Label-Noise Robust Generative Adversarial Networks

Training data

Noisy labeled

rcGAN

Conditioned on *clean* labels

Takuhiro Kaneko[1] Yoshitaka Ushiku[1] Tatsuya Harada[1, 2]

[1]The University of Tokyo [2]RIKEN

Talk

東京大学 THE UNIVERSITY OF TOKYO

AIP

CVPR LONG BEACH CALIFORNIA June 16-20, 2019

Code

Objective: Label-noise robust conditional image generation

Our goal is to construct a **label-noise robust conditional image generator**
that can reproduce ***clean* labeled data** (a)
even when ***noisy* labeled data** **(b)** are only available during the training.

(a) *Clean* labeled data

Unavailable

(b) *Noisy* labeled data

Available

Challenge: Limitation of naïve conditional generative models

Naïve **conditional generative models** (e.g., **AC-GAN** [1], **cGAN** [2, 3])
attempt to construct a generator conditioned on observable labels
(i.e., ***noisy* labels** in this case).

Training data **cGAN**

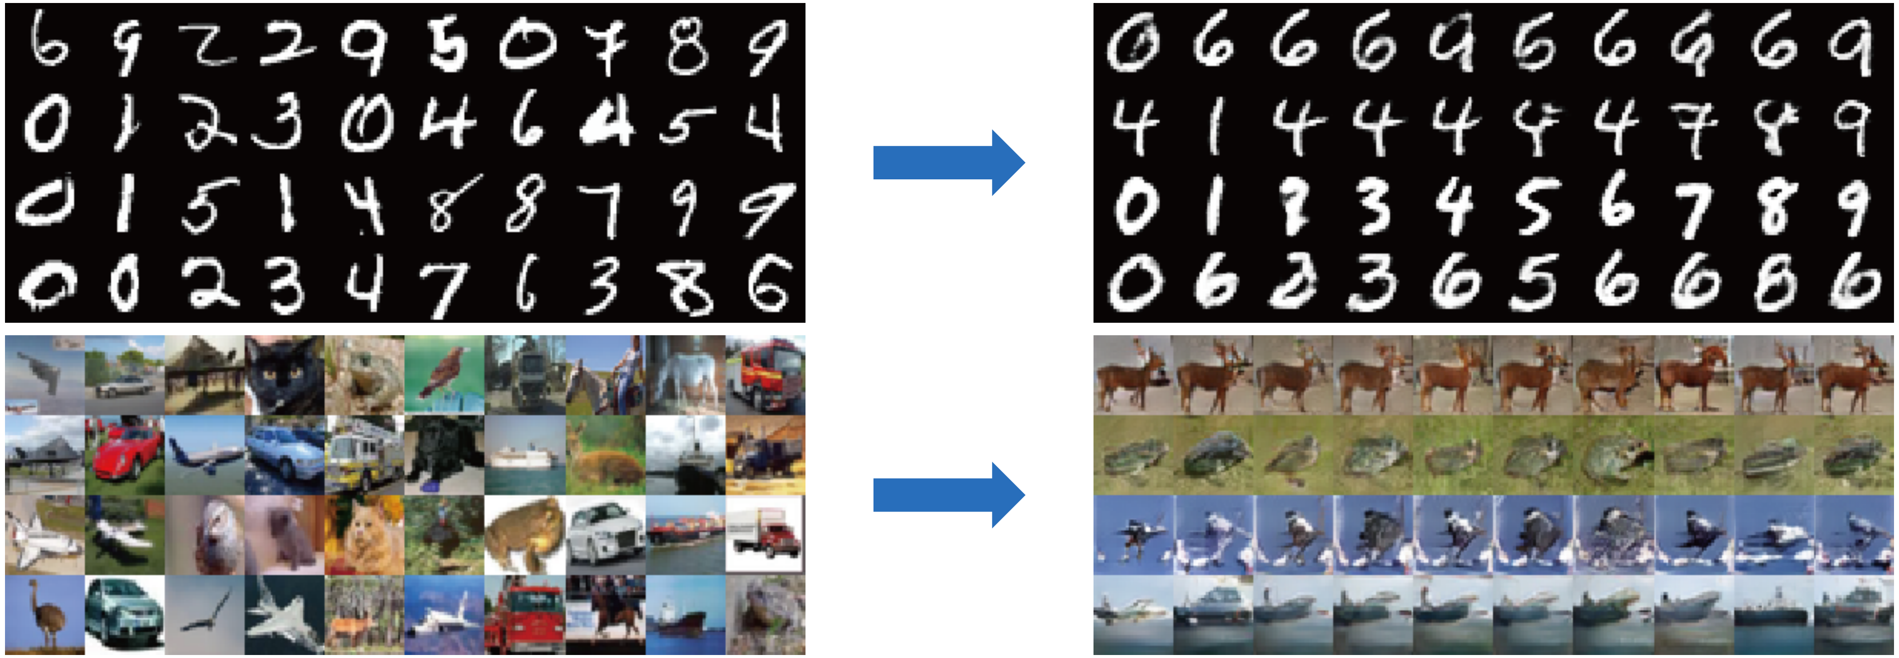

***Noisy* labeled** **Conditioned on *noisy* labels**

AC-GAN: auxiliary classifier GAN, cGAN: conditional GAN

[1] Odena et al. ICML 2017. [2] Mirza & Osindero. arXiv 2014. [3] Miyato & Koyama. ICLR 2018.

Contribution: Proposal of label-noise robust GANs

To overcome this limitation, we propose **label-noise robust GANs (rGANs)** that
can construct a generator conditioned on ***clean* labels**
even when trained with ***noisy* labeled data**.

