanding ConvNets

- Visualize the weights
- Visualize the last layer (via t-SNE)
- Visualize patches that maximally activate neurons
- Occlusion experiments
- Deconv approaches (single backward pass)
- Optimization over image approaches (optimization)

Question: Given a CNN **code**, is it possible to reconstruct the original image?



Find an image such that:

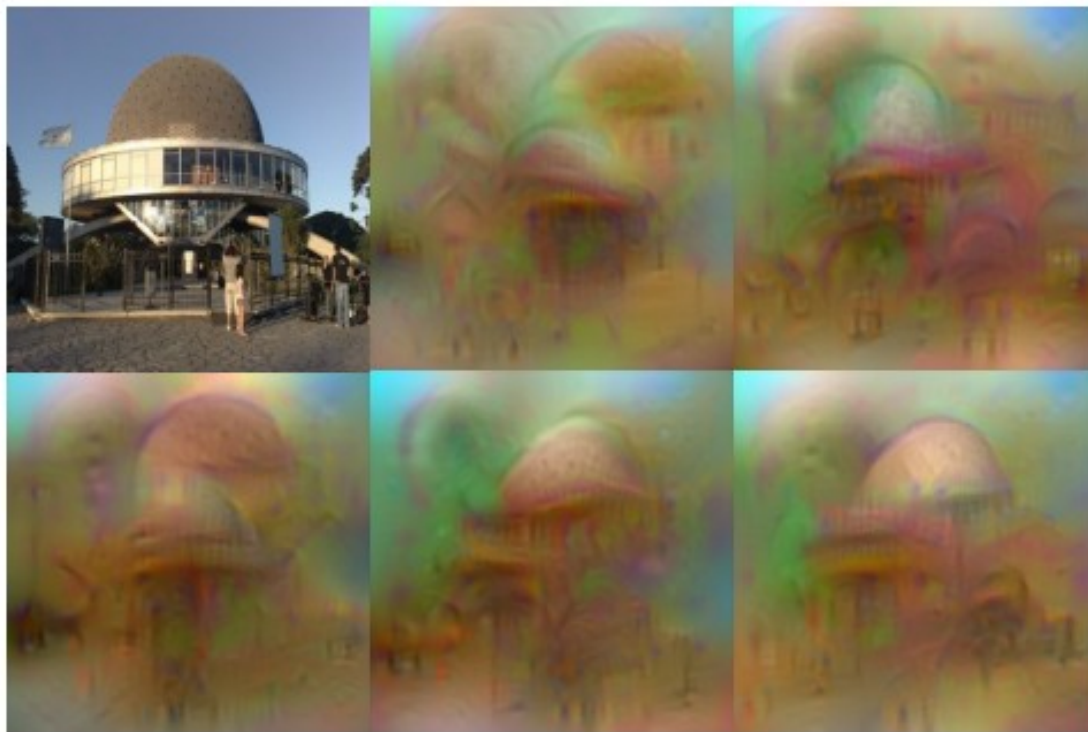
- Its code is similar to a given code
- It “looks natural” (image prior regularization)

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

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

Understanding Deep Image Representations by Inverting Them *[Mahendran and Vedaldi, 2014]*

original image

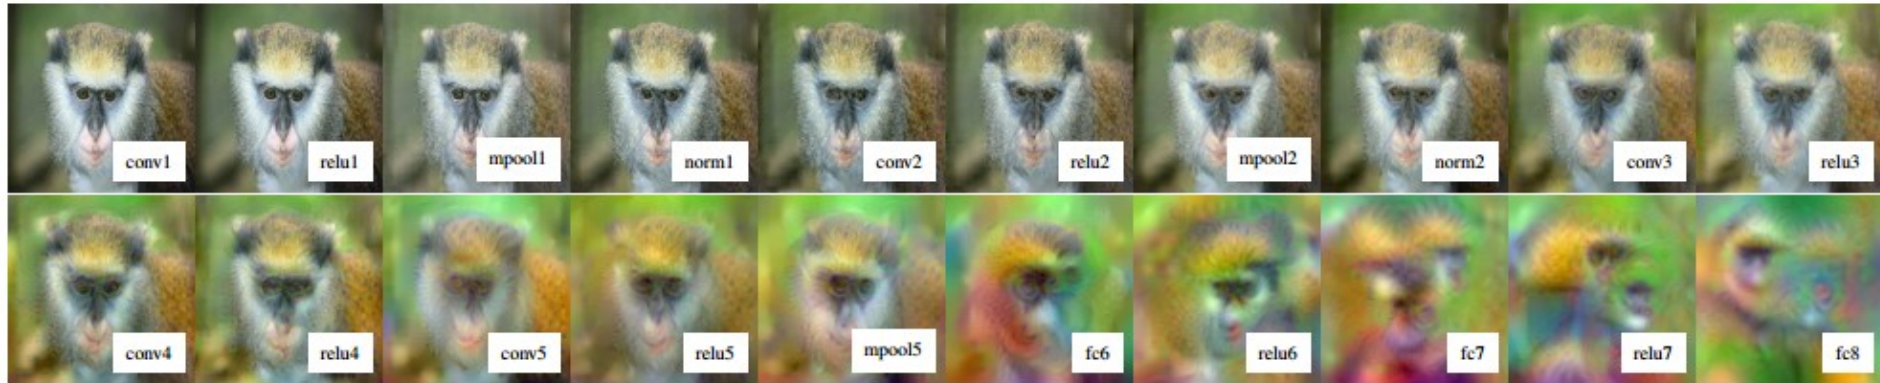


reconstructions
from the 1000
log probabilities
for ImageNet
(ILSVRC)
classes

Reconstructions from the representation after last last pooling layer (immediately before the first Fully Connected layer)



Reconstructions from intermediate layers



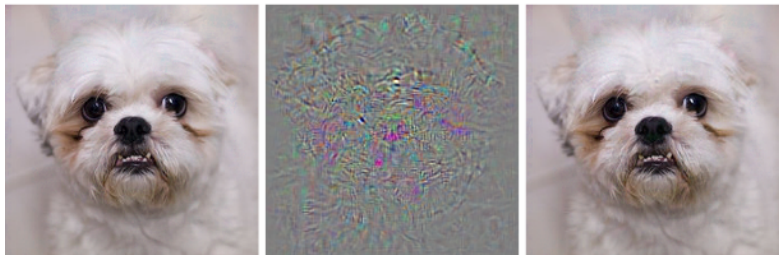
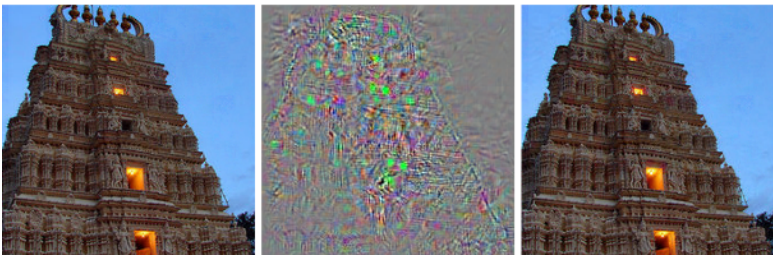
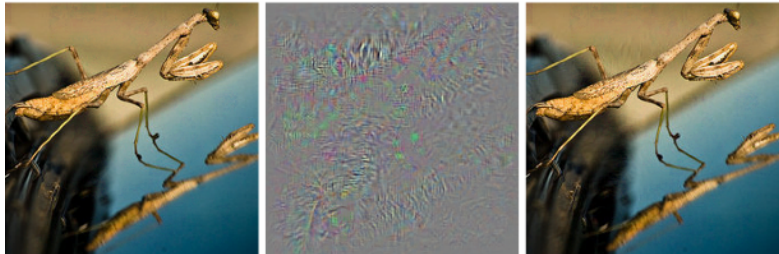
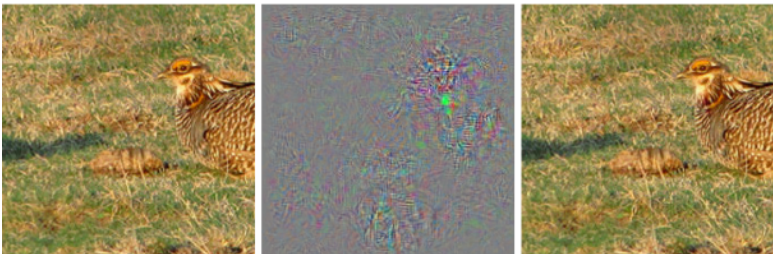
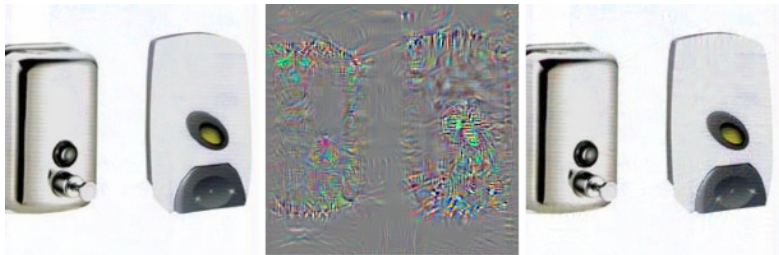
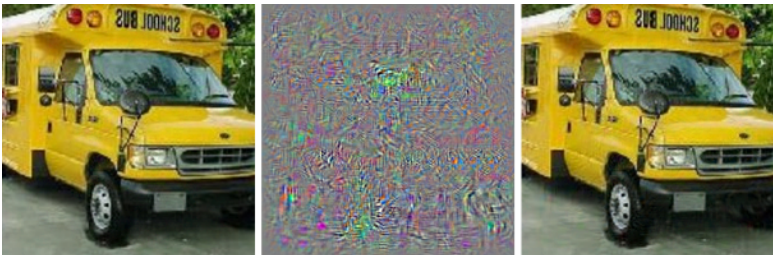
We can pose an optimization over the input image to maximize any class score.
That seems useful.

Question: Can we use this to “fool” ConvNets?

spoiler alert: yeah

- (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

[Intriguing properties of neural networks, Szegedy et al., 2013]



correct

+distort

ostrich

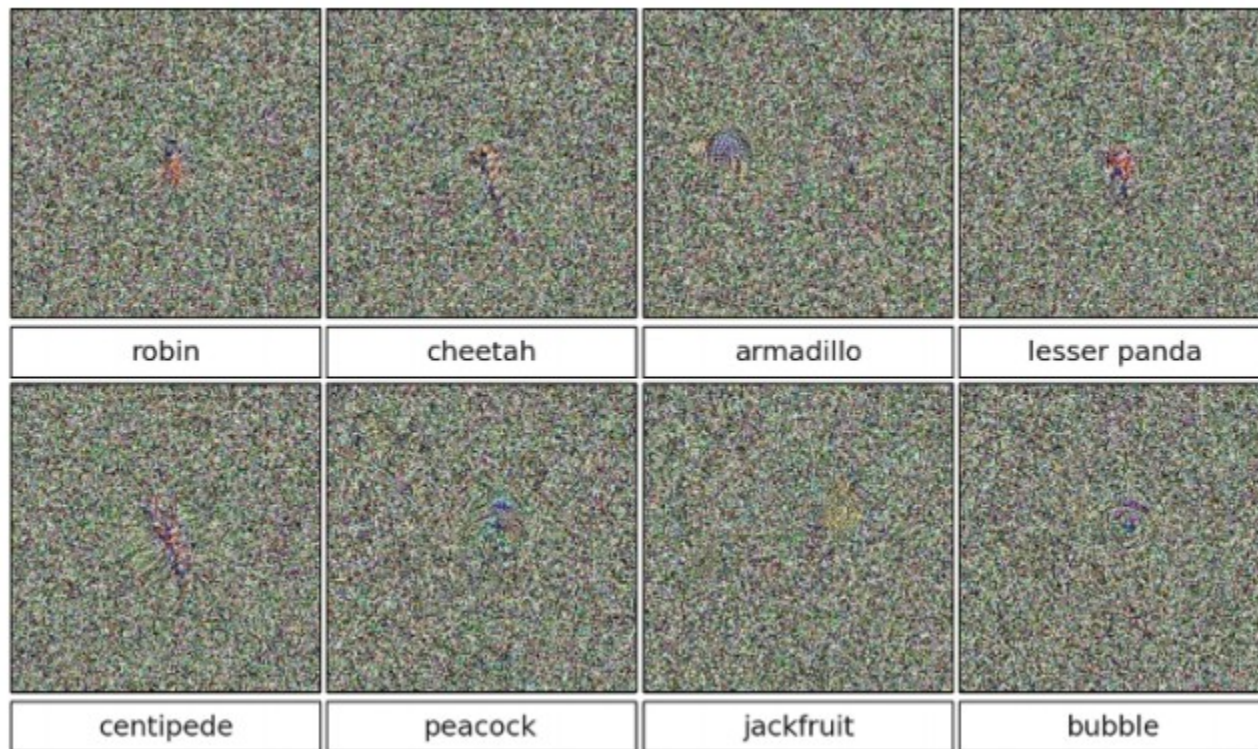
correct

+distort

ostrich

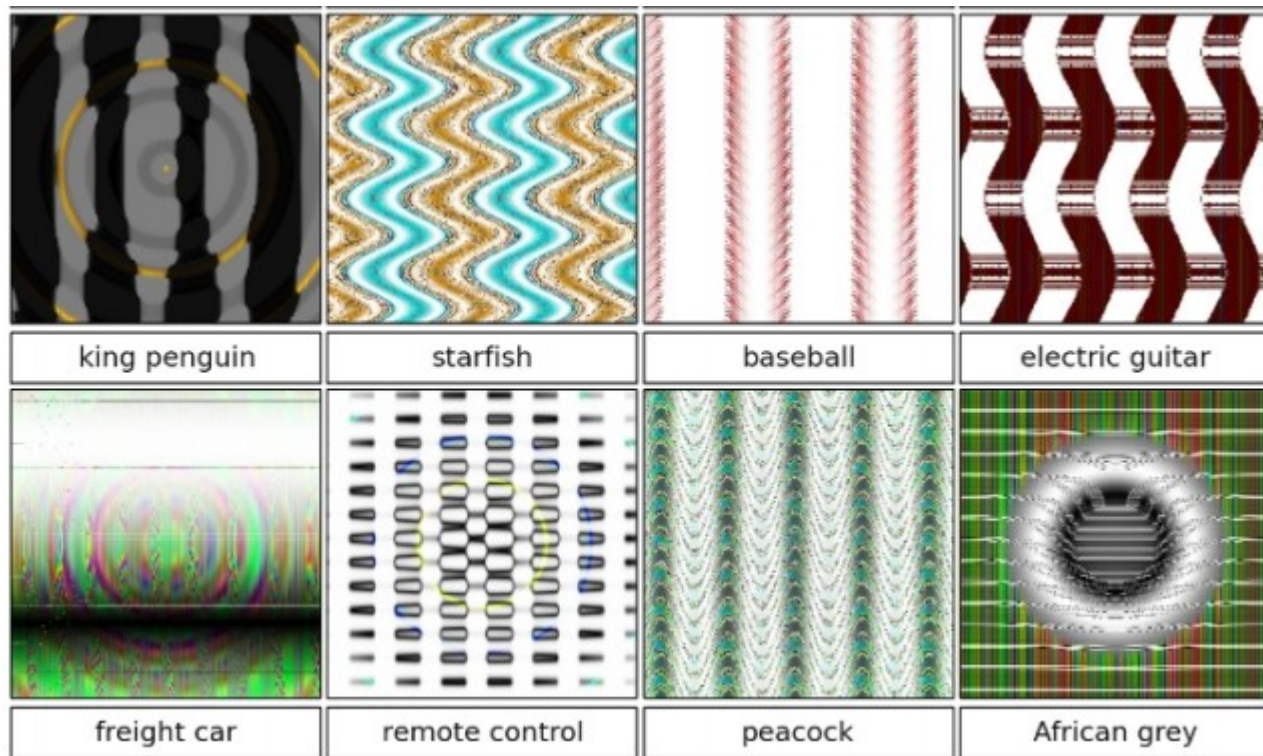
*[Deep Neural Networks are Easily Fooled: High Confidence Predictions for Unrecognizable Images
Nguyen, Yosinski, Clune, 2014]*

