

Image Generation and GANs

CSE P576

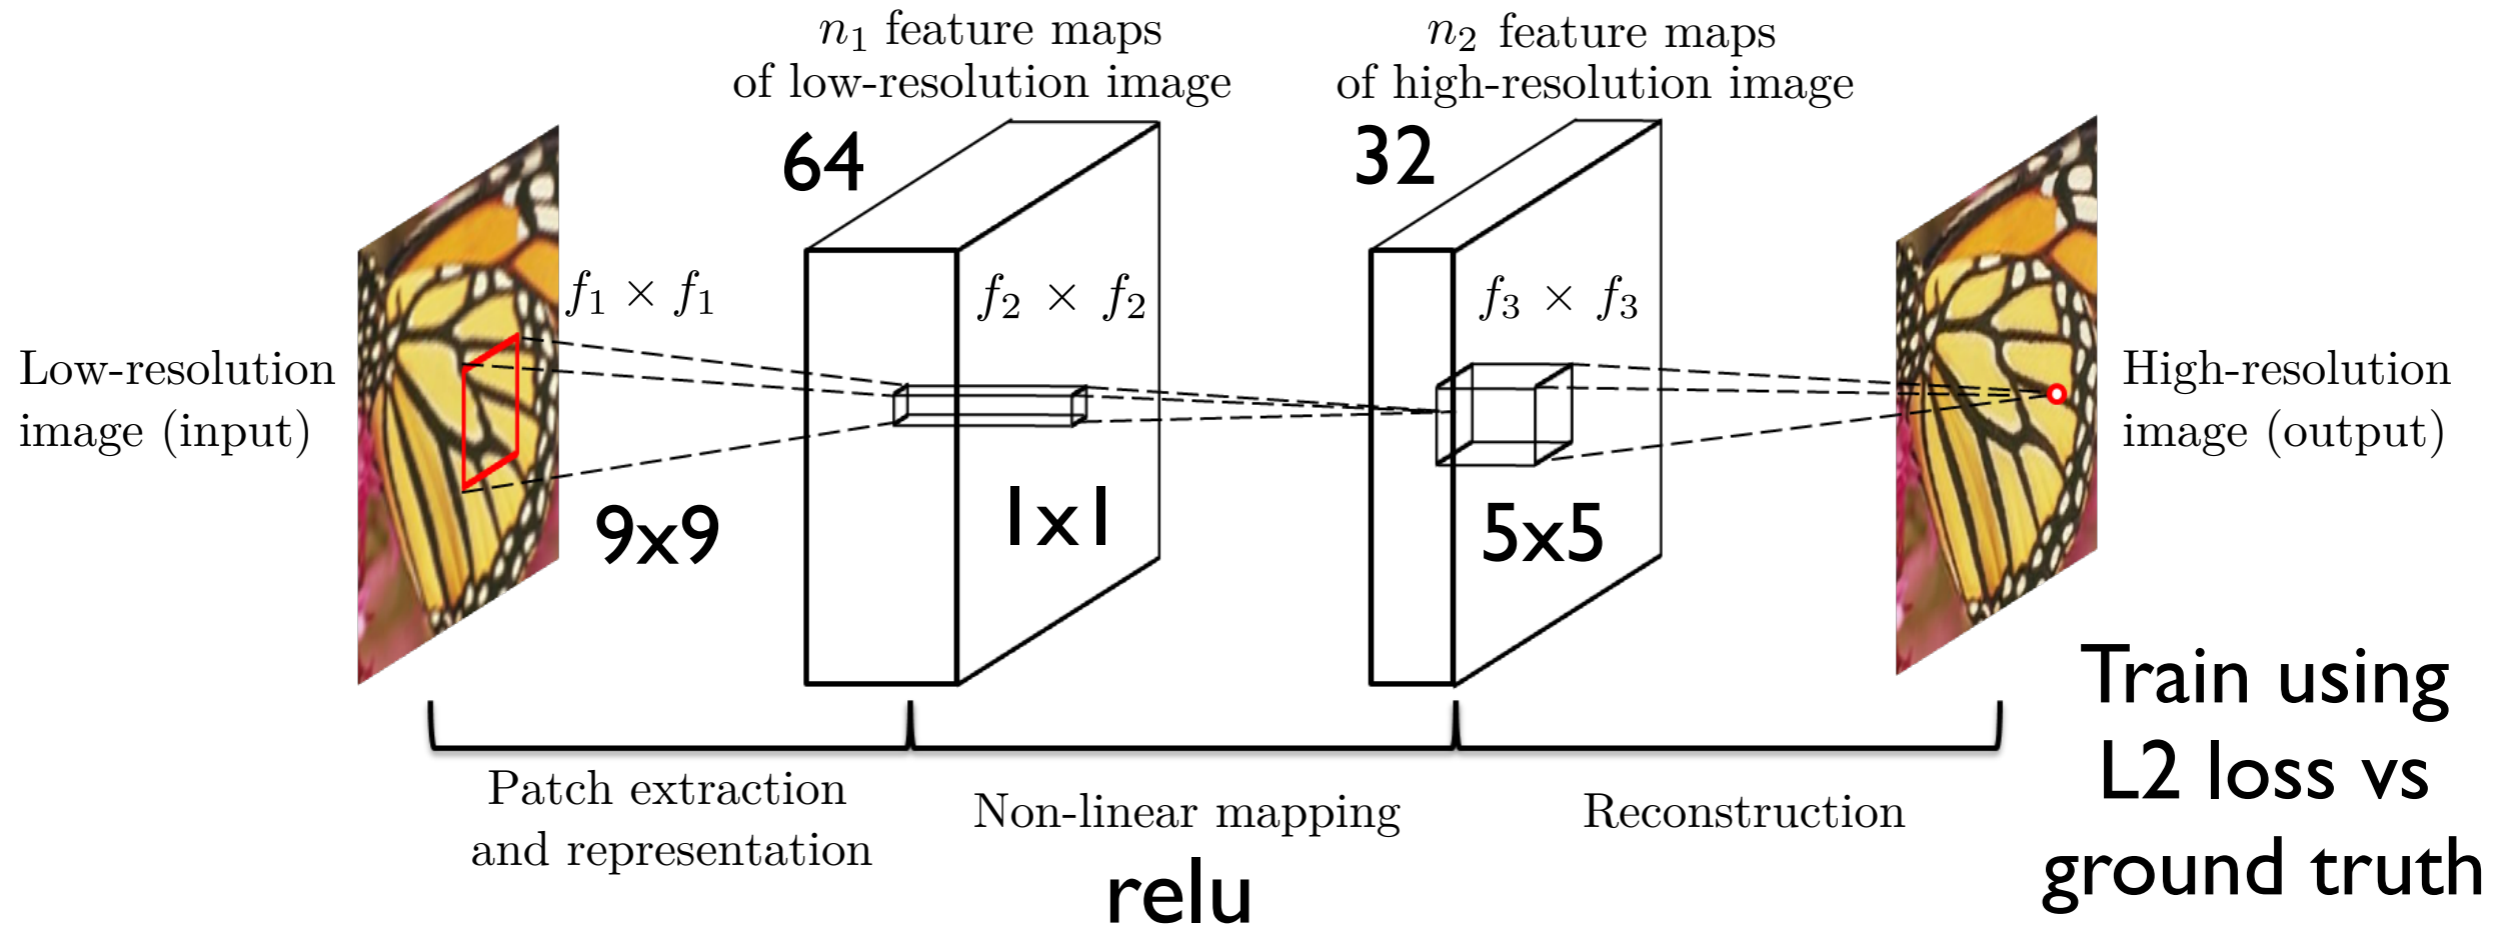
Dr. Matthew Brown

Image Generation + GANs

- Can a neural net define the loss function?
- Loss functions for Super-Resolution: L2, VGG, Adversarial
- Generative Adversarial Nets and Image Generation
- Conditional GANs, Image Translation, pix2pix

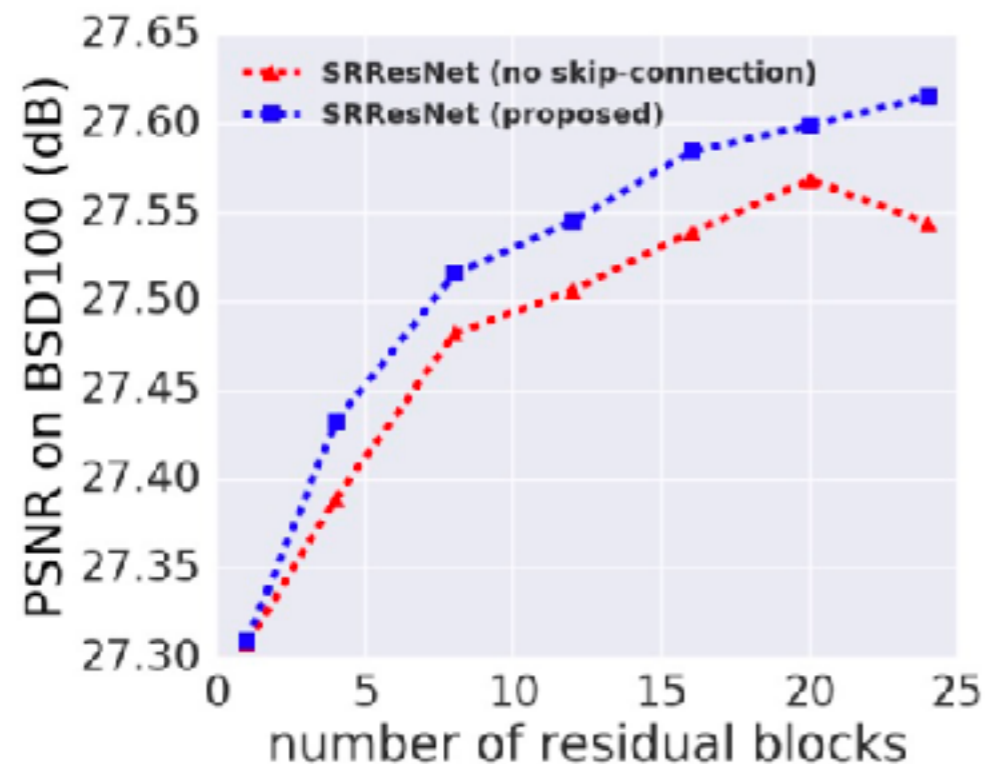
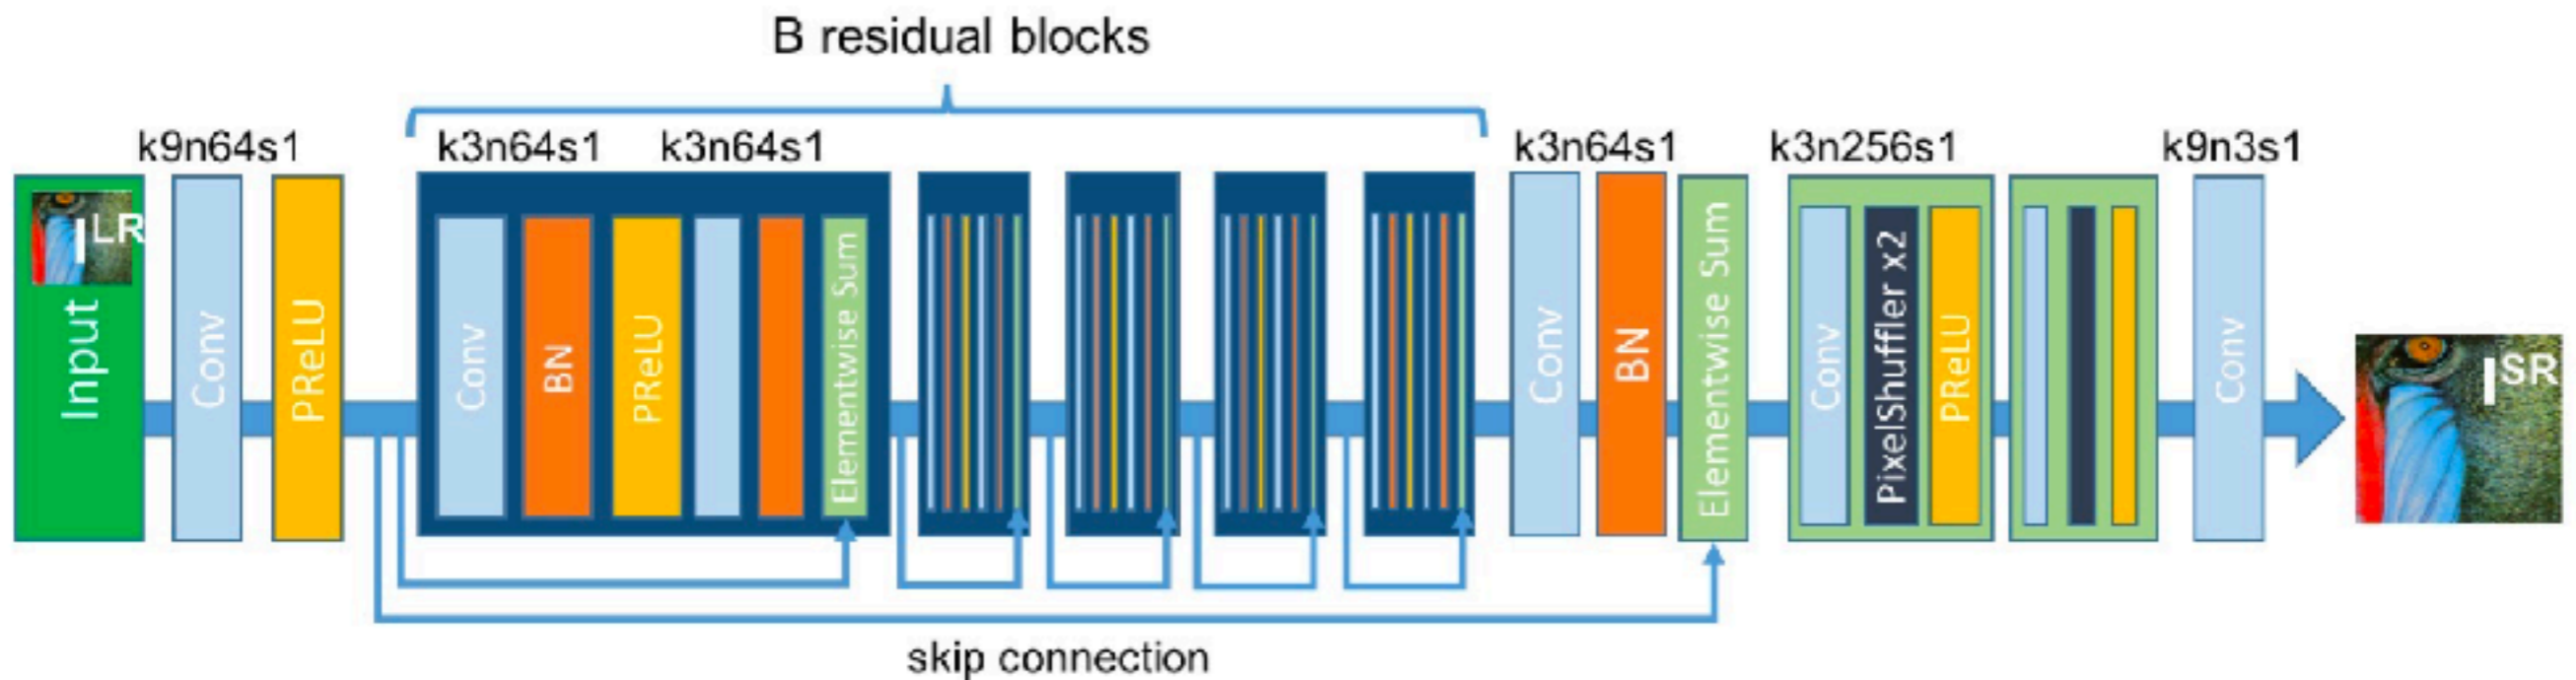
Super-Resolution: SRCNN

- Small network (3 layers) generates reasonable results



Super-Resolution: SRResNet

- Deeper networks generate better results, e.g., SRResNet



Trained with L2 loss, state-of-the-art PSNR in 2017



SRResNet: L2 Loss

bicubic
(21.59dB/0.6423)



SRResNet
(23.53dB/0.7832)



original



A state of the art super-res network trained with L2 loss is good at sharpening edges, but results lack realistic texture

SRResNet
(23.53dB/0.7832)



SRResNet
(23.53dB/0.7832)




More realistic, but L2
loss is worse
(21.15dB/0.6868)



Perceptual Metrics

- L2 loss does not match human perception in general



Human
preferred

Reference

L2
preferred


Texture Synthesis

- Which are the real radishes?