Training data **rcGAN**

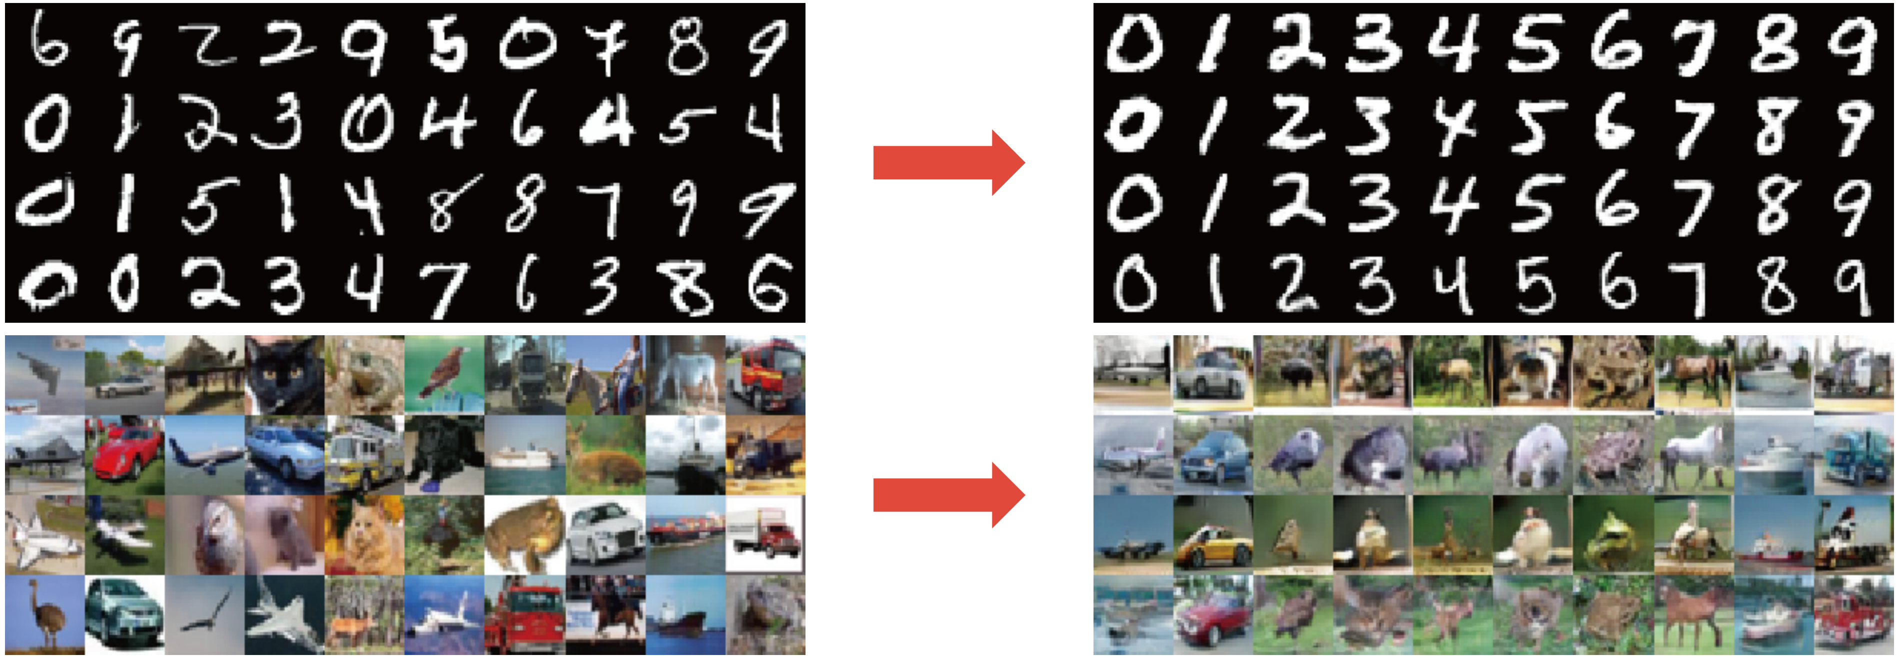

***Noisy* labeled** **Conditioned on *clean* labels**

rcGAN: label-noise robust conditional GAN

Main idea: Incorporation of noise transition model

We incorporate a **noise transition model** into naïve conditional GANs [1, 2, 3].