>99.6%
confidences

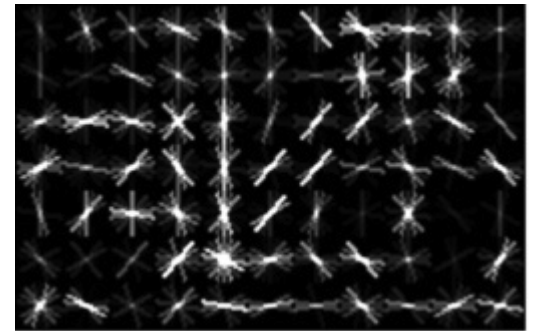


*[Deep Neural Networks are Easily Fooled: High Confidence Predictions for Unrecognizable Images
Nguyen, Yosinski, Clune, 2014]*

>99.6%
confidences



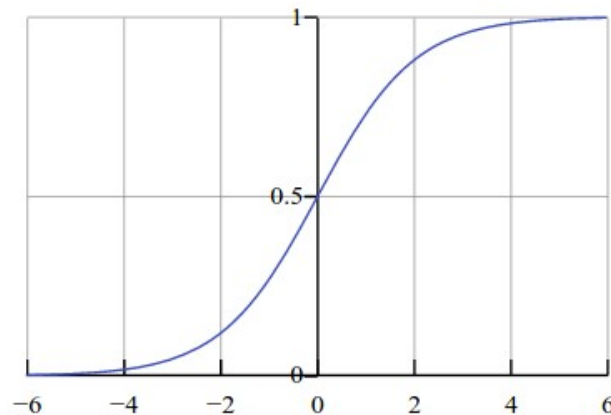
These kinds of results were around even before ConvNets...
[Exploring the Representation Capabilities of the HOG Descriptor, Tatu et al., 2011]



Identical HOG representation

Lets fool a binary linear classifier: (logistic regression)

$$P(y = 1 \mid x; w, b) = \frac{1}{1 + e^{-(w^T x + b)}} = \sigma(w^T x + b)$$



Since the probabilities of class 1 and 0 sum to one, the probability for class 0 is $P(y = 0 \mid x; w, b) = 1 - P(y = 1 \mid x; w, b)$. Hence, an example is classified as a positive example ($y = 1$) if $\sigma(w^T x + b) > 0.5$, or equivalently if the score $w^T x + b > 0$.

Lets fool a binary linear classifier:

X	2	-1	3	-2	2	2	1	-4	5	1	← input example
W	-1	-1	1	-1	1	-1	1	1	-1	1	← weights

$$P(y = 1 | x; w, b) = \frac{1}{1 + e^{-(w^T x + b)}} = \sigma(w^T x + b)$$

Lets fool a binary linear classifier:

X	2	-1	3	-2	2	2	1	-4	5	1	← input example
W	-1	-1	1	-1	1	-1	1	1	-1	1	← weights

class 1 score = dot product:

$$= -2 + 1 + 3 + 2 + 2 - 2 + 1 - 4 - 5 + 1 = -3$$

=> probability of class 1 is $1/(1+e^{(-(-3))}) = 0.0474$

i.e. the classifier is **95%** certain that this is class 0 example.

$$P(y = 1 | x; w, b) = \frac{1}{1 + e^{-(w^T x + b)}} = \sigma(w^T x + b)$$

Lets fool a binary linear classifier:

X	2	-1	3	-2	2	2	1	-4	5	1	← input example
W	-1	-1	1	-1	1	-1	1	1	-1	1	← weights
adversarial x	?	?	?	?	?	?	?	?	?	?	

class 1 score = dot product:

$$= -2 + 1 + 3 + 2 + 2 - 2 + 1 - 4 - 5 + 1 = -3$$

$$\Rightarrow \text{probability of class 1 is } 1/(1+e^{(-(-3))}) = 0.0474$$

i.e. the classifier is **95%** certain that this is class 0 example.

$$P(y = 1 | x; w, b) = \frac{1}{1 + e^{-(w^T x + b)}} = \sigma(w^T x + b)$$

Lets fool a binary linear classifier:

X	2	-1	3	-2	2	2	1	-4	5	1	← input example
W	-1	-1	1	-1	1	-1	1	1	-1	1	← weights
adversarial x	1.5	-1.5	3.5	-2.5	2.5	1.5	1.5	-3.5	4.5	1.5	

class 1 score before:

$$-2 + 1 + 3 + 2 + 2 - 2 + 1 - 4 - 5 + 1 = -3$$

$$\Rightarrow \text{probability of class 1 is } 1/(1+e^{(-(-3))}) = 0.0474$$

$$-1.5+1.5+3.5+2.5+2.5-1.5+1.5-3.5-4.5+1.5 = 2$$

$$\Rightarrow \text{probability of class 1 is now } 1/(1+e^{(-(-2))}) = 0.88$$

i.e. we improved the class 1 probability from 5% to 88%

$$P(y = 1 | x; w, b) = \frac{1}{1 + e^{-(w^T x + b)}} = \sigma(w^T x + b)$$

Lets fool a binary linear classifier:

X	2	-1	3	-2	2	2	1	-4	5	1	← input example
W	-1	-1	1	-1	1	-1	1	1	-1	1	← weights
adversarial x	1.5	-1.5	3.5	-2.5	2.5	1.5	1.5	-3.5	4.5	1.5	

class 1 score before:

$$-2 + 1 + 3 + 2 + 2 - 2 + 1 - 4 - 5 + 1 = -3$$

$$\Rightarrow \text{probability of class 1 is } 1/(1+e^{(-(-3))}) = 0.0474$$

$$-1.5+1.5+3.5+2.5+2.5-1.5+1.5-3.5-4.5+1.5 = 2$$

$$\Rightarrow \text{probability of class 1 is now } 1/(1+e^{(-(-2))}) = 0.88$$

i.e. we improved the class 1 probability from 5% to 88%

This was only with 10 input dimensions. A 224x224 input image has 150,528.

(It's significantly easier with more numbers, need smaller nudge for each)

Adversarial examples

x
 $y = \text{"panda"}$
w/ 57.7% confidence

+ .007 ×

$\text{sign}(\nabla_x J(\theta, x, y))$
"nematode"
w/ 8.2% confidence

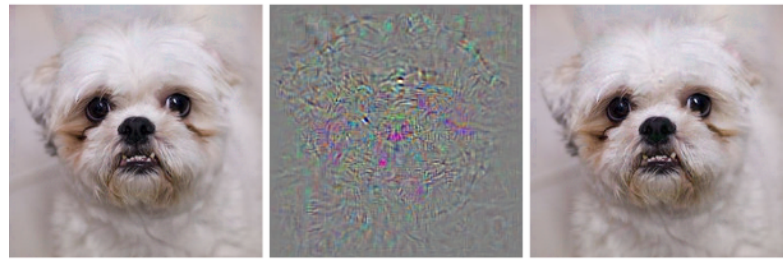
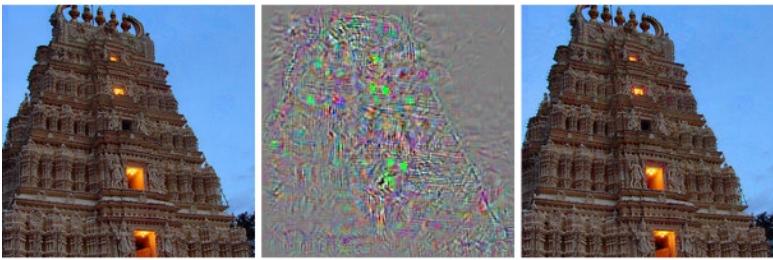
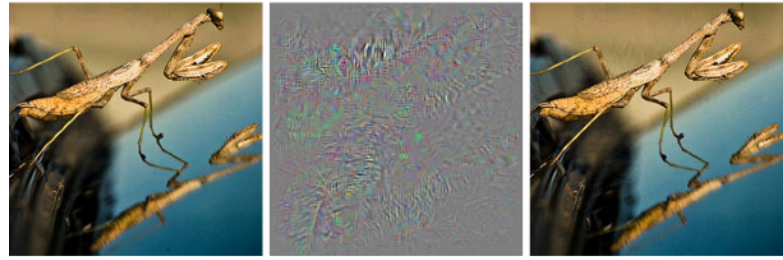
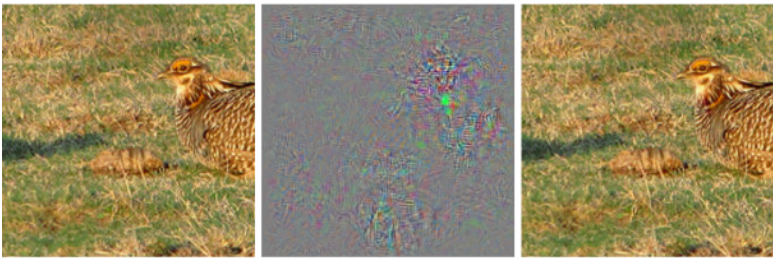
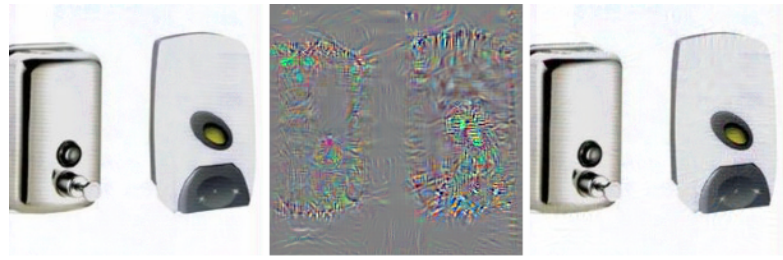
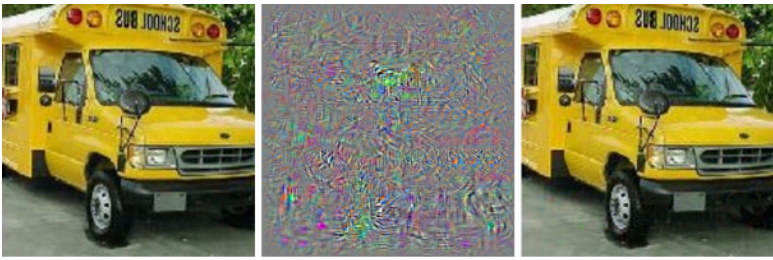
=

$x + \epsilon \cdot \text{sign}(\nabla_x J(\theta, x, y))$
"gibbon"
w/ 99.3 % confidence

Explaining and Harnessing Adversarial Examples

Ian J. Goodfellow, Jonathon Shlens, Christian Szegedy

[Intriguing properties of neural networks, Szegedy et al., 2013]



correct

+distort

ostrich

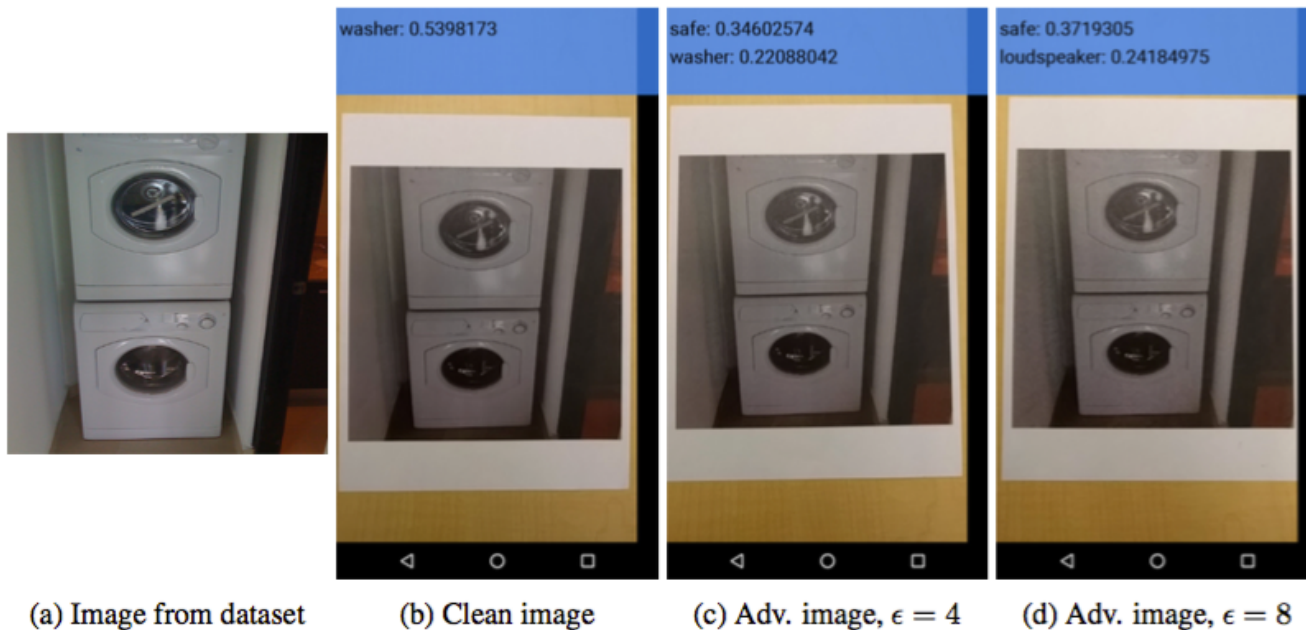
correct

+distort

ostrich

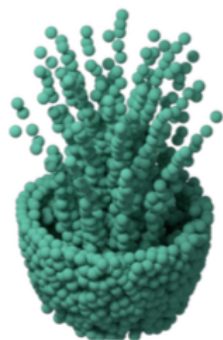
Adversarial examples

Can be printed on paper! Kurakin et al., 17

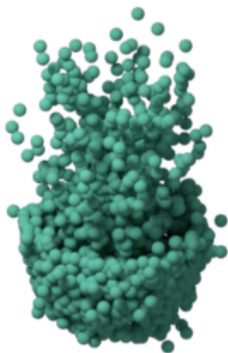


Adversarial examples

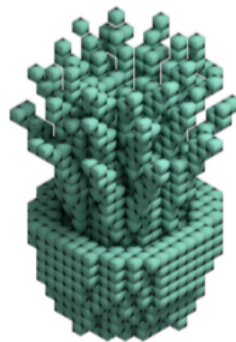
Also works for 3D models!
(though a little harder for point clouds)



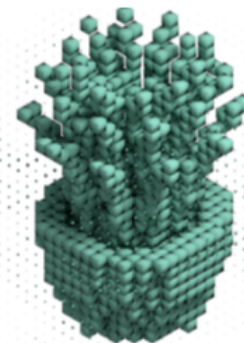
“plant”



“bench”



“plant”



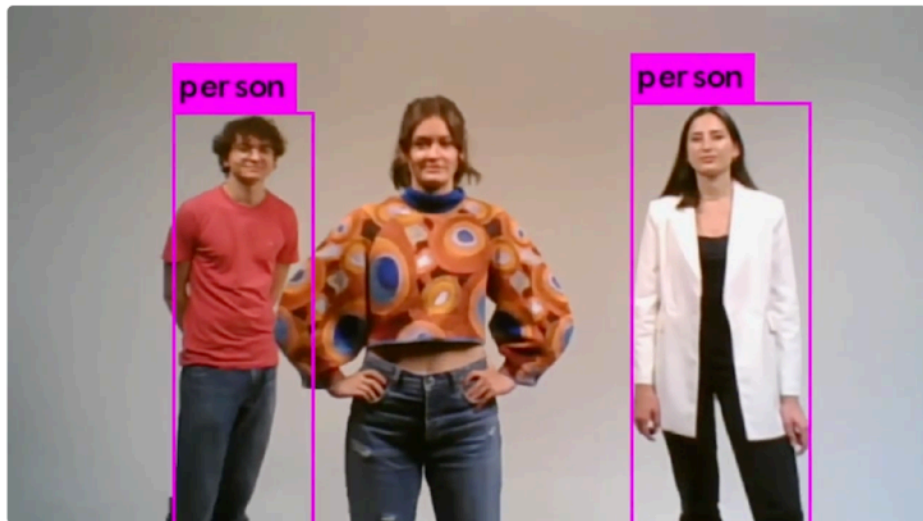
“bench”

point cloud

voxel

Su et al., ECCV 2018

Can use this to evade AI too!



