THE UNIVERSITY OF TENNESSEE KNOXVILLE

Decoupled Learning for Conditional Adversarial Networks

Zhifei Zhang, Yang Song, and Hairong Qi
Department of Electrical Engineering and Computer Science

WACV 2018

Motivation:

The conditional adversarial networks applied in existing works mainly consists of two parts: 1) the encoding-decoding nets (ED) and 2) the GANs, which are tied in the parts of decoder and generator. Therefore, the reconstruction loss and adversarial loss interact/ compete with each other, while potentially causes unstable results as shown in the follow.

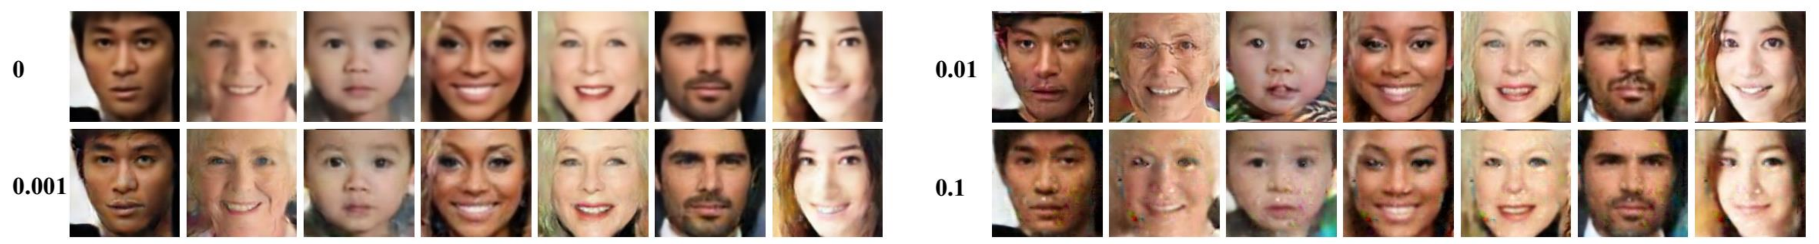

Existing works have to introduce a weighting factor (e.g., the values in the figure) to balance the effect of the two losses. How to adaptively set an appropriate weight or completely remove the necessity of the weighting factor is a problem unexplored.

Main Idea:

Decouple the interaction between the reconstruction loss and adversarial loss in backpropagation, avoiding the competition that may cause instability.

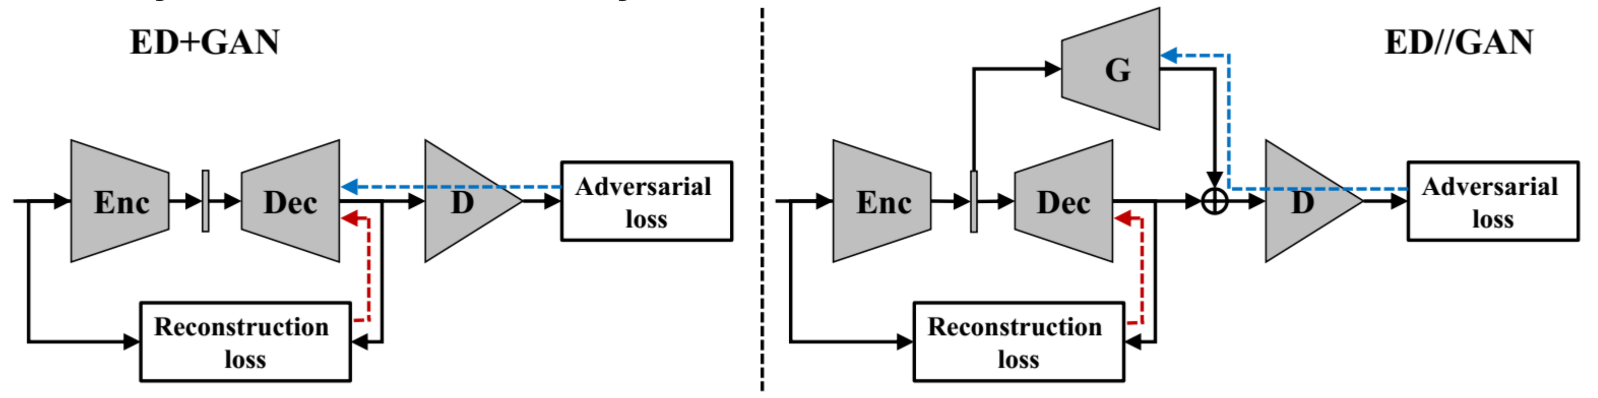

- ED+GAN: the traditional structure
- ED//GAN: the proposed structure(decoupled learning)
- Enc and Dec: the encoder and decoder networks
- G and D: the generator and discriminator
- Black arrows: feedforward path
- Red arrows: backpropagation of reconstruction loss
- Blue arrows: backpropagation of adversarial loss