Noise transition model

$$p(\tilde{y}|\hat{y})$$

***Noisy* label** ***Clean* label**

Two variants

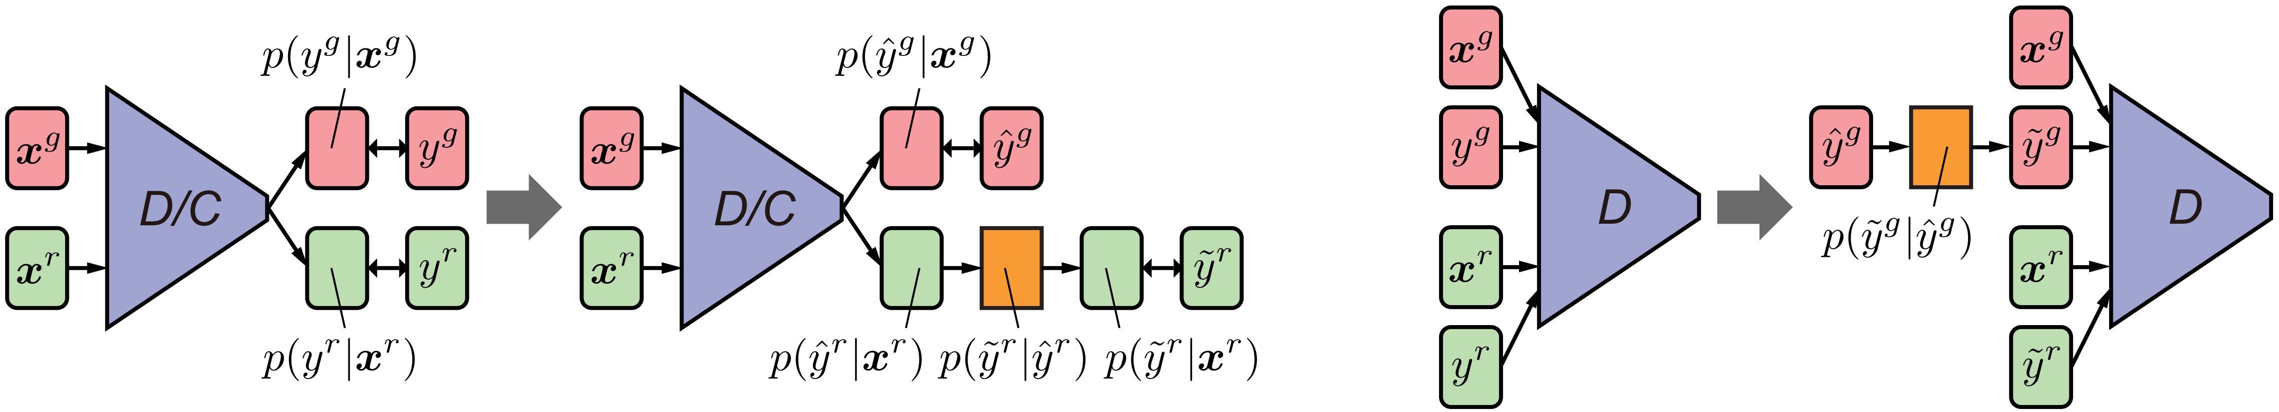

AC-GAN [1] **rAC-GAN** **cGAN** [2, 3] **rcGAN**

[1] Odena et al. ICML 2017. [2] Mirza & Osindero. arXiv 2014. [3] Miyato & Koyama. ICLR 2018.

Baseline: AC-GAN

Auxiliary classifier GAN [1]

***G**enerator*: $G(\boldsymbol{z}, y^g)$

***D**iscriminator/**C**lassifier*: $D(\boldsymbol{x})$ / $C(y|\boldsymbol{x})$

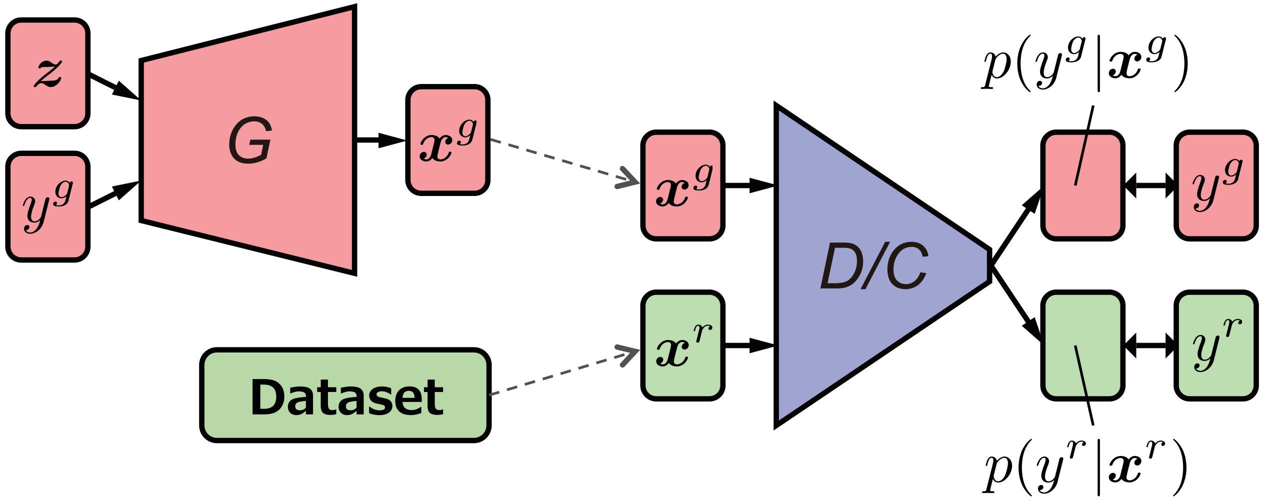

[1] Odena et al. ICML 2017.

