What we do in the Shadows GANs

Taras Lehinevych
taras.lehinevych@railsreactor.com

RAILS
REACTOR

Who am I

ML Engineer @ Rails Reactor

ML Engineer @ [redacted] (censored)

Disclaimer!
All the provided information based on open publication and datasets

Agenda

- Motivation
- Shadow detection
- Shadow removal
- Shadow generation
- Summary

What’s wrong with shadows?

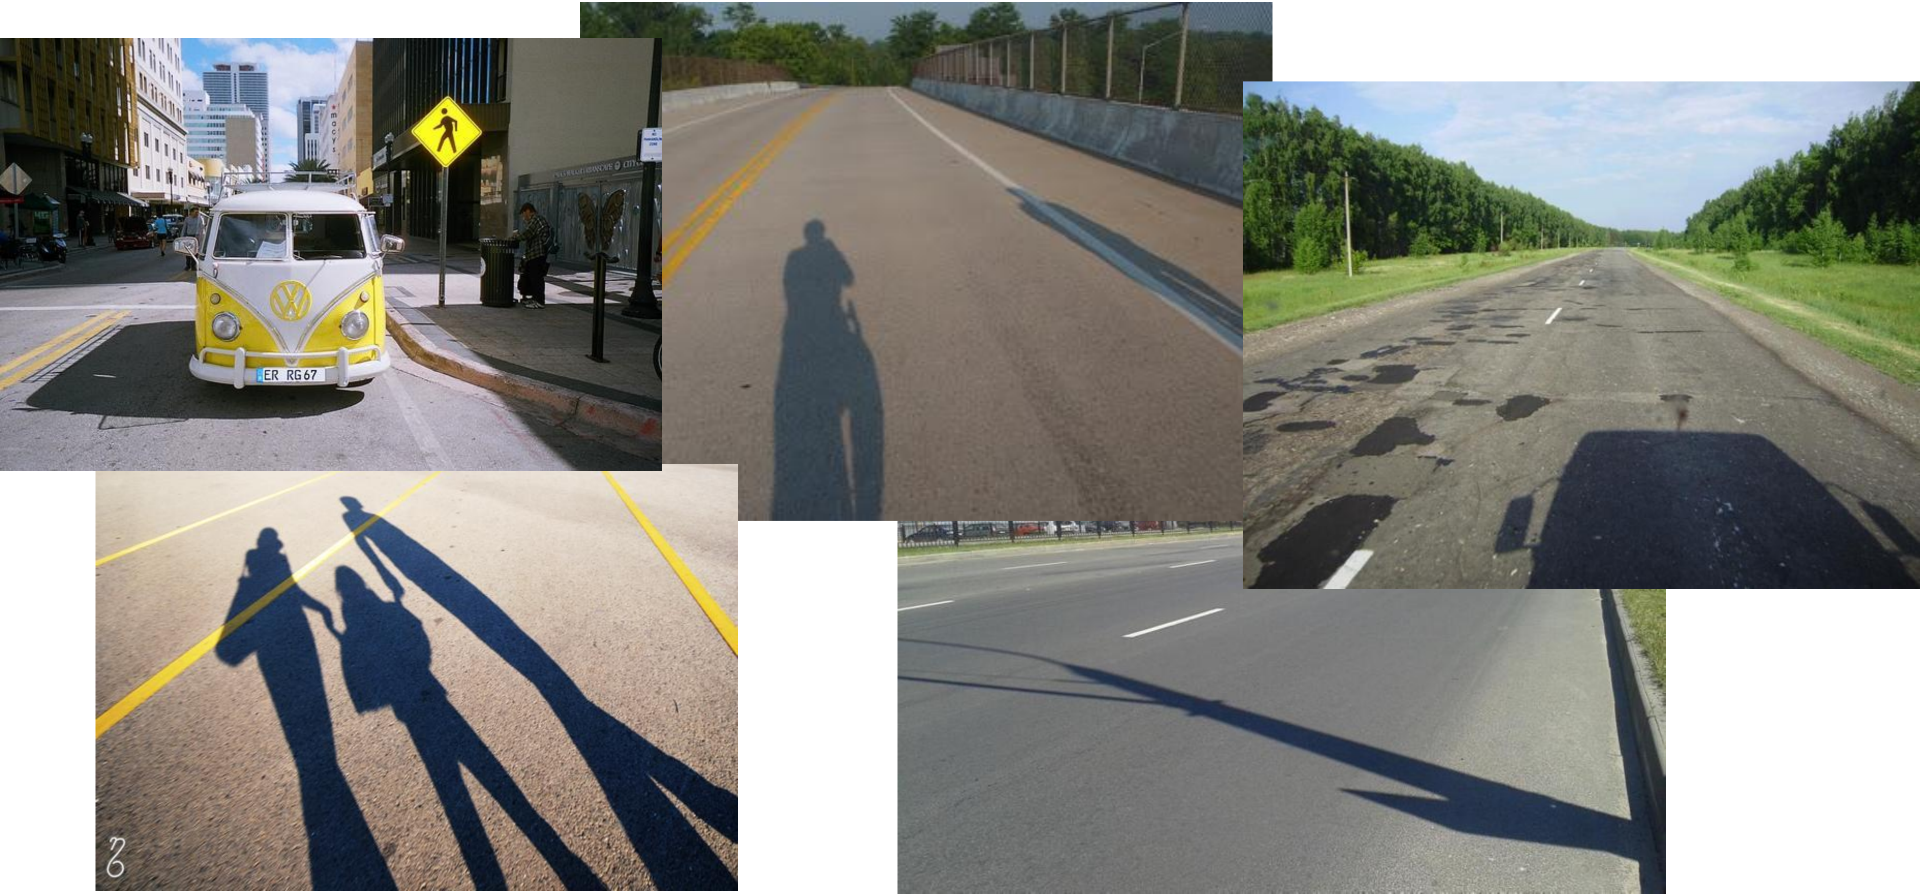

What do we want?

Datasets

SBU Dataset (link) - this new dataset contains 4,727 images (4,089 train images and 638 test images) with pixel based ground truth.

ISTD Dataset (link) - it contains 1870 triplets of shadow, shadow mask and shadow-free image under 135 different scenarios.

Approach #1 - Shadow Detection

Intrinsic image - get reflectance and illuminance constituents.

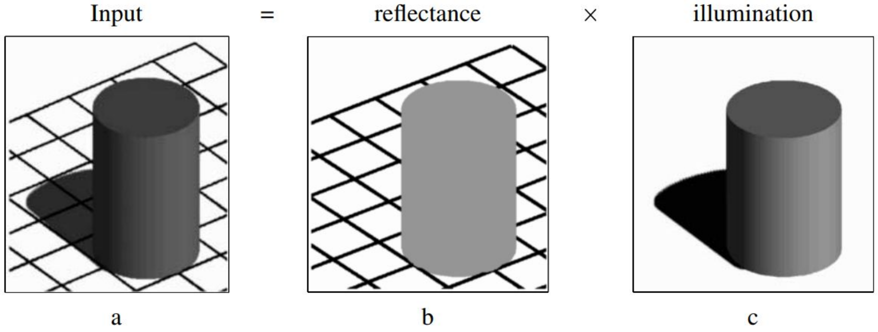

Cons:

- environment calibration
- twilight/night
- spectral properties of the camera
- modify the original color and intensity of images

Shadow Detection

[Fast Shadow Detection from a Single Image Using a Patched Convolutional Neural Network (2018)]