AI Camouflage for Data Privacy | Rachele Didero | MIT 2023



AI Camouflage for Data Privacy | Rachele Didero | MIT 2023

AI Camouflage for Data Privacy | Rachele Didero

https://www.youtube.com/watch?v=1UEQGBL_q4U

<https://www.capable.design/> — AI clothing for data privacy

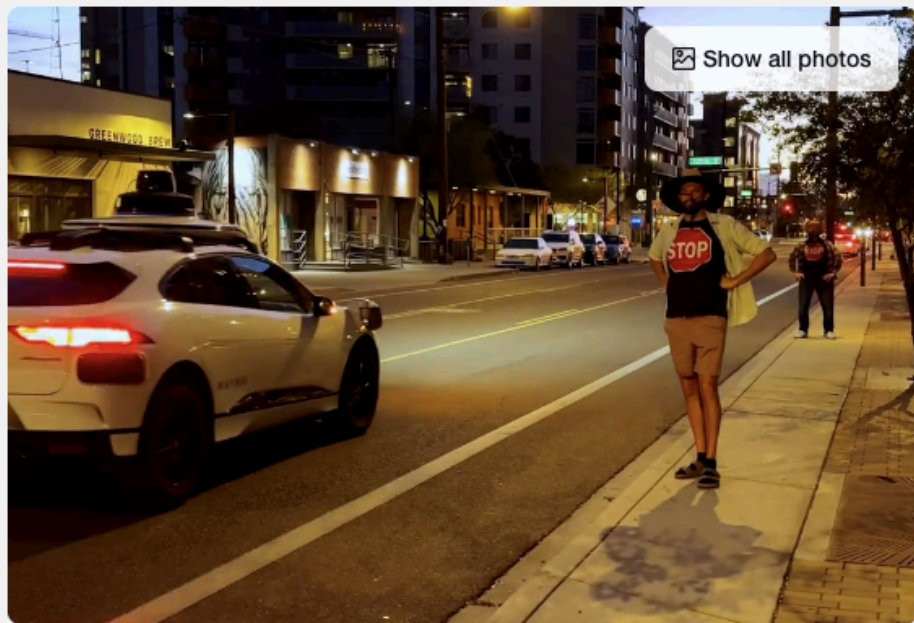
How a t-shirt stopped this autonomous car in its tracks

Autonomous cars are being touted as the future of travel with safety as a key focus – but an American has shown that in one way, their systems may be too sensitive.

🕒 6 May 2024 at 9:30 pm 💬 22 Comments



Jordan Mulach
Contributor



Neural Style Transfer and Texture Synthesis

Content Image



[This image](#) is licensed under [CC-BY 3.0](#)

+

Style Image



[Starry Night](#) by Van Gogh is in the public domain

=

Style Transfer!

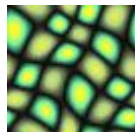


[This image](#) copyright Justin Johnson, 2015. Reproduced with permission.

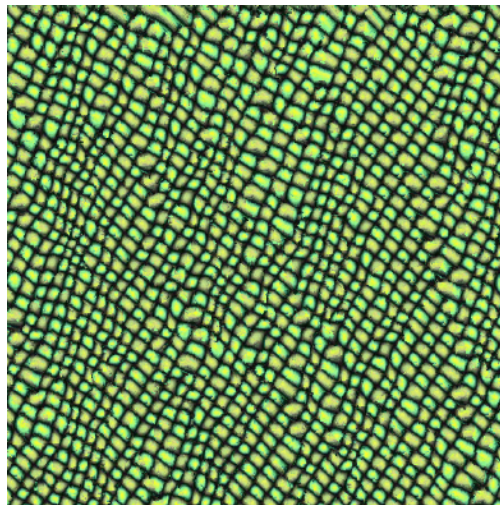
Gatys, Ecker, and Bethge, "Image style transfer using convolutional neural networks", CVPR 2016

Texture Synthesis

Given a sample patch of some texture, can we generate a bigger image of the same texture?



Input

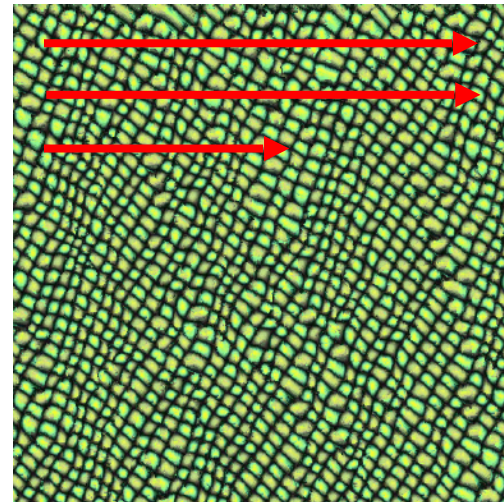
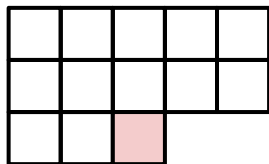
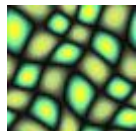


Output

[Output image](#) is licensed under the [MIT license](#)

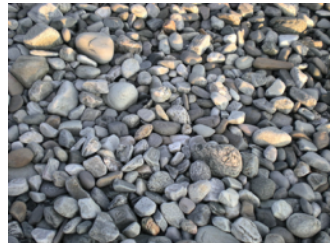
Texture Synthesis: Nearest Neighbor

Generate pixels one at a time in scanline order; form neighborhood of already generated pixels and copy nearest neighbor from input

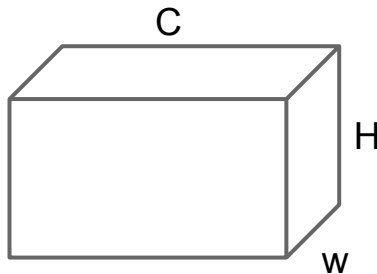
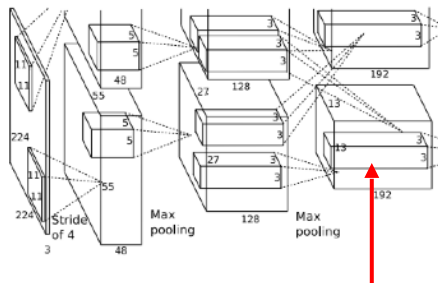


Wei and Levoy, "Fast Texture Synthesis using Tree-structured Vector Quantization", SIGGRAPH 2000
Efros and Leung, "Texture Synthesis by Non-parametric Sampling", ICCV 1999

Neural Texture Synthesis: Gram Matrix

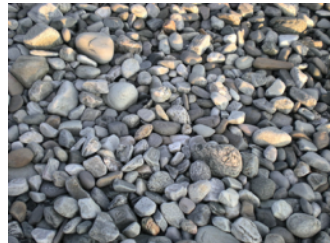


[This image](#) is in the public domain.

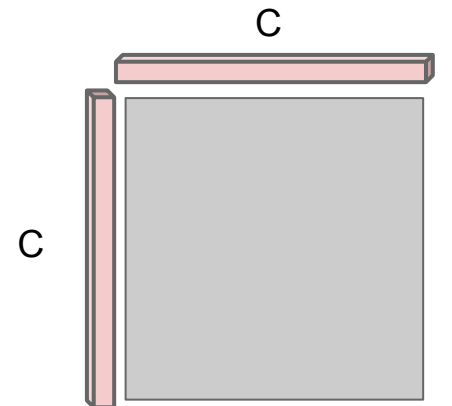
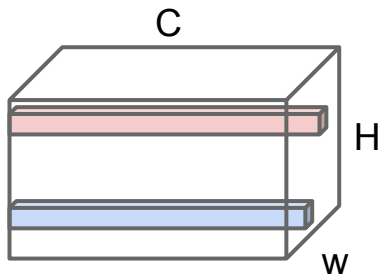
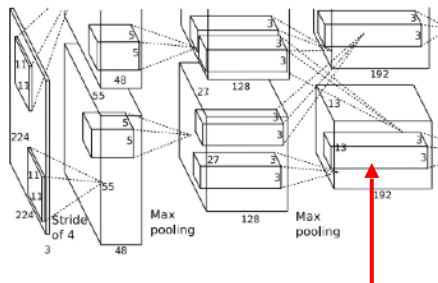


Each layer of CNN gives $C \times H \times W$ tensor of features; $H \times W$ grid of C -dimensional vectors

Neural Texture Synthesis: Gram Matrix

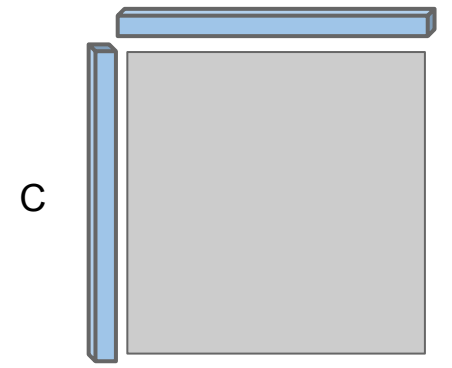


[This image](#) is in the public domain.

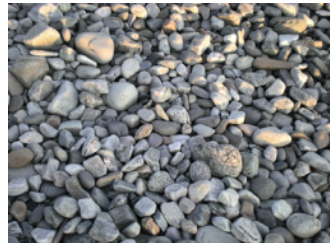


Each layer of CNN gives $C \times H \times W$ tensor of features; $H \times W$ grid of C -dimensional vectors

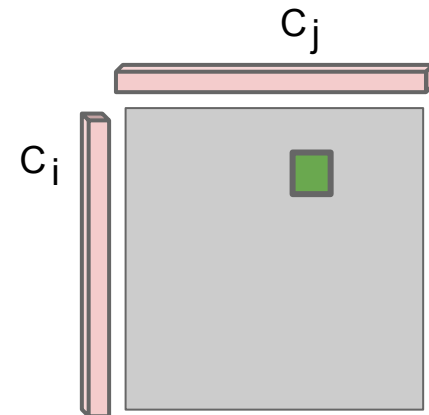
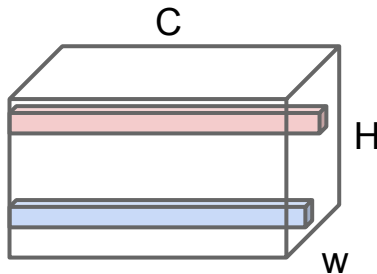
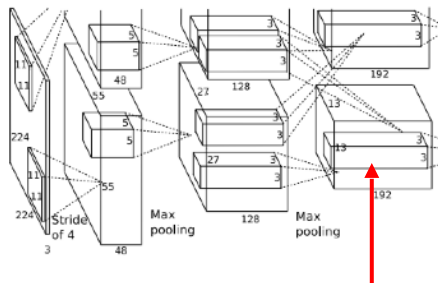
Outer product of C -dimensional vector with itself gives $C \times C$ matrix measuring co-occurrence



Neural Texture Synthesis: Gram Matrix



[This image](#) is in the public domain.



Each layer of CNN gives $C \times H \times W$ tensor of features; $H \times W$ grid of C -dimensional vectors

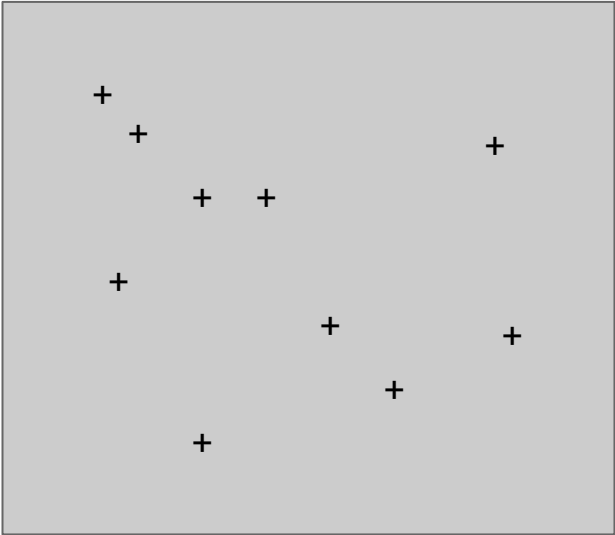
Outer product of C -dimensional vector with itself gives $C \times C$ matrix measuring co-occurrence

■ The green box $G(i,j)$ represents the AVERAGE, over image positions in the image, of the product of features i and j .

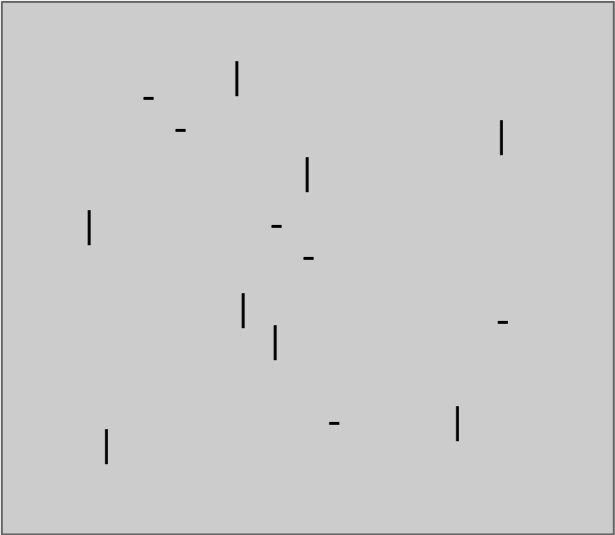
Suppose “ i ” represents horizontal lines and “ j ” represents vertical lines.

Gram Matrix captures co-occurrences

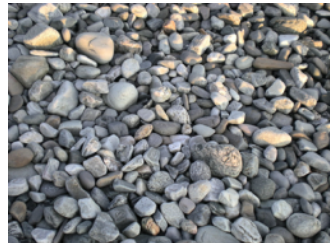
Vertical and horizontal co-occur



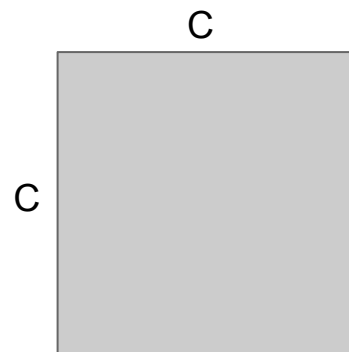
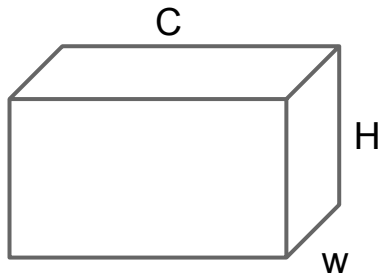
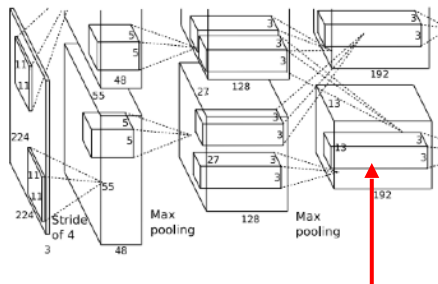
Vertical and horizontal DO NOT co-occur



Neural Texture Synthesis: Gram Matrix



[This image](#) is in the public domain.