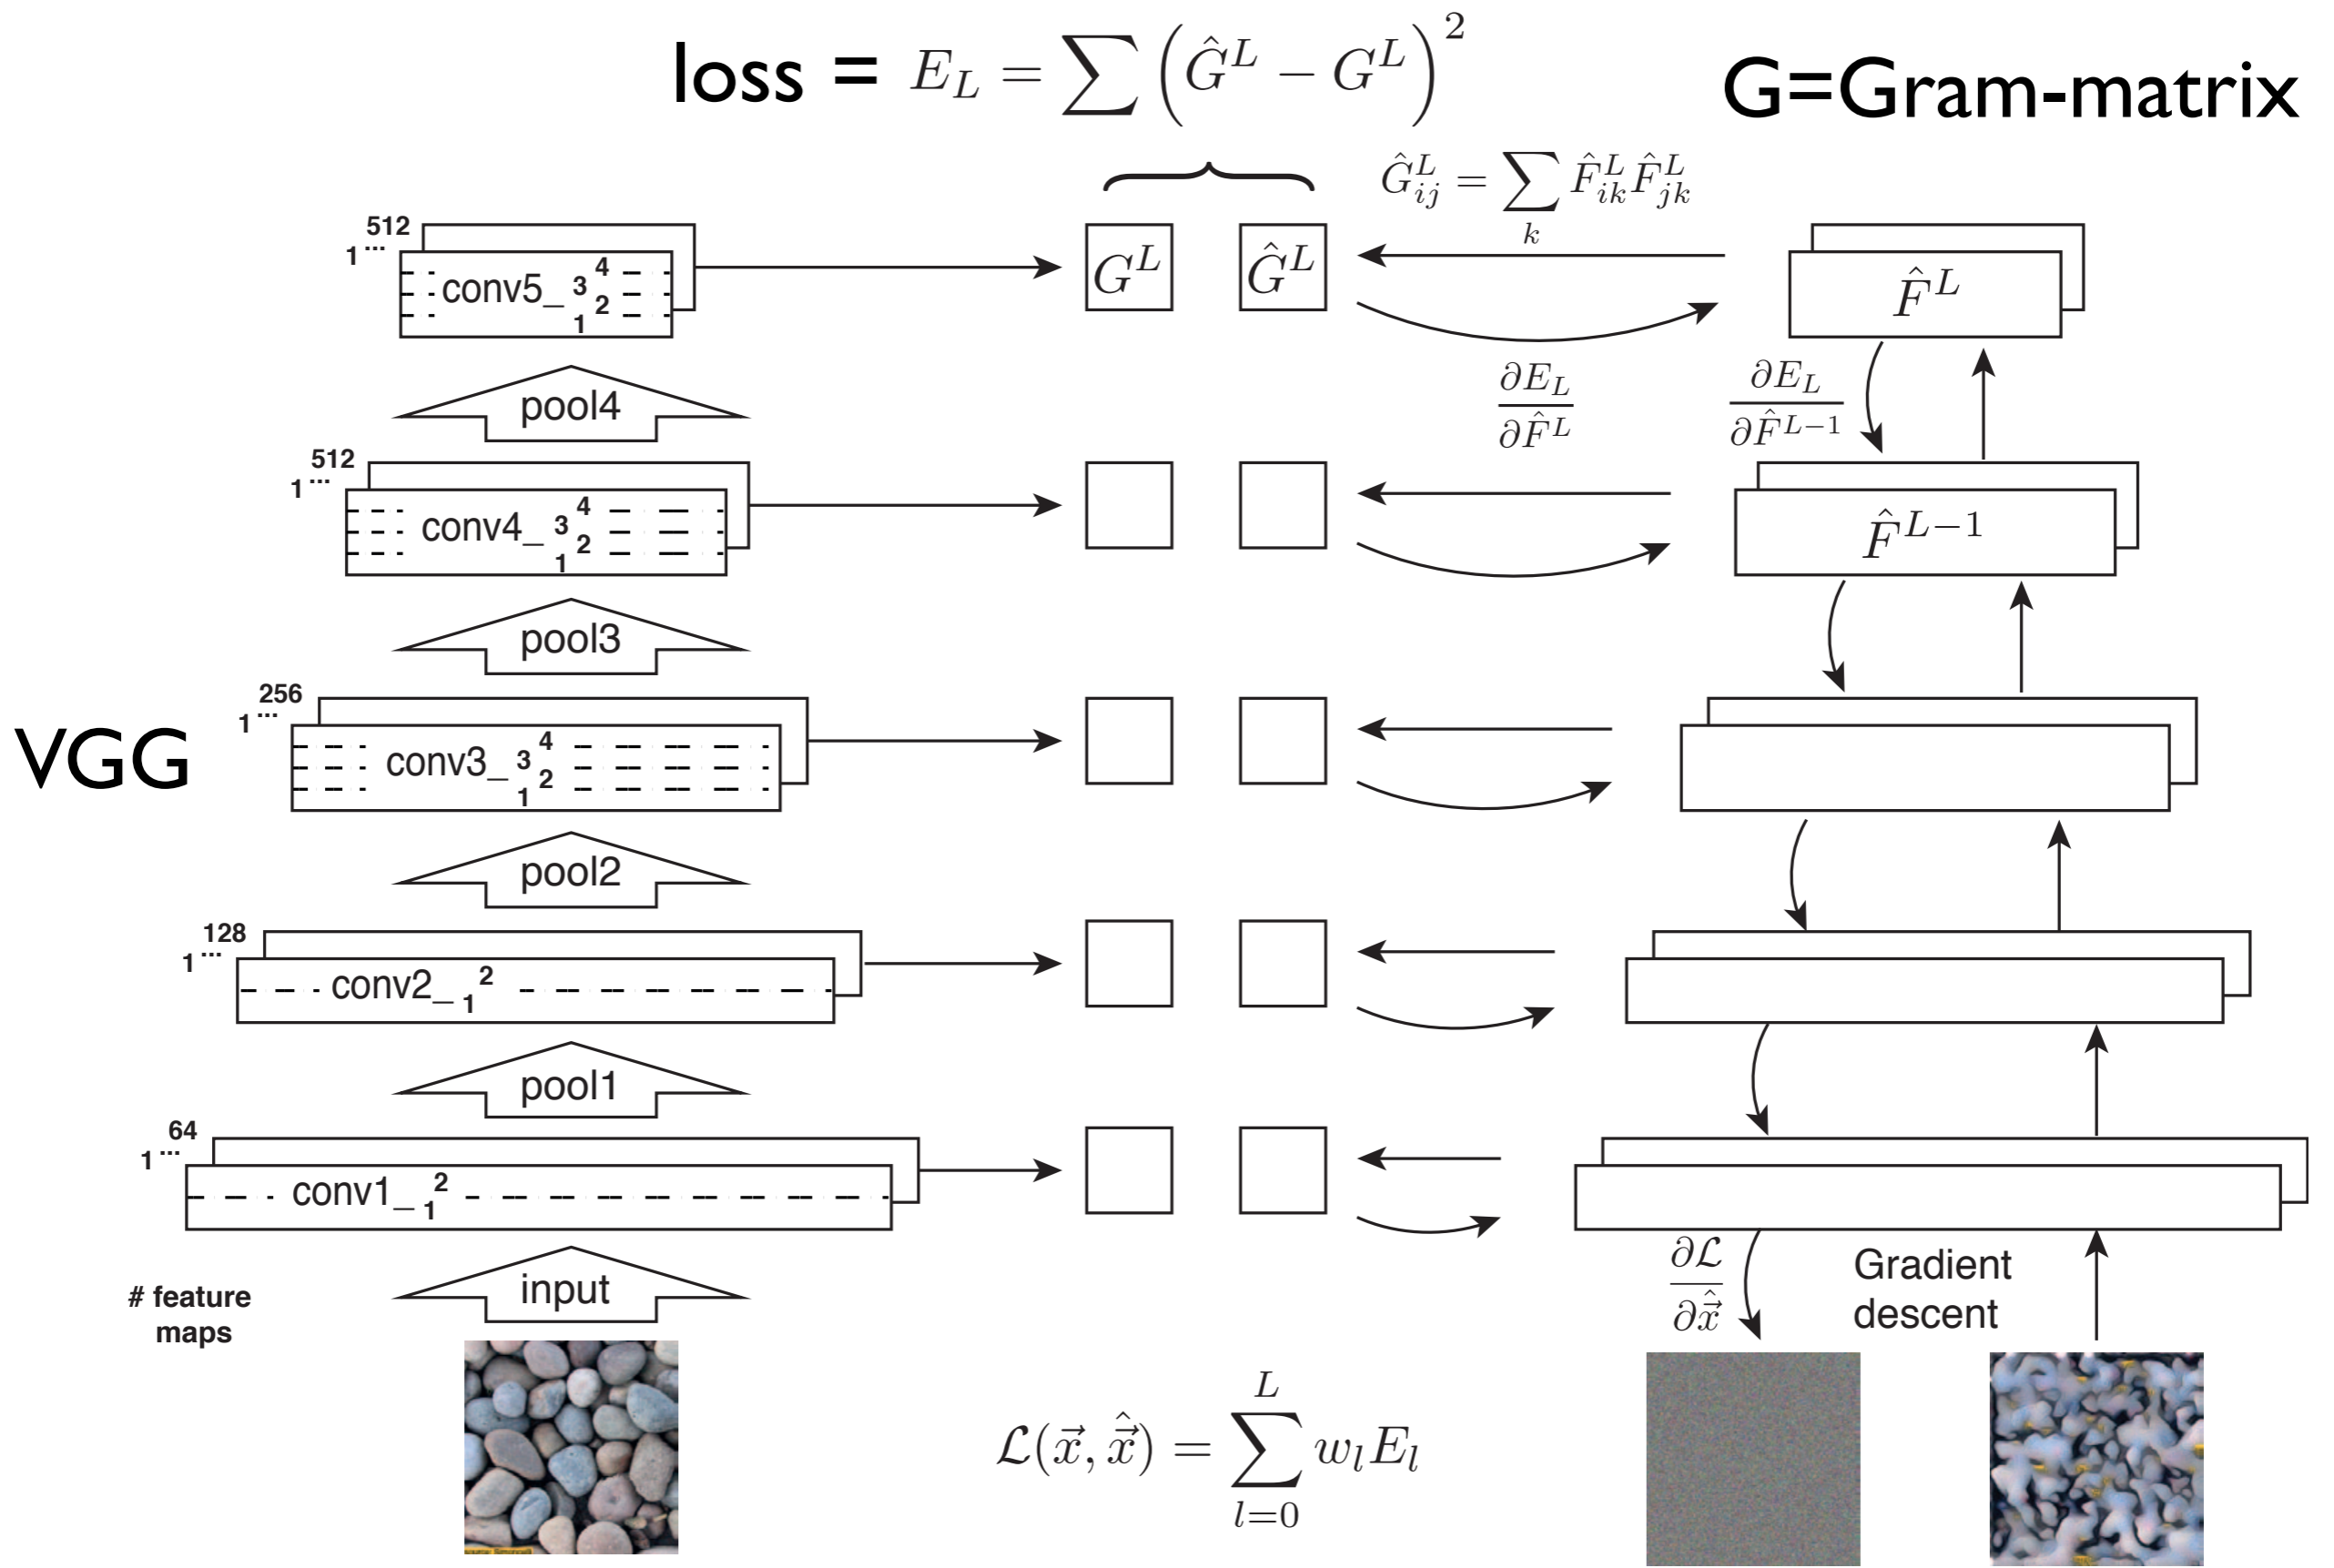
Texture Synthesis

- Which are the real rocks?



$$\text{loss} = E_L = \sum \left(\hat{G}^L - G^L \right)^2$$

G=Gram-matrix





Match Σ of outer products of feature activations

$$\hat{\vec{x}} := \vec{x} - \alpha \frac{\partial \mathcal{L}}{\partial \hat{\vec{x}}}$$

Texture Synthesis


Original



Portilla
Simoncelli




Gatys




- Portilla Simoncelli 1999 texture model also used correlation of filter responses (though shallow features / hand tuned)

simple features




Match texture statistics up to level =



more complex features



Style Transfer

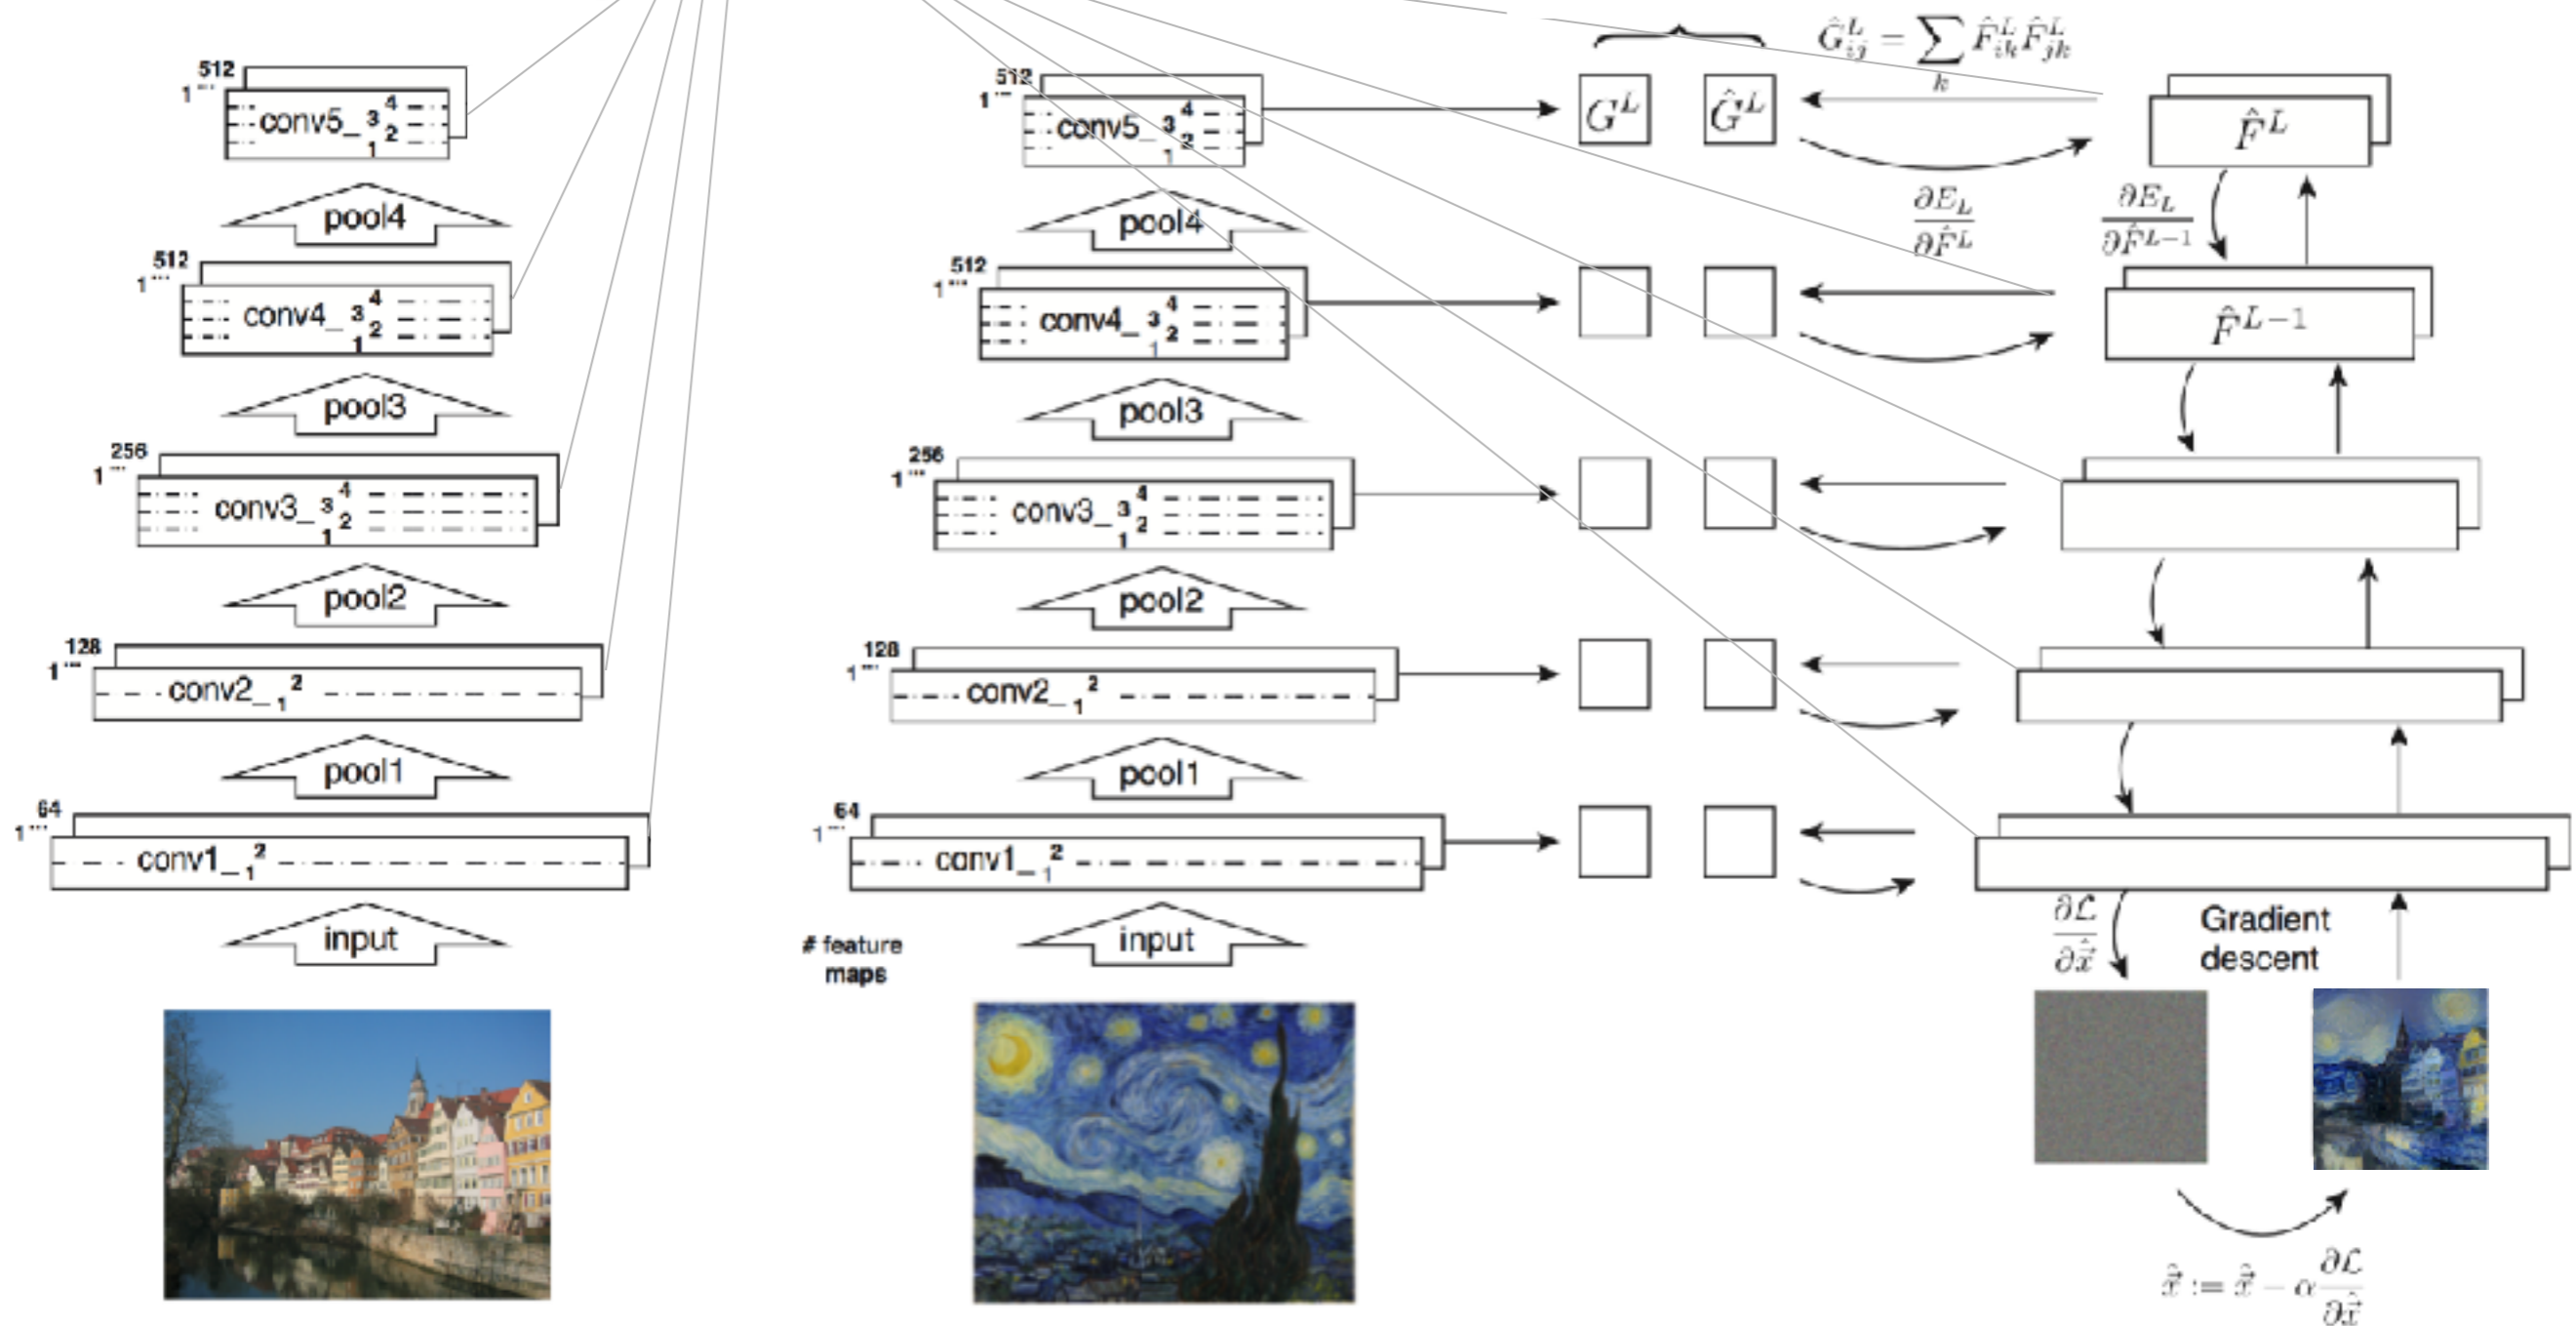


Re-render an image given the “style” of an artist [Gatys et al 2015]