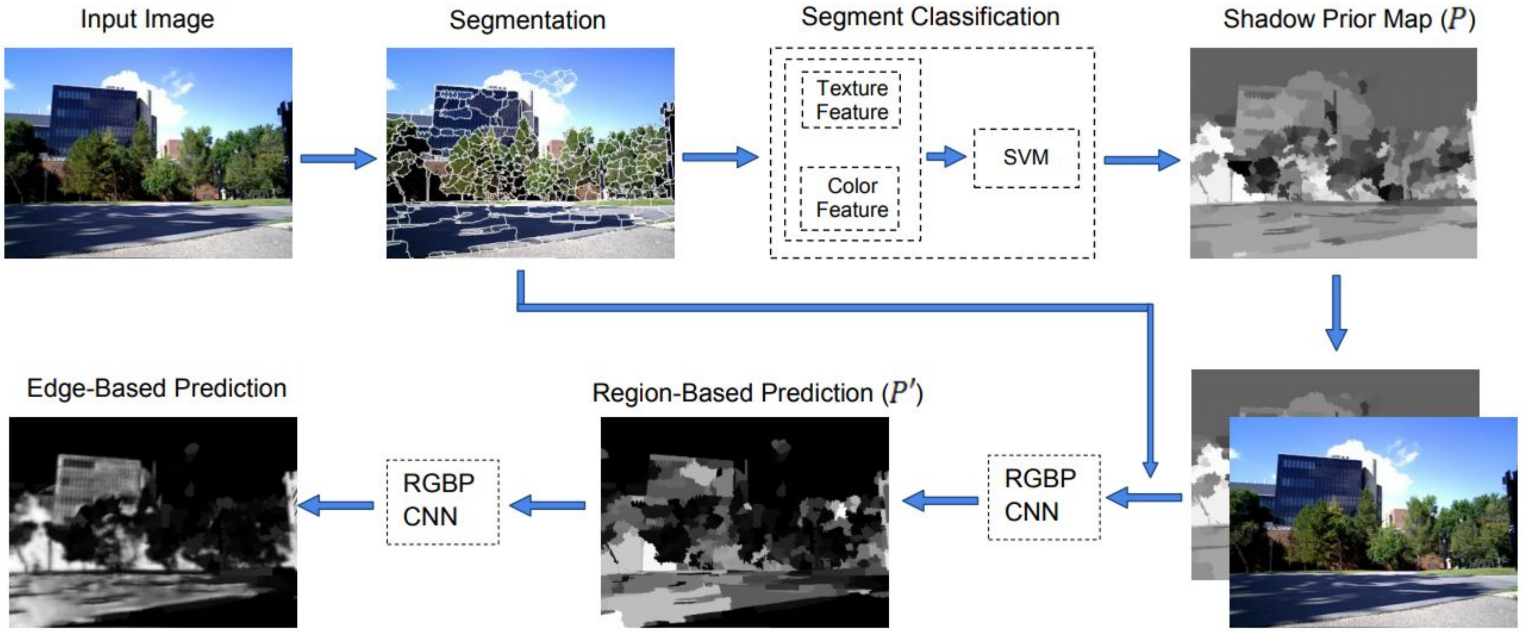

Patched-CNN

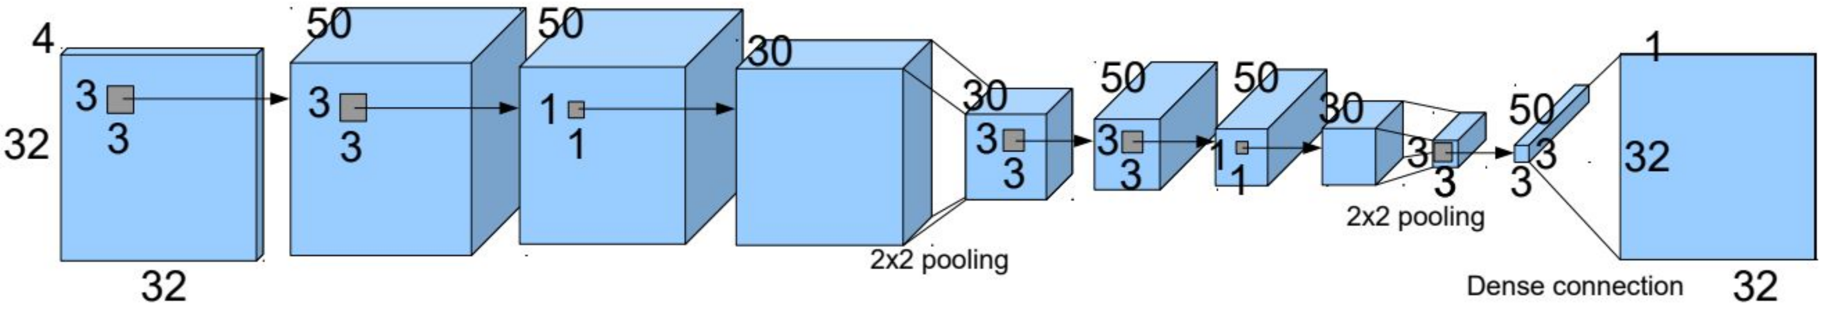

Input: 32x32 RGBP path

NN: 6 conv 2 pooling and one fully connected layers

Output: probability map of patch

Loss: negative log-likelihood of the prediction of all pixels

http://chenpingyu.org/docs/yago_eccv2016.pdf

Evaluation Metric

$$ShadowAccuracy = \frac{TP}{all\ shadow\ pixels}$$

$$Non-shadowAccuracy = \frac{TN}{all\ non-shadow\ pixels}$$

$$TotalAccuracy = \frac{TP+TN}{all\ pixels}$$

Fig. 2. Comparison of our qualitative results with the results of other methods. Rows from top to bottom: input images, ground truths, results of unary-pairwise method, results of stacked-CNN, obtained probanility map of our method, binary mask of shadows based on the probability map of our method.

https://arxiv.org/pdf/1709.09283.pdf

Shadow detected

What next?

