Balanced Unsupervised Style Transfer using Generative Adversarial Networks

Dr. Michal Holtzman Gazit

End-to-End Learning for Many Tasks

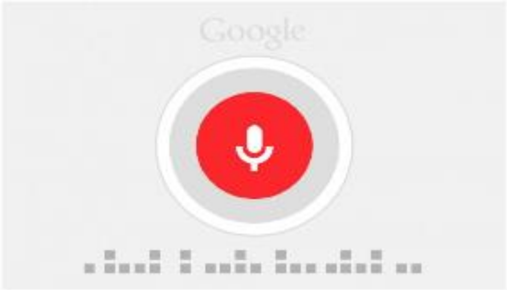

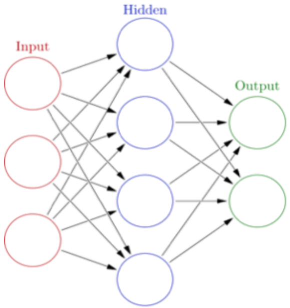

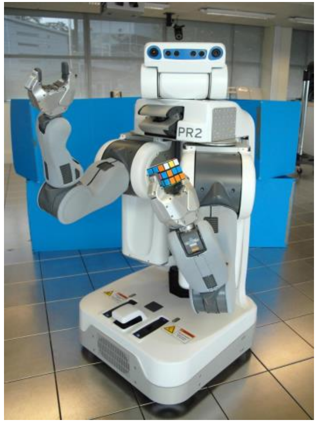

Computer vision tasks

- Segmentation
- Tracking
- Detection
- Registration
- Classification
- All need data

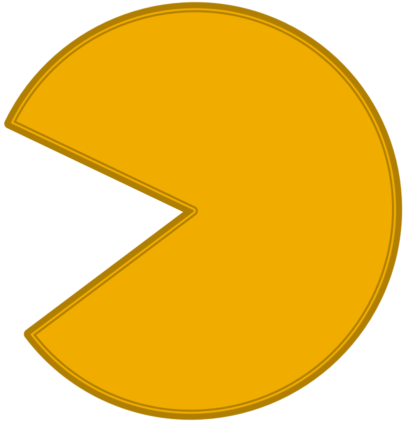

Dataset

- No real images (or maybe just a few)
- Can use synthetic images (or images from a different domain)
- Synthesized images try to simulate the relevant image domain
  - Big labeled database for the relevant task

Problem

- “Simulated” images are not similar enough to the “real” images
- Do not represent the true statistics of the real domain
- An algorithm trained on these images is bound to succeed less in real scenario

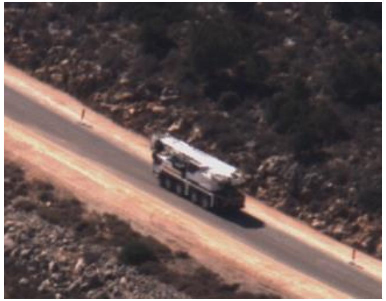

Solution

- Better mimic the style of the real domain – allowing better training hence deployment of the computer vision task

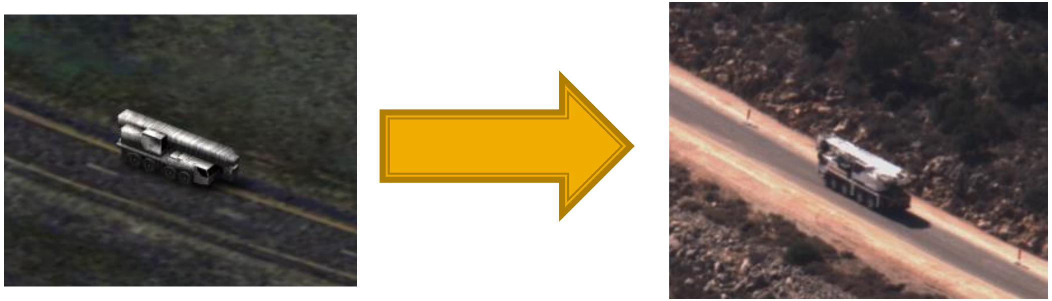

Generative Adversarial Networks

- Minimax Game - alternates between Generator and Discriminator
- Generator – trains to generate samples that fool the discriminator
- Discriminator – trains to distinguish – real or fake