Content and Style Losses

$$E(\text{content}) = \frac{1}{2} \sum_{i,j} (F_{ij}^l - P_{ij}^l)^2$$

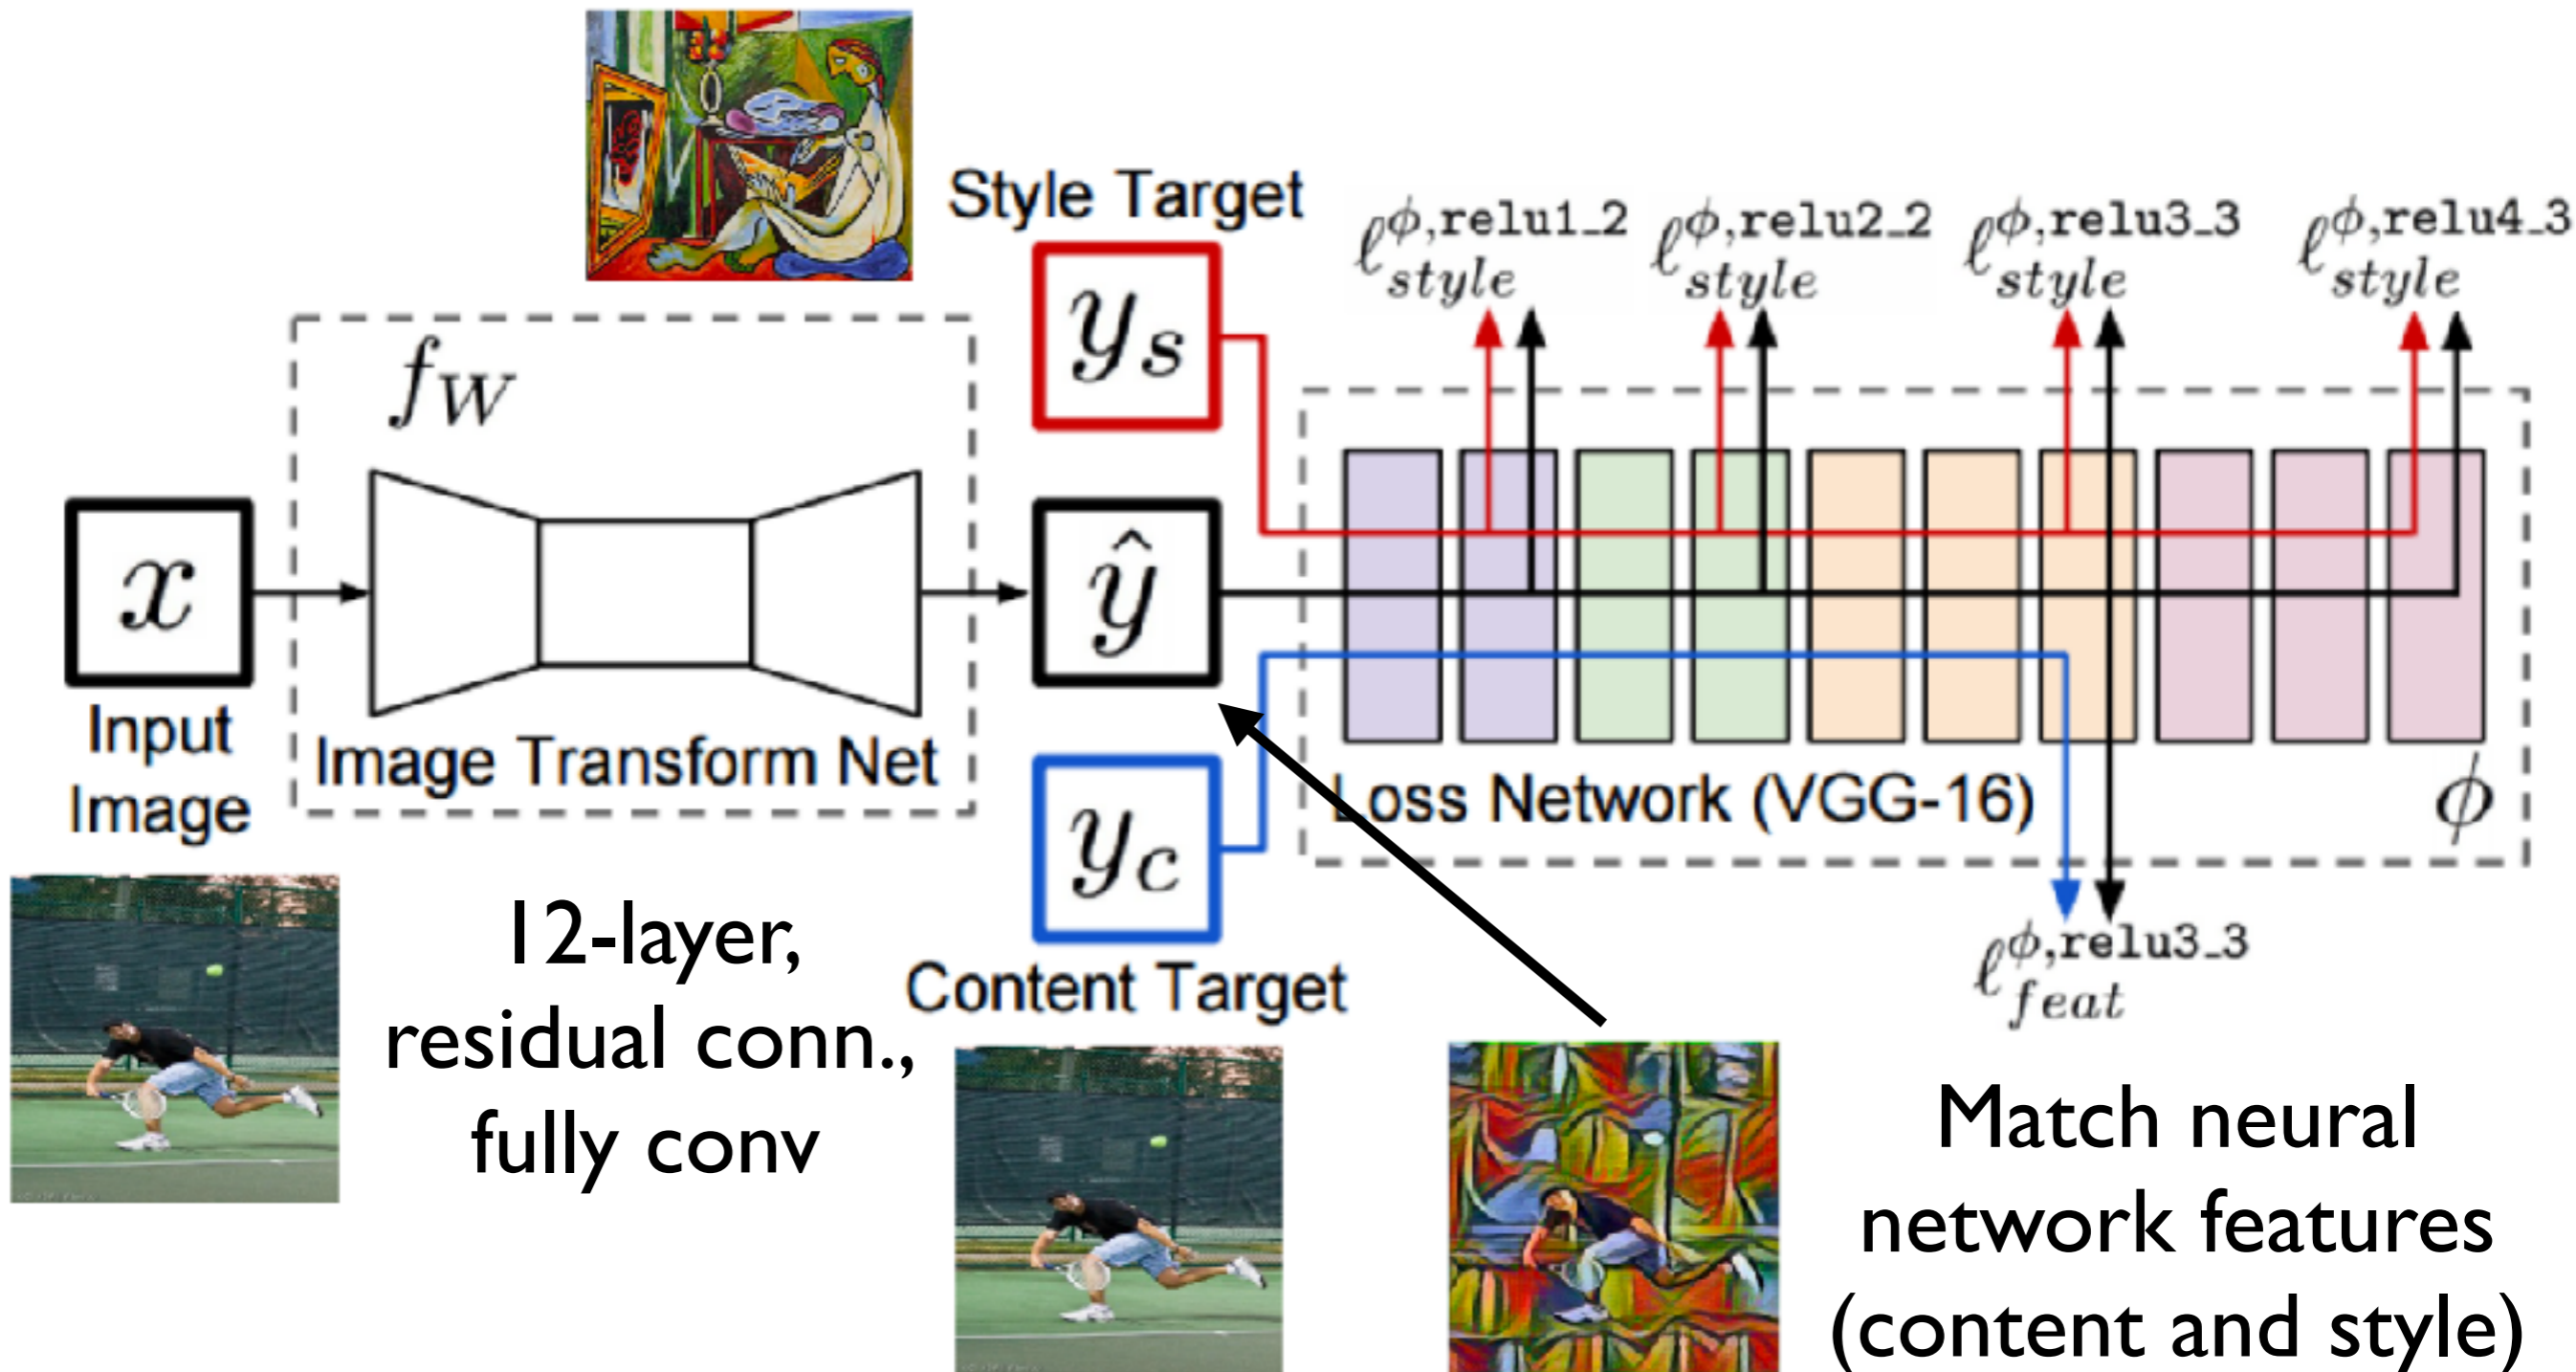
$$E(\text{style}) = E_L = \sum (\hat{G}^L - G^L)^2$$



Match content

Match style

Feedforward Style Transfer



Feedforward Style Transfer





Style

The Starry Night,
Vincent van Gogh,
1889



Style

The Muse,
Pablo Picasso,
1935



Content

Gatys

Johnson

Content

Gatys

Johnson

Comparable results to with much lower computational costs
(single feedforward pass vs 1000s of backprop iterations)

Super-Resolution

- Small networks are generally good at sharpening edges and can work well for small factor (e.g., 2) super-resolution
- Better results can be achieved by using deeper networks, + more sophisticated loss functions (perceptual loss, GANs)