Generative Adversarial Networks

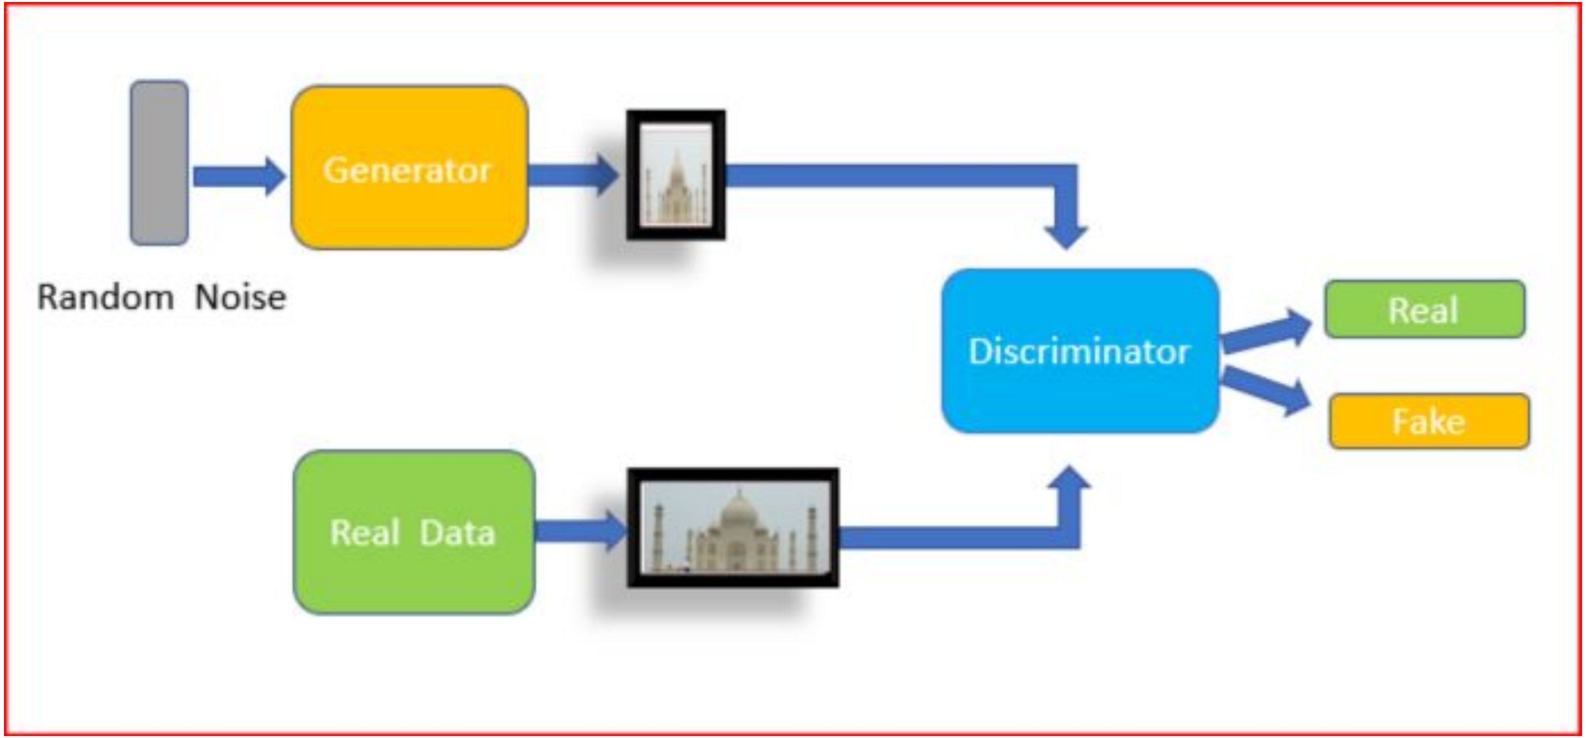

https://medium.com/datadriveninvestor/generative-adversarial-network-gan-using-keras-ce1c05cfdfd3

GAN objective

$$\min_G \max_D V(D, G) = \mathbb{E}_{\boldsymbol{x} \in p_{\text{data}}(\boldsymbol{x})}[\log D(\boldsymbol{x})] + \mathbb{E}_{\boldsymbol{z} \in p_{\boldsymbol{z}}(\boldsymbol{z})}[\log(1 - D(G(\boldsymbol{z}))]$$

https://arxiv.org/abs/1406.2661

Conditional GAN

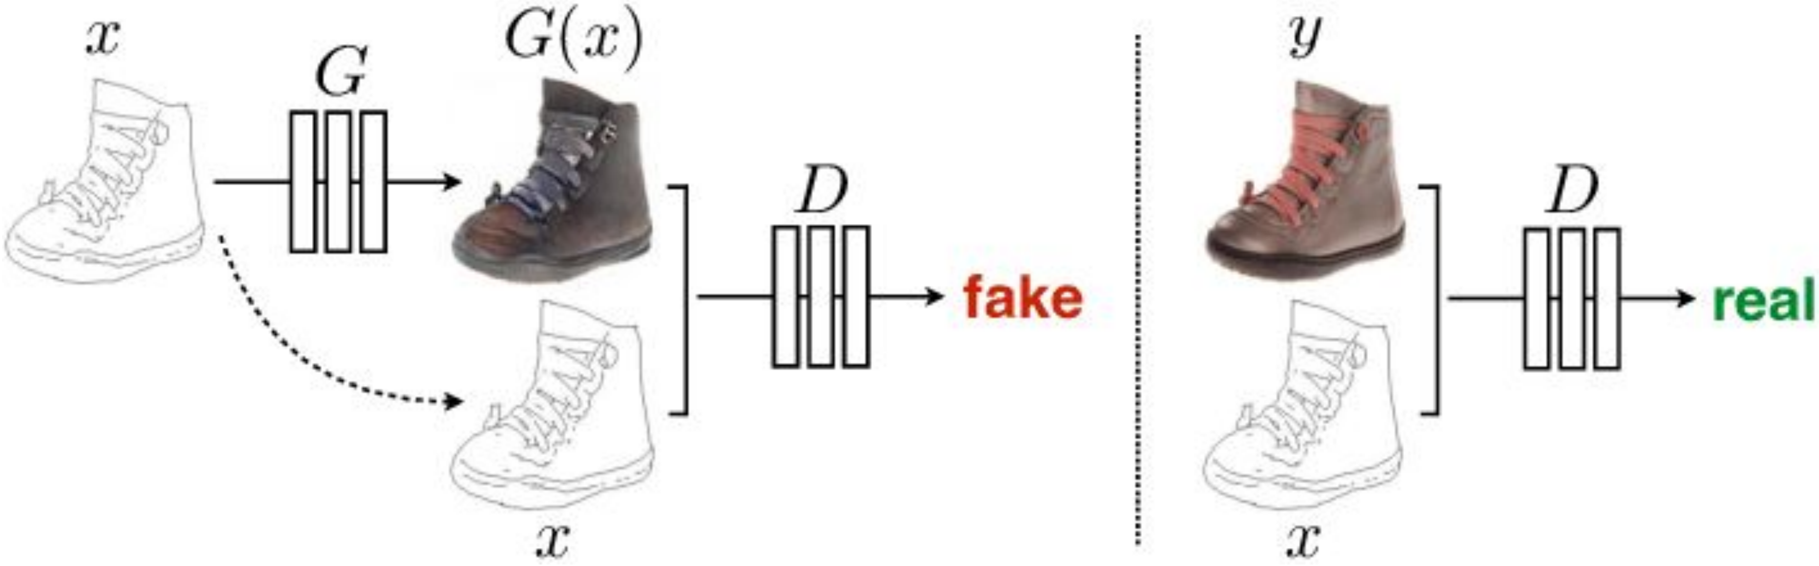

https://arxiv.org/abs/1611.07004

Conditional GAN

$$\min_G \max_D V(D, G) = \mathbb{E}_{\boldsymbol{x} \sim p_{\text{data}}(\boldsymbol{x})}[\log D(\boldsymbol{x}|\boldsymbol{y})] + \mathbb{E}_{\boldsymbol{z} \sim p_z(\boldsymbol{z})}[\log(1 - D(G(\boldsymbol{z}|\boldsymbol{y})))].$$