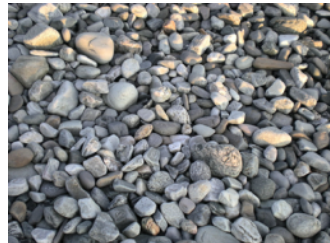
Gram Matrix

Each layer of CNN gives $C \times H \times W$ tensor of features; $H \times W$ grid of C -dimensional vectors

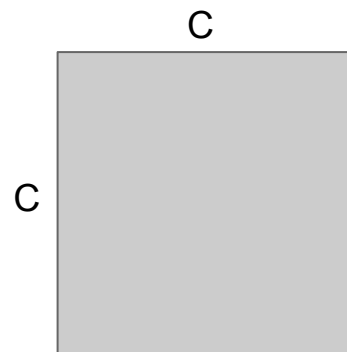
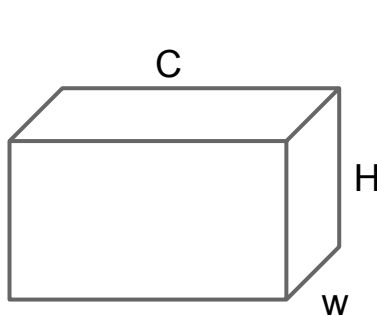
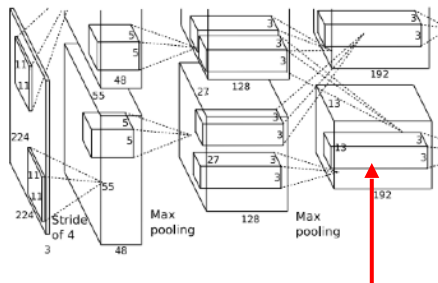
Outer product of C -dimensional vector with itself gives $C \times C$ matrix measuring co-occurrence

Average over all HW outer products, giving **Gram matrix** of shape $C \times C$

Neural Texture Synthesis: Gram Matrix



[This image](#) is in the public domain.



Each layer of CNN gives $C \times H \times W$ tensor of features; $H \times W$ grid of C -dimensional vectors

Outer product of C -dimensional vector with itself gives $C \times C$ matrix measuring co-occurrence

Average over all $H \times W$ outer products, giving **Gram matrix** of shape $C \times C$

Efficient to compute; reshape features from

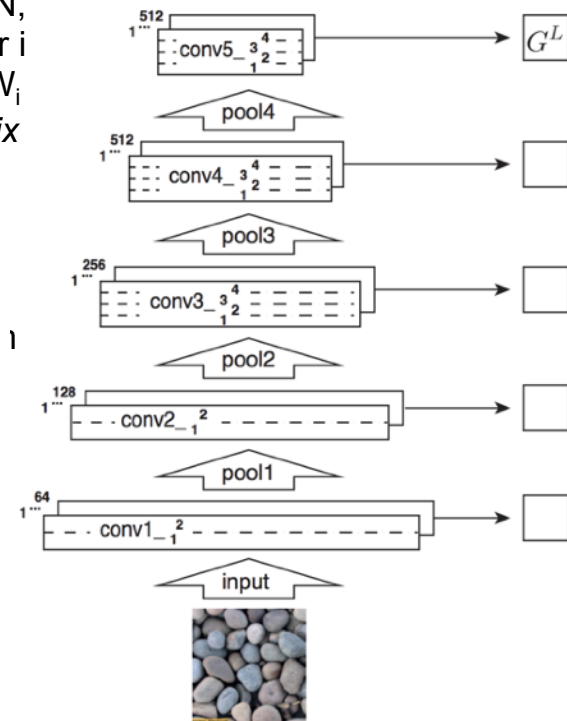
$C \times H \times W$ to $= C \times HW$

then compute $G = FF^T$

Neural Texture Synthesis

1. Pretrain a CNN on ImageNet (VGG-19)
2. Run input texture forward through CNN, record activations on every layer; layer i gives feature map of shape $C_i \times H_i \times W_i$
3. At each layer compute the *Gram matrix* giving outer product of features:

$$G_{ij}^l = \sum_k F_{ik}^l F_{jk}^l \quad (\text{shape } C_i \times C_j)$$



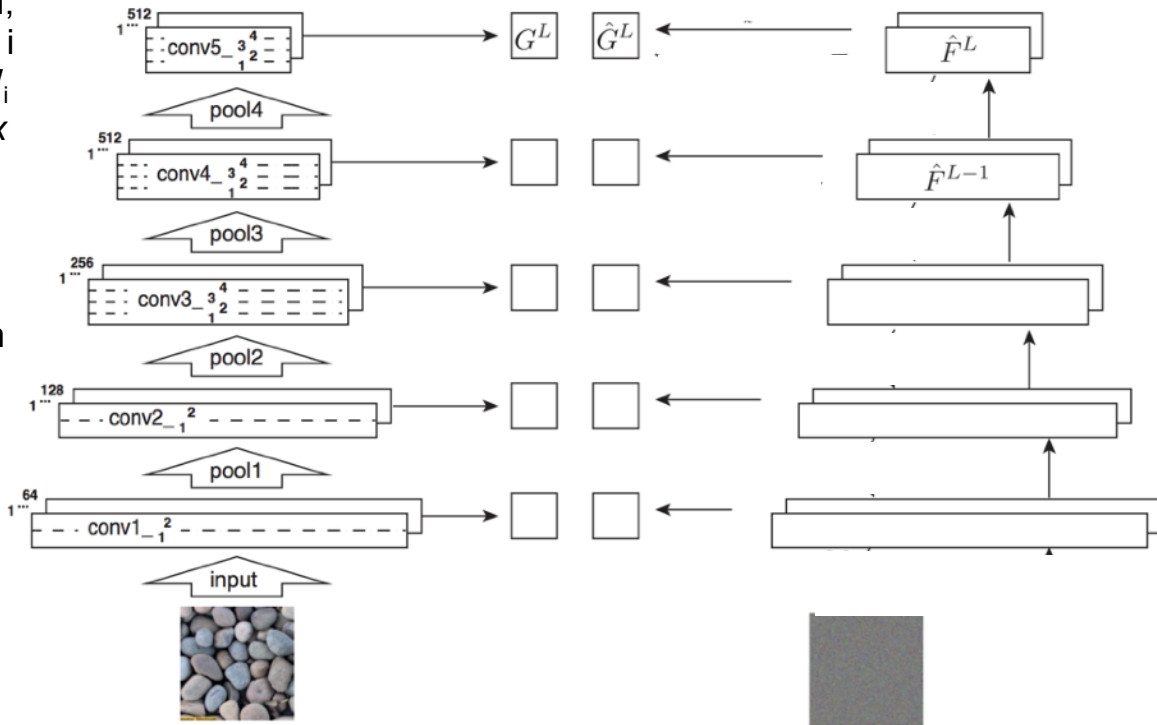
Gatys, Ecker, and Bethge, "Texture Synthesis Using Convolutional Neural Networks", NIPS 2015
Figure copyright Leon Gatys, Alexander S. Ecker, and Matthias Bethge, 2015. Reproduced with permission.

Neural Texture Synthesis

1. Pretrain a CNN on ImageNet (VGG-19)
2. Run input texture forward through CNN, record activations on every layer; layer i gives feature map of shape $C_i \times H_i \times W_i$
3. At each layer compute the *Gram matrix* giving outer product of features:

$$G_{ij}^l = \sum_k F_{ik}^l F_{jk}^l \quad (\text{shape } C_i \times C_i)$$

4. Initialize generated image from random noise
5. Pass generated image through CNN, compute Gram matrix on each layer



Gatys, Ecker, and Bethge, "Texture Synthesis Using Convolutional Neural Networks", NIPS 2015
 Figure copyright Leon Gatys, Alexander S. Ecker, and Matthias Bethge, 2015. Reproduced with permission.

Neural Texture Synthesis

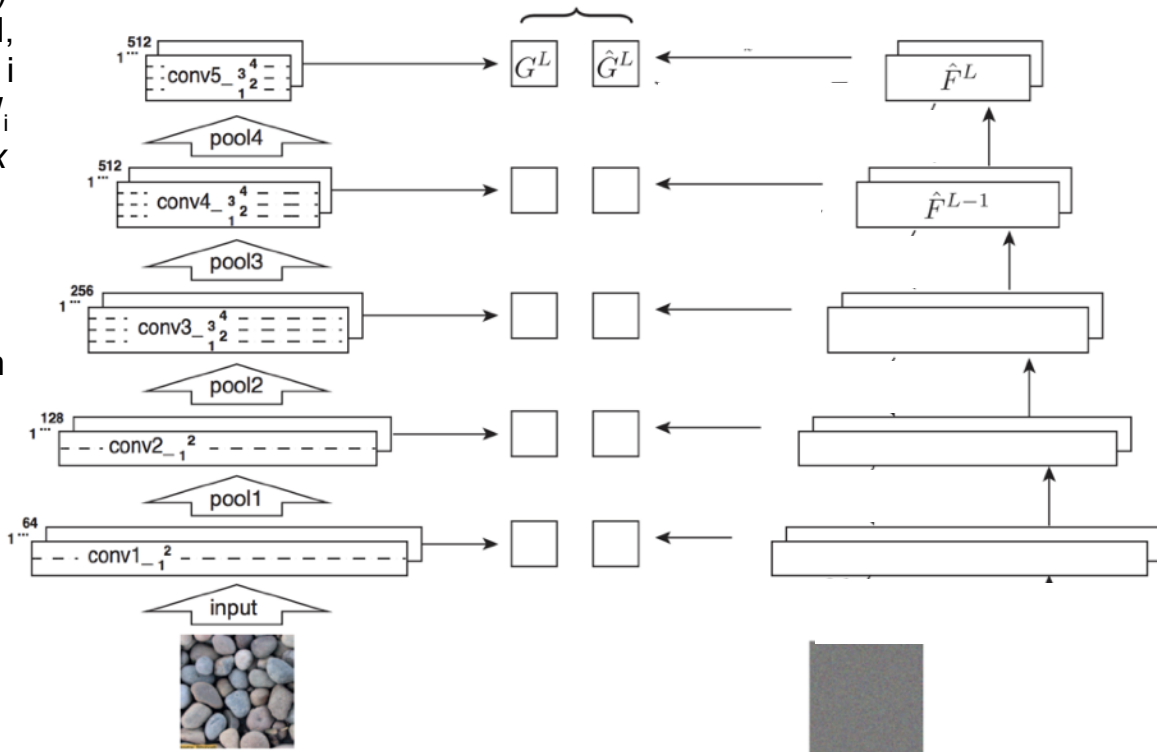
1. Pretrain a CNN on ImageNet (VGG-19)
2. Run input texture forward through CNN, record activations on every layer; layer i gives feature map of shape $C_i \times H_i \times W_i$
3. At each layer compute the *Gram matrix* giving outer product of features:

$$G_{ij}^l = \sum_k F_{ik}^l F_{jk}^l \quad (\text{shape } C_i \times C_i)$$

4. Initialize generated image from random noise
5. Pass generated image through CNN, compute Gram matrix on each layer
6. Compute loss: weighted sum of L2 distance between Gram matrices

$$E_l = \frac{1}{4N_l^2 M_l^2} \sum_{i,j} (G_{ij}^l - \hat{G}_{ij}^l)^2$$

$$\mathcal{L}(\vec{x}, \hat{\vec{x}}) = \sum_{l=0}^L w_l E_l$$



Gatys, Ecker, and Bethge, "Texture Synthesis Using Convolutional Neural Networks", NIPS 2015
 Figure copyright Leon Gatys, Alexander S. Ecker, and Matthias Bethge, 2015. Reproduced with permission.

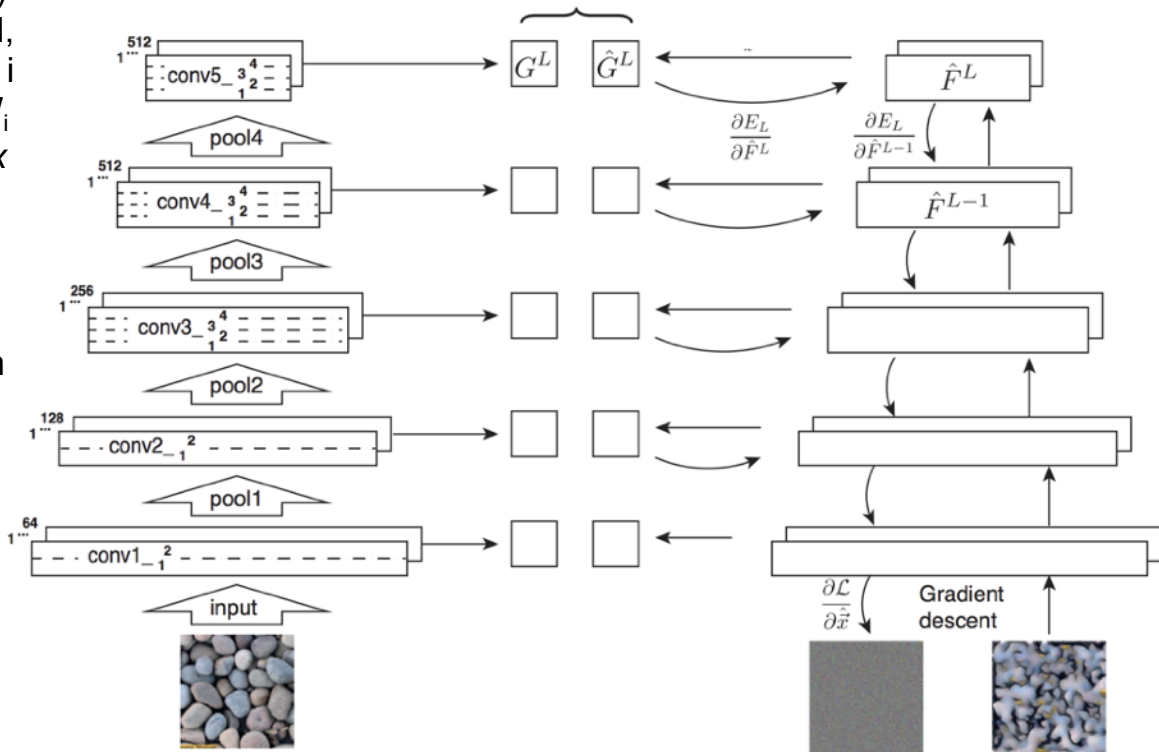
Neural Texture Synthesis

1. Pretrain a CNN on ImageNet (VGG-19)
2. Run input texture forward through CNN, record activations on every layer; layer i gives feature map of shape $C_i \times H_i \times W_i$
3. At each layer compute the *Gram matrix* giving outer product of features:

$$G_{ij}^l = \sum_k F_{ik}^l F_{jk}^l \quad (\text{shape } C_i \times C_i)$$

4. Initialize generated image from random noise
5. Pass generated image through CNN, compute Gram matrix on each layer
6. Compute loss: weighted sum of L2 distance between Gram matrices
7. Backprop to get gradient on image
8. Make gradient step on image
9. GOTO 5

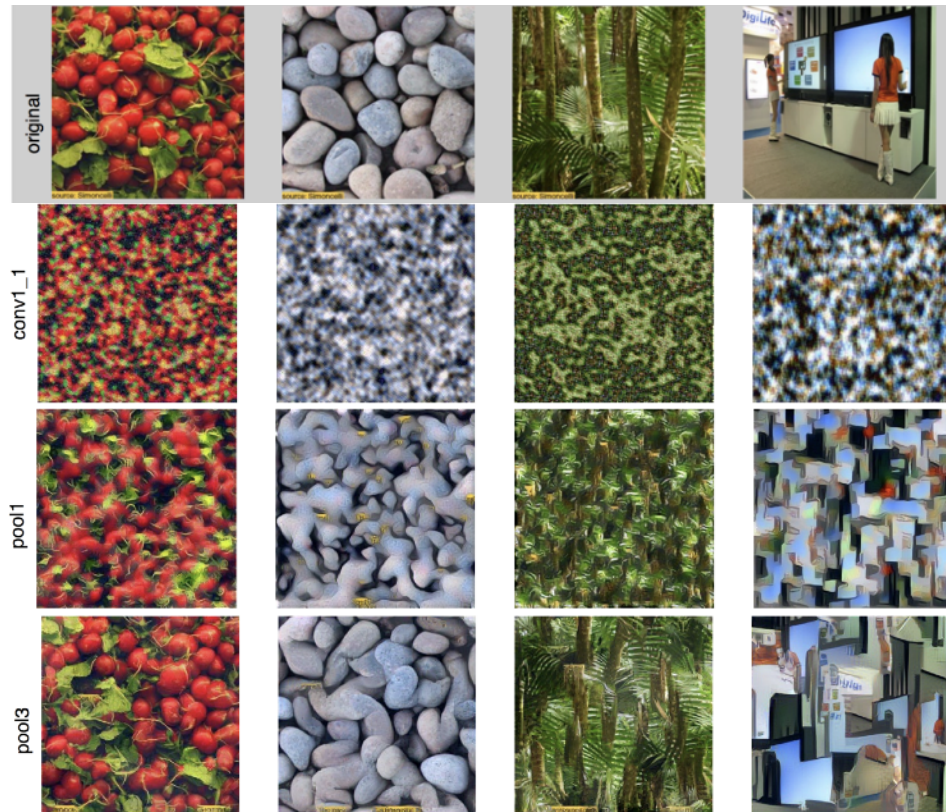
$$E_l = \frac{1}{4N_l^2 M_l^2} \sum_{i,j} (G_{ij}^l - \hat{G}_{ij}^l)^2 \quad \mathcal{L}(\vec{x}, \hat{x}) = \sum_{l=0}^L w_l E_l$$



Gatys, Ecker, and Bethge, "Texture Synthesis Using Convolutional Neural Networks", NIPS 2015
Figure copyright Leon Gatys, Alexander S. Ecker, and Matthias Bethge, 2015. Reproduced with permission.