Original

Bicubic

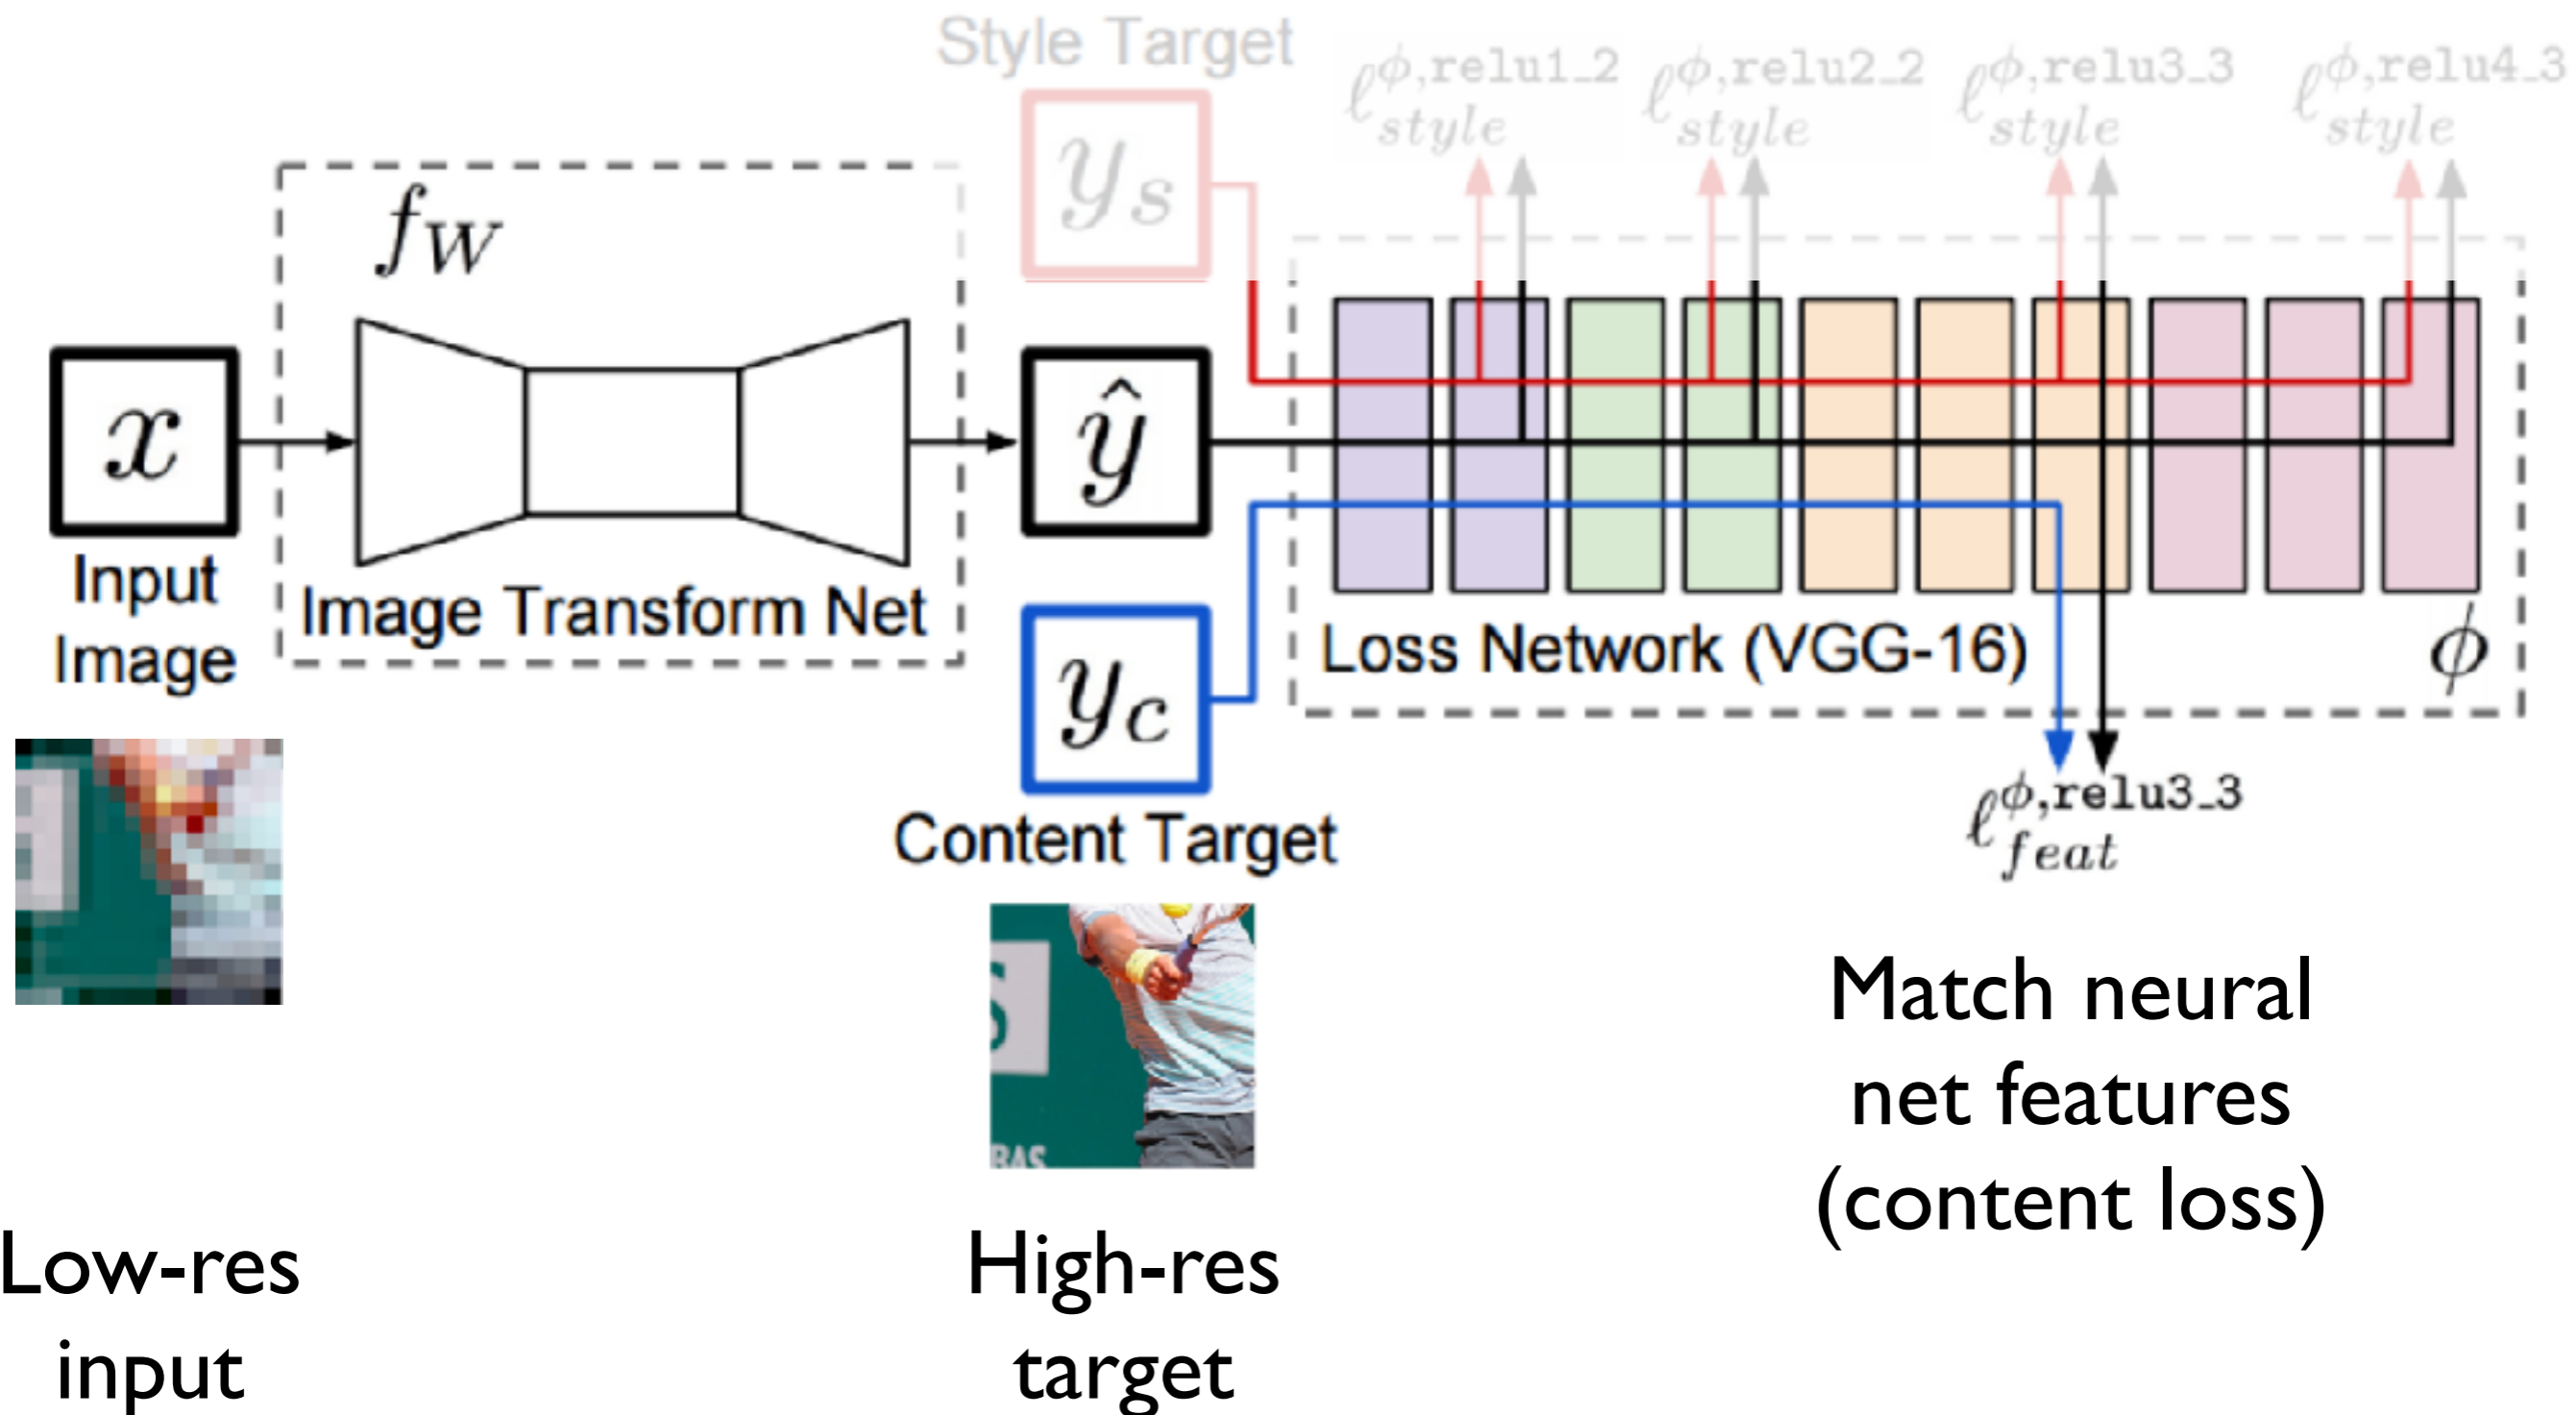
SRCNN

Johnson et al*

*12-layer, residual conn., fully conv, VGG loss [Johnson et al. 2016] 18

Super-Resolution with VGG loss

Note: style loss not used for super-res



Super-Resolution with VGG loss

- Results



Ground
Truth

Bicubic

SRCNN

Johnson et al.
VGG loss

Super-Resolution with VGG loss



SRCNN



Johnson

Super-Resolution with VGG loss



Ground Truth

This image

Set5 mean

Set14 mean

BSD100 mean

Bicubic

22.75 / 0.5946

23.80 / 0.6455

22.37 / 0.5518

22.11 / 0.5322

Ours (ℓ_{pixel})

23.42 / 0.6168

24.77 / 0.6864

23.02 / 0.5787

22.54 / 0.5526

Ours (ℓ_{feat})

21.90 / 0.6083

23.26 / 0.7058

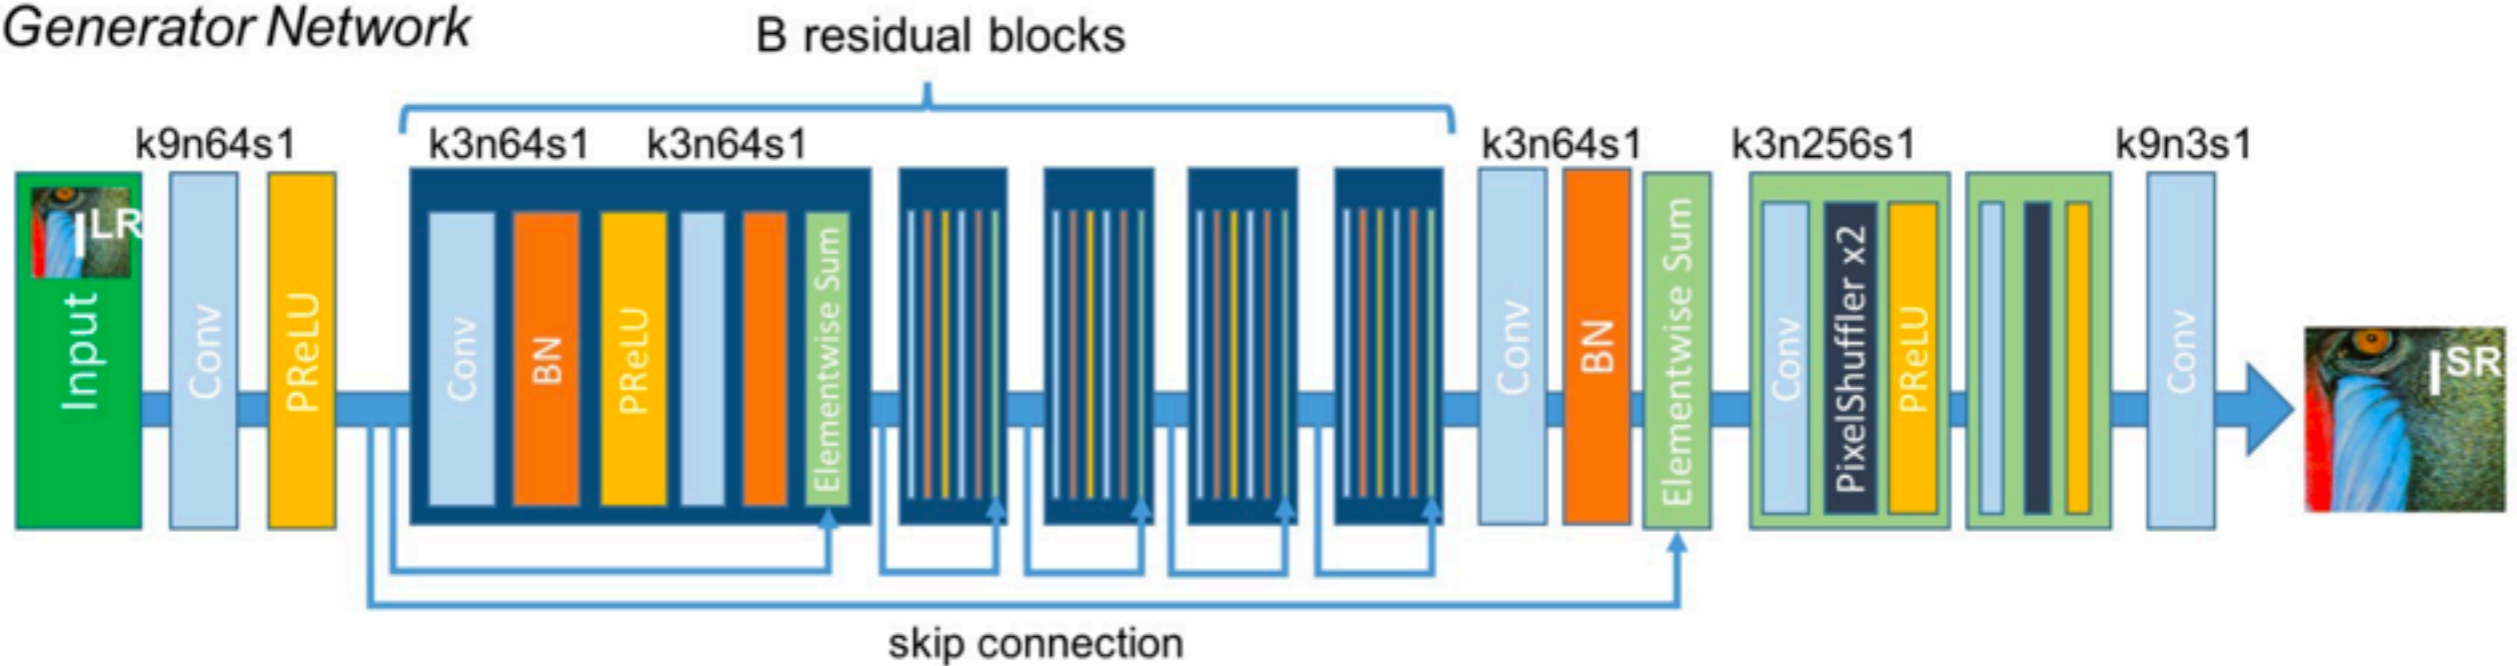
21.64 / 0.5837

21.35 / 0.5474


SRGAN

- Can we train a neural network to define the loss?

Generator Network



Discriminator Network



SRResNet
(23.53dB/0.7832)