https://arxiv.org/abs/1611.07004

Approach #2 - Shadow Detection & Removing

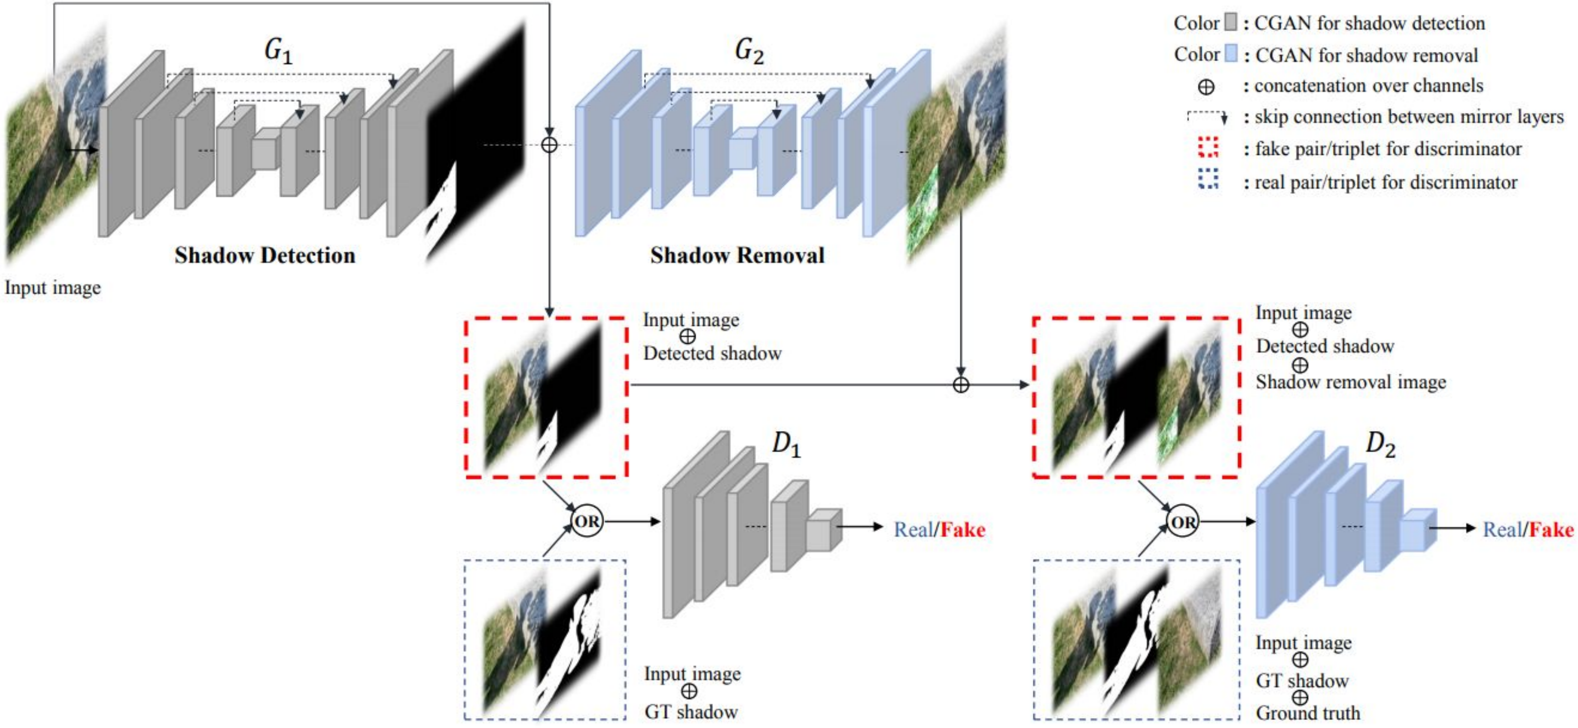

[Stacked Conditional Generative Adversarial Networks for Jointly Learning Shadow Detection and Shadow Removal (2017)]

Shadow Detection & Removing

$$\min_{G_1,G_2}\max_{D_1,D_2} \mathcal{L}_{data_1}(G_1) + \lambda_1\mathcal{L}_{data_2}(G_2|G_1) +$$
$$\lambda_2\mathcal{L}_{CGAN_1}(G_1, D_1) + \lambda_3\mathcal{L}_{CGAN_2}(G_2, D_2|G_1).$$

$$\mathrm{BER} = 1 - \frac{1}{2}(\frac{TP}{TP+FN} + \frac{TN}{TN+FP})$$

Detection: Balance Error Rate (BER) for ground-truth masks and the predicted ones.

Removal: Root Mean Square Error (RMSE) in LAB color space between the ground truth shadow-free image and the recovered image.

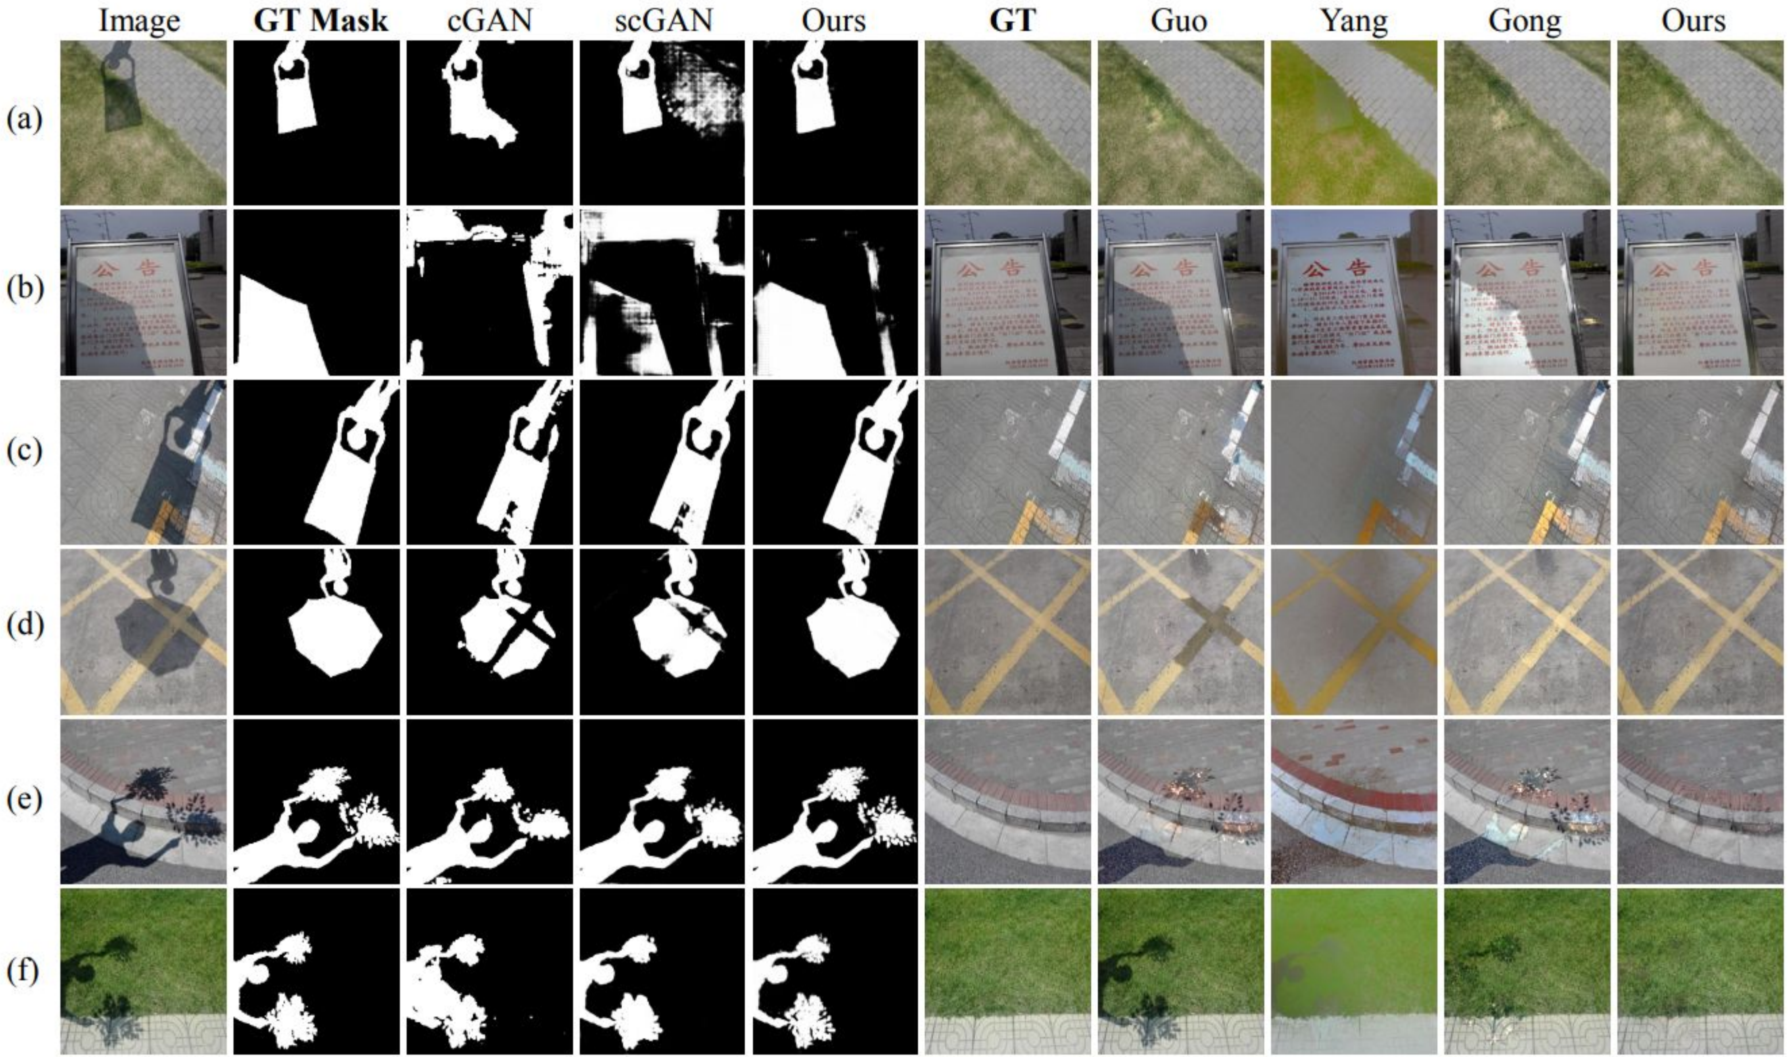
Image
GT Mask
cGAN
scGAN
Ours
GT
Guo
Yang
Gong
Ours
(a)
(b)
(c)
(d)
(e)
(f)
公 告