Neural Texture Synthesis

Reconstructing texture
from higher layers recovers
larger features from the
input texture



Gatys, Ecker, and Bethge, "Texture Synthesis Using Convolutional Neural Networks", NIPS 2015
Figure copyright Leon Gatys, Alexander S. Ecker, and Matthias Bethge, 2015. Reproduced with permission.

Neural Texture Synthesis: Texture = Artwork

Texture synthesis
(Gram
reconstruction)

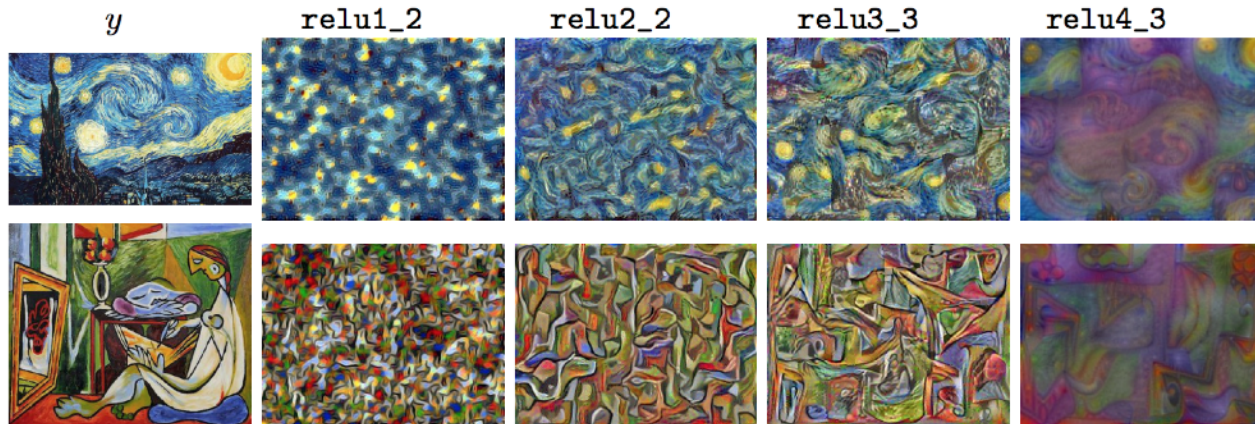
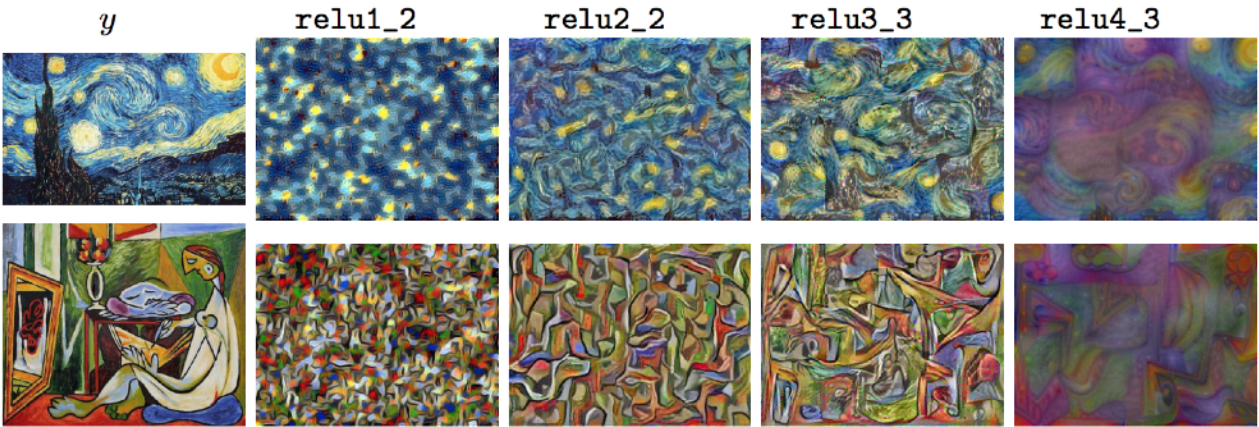


Figure from Johnson, Alahi, and Fei-Fei, "Perceptual Losses for Real-Time Style Transfer and Super-Resolution", ECCV 2016. Copyright Springer, 2016. Reproduced for educational purposes.

Neural Style Transfer: Feature + Gram Reconstruction

Texture synthesis
(Gram reconstruction)



Feature reconstruction

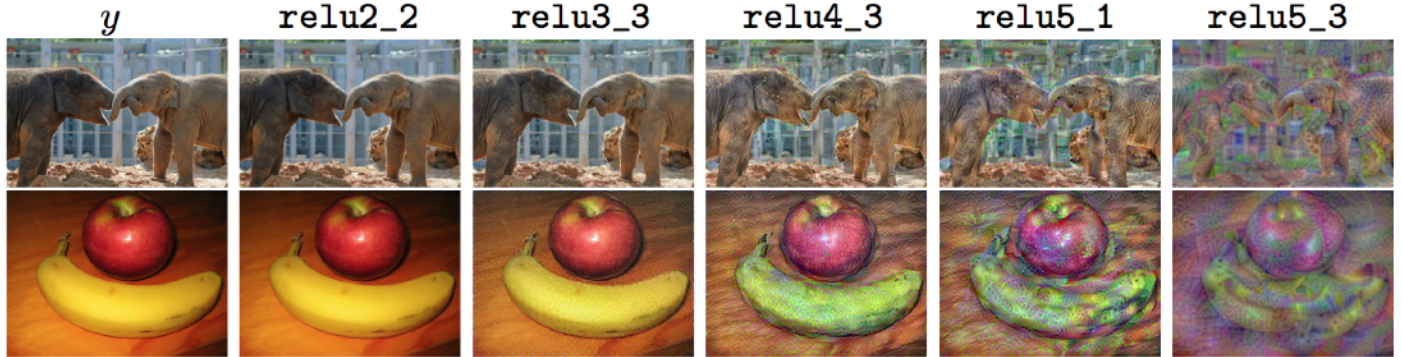


Figure from Johnson, Alahi, and Fei-Fei, "Perceptual Losses for Real-Time Style Transfer and Super-Resolution", ECCV 2016. Copyright Springer, 2016. Reproduced for educational purposes.

Neural Style Transfer

Content Image



[This image](#) is licensed under [CC-BY 3.0](#)

Style Image



[Starry Night](#) by Van Gogh is in the public domain

+

Gatys, Ecker, and Bethge, "Texture Synthesis Using Convolutional Neural Networks", NIPS 2015

Neural Style Transfer

Content Image



[This image](#) is licensed under [CC-BY 3.0](#)

+

Style Image



[Starry Night](#) by Van Gogh is in the public domain

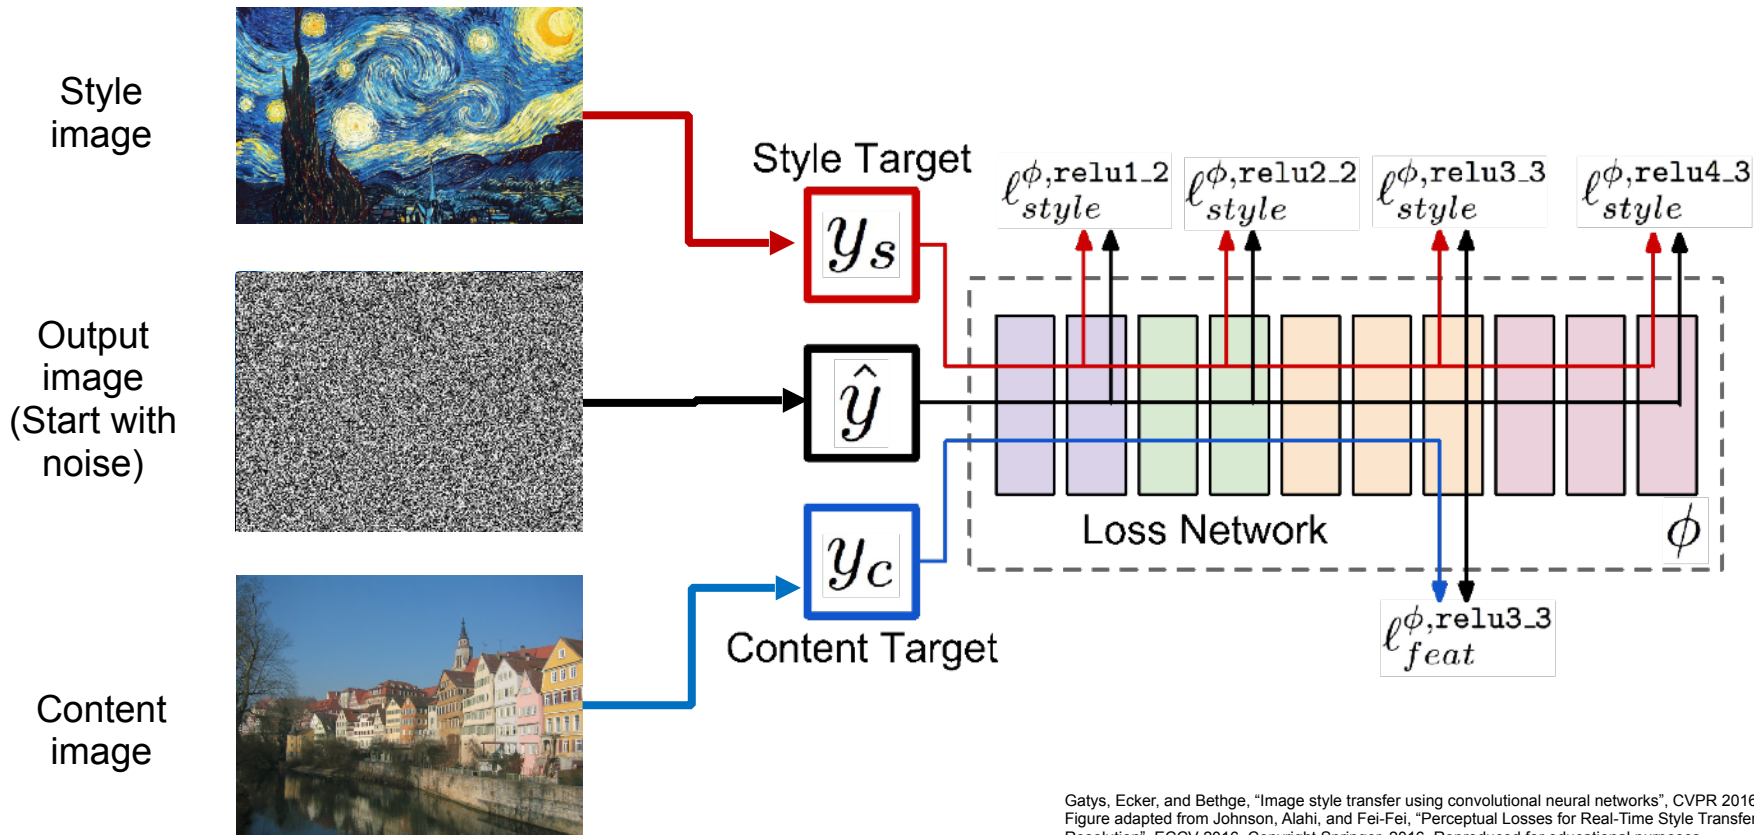
=

Style Transfer!

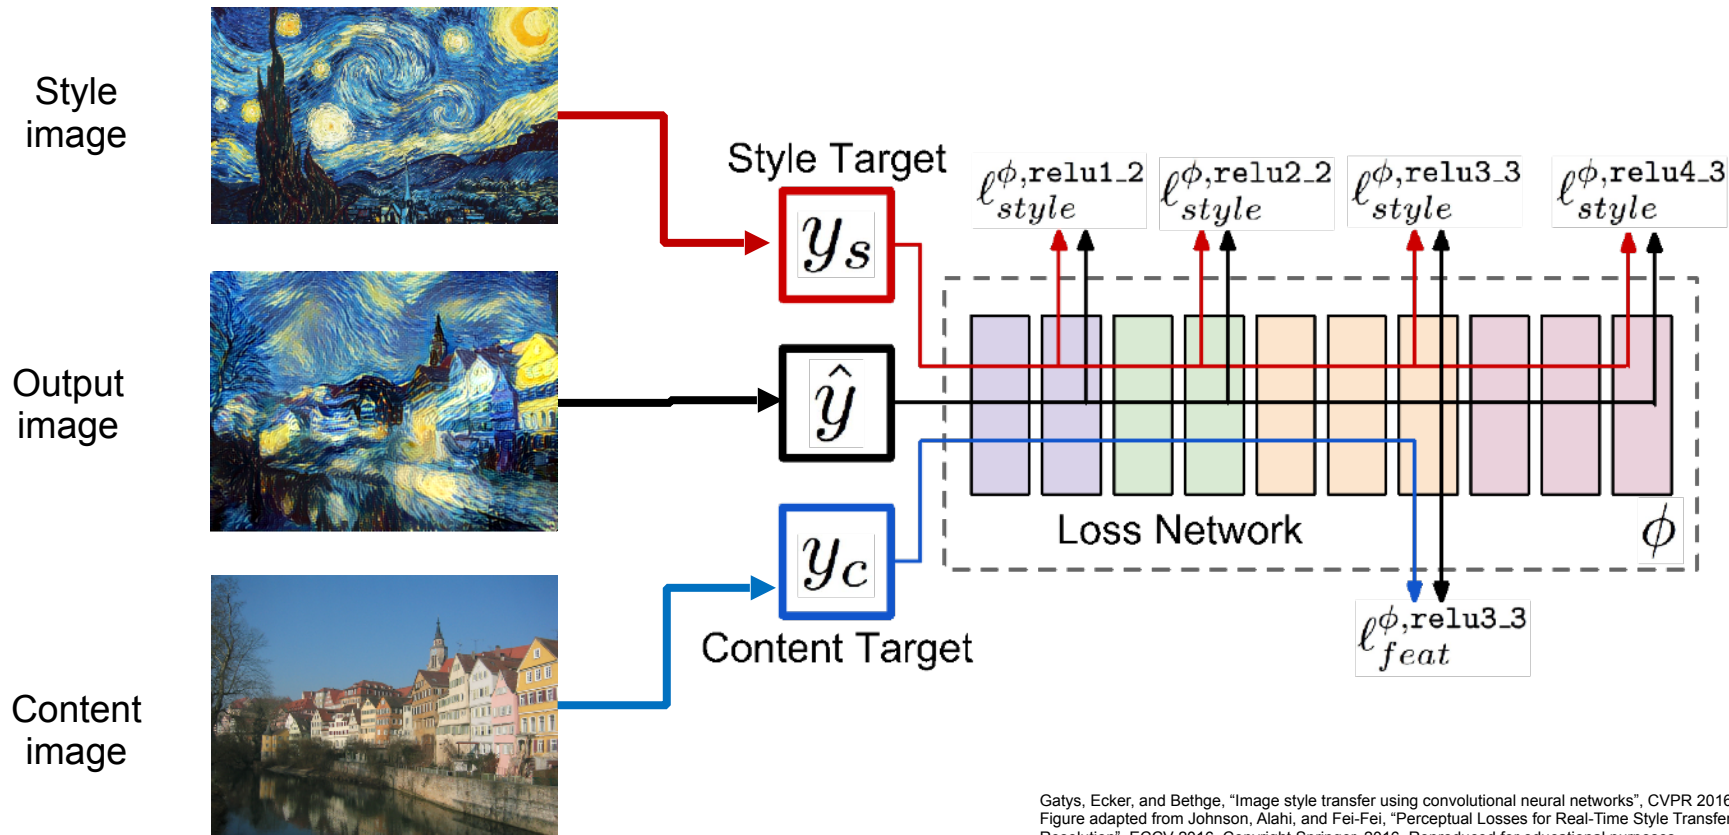


[This image](#) copyright Justin Johnson, 2015. Reproduced with permission.