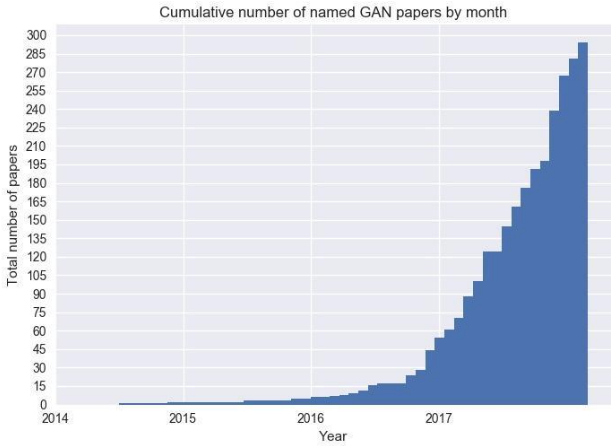

]Goodfellow et al. [2014]www.medium.com/@devnag/generative-adversarial-networks-gans –in-50-lines-of-code[

Generative Adversarial networks

$$\min_{G}\max_{D}\mathbb{E}_{x,y}[\log D(G(x))+\log(1-D(y))]$$

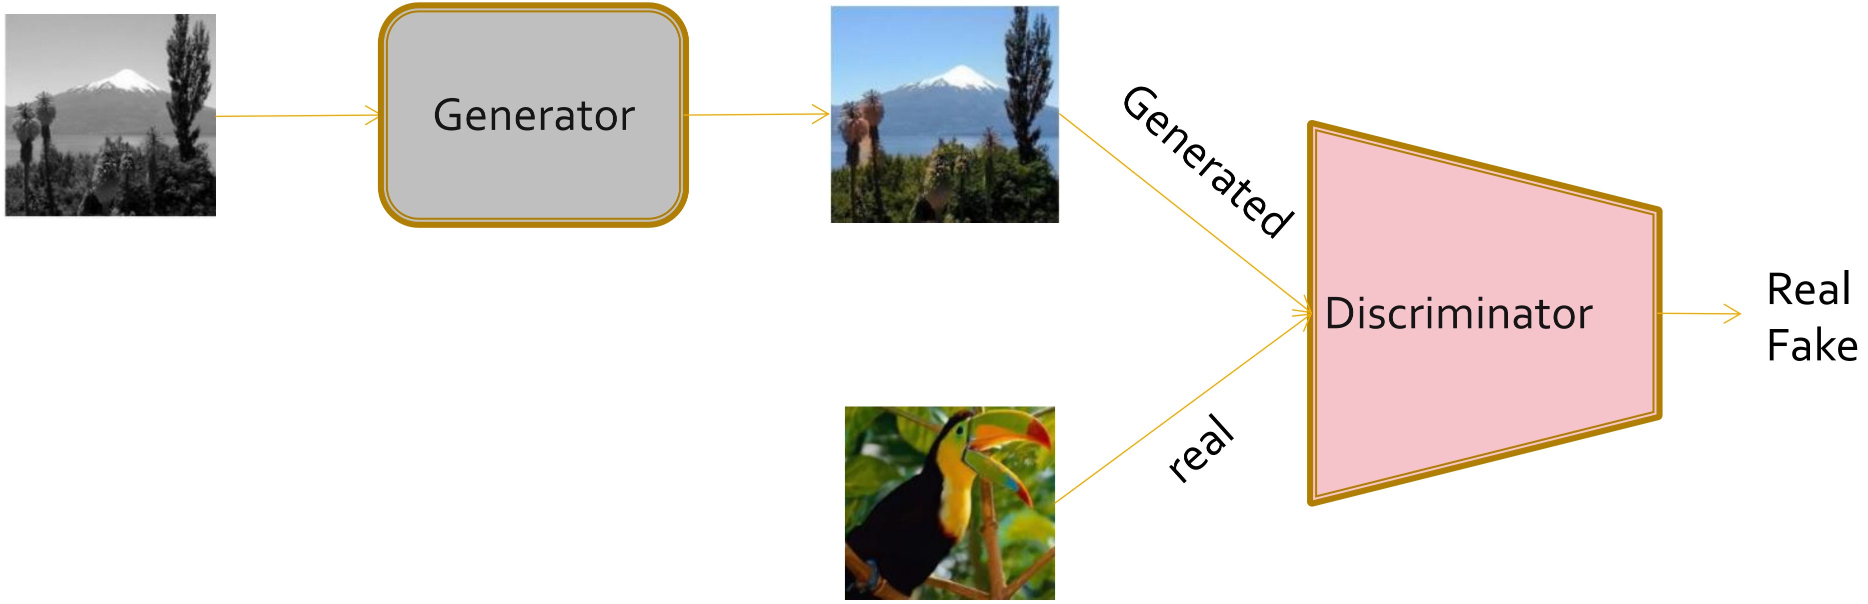

]Goodfellow et al. [2014

Cycle GAN

- How does it work?