Shadow Detection & Removing

Approach #3 - Shadow Generation/Augmentation

Cycle GAN - Datasets

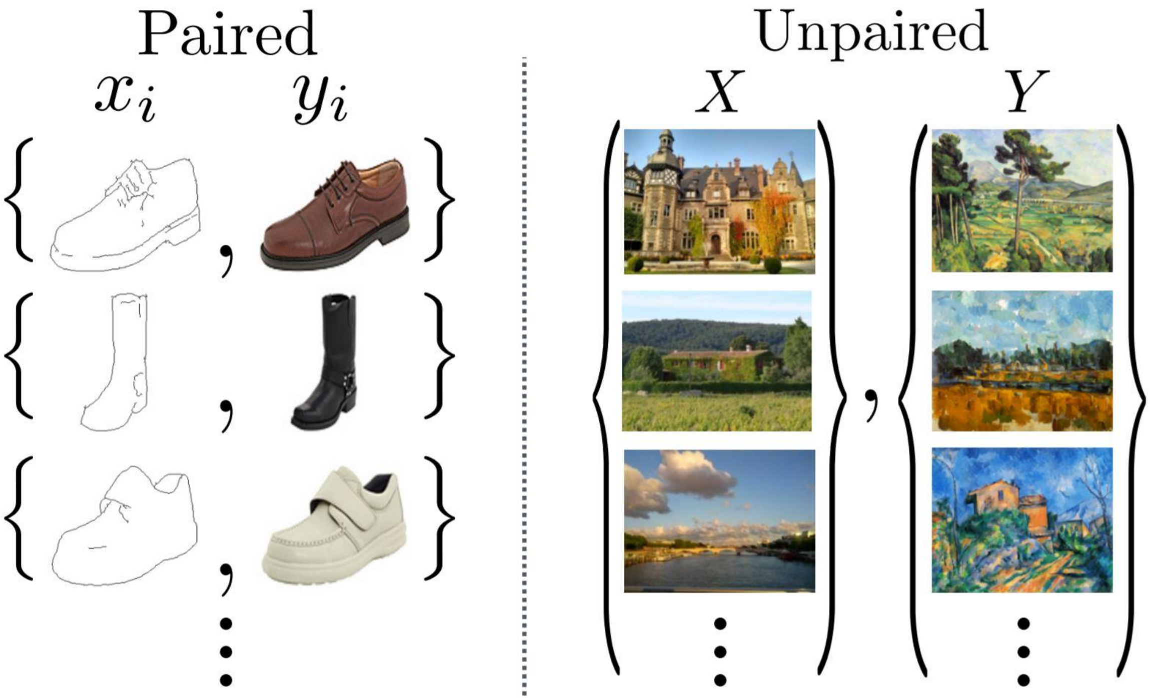

https://hardikbansal.github.io/CycleGANBlog/

Cycle GAN

Unpaired Image-to-Image Translation

using

Cycle-Consistent Adversarial Networks

(2017)