Baseline: AC-GAN

Auxiliary classifier GAN [1]

***G**enerator*: $G(\boldsymbol{z}, y^g)$

***D**iscriminator/**C**lassifier*: $D(\boldsymbol{x})$ / $C(y|\boldsymbol{x})$

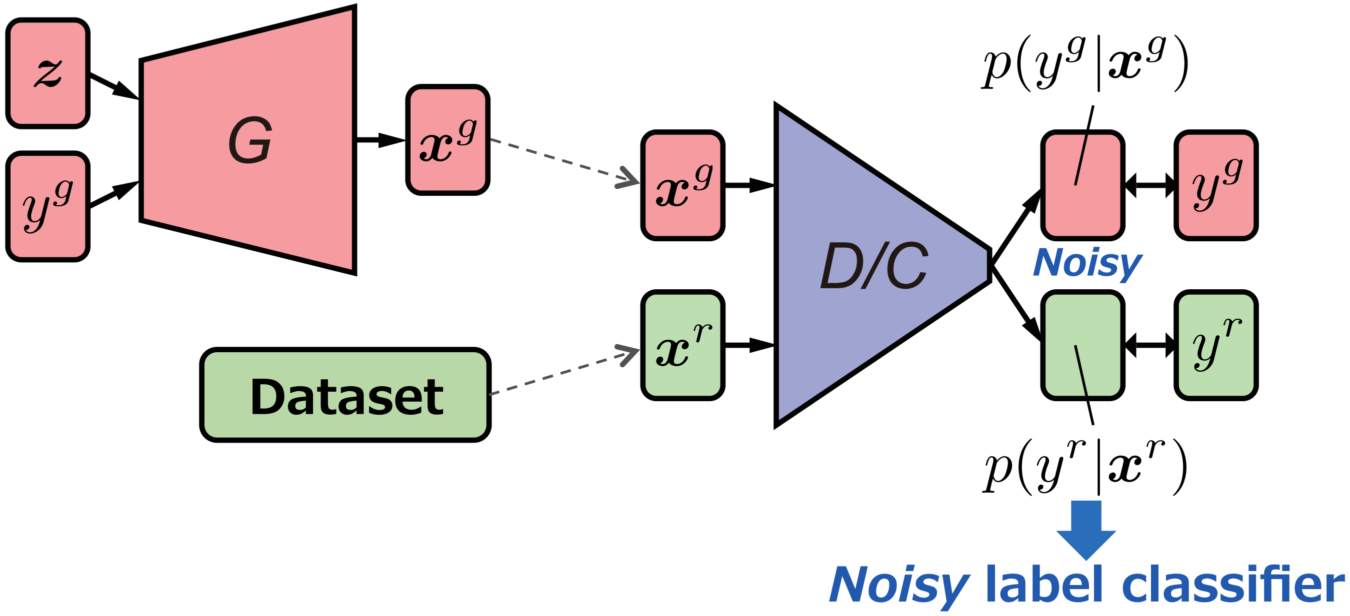

Limitation

C can fit ***noisy* labels** when trained with ***noisy* labeled data**.

[1] Odena et al. ICML 2017.