Approach:

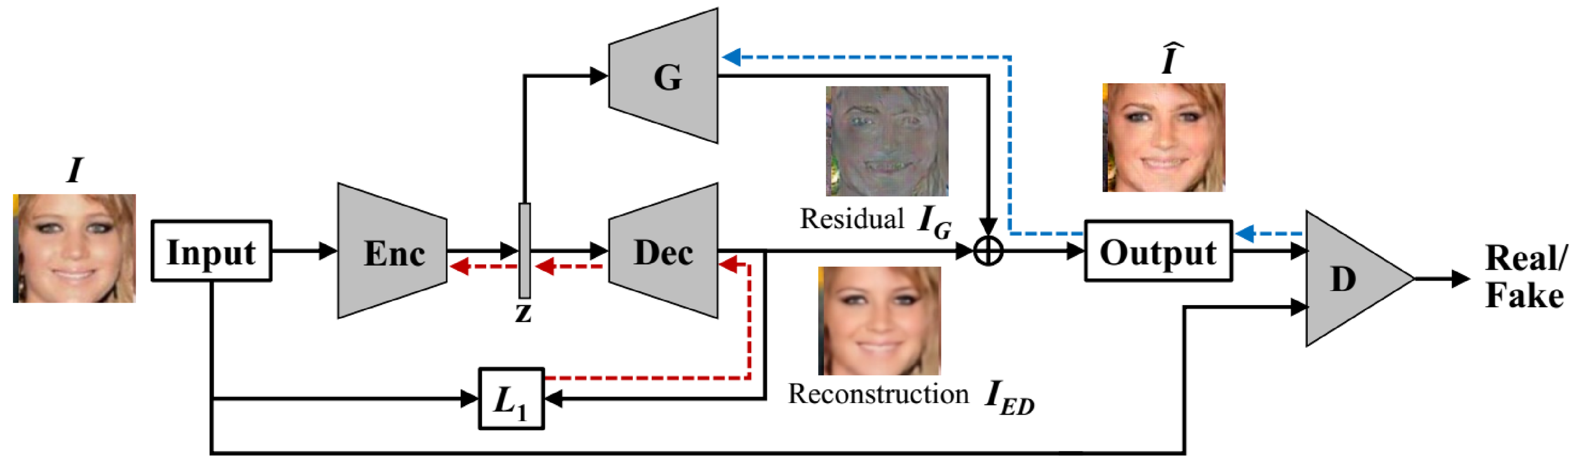

Reconstruction loss:

$$\mathcal{L}_{const}(Enc, Dec) = \|I - Dec(Enc(I))\|_1$$

Adversarial loss:

$$\mathcal{L}_{adv}(D) = \mathbb{E}\left[\log\left(1 - D(I)\right)\right] + \mathbb{E}\left[\log D\left(I_{ED} + I_G\right)\right],$$
$$\mathcal{L}_{adv}(G) = \mathbb{E}\left[\log\left(1 - D\left(I_{ED} + G(z)\right)\right)\right],$$

where $I_{ED} = Dec(z)$ and $z = Enc(I)$.

The objective function:

Above all, the objective function of the proposed decoupled learning (ED//GAN) is

$$\min_{Enc,Dec} \mathcal{L}_{const}(Enc, Dec) + \min_{G} \mathcal{L}_{adv}(G) + \min_{D} \mathcal{L}_{adv}(D).$$

There are no weighting parameters between the losses in the objective function, which relaxes the manual tuning.

Visualizing the Boost from Adversarial Learning:

Left: output from ED. Middle: output from G (trained by the adversarial loss). Right: the final out of ED//GAN, adding the output from G to that from ED.