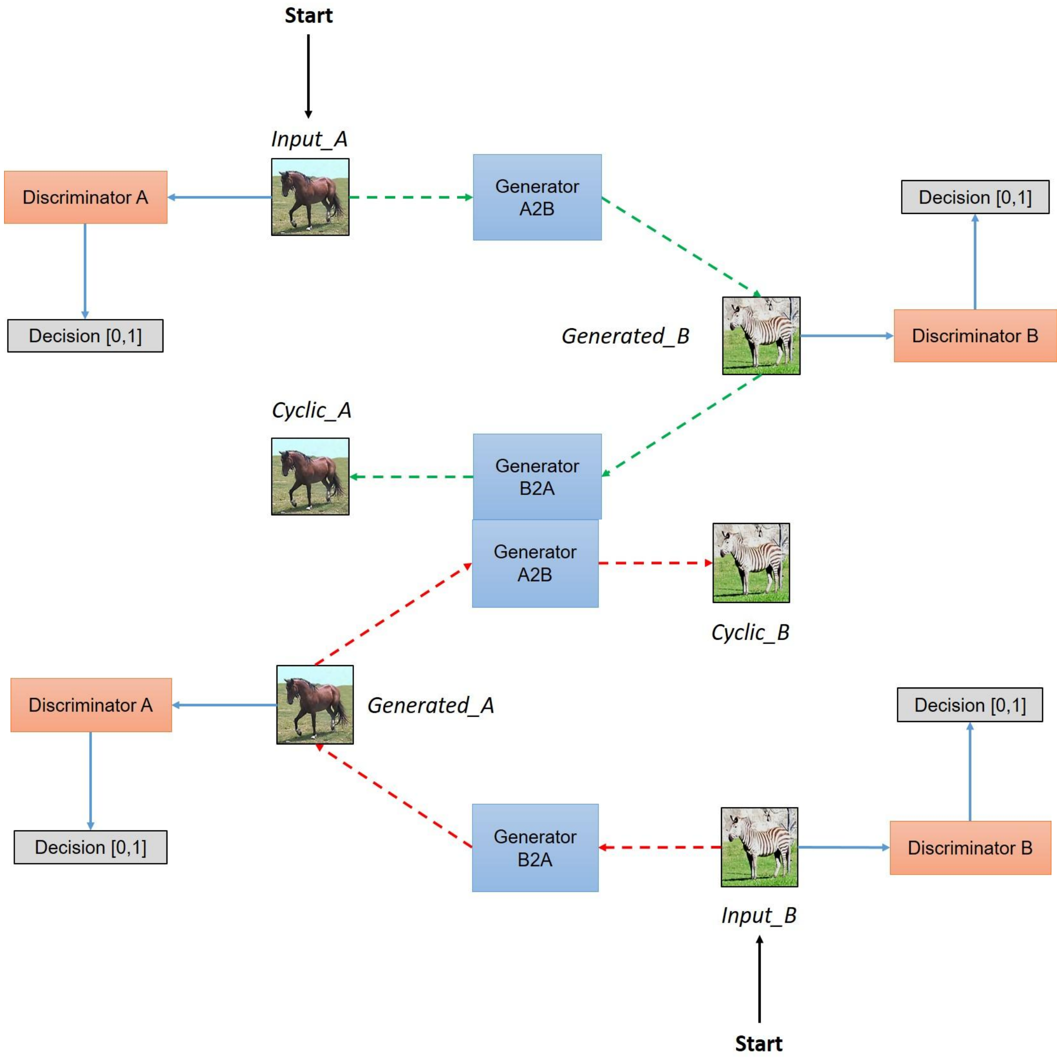

Cycle GAN objective

$$\begin{aligned}\mathcal{L}(G, F, D_X, D_Y) =& \mathcal{L}_{\text{GAN}}(G, D_Y, X, Y) \\ &+ \mathcal{L}_{\text{GAN}}(F, D_X, Y, X) \\ &+ \lambda \mathcal{L}_{\text{cyc}}(G, F),\end{aligned}$$

Cycle GAN

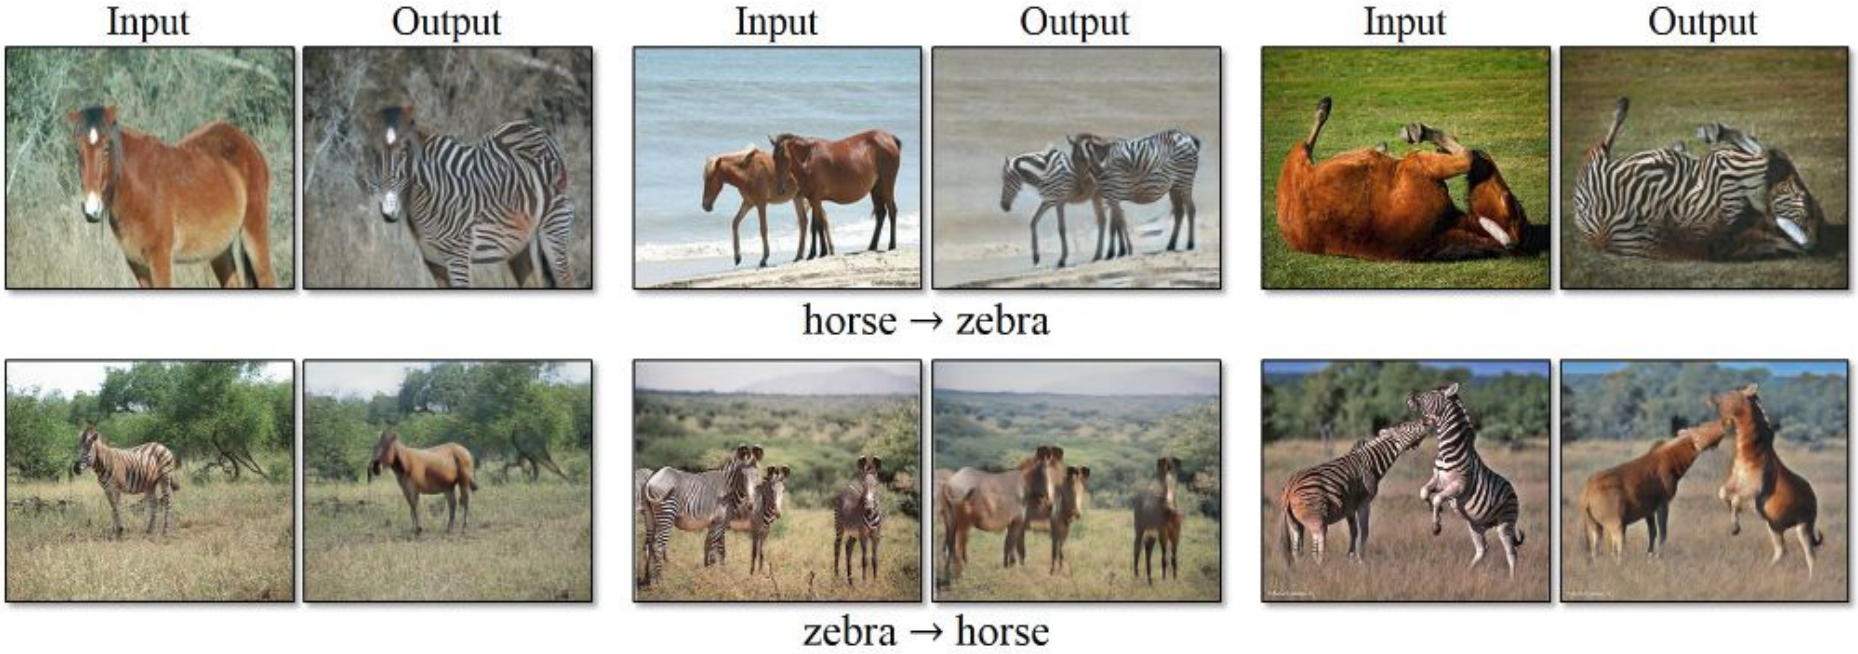
Input
Output
Input
Output
Input
Output
horse → zebra
zebra → horse