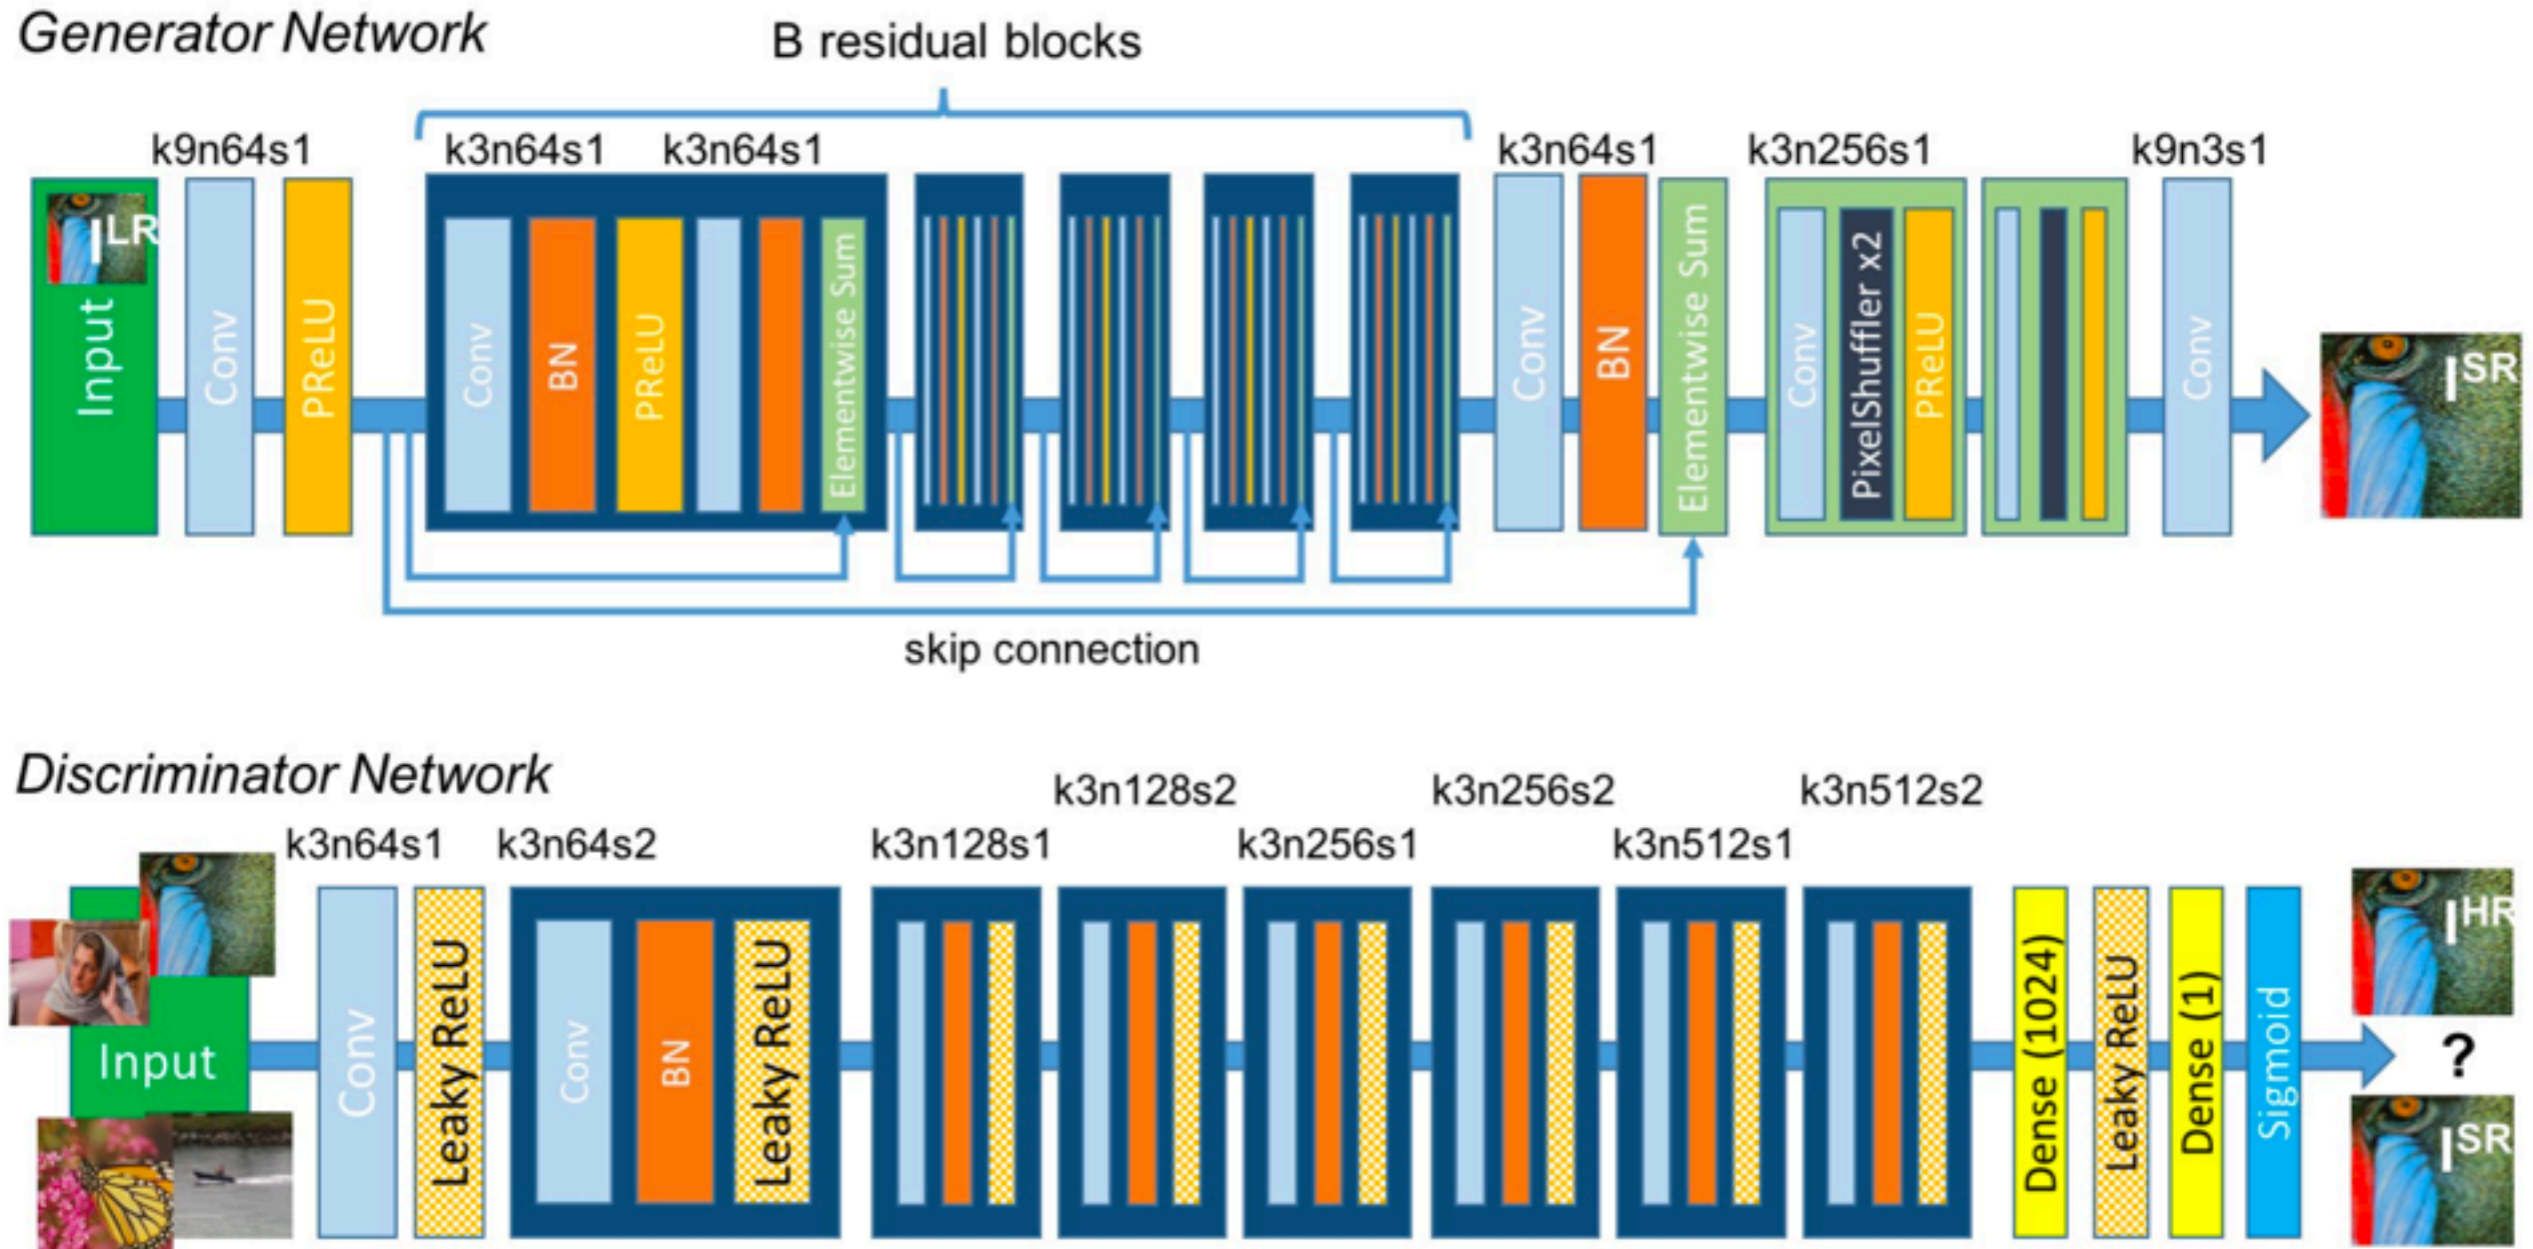


SRGAN
(21.15dB/0.6868)



SRGAN

Generator performs super-resolution, tries to fool discriminator



Discriminator tries to detect real vs super-resolved images

Generative Adversarial Networks

Setup: Assume we have data x_i drawn from distribution $p_{\text{data}}(x)$. Want to sample from p_{data} .

Generative Adversarial Networks

Setup: Assume we have data x_i drawn from distribution $p_{\text{data}}(x)$. Want to sample from p_{data} .

Idea: Introduce a latent variable z with simple prior $p(z)$.

Sample $z \sim p(z)$ and pass to a **Generator Network** $x = G(z)$

Then x is a sample from the **Generator distribution** p_G . Want $p_G = p_{\text{data}}$!


Generative Adversarial Networks

Setup: Assume we have data x_i drawn from distribution $p_{\text{data}}(x)$. Want to sample from p_{data} .

Idea: Introduce a latent variable z with simple prior $p(z)$.

Sample $z \sim p(z)$ and pass to a **Generator Network** $x = G(z)$

Then x is a sample from the **Generator distribution** p_G . Want $p_G = p_{\text{data}}$!



Train **Generator Network** G to convert z into fake data x sampled from p_G

Generative Adversarial Networks

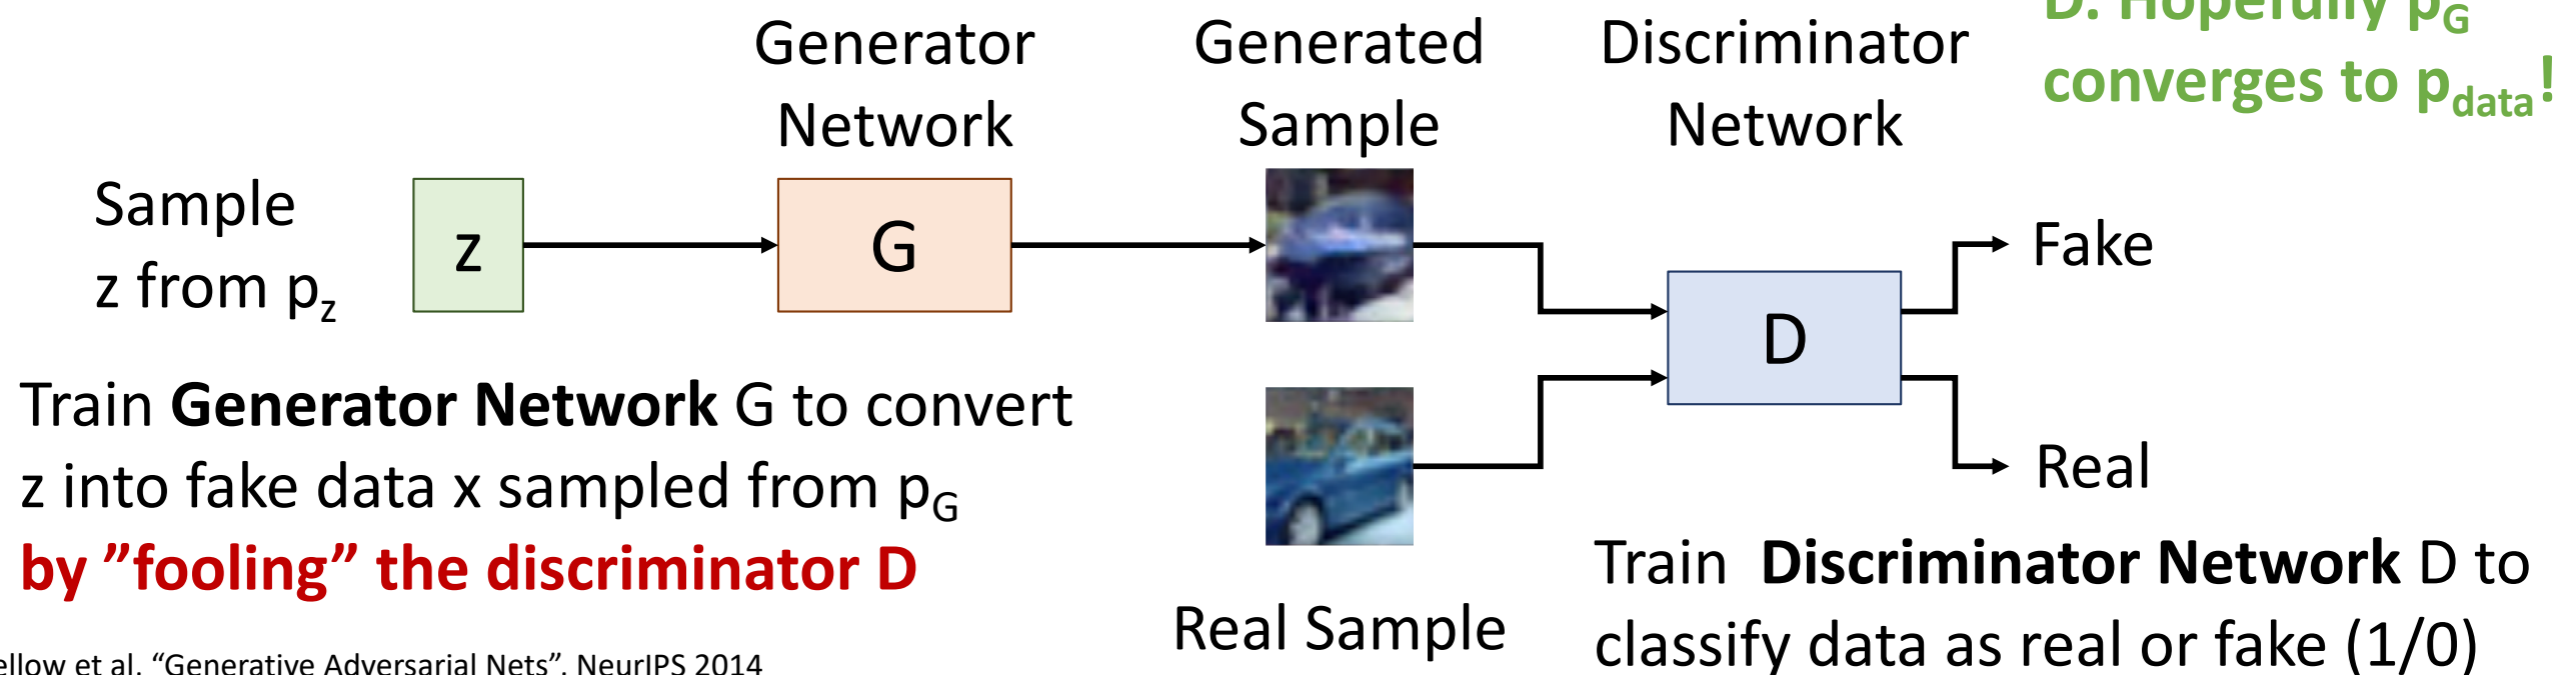
Setup: Assume we have data x_i drawn from distribution $p_{\text{data}}(x)$. Want to sample from p_{data} .

Idea: Introduce a latent variable z with simple prior $p(z)$.

Sample $z \sim p(z)$ and pass to a **Generator Network** $x = G(z)$

Then x is a sample from the **Generator distribution** p_G . Want $p_G = p_{\text{data}}$!

Jointly train G and D. Hopefully p_G converges to p_{data} !



Goodfellow et al, "Generative Adversarial Nets", NeurIPS 2014

Generative Adversarial Networks: Training Objective

Jointly train generator G and discriminator D with a **minimax game**

$$\min_G \max_D \left(E_{x \sim p_{data}} [\log D(x)] + E_{z \sim p(z)} \left[\log \left(1 - D(G(z)) \right) \right] \right)$$


Goodfellow et al, "Generative Adversarial Nets", NeurIPS 2014

Generative Adversarial Networks: Training Objective

Jointly train generator G and discriminator D with a **minimax game**

Discriminator wants
 $D(x) = 1$ for real data

$$\min_G \max_D \left(E_{x \sim p_{data}} [\log D(x)] + E_{z \sim p(z)} [\log (1 - D(G(z)))] \right)$$



Goodfellow et al, "Generative Adversarial Nets", NeurIPS 2014

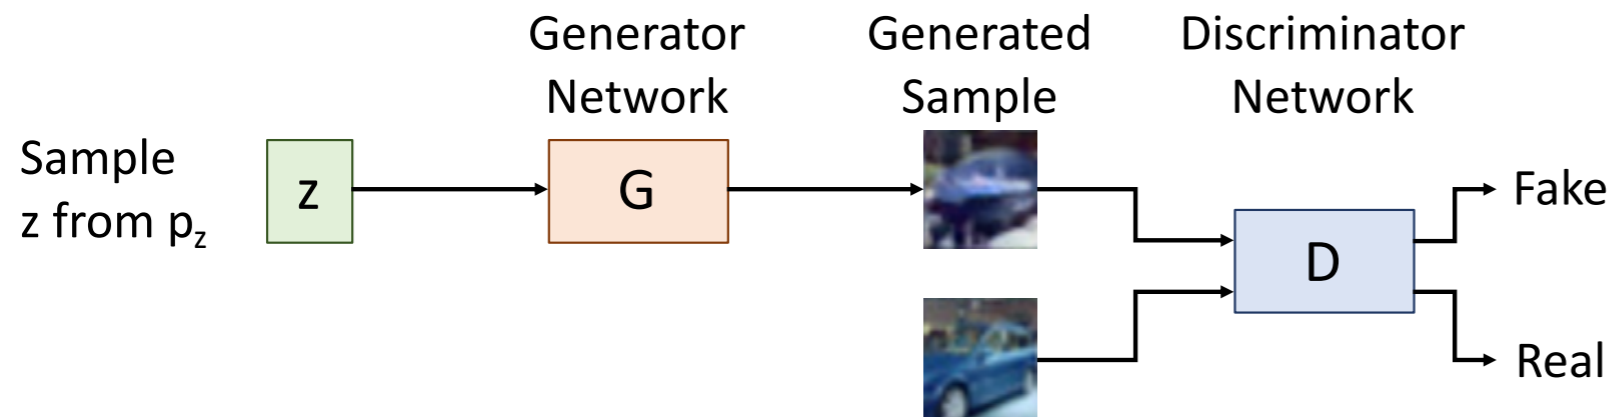
Generative Adversarial Networks: Training Objective

Jointly train generator G and discriminator D with a **minimax game**

Discriminator wants
 $D(x) = 1$ for real data

Discriminator wants
 $D(x) = 0$ for fake data

$$\min_G \max_D \left(E_{x \sim p_{data}} [\log D(x)] + E_{z \sim p(z)} [\log (1 - D(G(z)))] \right)$$



Goodfellow et al, "Generative Adversarial Nets", NeurIPS 2014


Generative Adversarial Networks: Training Objective

Jointly train generator G and discriminator D with a **minimax game**

Discriminator wants
 $D(x) = 1$ for real data

Discriminator wants
 $D(x) = 0$ for fake data

$$\min_G \max_D \left(E_{x \sim p_{data}} [\log D(x)] + E_{z \sim p(z)} \left[\log \left(1 - D(G(z)) \right) \right] \right)$$



Generator wants
 $D(x) = 1$ for fake data

Goodfellow et al, "Generative Adversarial Nets", NeurIPS 2014

Generative Adversarial Networks: Training Objective

Jointly train generator G and discriminator D with a **minimax game**

Train G and D using alternating gradient updates

$$\min_G \max_D \left(E_{x \sim p_{data}} [\log D(x)] + E_{z \sim p(z)} \left[\log \left(1 - D(G(z)) \right) \right] \right)$$
$$= \min_G \max_D V(G, D)$$

Goodfellow et al, "Generative Adversarial Nets", NeurIPS 2014

Generative Adversarial Networks: Training Objective

Jointly train generator G and discriminator D with a **minimax game**

Train G and D using alternating gradient updates

$$\min_G \max_D \left(E_{x \sim p_{data}} [\log D(x)] + E_{z \sim p(z)} \left[\log \left(1 - D(G(z)) \right) \right] \right)$$

$$= \min_G \max_D V(G, D)$$

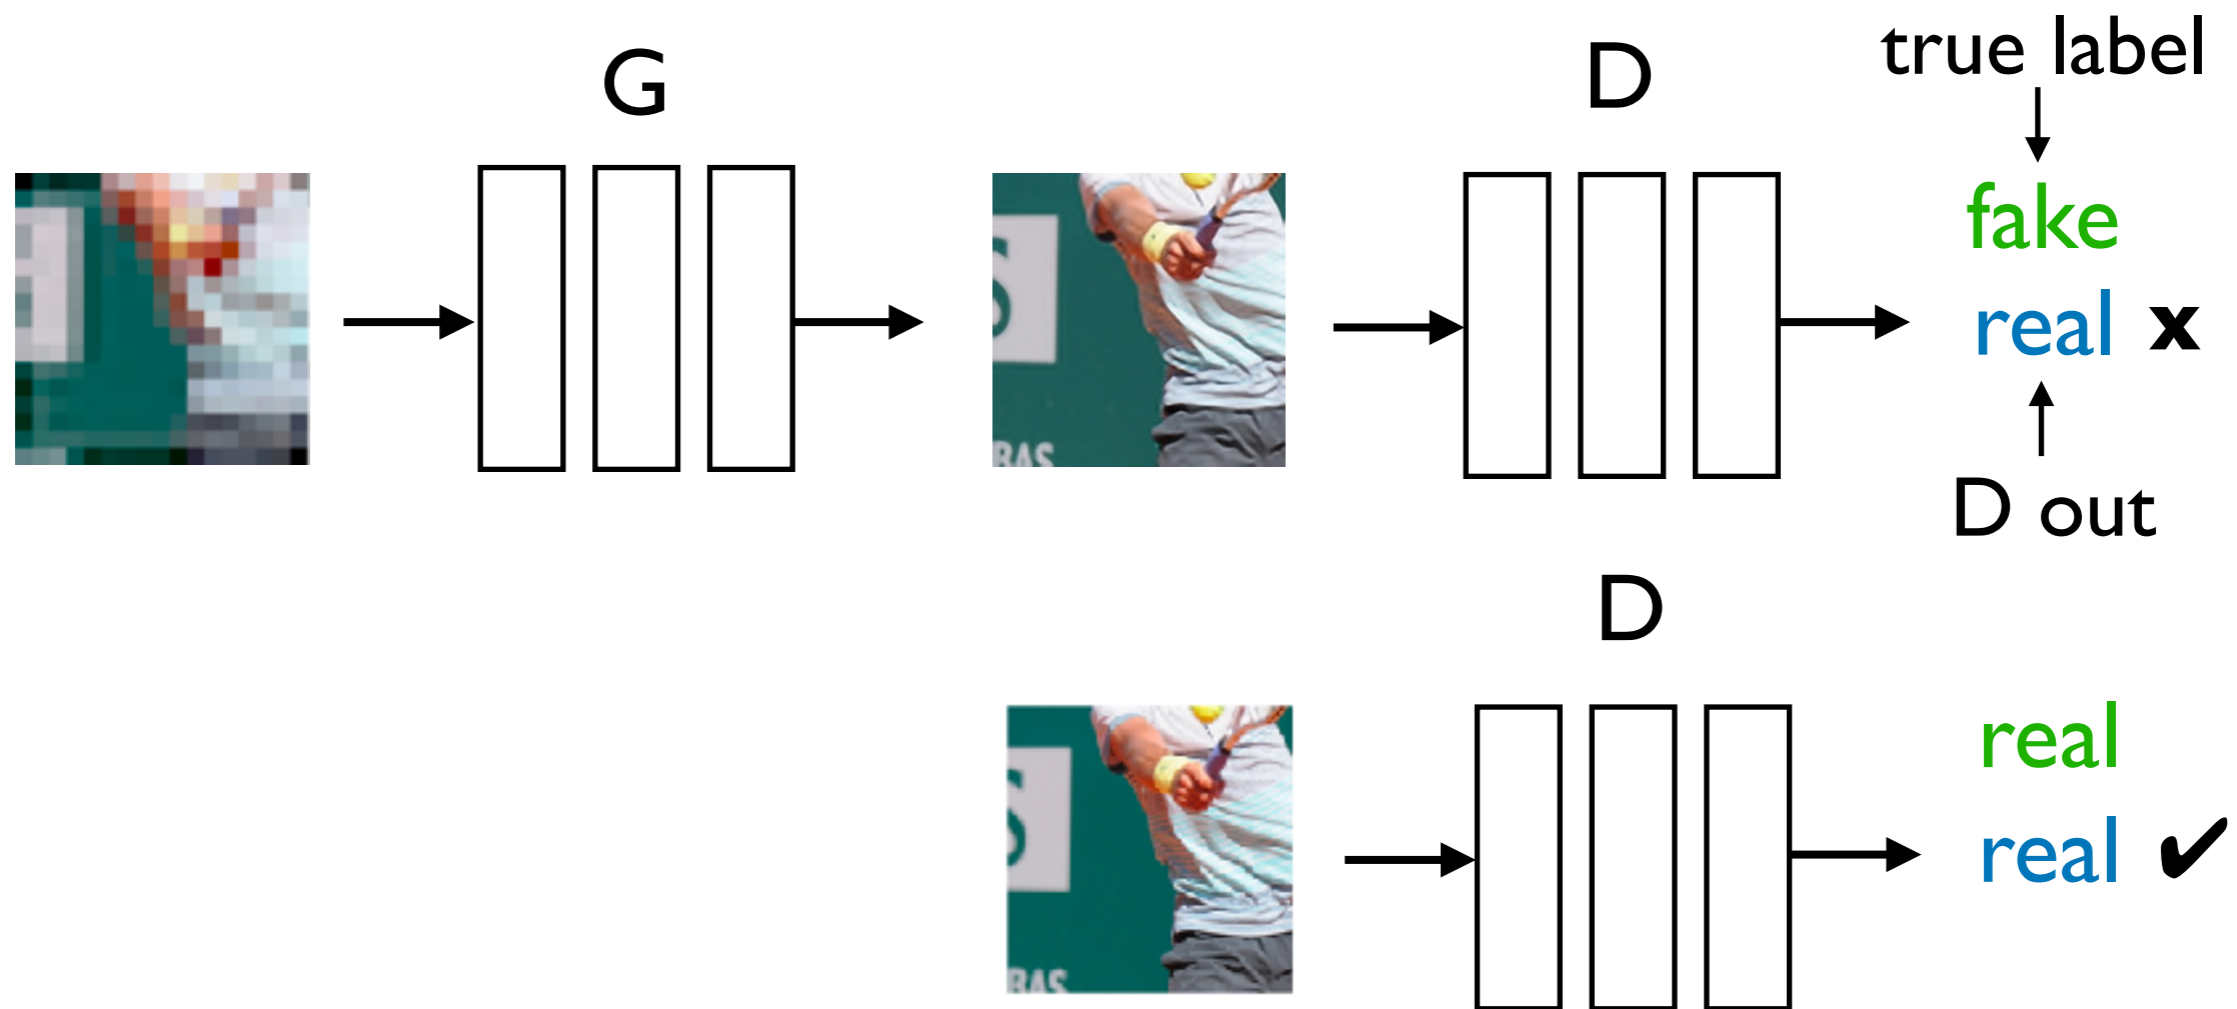
We are not minimizing any overall loss! No training curves to look at!