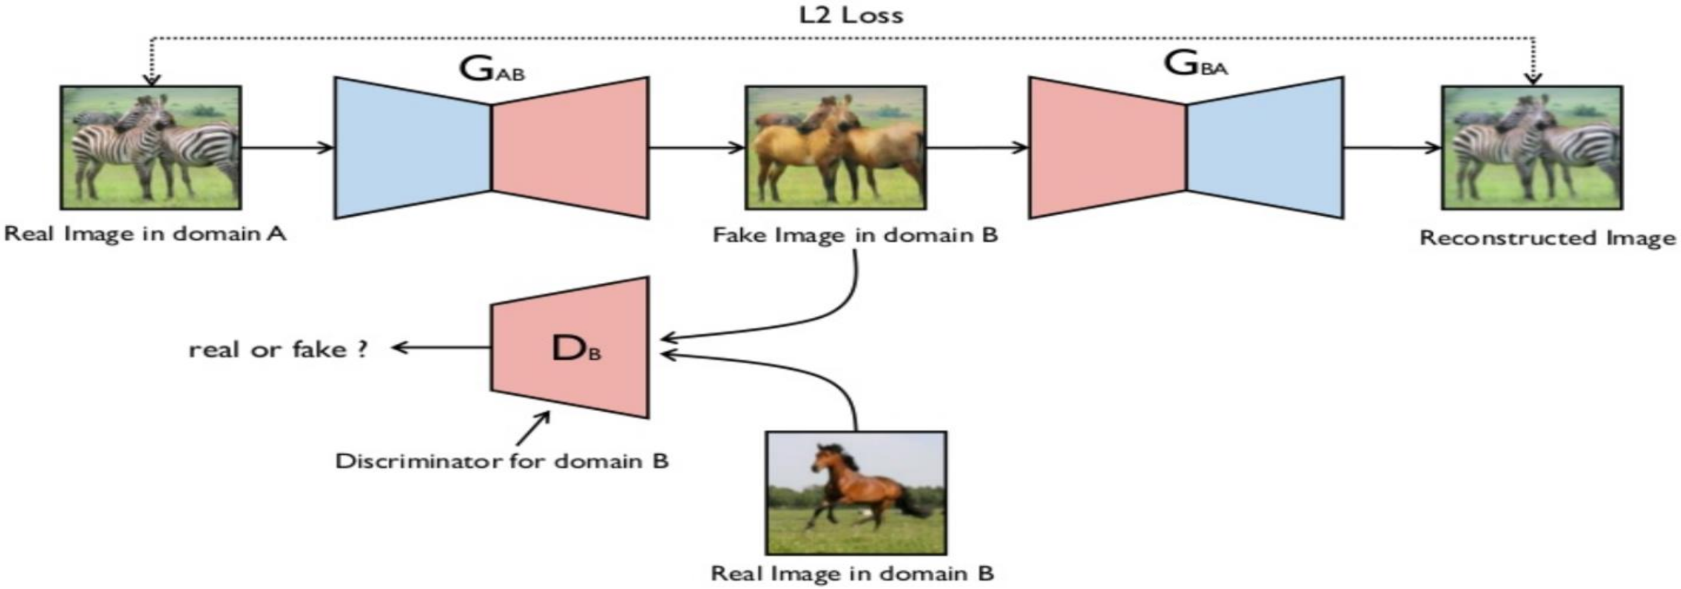

Upaired Image-to-Image Translation using Cycle Consistent Adversarial Networks – J.Y Zhu, T. Park, P. Isola, A Efros – ICCV 17

Cycle consistency

$$\mathcal{L}_{\text{cyc}}(G, F) = \mathbb{E}_{x \sim p_{\text{data}}(x)}[\|F(G(x)) - x\|_1] + \mathbb{E}_{y \sim p_{\text{data}}(y)}[\|G(F(y)) - y\|_1].$$

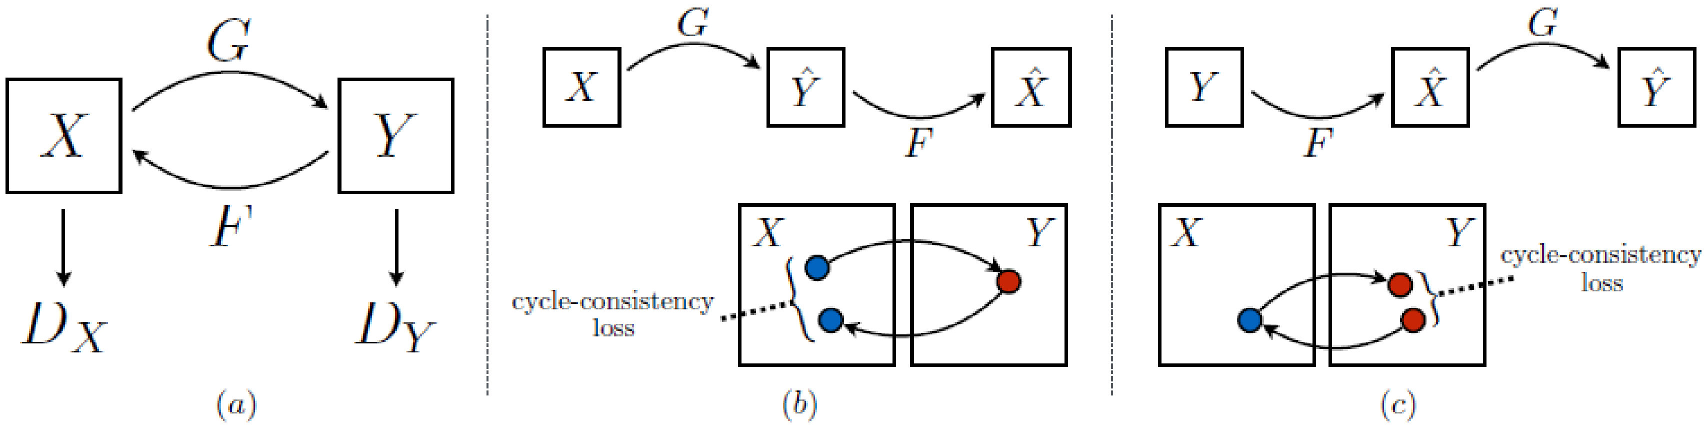

Objective

$$\begin{aligned}\mathcal{L}_{\text{GAN}}(G, D_Y, X, Y) = &\mathbb{E}_{y \sim p_{\text{data}}(y)}[\log D_Y(y)] \\ &+ \mathbb{E}_{x \sim p_{\text{data}}(x)}[\log(1 - D_Y(G(x))],\end{aligned}$$