Cycle GAN

Identity Loss

$$\mathcal{L}_{\text{identity}}(G, F) =$$

$$\mathbb{E}_{y \sim p_{\text{data}}(y)}[\|G(y) - y\|_1] +$$

$$\mathbb{E}_{x \sim p_{\text{data}}(x)}[\|F(x) - x\|_1].$$

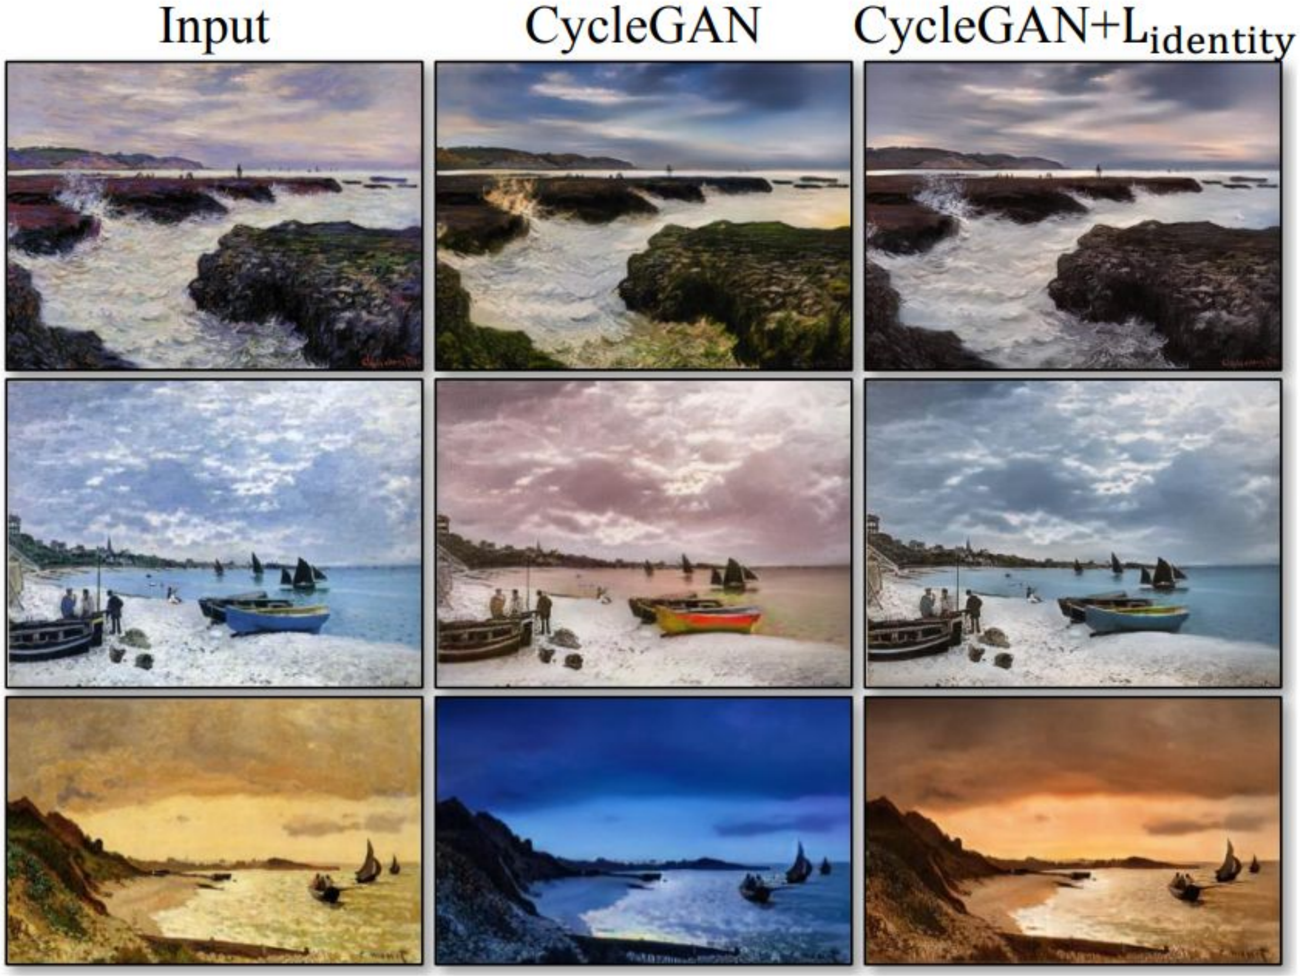

Cycle GAN

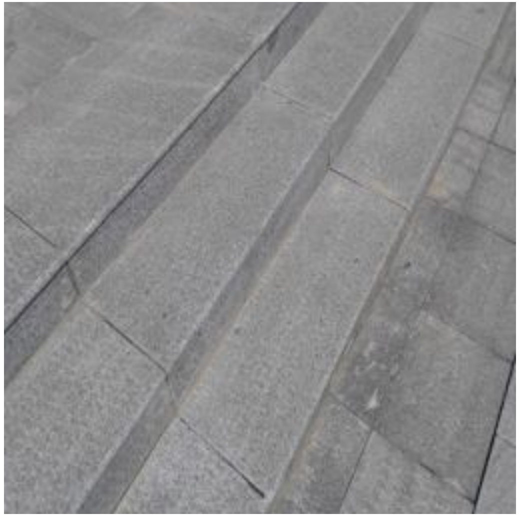

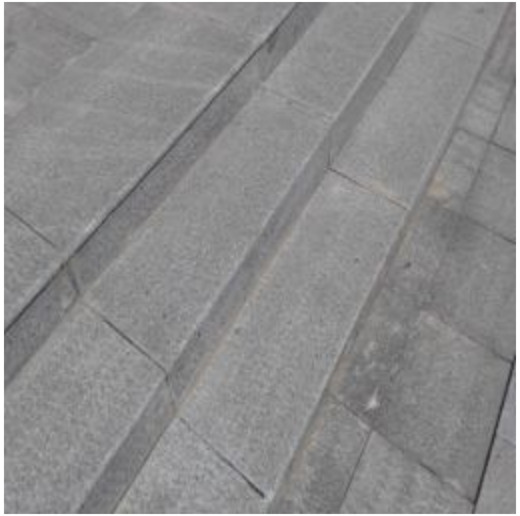

Mask-ShadowGAN

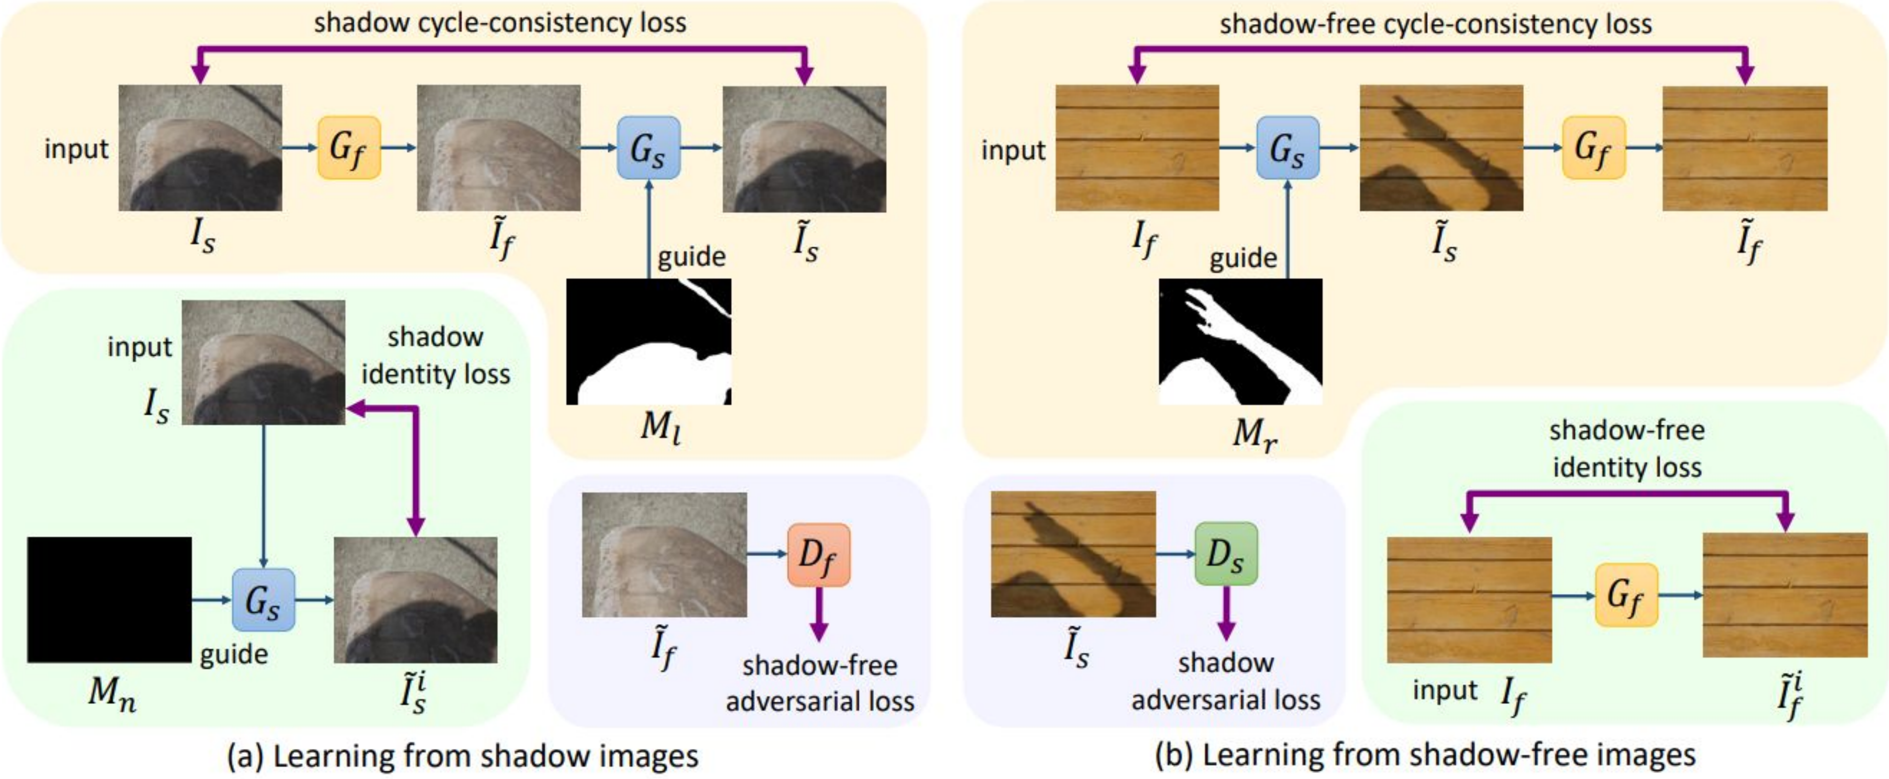

https://arxiv.org/pdf/1903.10683.pdf

Mask-ShadowGAN

$$
\begin{aligned}