Proposal: rAC-GAN

Label-noise robust auxiliary classifier GAN

***G**enerator*: $G(\boldsymbol{z}, \hat{y}^g)$

***D**iscriminator/**C**lassifier*: $D(\boldsymbol{x})$ / $\sum p(\tilde{y}|\hat{y})\hat{C}(\hat{y}|\boldsymbol{x})$

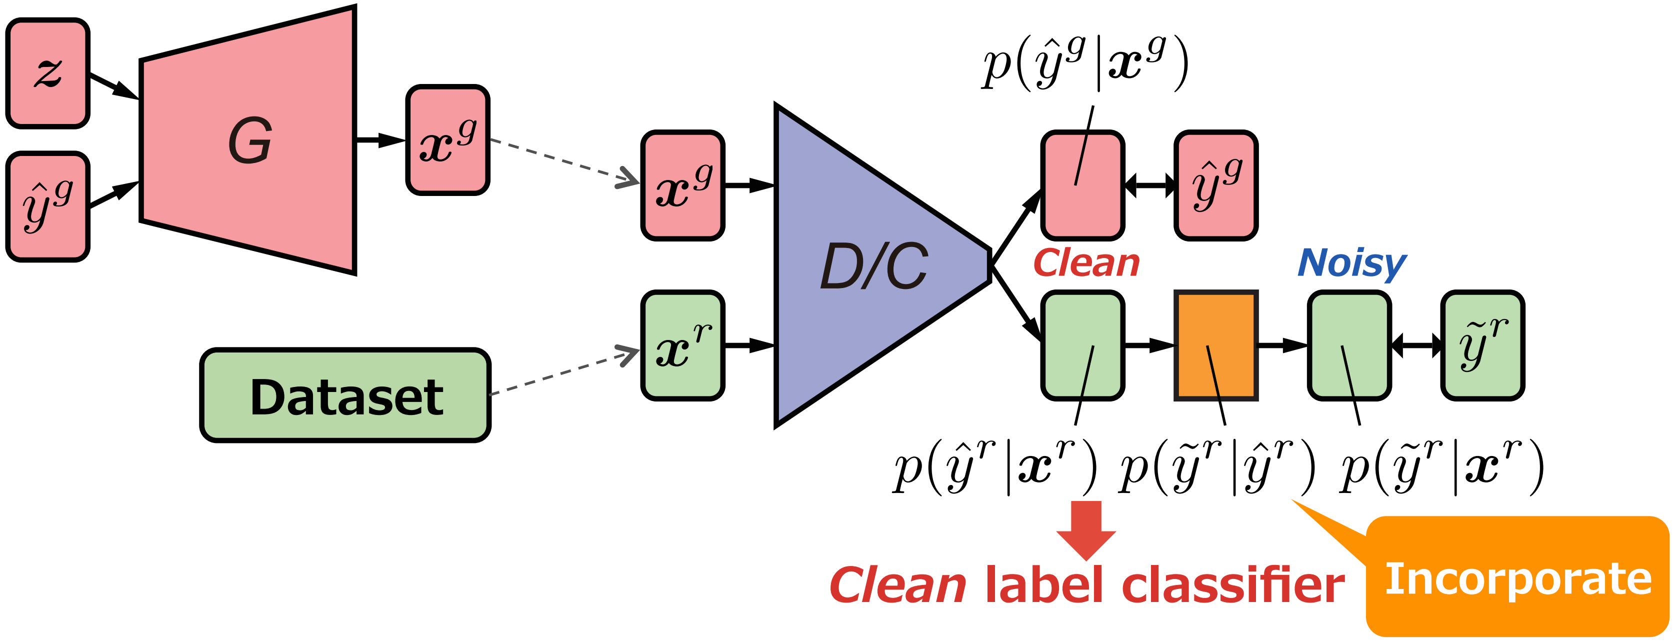

Solution

We correct the ***C***'s prediction using the **noise transition model** (forward correction [4]).

[4] Patrini et al. CVPR 2017.

Baseline: cGAN

Conditional GAN [2, 3]