Experimental Results:

Normalized Relative Discriminative Score (NRDS):

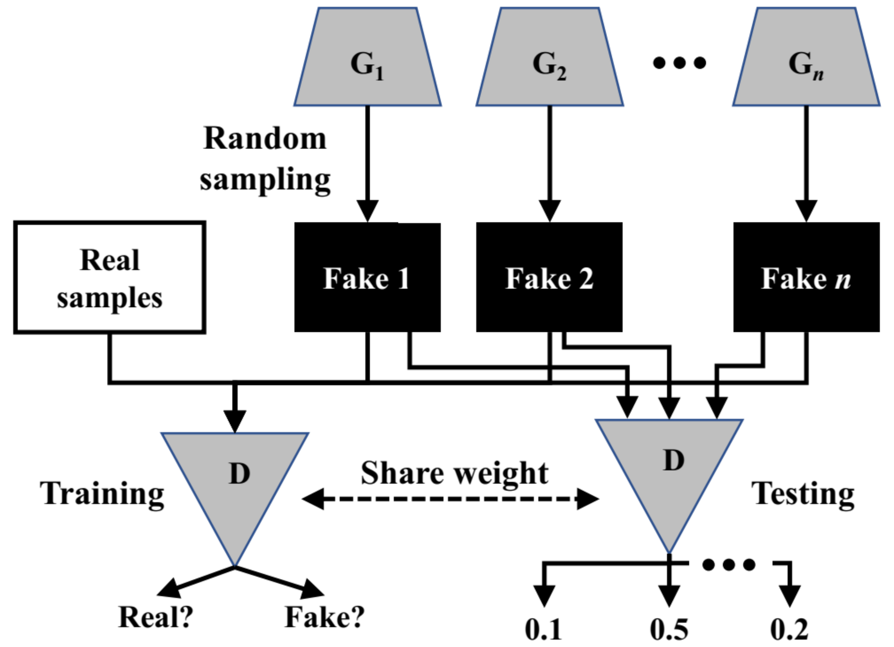

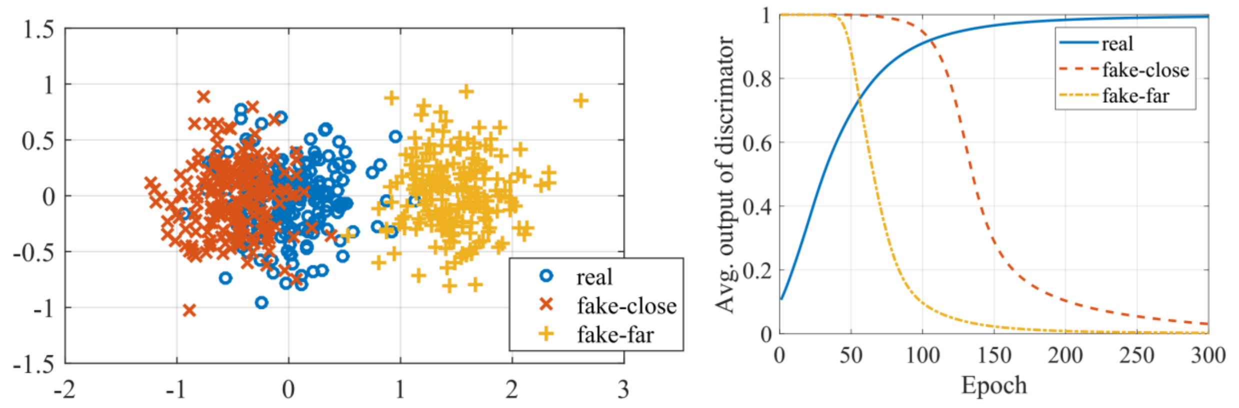

A toy example of computing NRDS. Left: the real and fake samples randomly sampled from 2-D normal distributions with different means but with the same covariance. Right: the curves of epoch vs. averaged output of discriminator on corresponding sets (colors) of samples.