& L_{final}(G_s, G_f, D_s, D_f) \\
& = \omega_1(L^a_{GAN}(G_f, D_f) + L^b_{GAN}(G_s, D_s)) \\
& + \omega_2(L^a_{cycle}(G_f, G_s) + L^b_{cycle}(G_s, G_f)) \\
& + \omega_3(L^a_{identity}(G_s) + L^b_{identity}(G_f)) \; .
\end{aligned}
$$

Mask-ShadowGAN

  

Mask-ShadowGAN

Mask-ShadowGAN

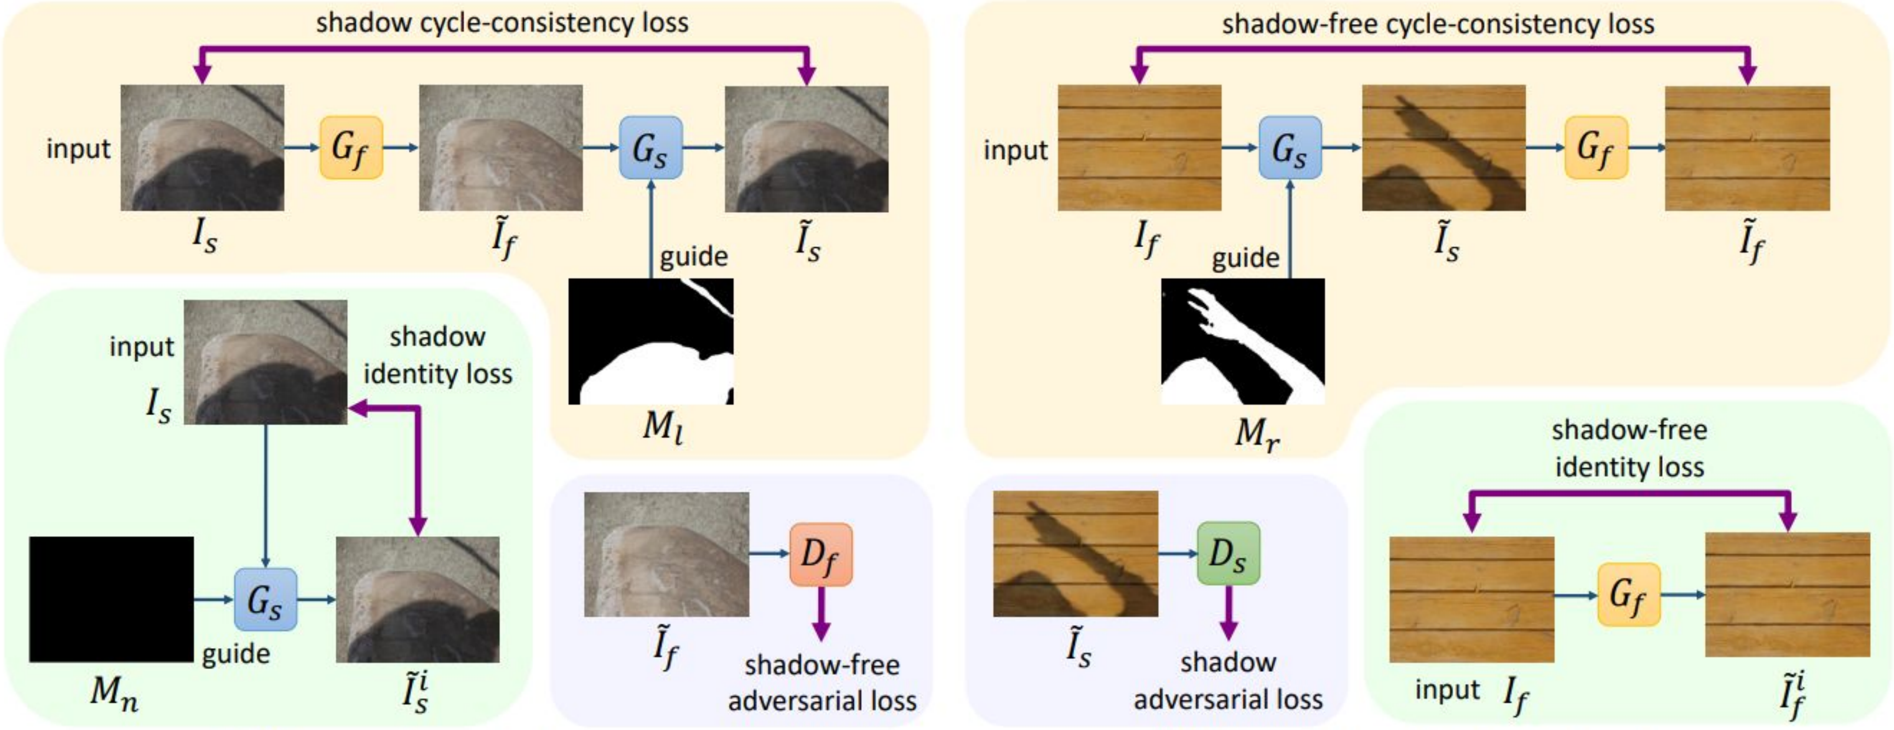

(a) Learning from shadow images

(b) Learning from shadow-free images

https://arxiv.org/pdf/1903.10683.pdf

Summary

Thanks for your time and attention!

Questions?