For t in $1, \dots, T$:

1. (Update D) $D = D + \alpha_D \frac{\partial V}{\partial D}$
2. (Update G) $G = G - \alpha_G \frac{\partial V}{\partial G}$

SRGAN

- Generator performs superres, discriminator attempts to detect real images vs super-resolved low res images




$$\min_{\theta_G} \max_{\theta_D} \mathbb{E}_{I^{HR} \sim p_{\text{train}}(I^{HR})} [\log D_{\theta_D}(I^{HR})] +$$

$$\mathbb{E}_{I^{LR} \sim p_G(I^{LR})} [\log(1 - D_{\theta_D}(G_{\theta_G}(I^{LR})))]$$

Generative Adversarial Networks: Results


Generated samples



Nearest neighbor from training set

Goodfellow et al, "Generative Adversarial Nets", NeurIPS 2014

Generative Adversarial Networks: DC-GAN



Generator

Radford et al, "Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks", ICLR 2016

Generative Adversarial Networks: DC-GAN

Samples from the model look much better!



Radford et al,
ICLR 2016

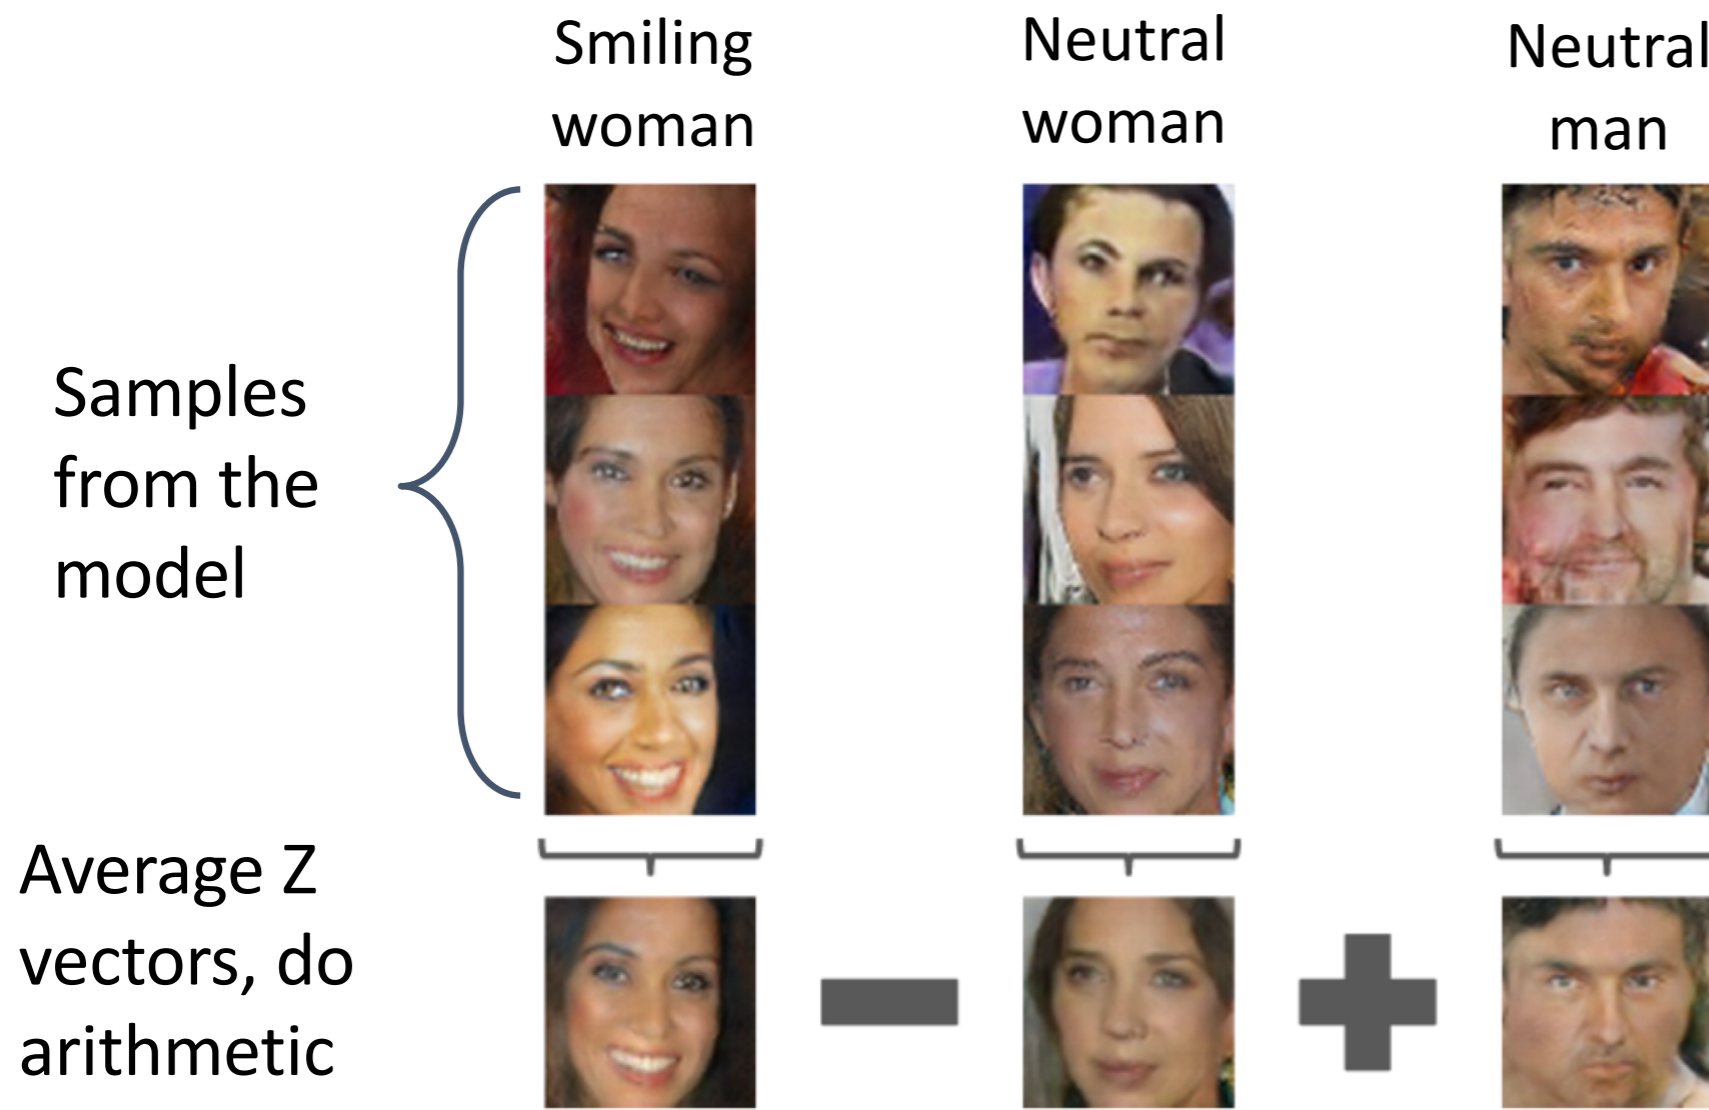
Generative Adversarial Networks: Interpolation