***G**enerator*: $G(\boldsymbol{z}, y^g)$

***D**iscriminator*: $D(\boldsymbol{x}, y)$

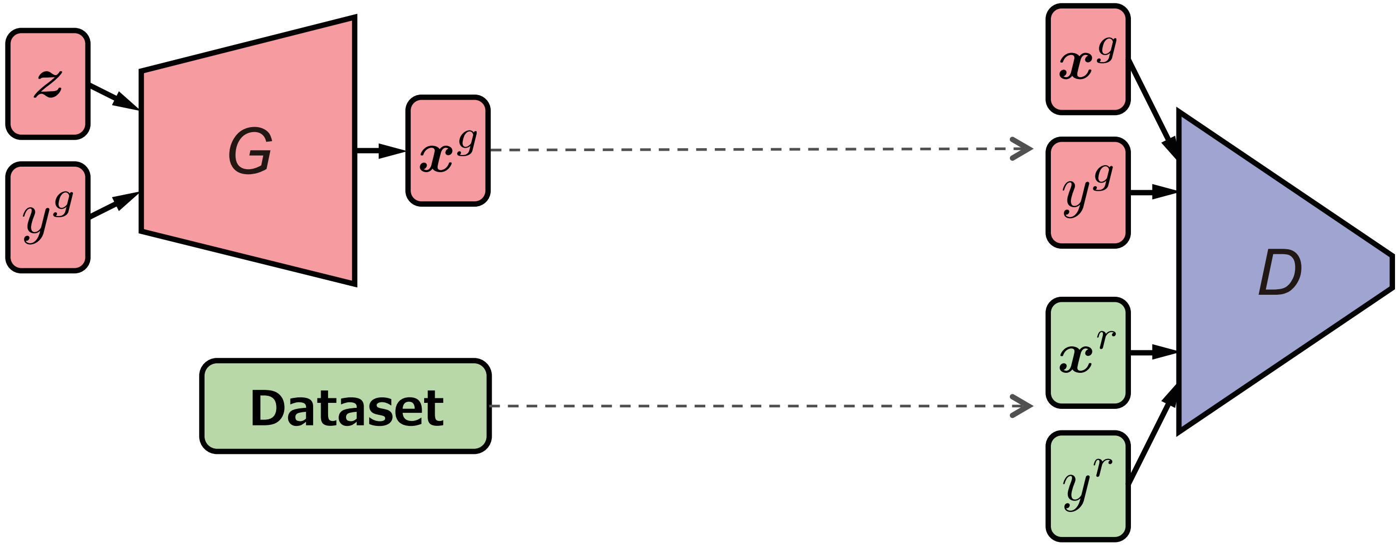

[2] Mirza & Osindero. arXiv 2014. [3] Miyato & Koyama. ICLR 2018.

Baseline: cGAN

Conditional GAN [2, 3]