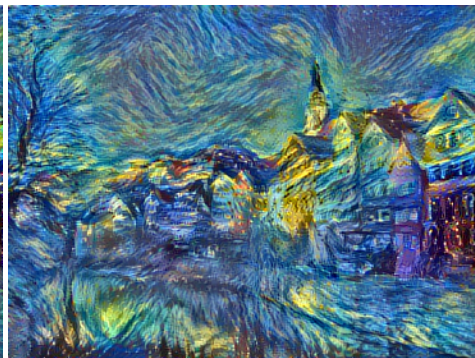
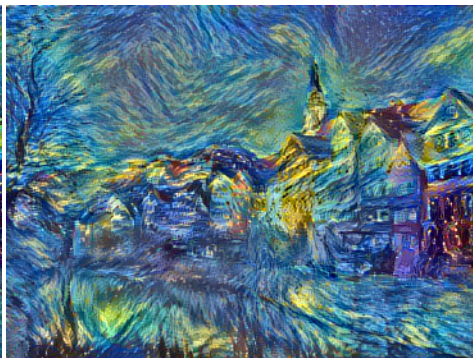
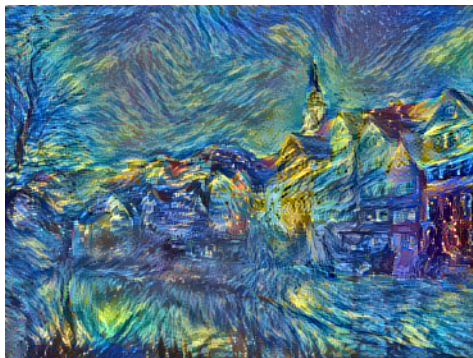
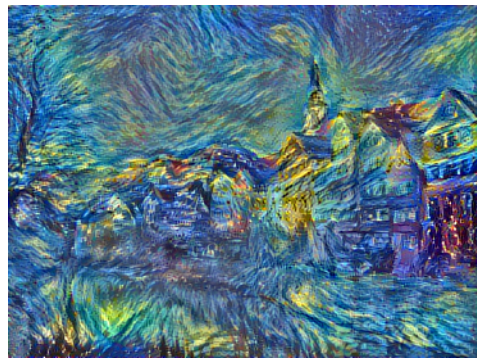
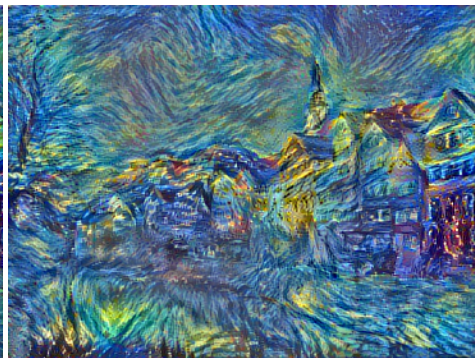
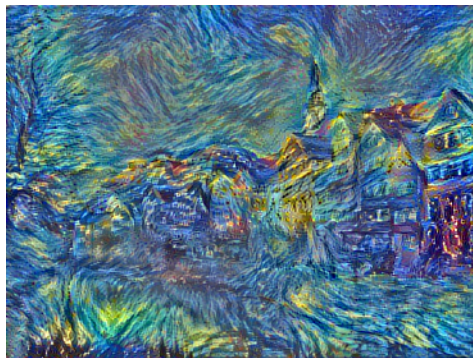
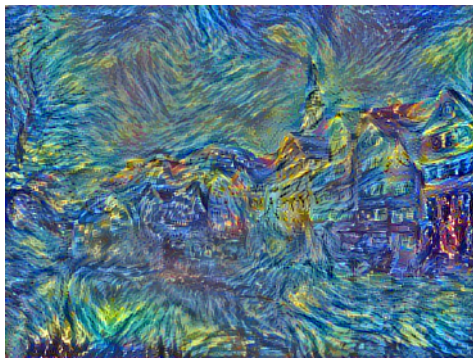
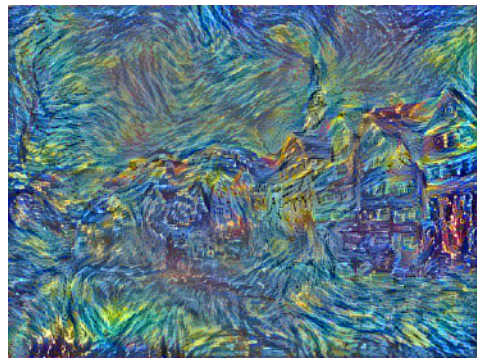
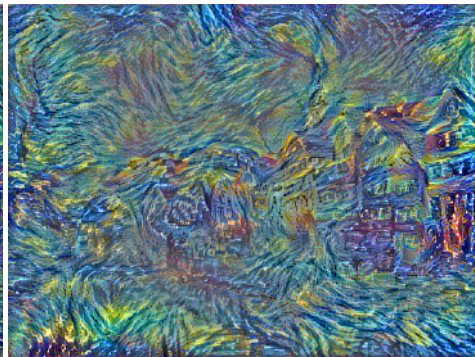
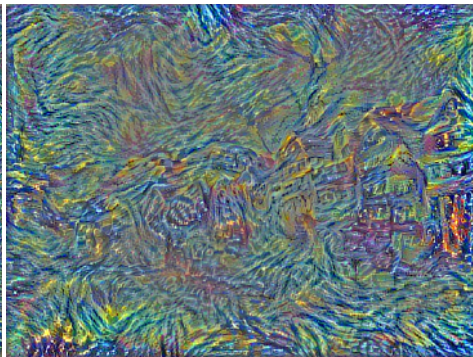
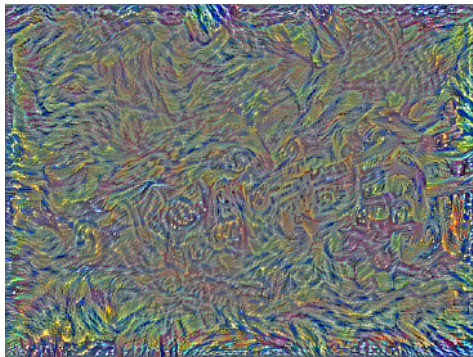
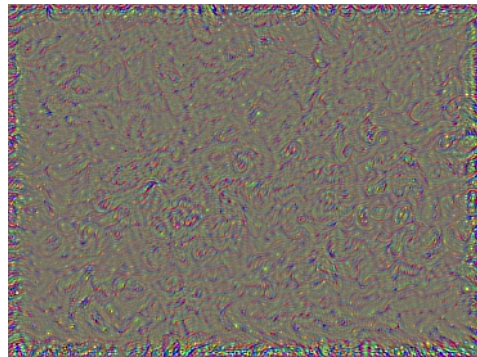
Gatys, Ecker, and Bethge, "Image style transfer using convolutional neural networks", CVPR 2016



Gatys, Ecker, and Bethge, "Image style transfer using convolutional neural networks", CVPR 2016
 Figure adapted 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.



Gatys, Ecker, and Bethge, "Image style transfer using convolutional neural networks", CVPR 2016
 Figure adapted 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.



Neural Style Transfer

Example outputs from [implementation](#) (in Torch)



Gatys, Ecker, and Bethge, "Image style transfer using convolutional neural networks", CVPR 2016
Figure copyright Justin Johnson, 2015.

Neural Style Transfer



More weight to
content loss



More weight to
style loss

Neural Style Transfer

Resizing style image before running style transfer algorithm can transfer different types of features



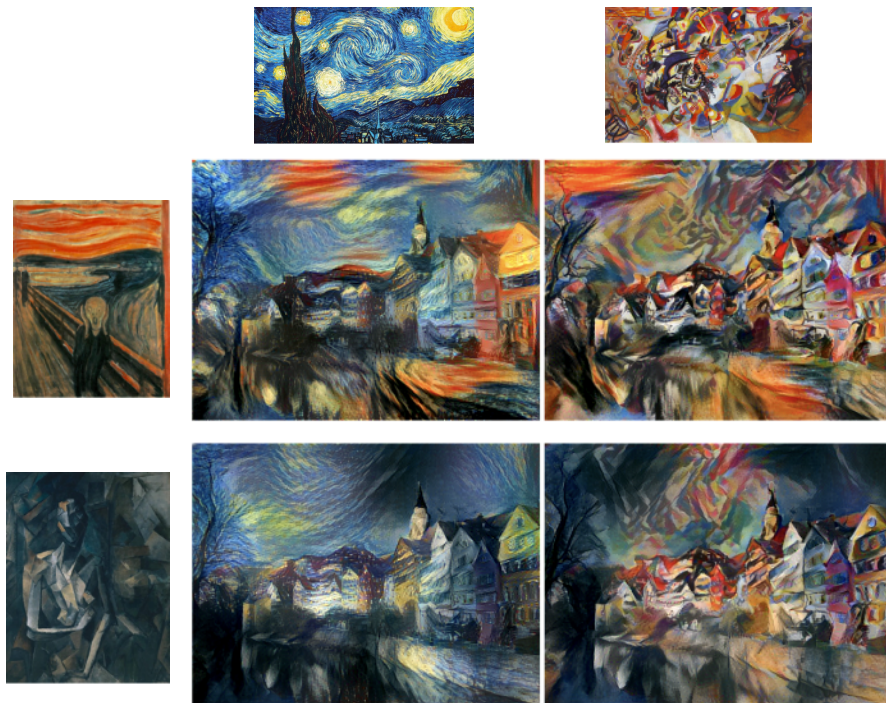
Larger style
image

Smaller style
image

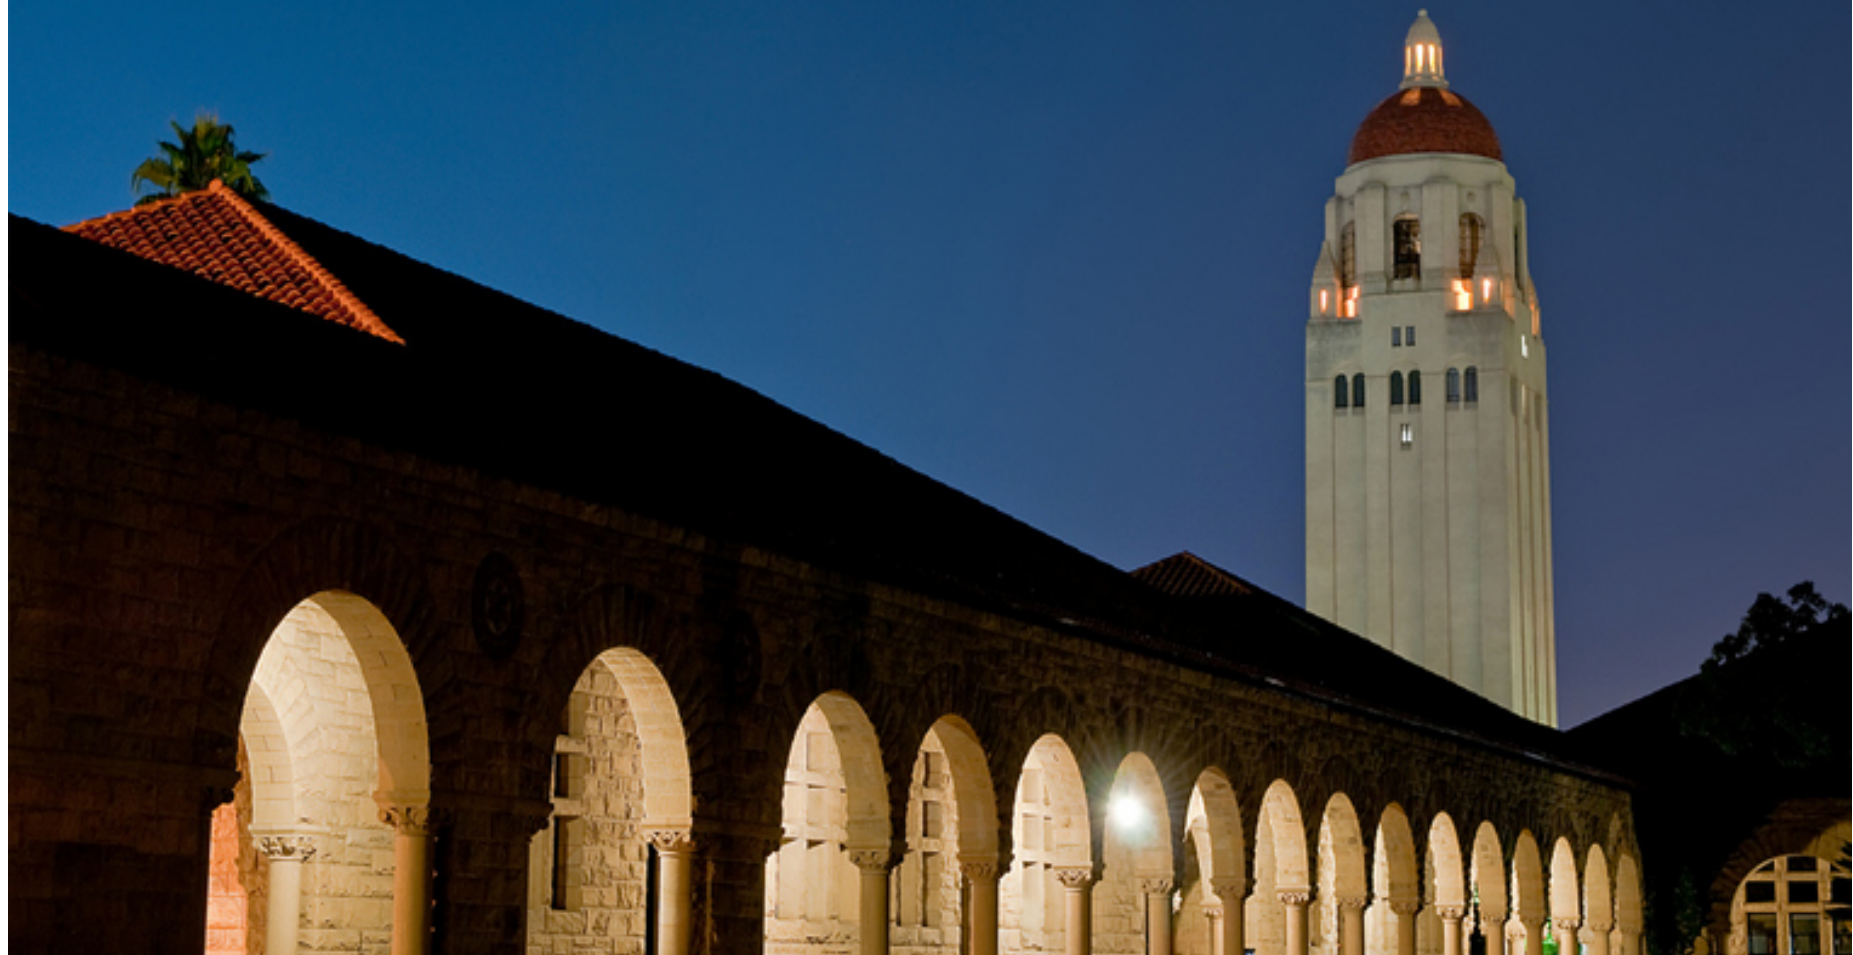
Gatys, Ecker, and Bethge, "Image style transfer using convolutional neural networks", CVPR 2016
Figure copyright Justin Johnson, 2015.