Interpolating
between
points in
latent z
space

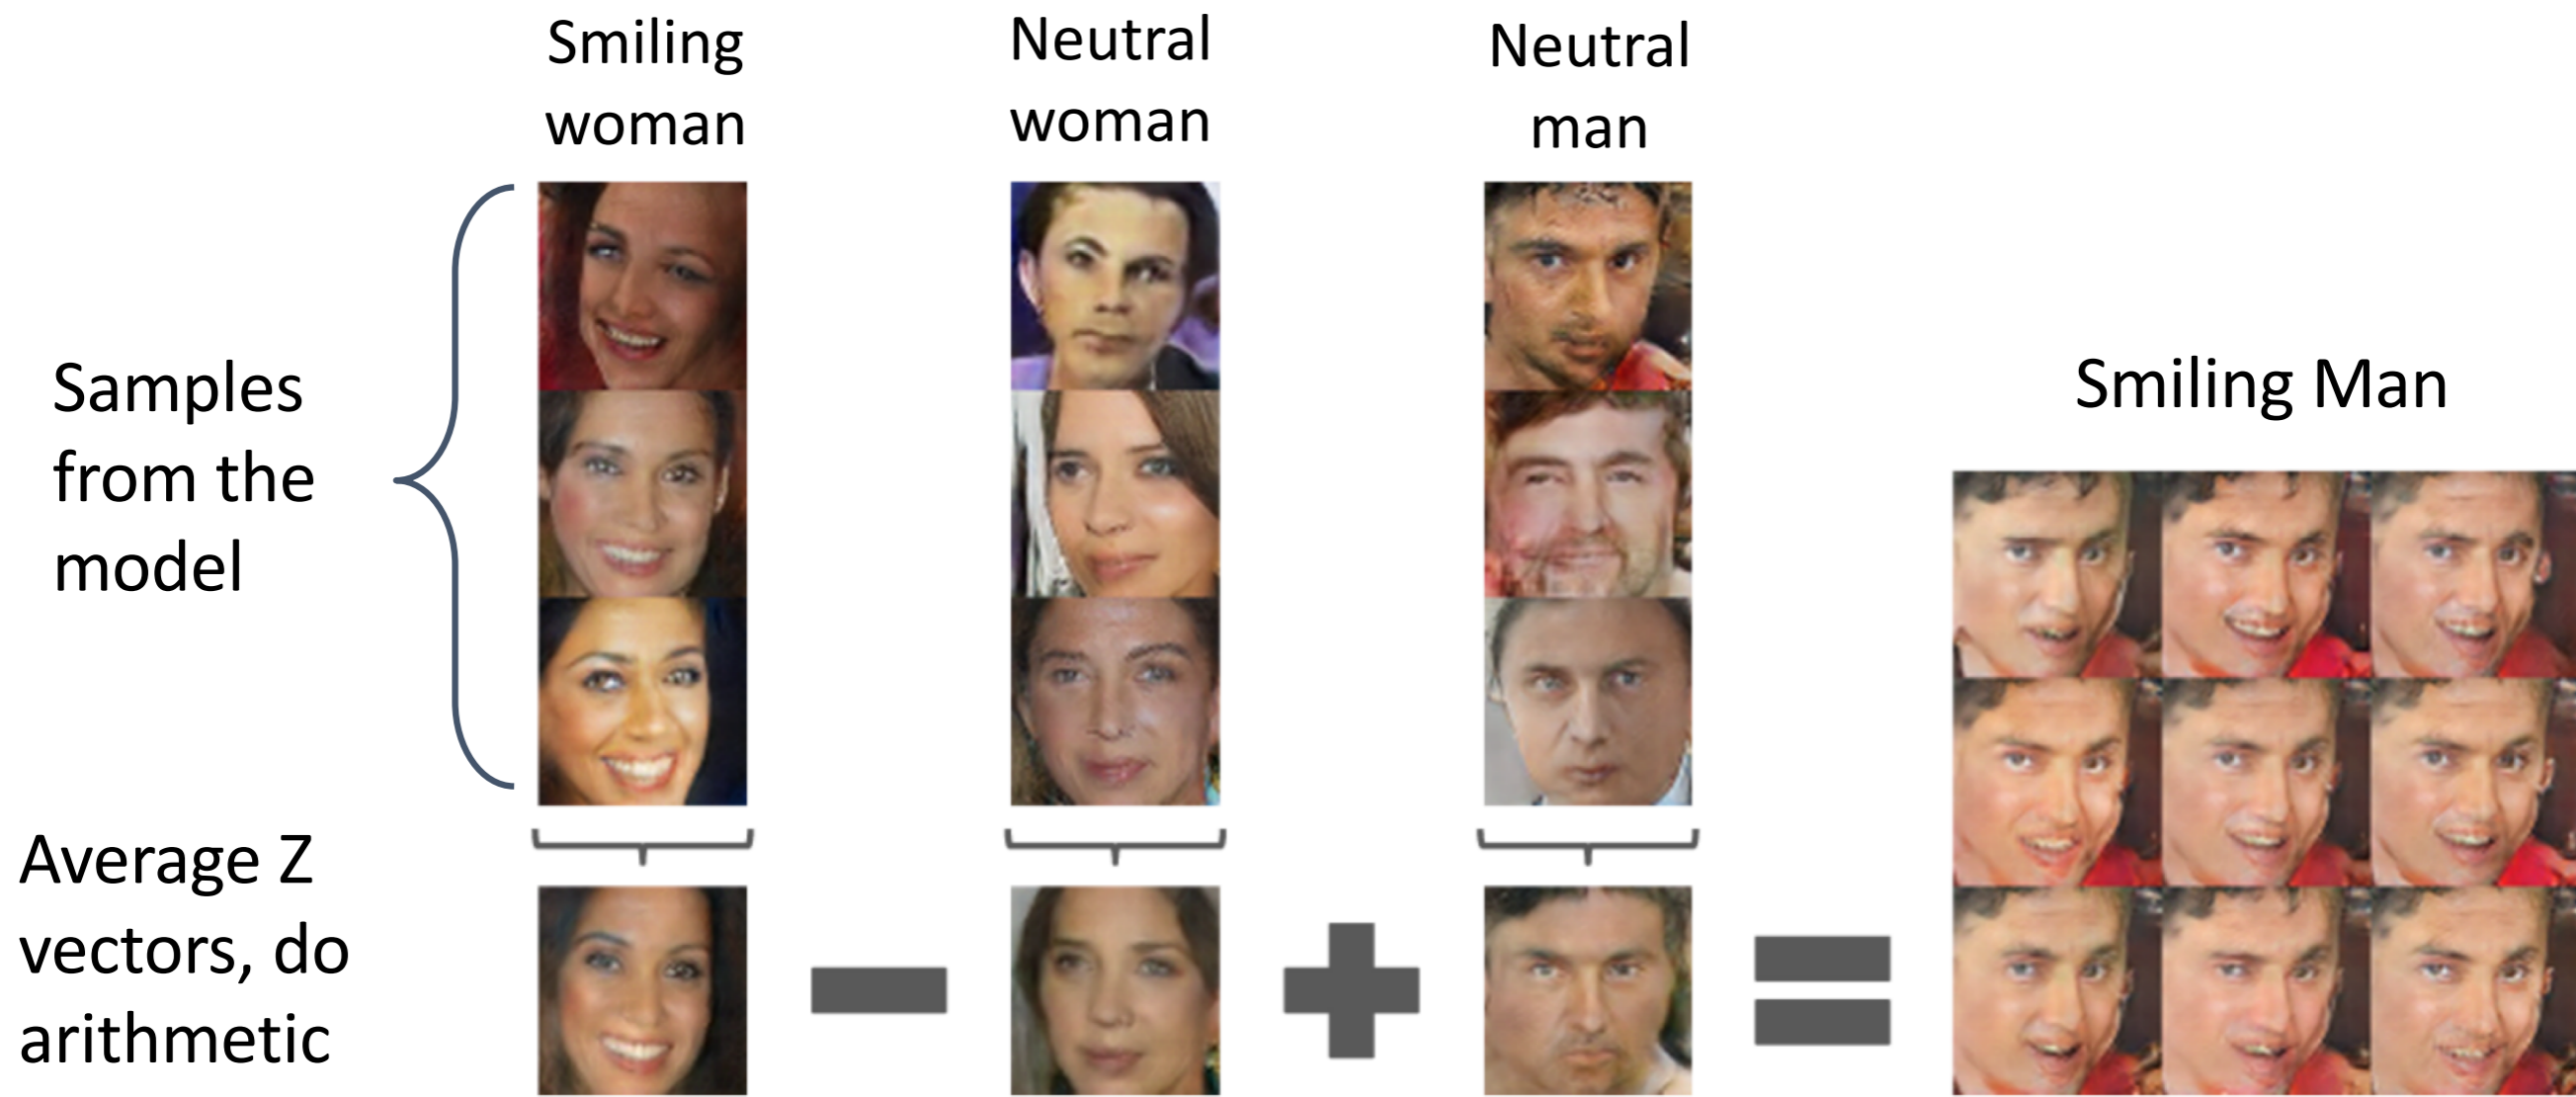
Radford et al,
ICLR 2016

Generative Adversarial Networks: Vector Math




Radford et al, ICLR 2016

Generative Adversarial Networks: Vector Math



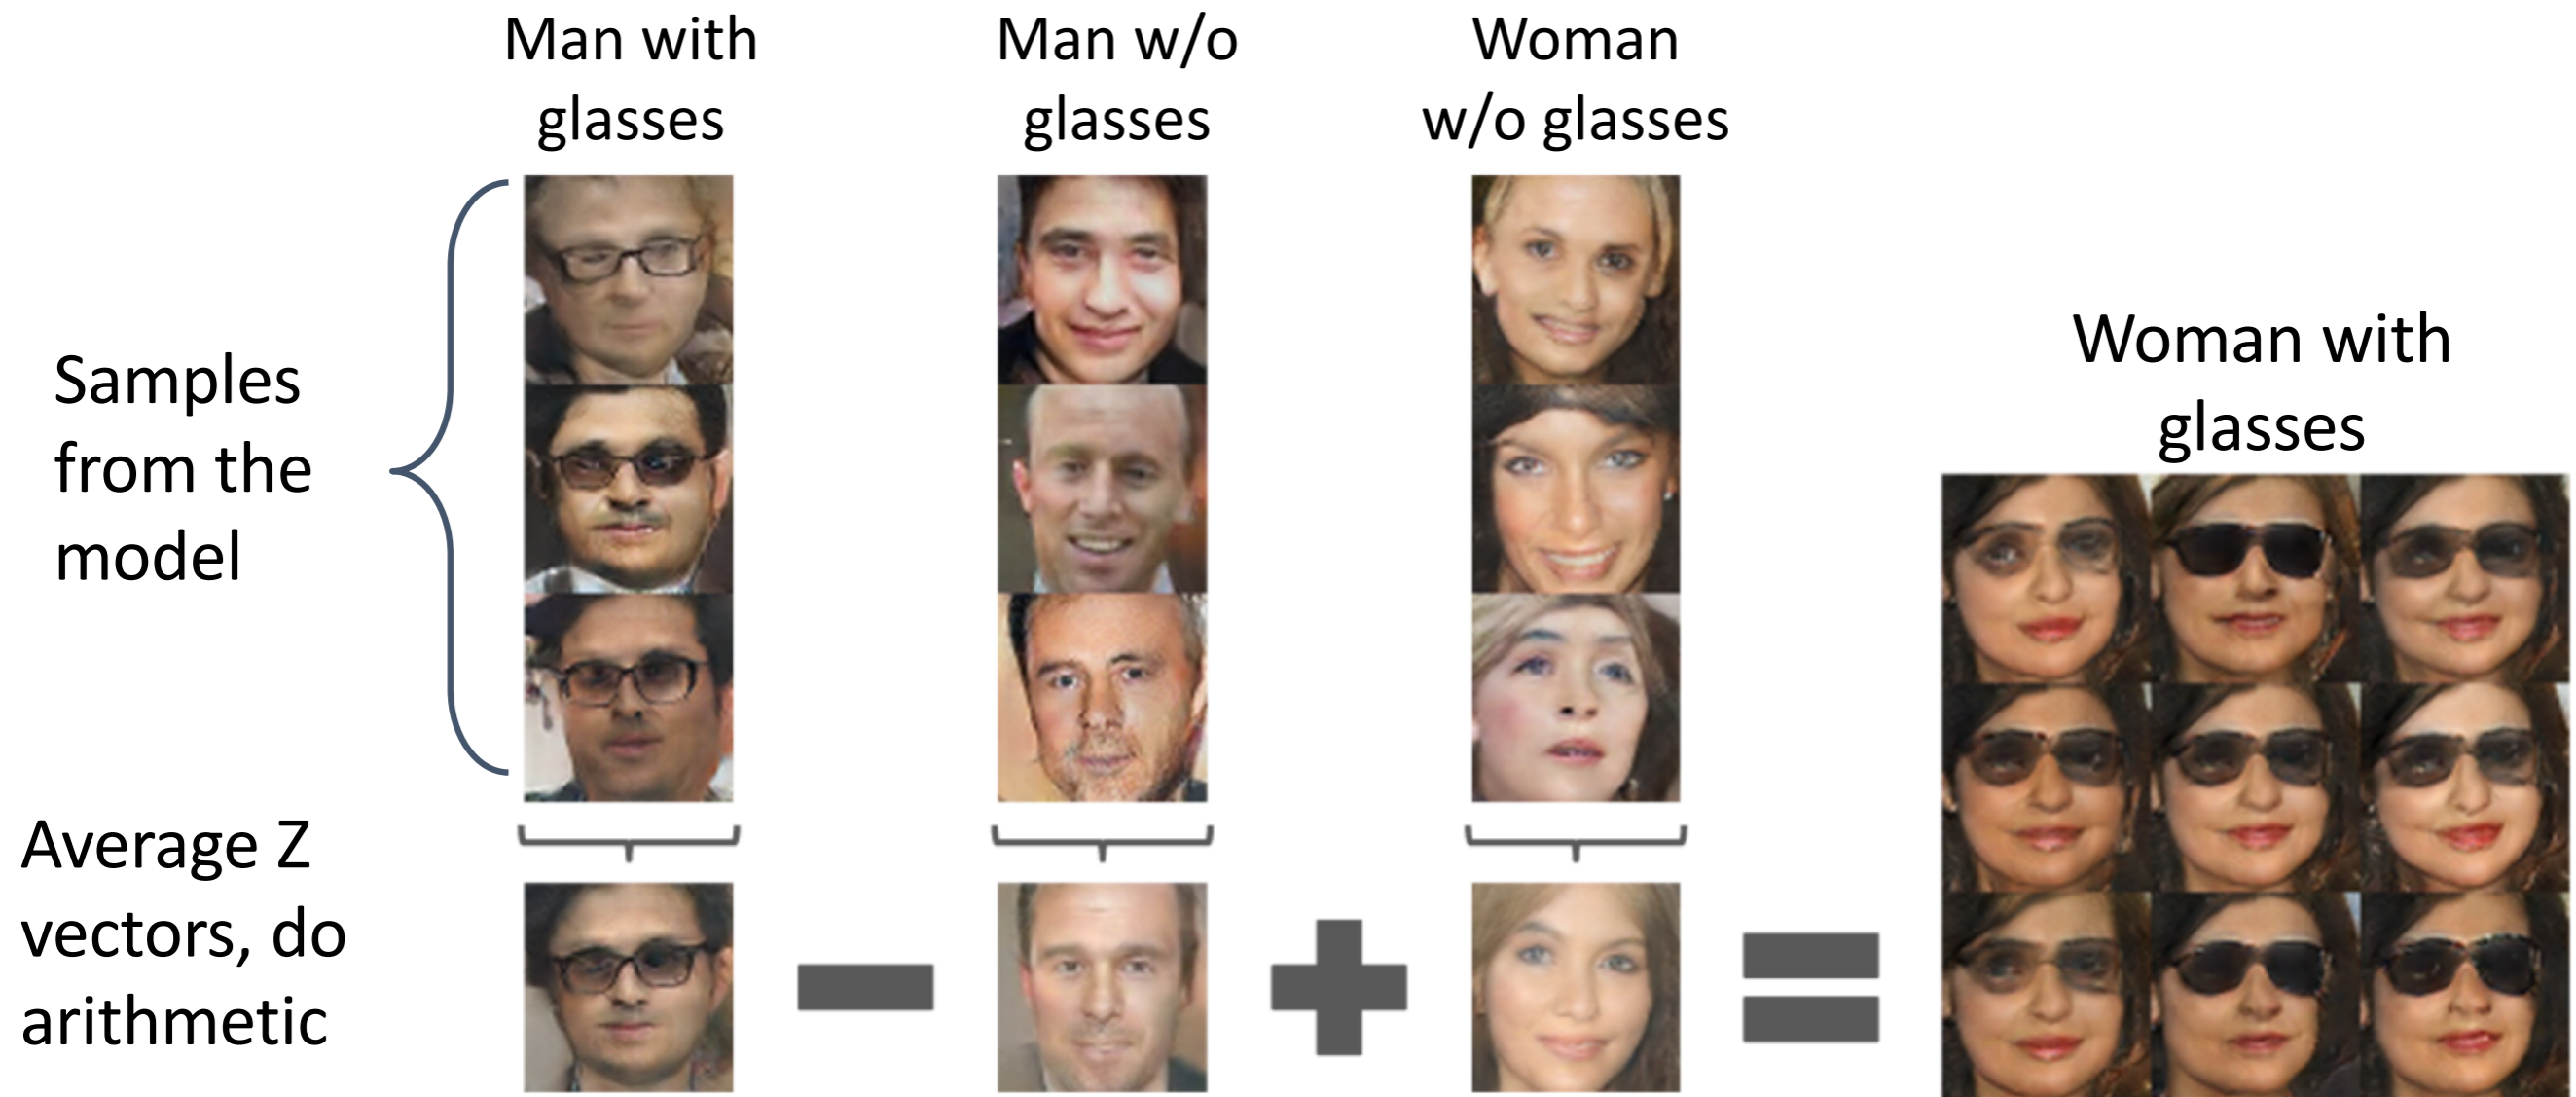
Radford et al, ICLR 2016

Generative Adversarial Networks: Vector Math




Radford et al, ICLR 2016


Generative Adversarial Networks: Vector Math




Radford et al, ICLR 2016




Odena et al.,
2016 [1]



Miyato et al.,
2017 [3]



Zhang et al.,
2018 [2]



Brock et al.,
2018 [4]

BigGAN

[Odena 2019]

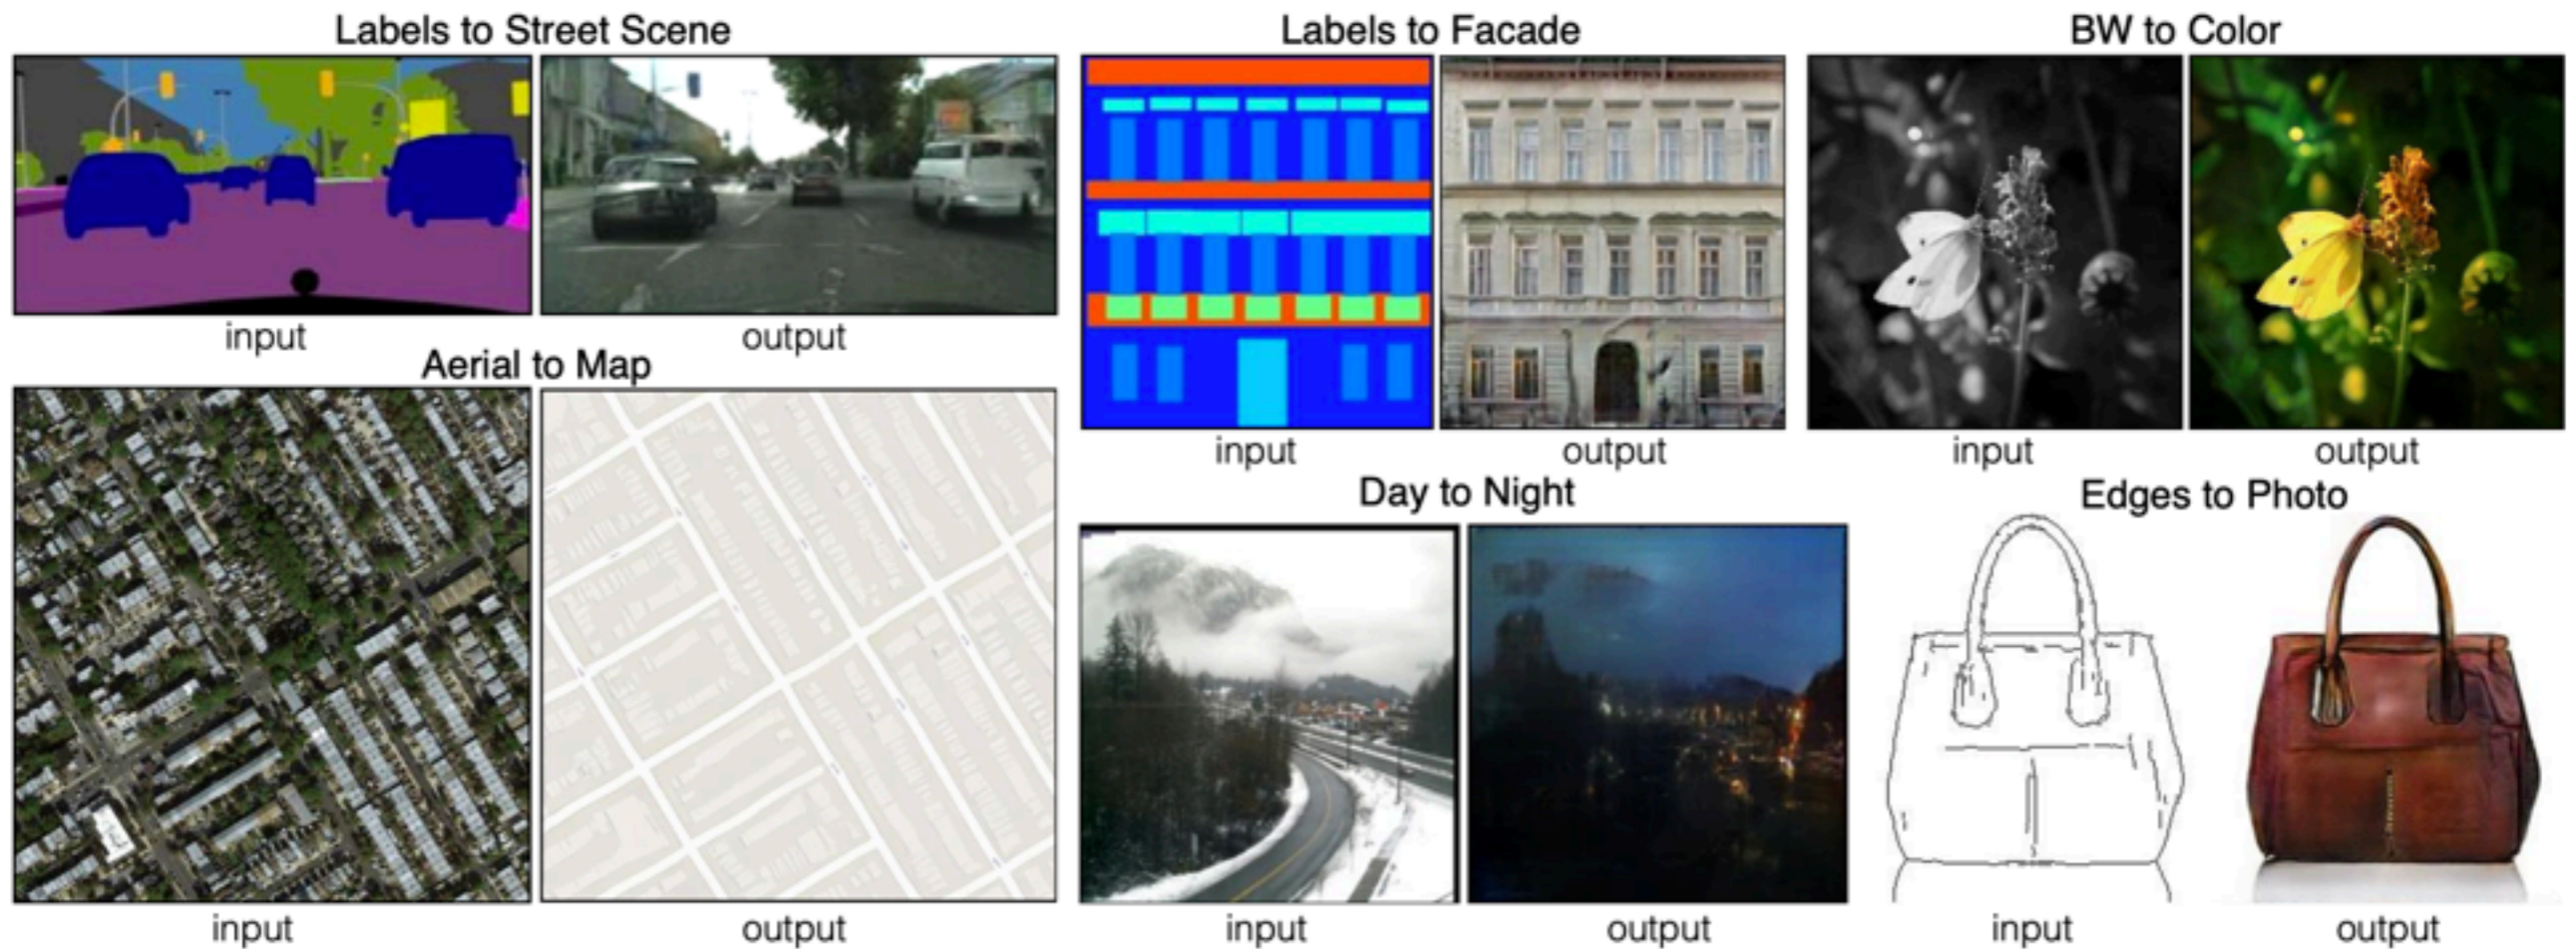


GAN Questions

- Problem 1 What are the trade-offs between GANs and other generative models?
- Problem 2 What sorts of distributions can GANs model?
- Problem 3 How can we Scale GANs beyond image synthesis?
- Problem 4 What can we say about the global convergence of the training dynamics?
- Problem 5 How should we evaluate GANs and when should we use them?
- Problem 6 How does GAN training scale with batch size?
- Problem 7 What is the relationship between GANs and adversarial examples?