***G**enerator*: $G(\boldsymbol{z}, y^g)$

***D**iscriminator*: $D(\boldsymbol{x}, y)$

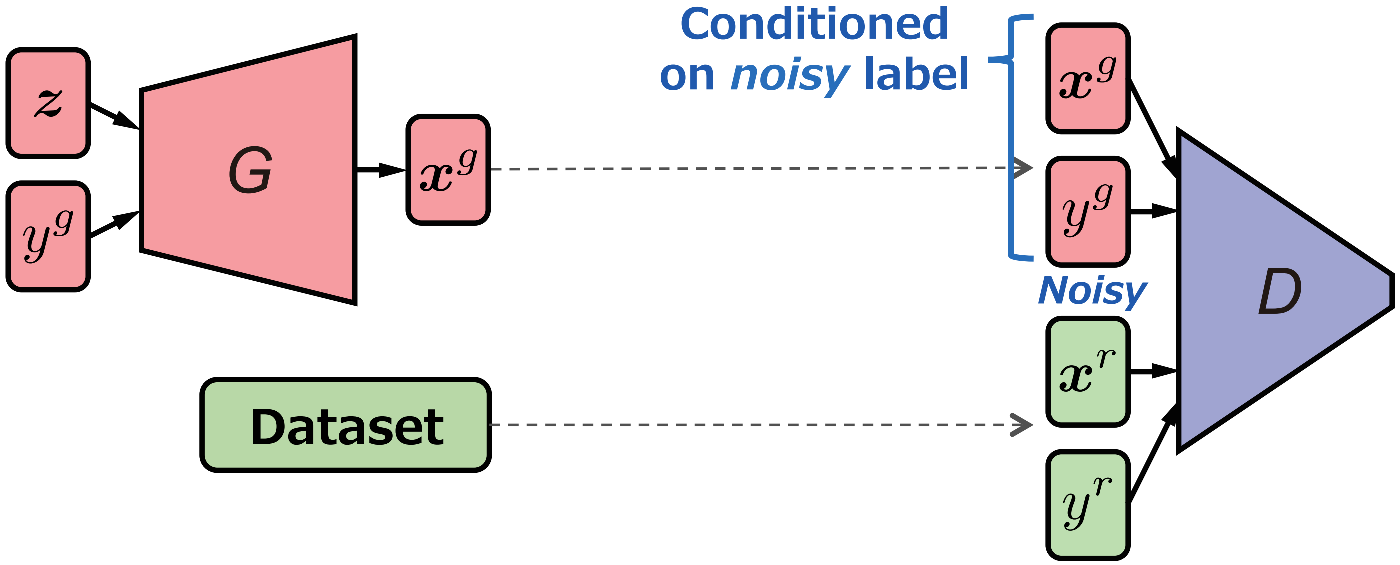

Limitation

G is optimized conditioned on ***noisy* labels** when trained with ***noisy* labeled data**.

[2] Mirza & Osindero. arXiv 2014. [3] Miyato & Koyama. ICLR 2018.

Proposal: rcGAN

Label-noise robust conditional GAN