Neural Style Transfer: Multiple Style Images

Mix style from multiple images by taking a weighted average of Gram matrices

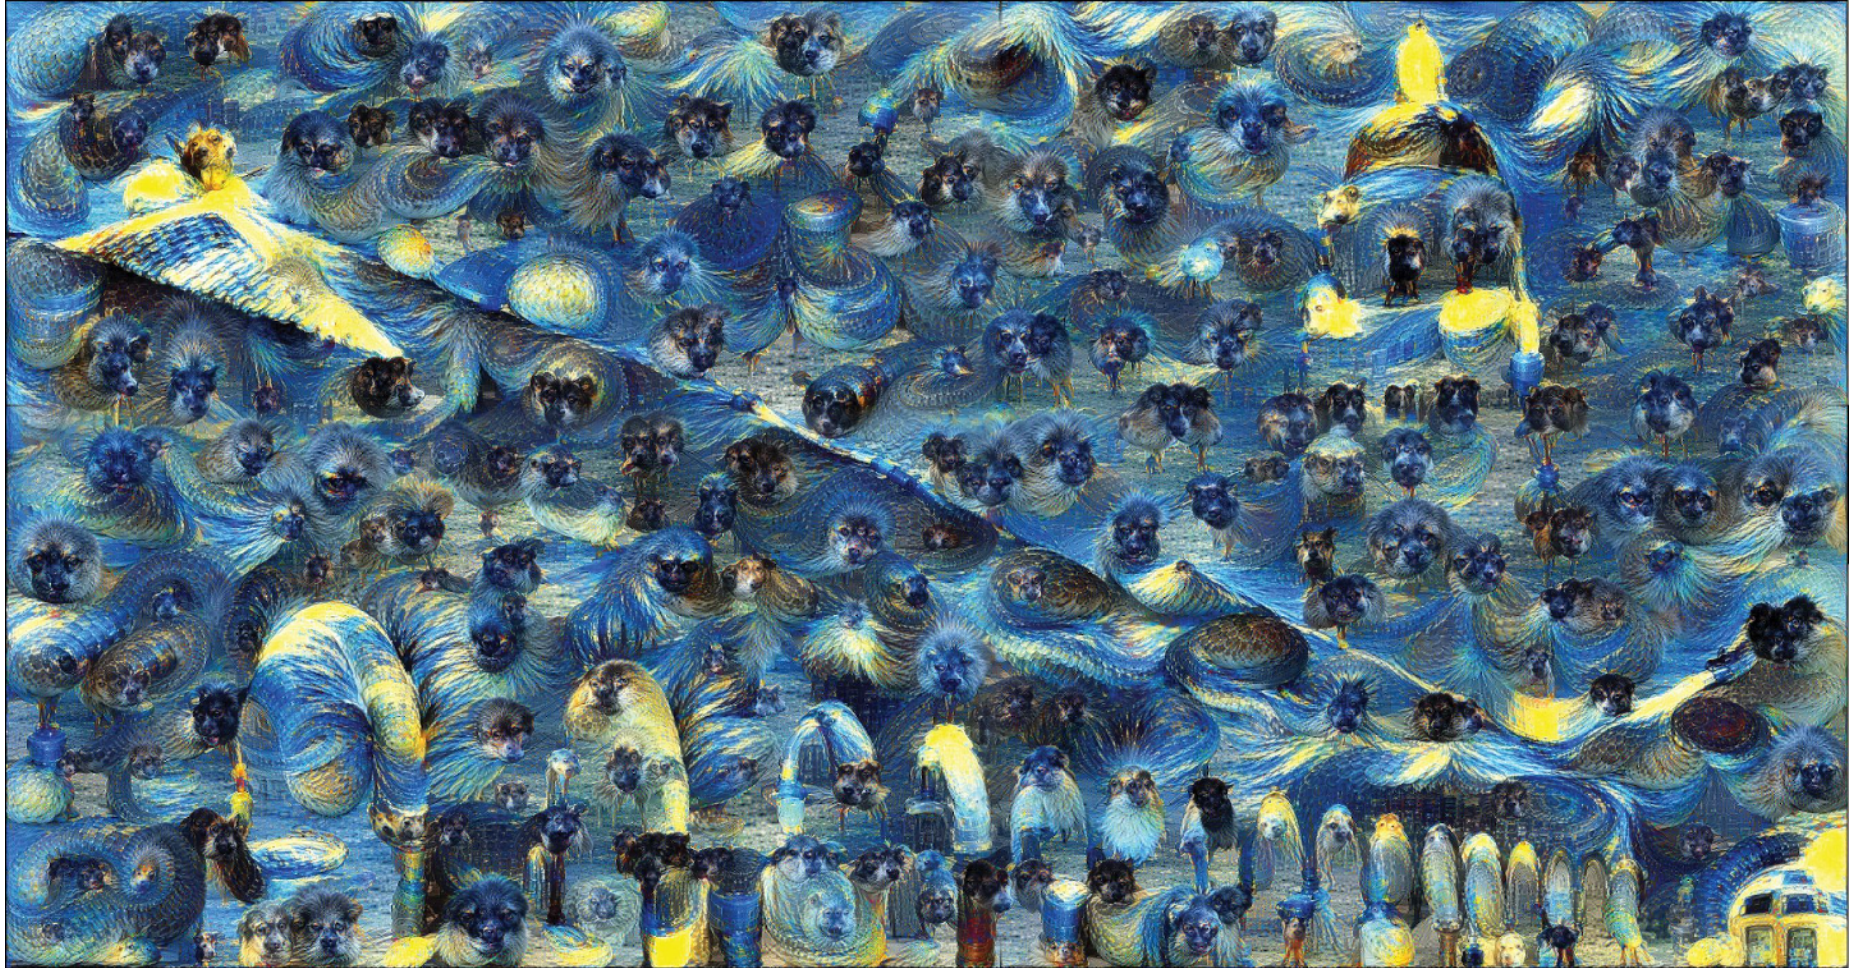


Gatys, Ecker, and Bethge, "Image style transfer using convolutional neural networks", CVPR 2016
Figure copyright Justin Johnson, 2015.









WHAT DOES IT TAKE TO GENERATE NATURAL TEXTURES?

Ivan Ustyuzhaninov^{*,1,2,3}, Wieland Brendel^{*,1,2}, Leon Gatys^{1,2,3}, Matthias Bethge^{1,2,3,4}

* contributed equally

¹Centre for Integrative Neuroscience, University of Tübingen, Germany

²Bernstein Center for Computational Neuroscience, Tübingen, Germany

³Graduate School of Neural Information Processing, University of Tübingen, Germany

⁴Max Planck Institute for Biological Cybernetics, Tübingen, Germany

Original



single-scale
 $0.195 \cdot 10^{-3}$



$0.194 \cdot 10^{-3}$



$0.283 \cdot 10^{-3}$



$0.089 \cdot 10^{-3}$



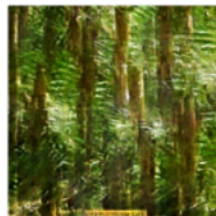
multi-scale
 $0.094 \cdot 10^{-3}$



$0.157 \cdot 10^{-3}$



$0.212 \cdot 10^{-3}$



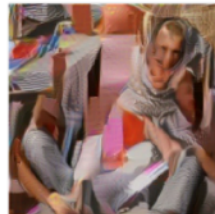
$0.077 \cdot 10^{-3}$



Gatys et al. [1]
 $0.128 \cdot 10^{-3}$



$0.089 \cdot 10^{-3}$



$0.187 \cdot 10^{-3}$

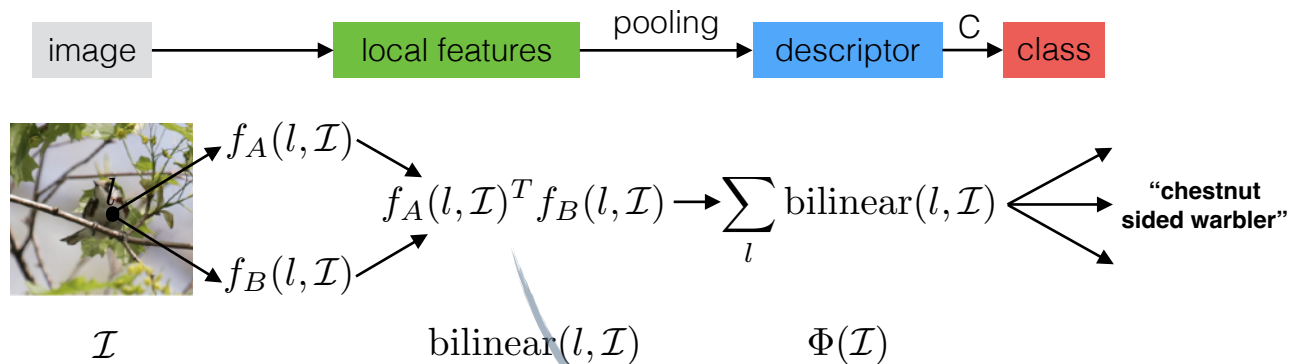


$0.022 \cdot 10^{-3}$



Bilinear/Second-Order Pooling

CNN activations pooled after outer-product encoding



Generalizes texture encoders

- ▶ Fisher vectors, Bag of Visual Words, VLAD
- ▶ Gram-matrix (when $f_A=f_B$)
- ▶ Excellent transfer from ImageNet to fine-grained domains (e.g., birds, cars, airplanes)
- ▶ Looks a lot like attention, but we didn't call it that!

f_A

	bea	tail	bell	legs	bell
red					
blue					
gray					
blue					
black					

f_B

"gray belly"

Bilinear/Second-Order Pooling

Fine-grained classification (VGG-D + VGG-M networks)



CUB 200-2011
200 species, 11,788 images



FGVC Aircraft
100 variants, 10,000 images



Stanford cars
196 models, 16,185 images

Method	Birds	Aircraft	Cars
Fully connected [D]	70.4	76.6	79.8
Fisher vector [D]	74.7	78.7	85.7
Bilinear [D+D]	84.0	83.9	90.6
Bilinear [D+M]	84.1	84.5	91.3
Previous work	84.1 [1]	80.7 [2]	92.6 [3]

Method	NABirds
Inception-BN	73.1 [4]
B-CNN [D+M]	79.4

48,562 images of 555 categories

[1] Spatial Transformer Networks, Jaderberg et al., NIPS 15

[2] Revisiting the Fisher vector for Fine-grained Classification, Gosselin et al., PR Letters 14

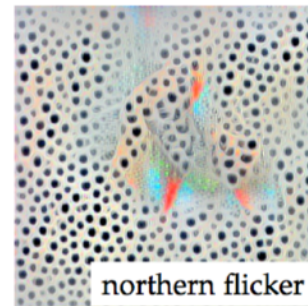
[3] Fine-Grained Rec. w/o Part Annotations, Krause et al., CVPR 15

[4] Batch-normalized Inception Architectures, Szegedy et al., CVPR 15

Visualizing Texture Classifiers

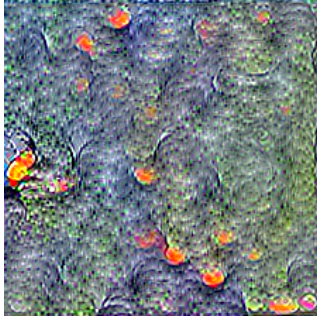
“inverse” images for bilinear CNNs

Maximal images: $\arg \max_{\mathcal{I}} \log P(c|\mathcal{I}, \mathbf{W}) + \log \Gamma(\mathcal{I})$

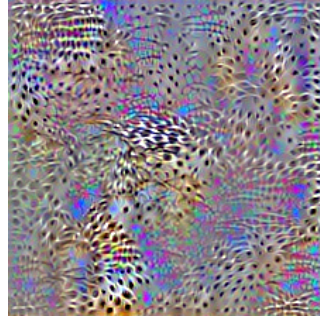


Visualizing Texture Classifiers

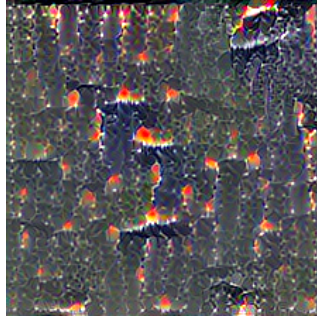
What texture are birds?



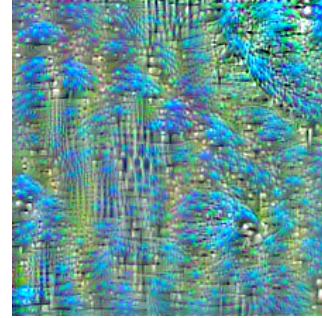
Crested auklet



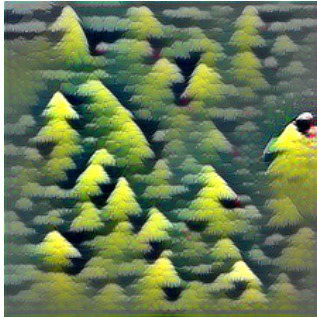
Cactus wren



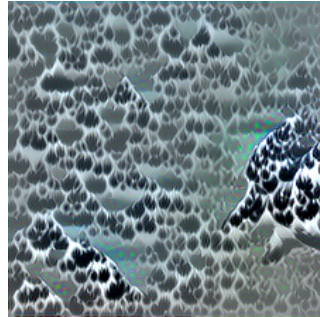
Red winged blackbird



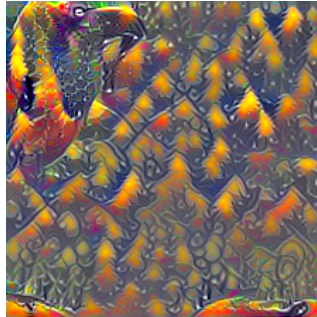
Indigo bunting



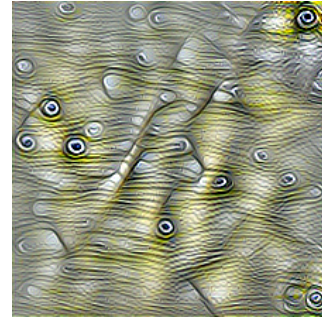
American goldfinch



Pied kingfisher



Hooded oriole

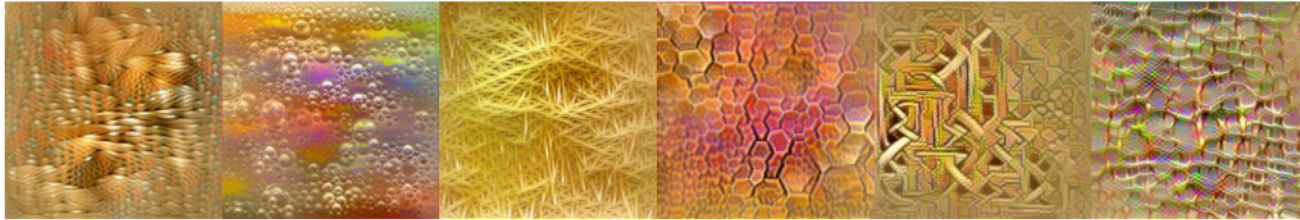


White eyed vireo

Visualizing Texture Classifiers

What texture are bookstores?