Compare ED+GAN and ED//GAN:

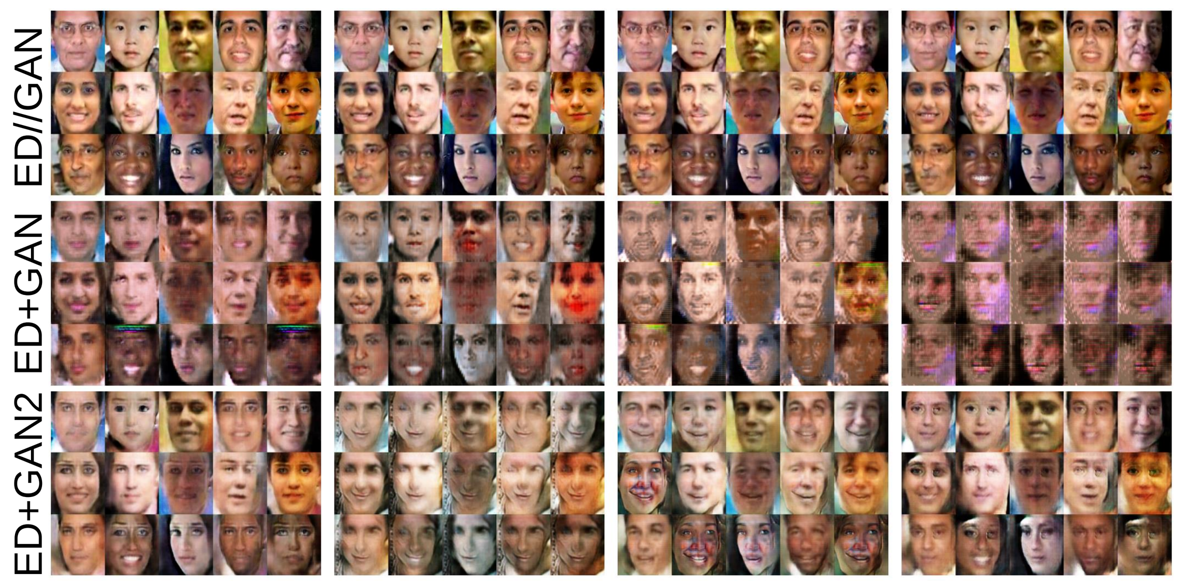

Table 2. NRDS with different weight settings and their std.

	0.001	0.01	0.1	1	std
ED+GAN	.1172	.1143	.1163	.0731	.0215
ED+GAN2	.1066	.1143	.1268	.1267	.0099
ED//GAN	.1432	.1434	.1458	.1466	.0017

ED+GAN2 denote the structure with batch normalization

ED+GAN is sensitive to weight variation. By contrast, ED//GAN is robust to weight variation, relaxing the weight tuning.

Adapt Existing Works to ED//GAN:

Pix2Pix [Isola et al., 2017] vs ED//GAN version

Method	ED+GAN			ED//GAN
	1:1	100:1	1000:1	
NRDS	.2190	**.2641**	.2572	.2597

CAAE [Zhang et al., 2017] vs ED//GAN version

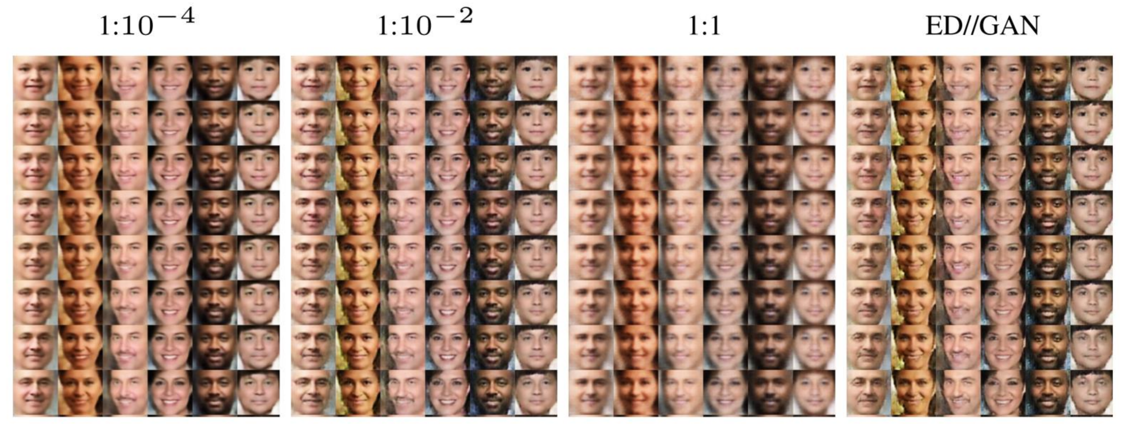

ED+GAN			ED//GAN
$1{:}10^{-4}$	$1{:}10^{-2}$	1:1	
.2527	.2496	.2430	**.2547**