Image Translation

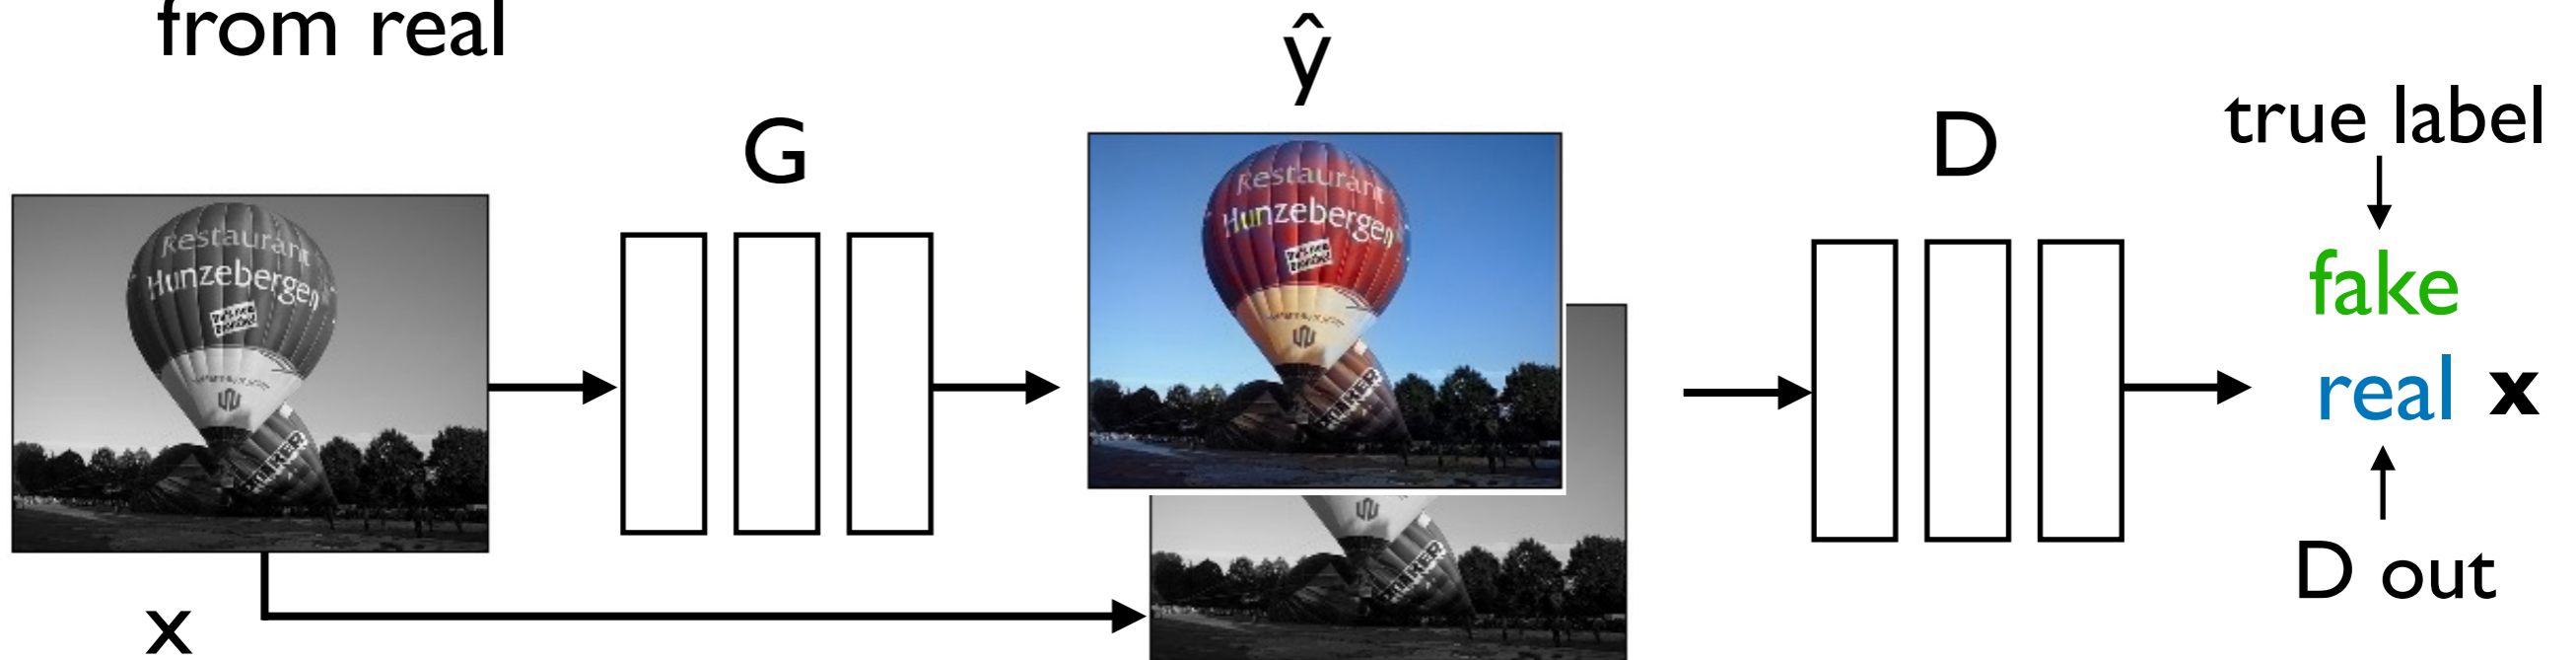
- Many problems in vision/graphics can be viewed as image translation problems



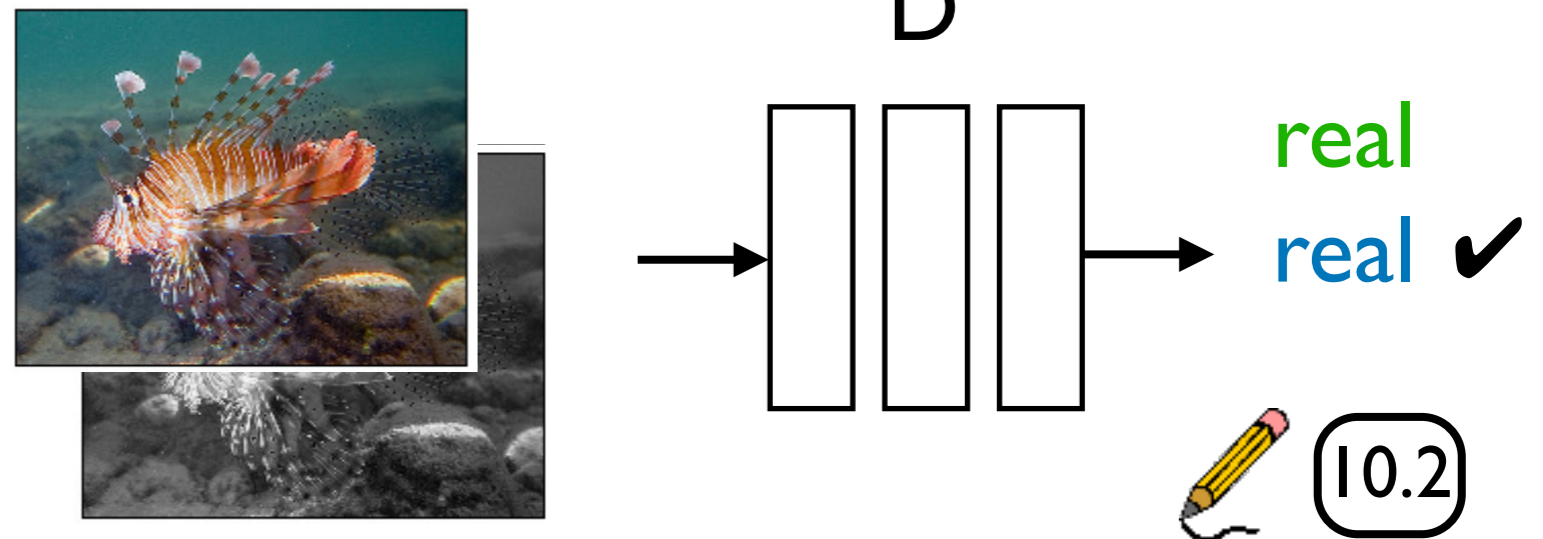
Can we build a general machine to translate images?

Image Translation

- e.g., translation from grey to color should be indistinguishable from real

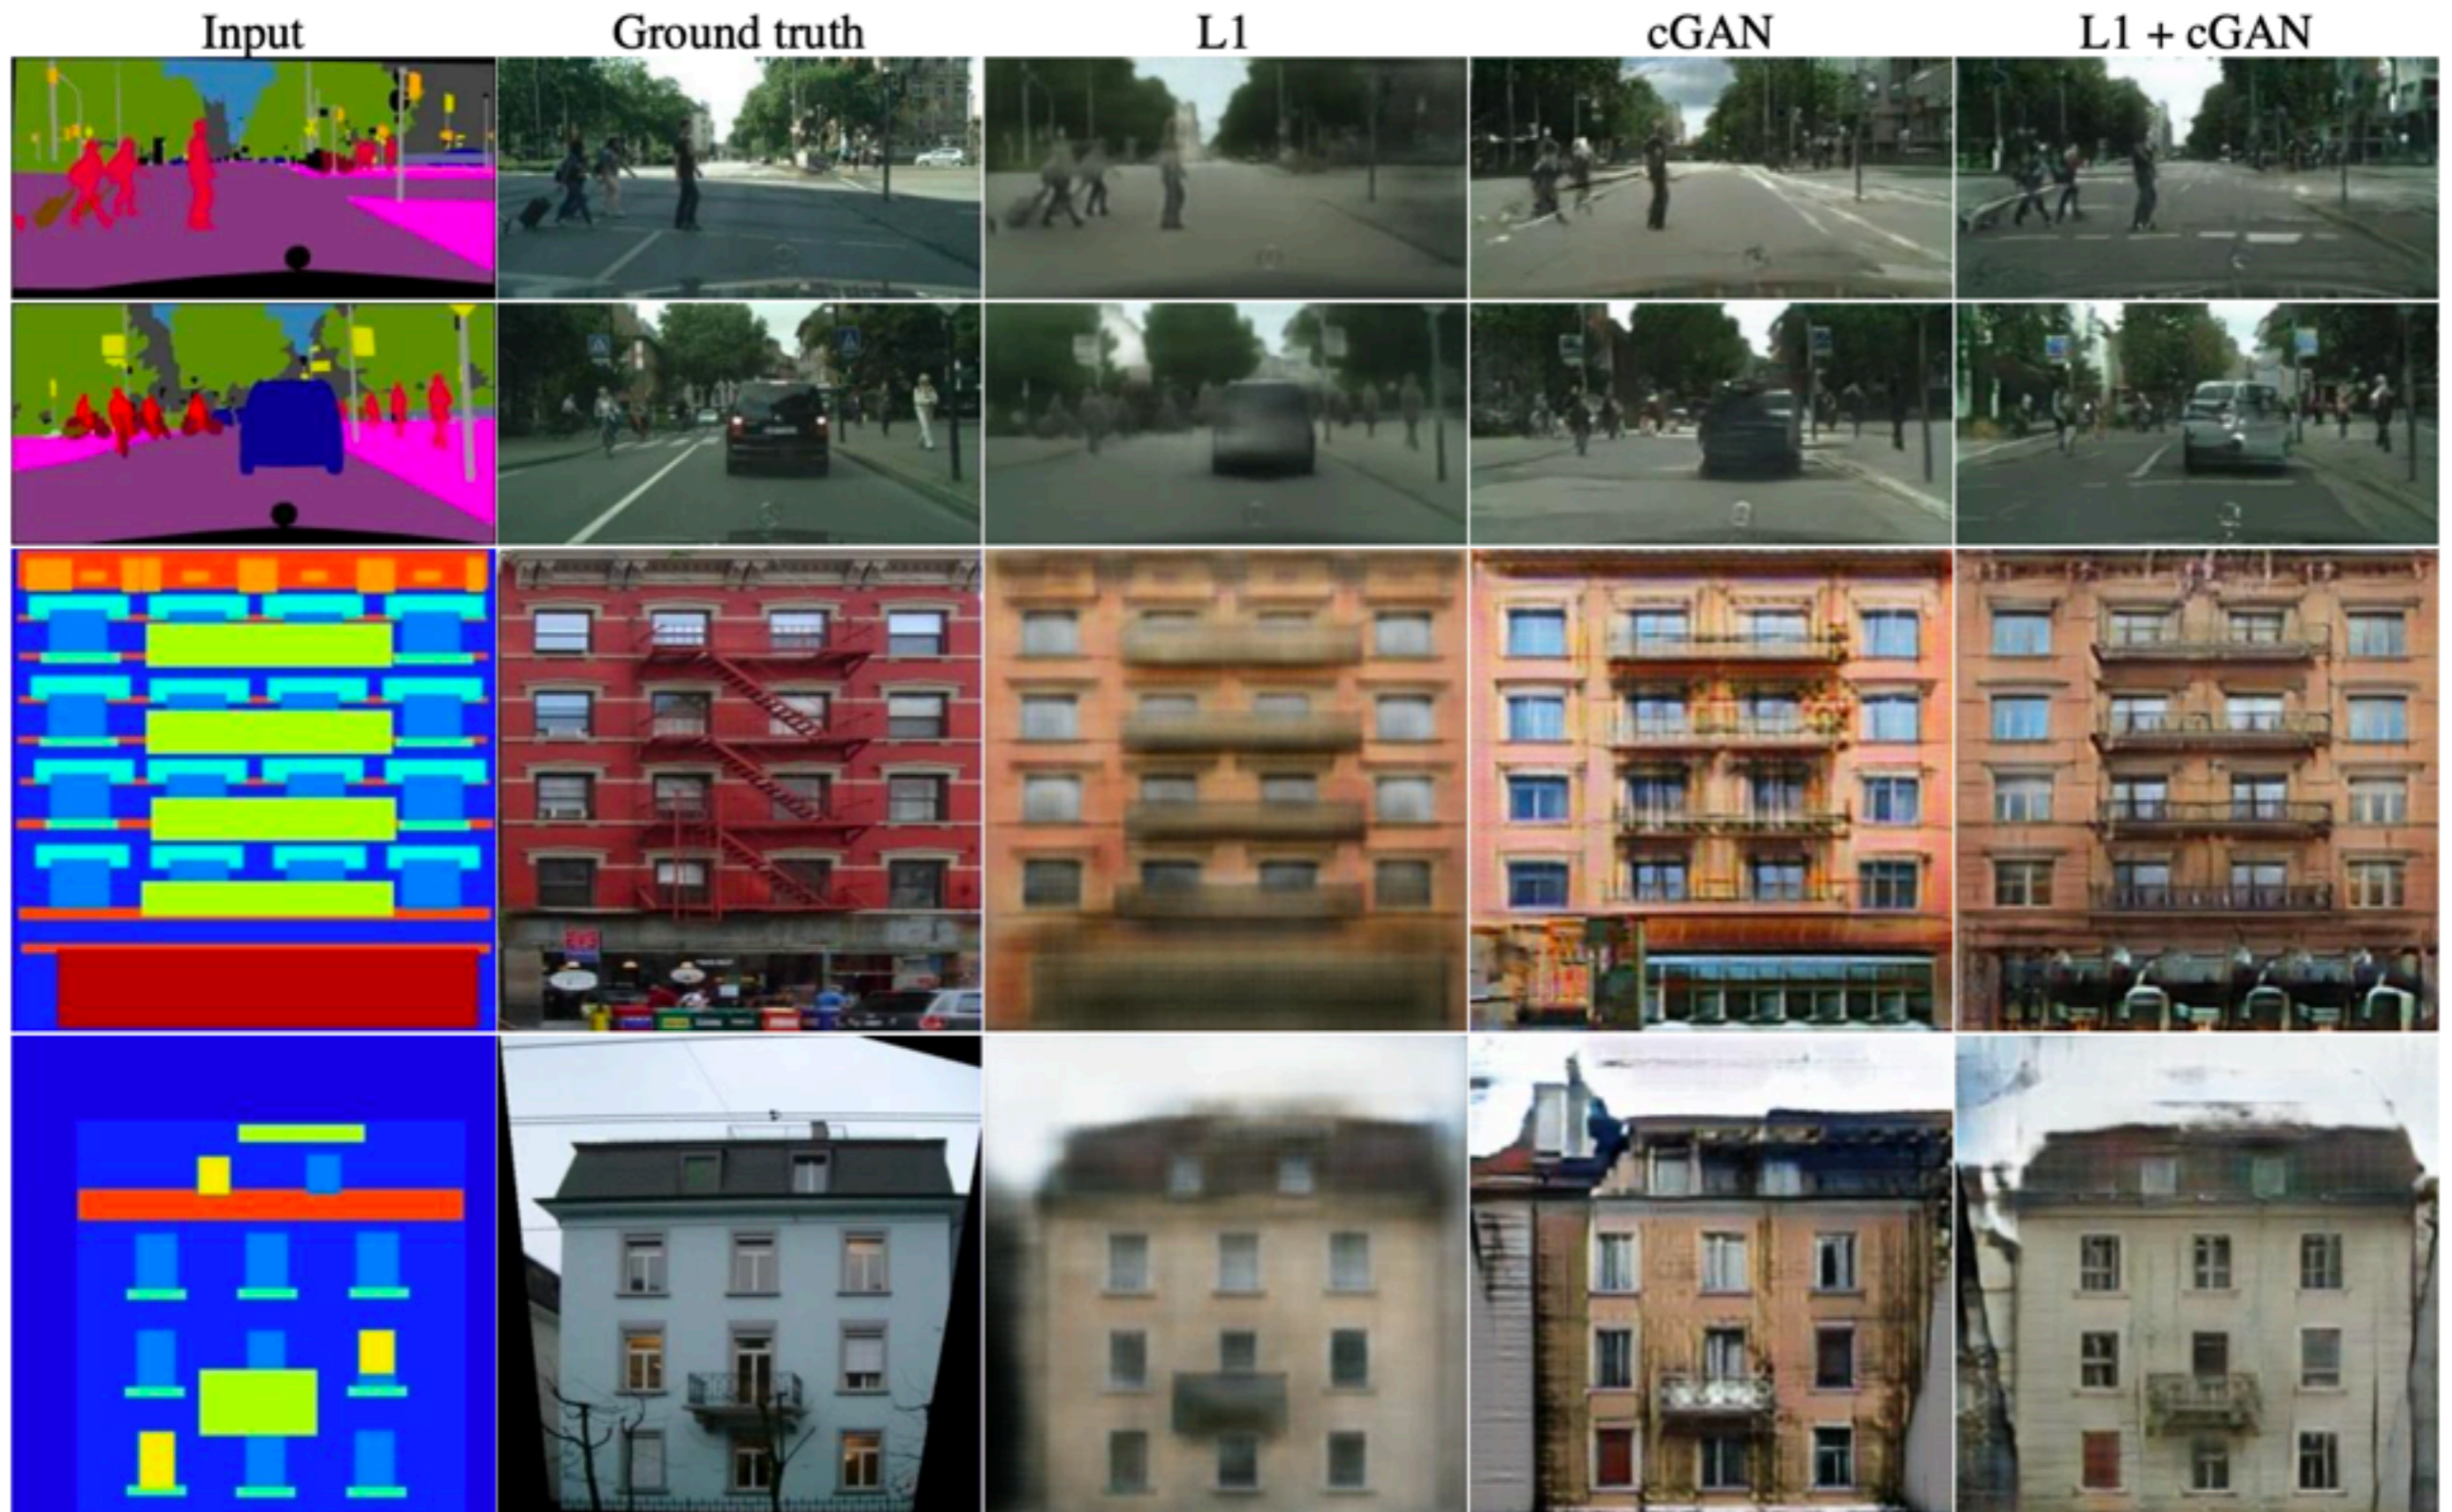


Note: pix2pix has an additional supervisory
L1 loss = $|y - \hat{y}|$



This is a (conditional) Generative Adversarial Network

pix2pix: Segmentation \rightarrow Image




pix2pix: Edges \rightarrow Image



Model trained using edge detection, but works for hand drawings:



#edges2cats by Christopher Hesse



sketch by Ivy Tsai



<https://affinelayer.com/pixsrv/>

THANKS