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

Additions to the original objective

- Similarity to the original image
- Make the result valuable to the task
- Balance the power between G and D via regularization

Similarity loss

- GAN is unpaired
- We want only style transfer
- Preserve image shape
- Enforce similarity to some extent between input image and generated one

$$Ls = \left\| x - G(x) \right\|_{L_1}$$

Task loss

- Insure output is valuable to the required task
- Example – Classification
  - Add constraint on image class
  - Add loss – before and after generation

$$L_T = \|T(x) - t\| + \|T(G(x)) - t\|$$

Bousmalis K., et al. (CVPR 2017)

Nash Equilibrium?

- Problem: Discriminator always wins
- Regularization to balance the power between G and D

$$L_{WGANLP} = \max\{0, \|\nabla D(\hat{x})\| - 1\}^2$$

$$\hat{x} = \alpha G(x) + (1 - \alpha) x$$

$$\alpha \sim U[0,1]$$

- Penalize where gradient norm of D is larger than 1

ON THE REGULARIZATION OF WASSERSTEIN GANS – H . Petzka, A Fischer, D Lukovnicov – ICLR 18

Results

Results

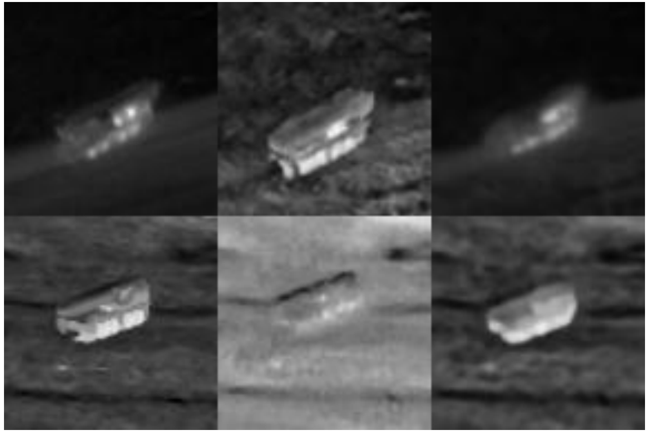

Task success rate w/o GAN images	Task success rate w GAN images
0.743	0.777

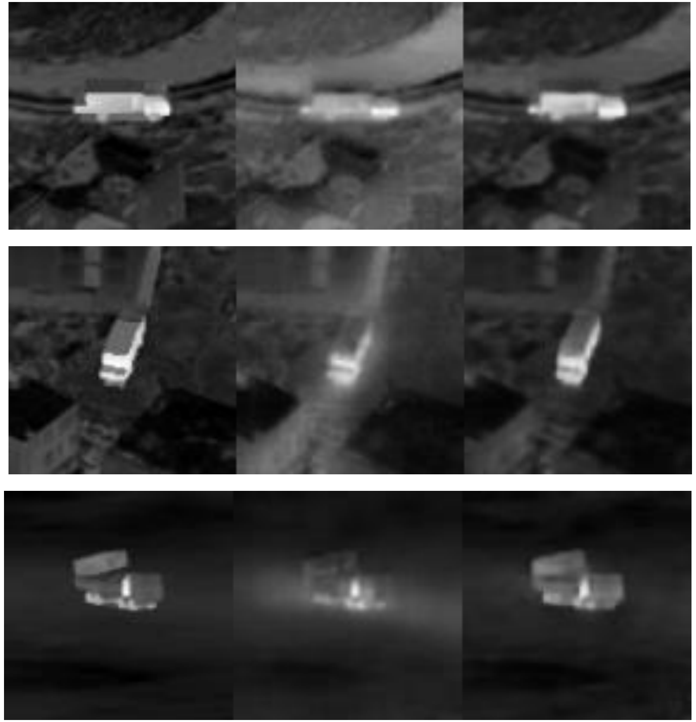

michalhg@rafael.co.il

THANK YOU