Describable Texture Dataset



braided

bubbly

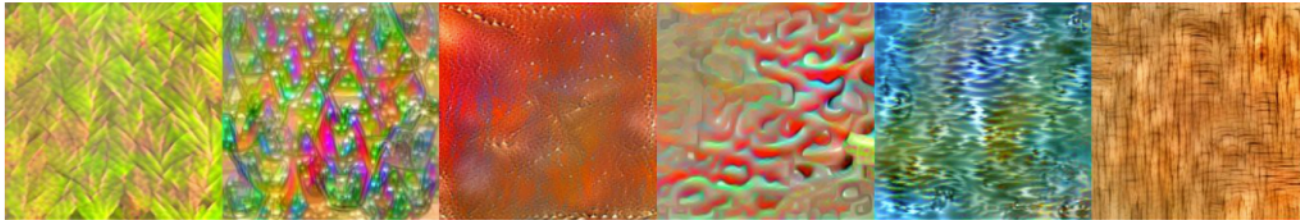
fibrous

honeycombed

interlaced

meshed

Flickr Material Dataset



foliage

glass

leather

plastic

water

wood

MIT Indoor



bathroom

bookstore

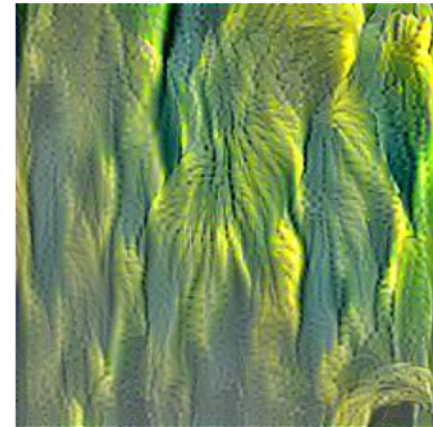
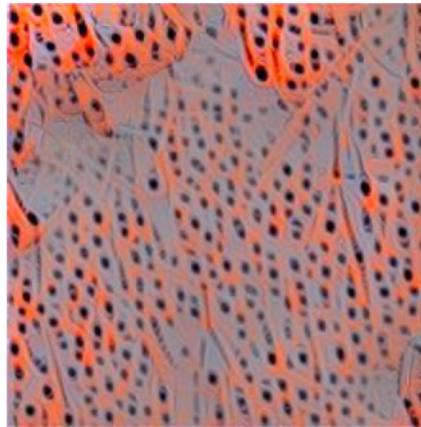
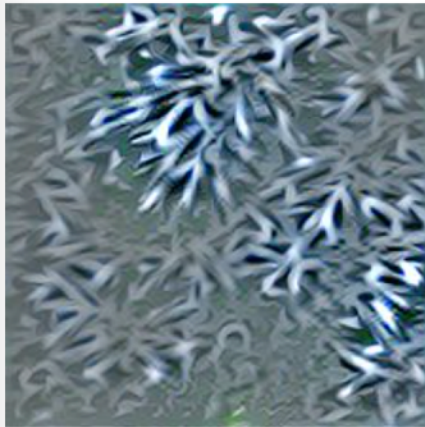
bowling

closet

classroom

laundromat

Oxford flowers



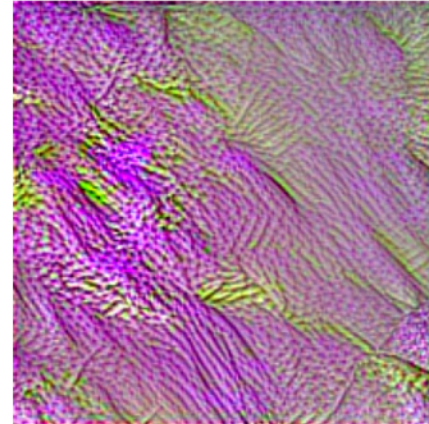
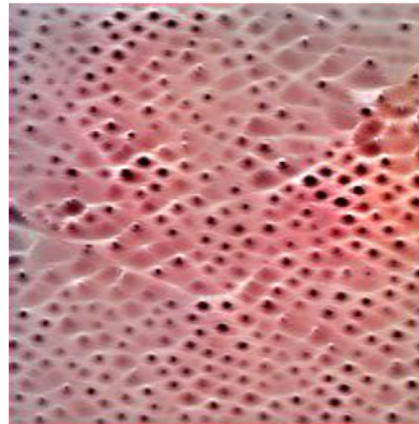
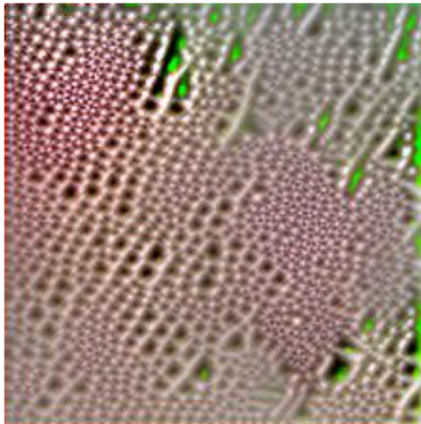
Visualizing Texture Classifiers

FGVC butterflies and moths

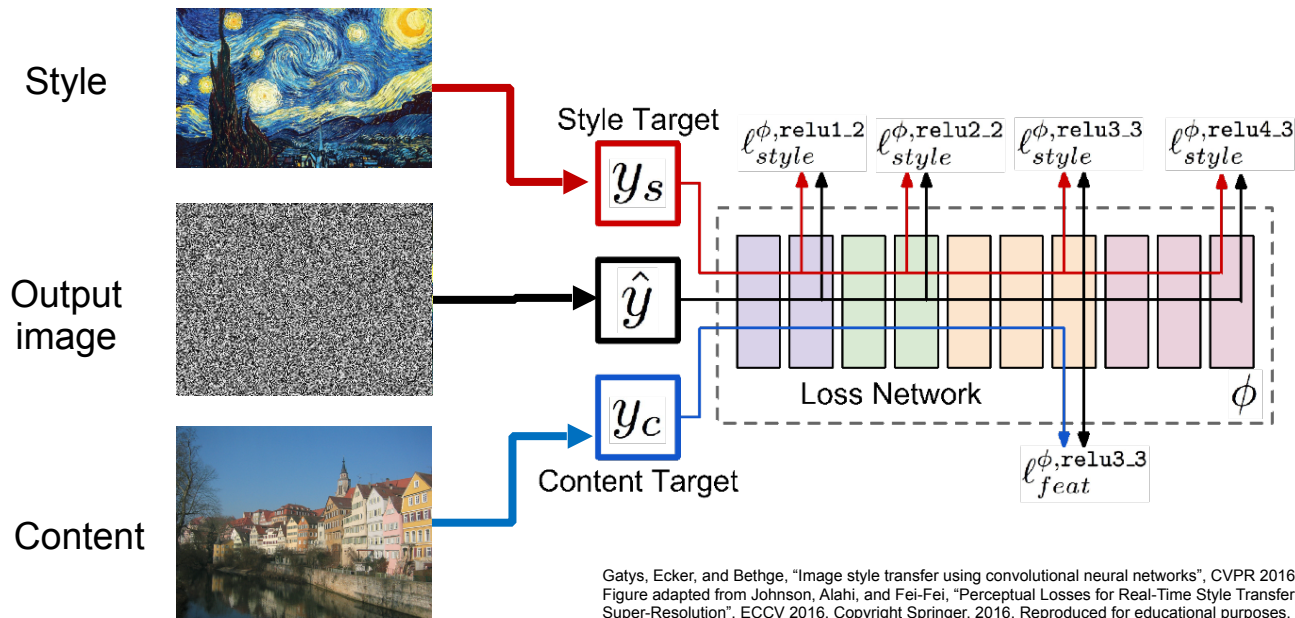


Visualizing Texture Classifiers

FGVC fungi



Can we get rid of the optimization process?

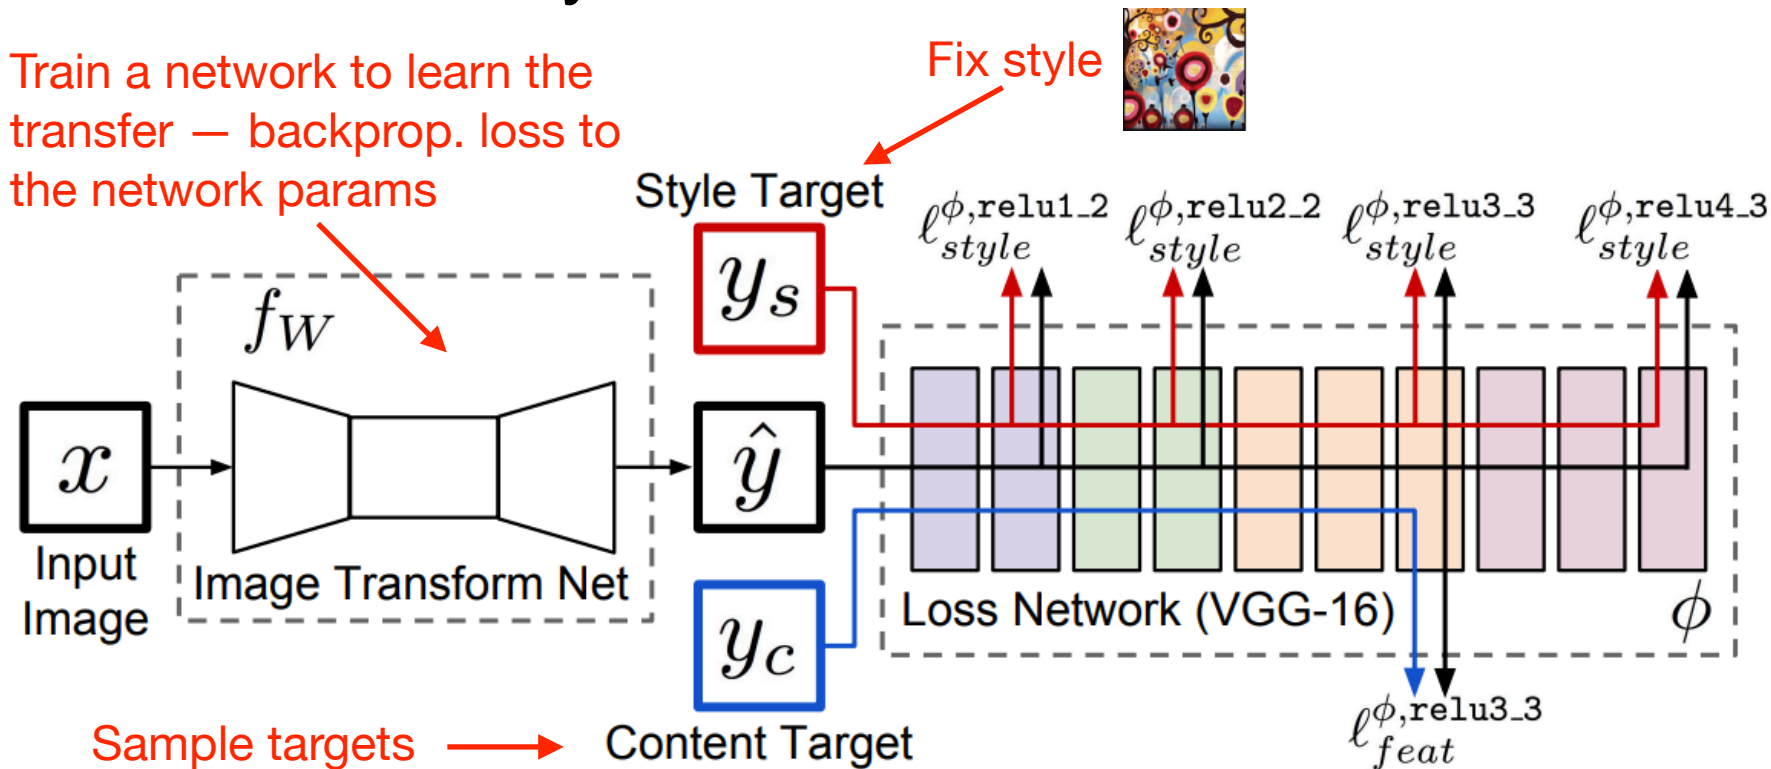


Needs many forward and backward passes!

Gatys, Ecker, and Bethge, "Image style transfer using convolutional neural networks", CVPR 2016
Figure adapted 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.

Feed-forward style transfer

Train a network to learn the transfer — backprop. loss to the network params



Perceptual Losses for Real-Time Style Transfer and Super-Resolution, Johnson et al., ECCV'26

Feed-forward style transfer — almost real-time!

Style
The Starry Night,
Vincent van Gogh,
1889



Style
The Muse,
Pablo Picasso,
1935



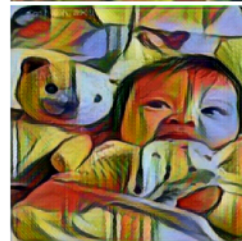
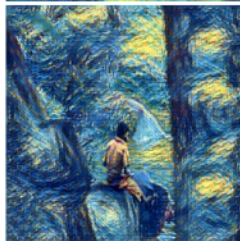
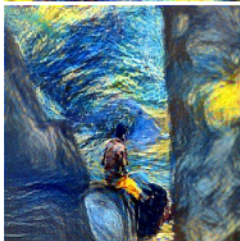
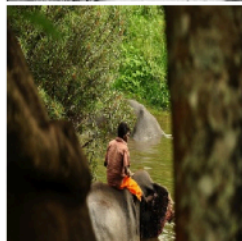
Gatys et al.

Johnson et al.



Gatys et al.

Johnson et al.



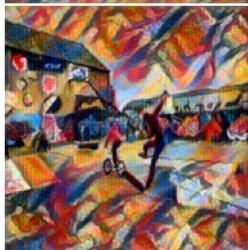
Perceptual Losses for Real-Time Style Transfer and Super-Resolution, Johnson et al., ECCV'26

Feed-forward style transfer — almost real-time!

Style
Composition VII,
Wassily
Kandinsky, 1913



Style
The Great Wave off Kanagawa,
Hokusai,
1829-1832



Gatys et al.

Johnson et al.

Gatys et al.

Johnson et al.

Perceptual Losses for Real-Time Style Transfer and Super-Resolution, Johnson et al., ECCV'26

Fast style transfer — concurrent / follow up papers

- Concurrently work on fast style transfer
 - Ulyanov, D., Lebedev, V., Vedaldi, A., Lempitsky, V.: Texture networks: Feedforward synthesis of textures and stylized images. ICML. (2016)
 - Li, C., Wand, M.: Precomputed real-time texture synthesis with markovian generative adversarial networks. ECCV. (2016)
- Dramatically improved the quality of the above methods
 - Instance Normalization: The Missing Ingredient for Fast Stylization, Ulyanov et al., 2017
- A single network to generate multiple styles though conditional instance normalization layers
 - A learned representation for artistic style. Dumoulin et al., ICLR 2017.



A learned representation for artistic style. Dumoulin et al., ICLR 2017

Summary

Many methods for understanding CNN representations

Activations: Nearest neighbors, Dimensionality reduction, maximal patches, occlusion

Gradients: Saliency maps, class visualization, feature inversion

Fun: DeepDream, Texture Synthesis, Style Transfer

Bonus: Works for fine-grained categorization too!

Future work: Generative models!