***G**enerator*: $G(\boldsymbol{z}, \hat{y}^g)$

***D**iscriminator*: $D(\boldsymbol{x}, \tilde{y})$ $(\tilde{y}^g \sim p(\tilde{y}|\hat{y}^g))$

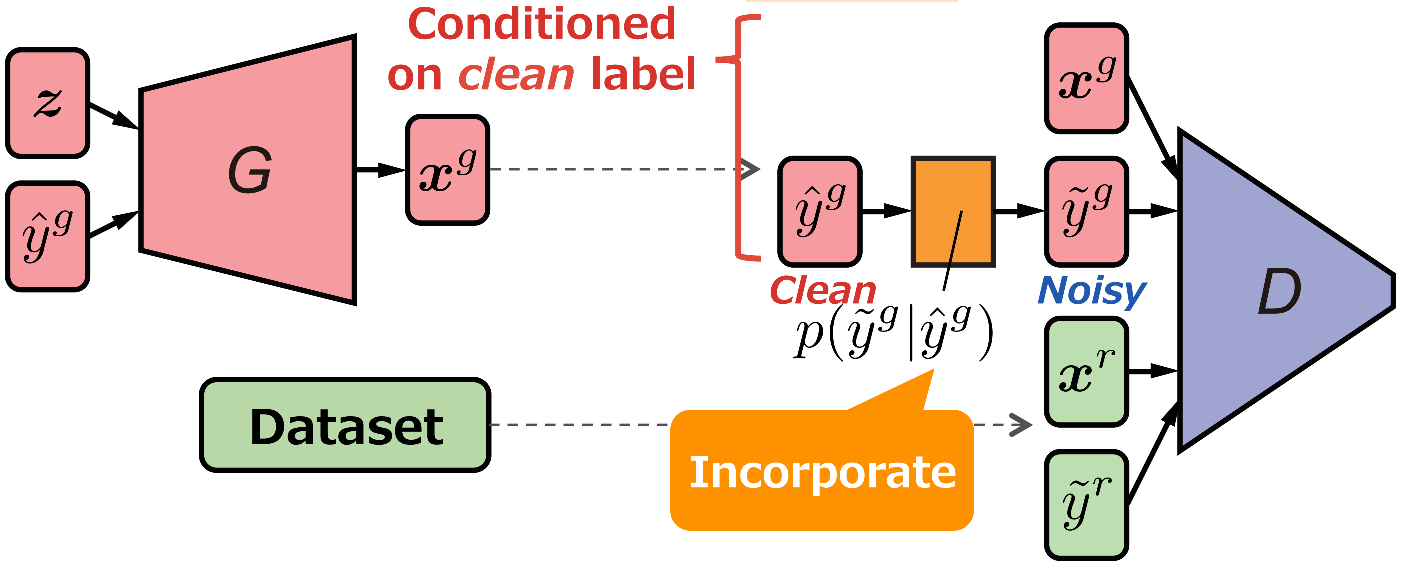

Solution

We correct the ***D***'s input using the **noise transition model**.

Experiment: Comprehensive study

Experimental conditions

Dataset: CIFAR-10 [5], CIFAR-100 [5]

Noise: Symmetric [6], Asymmetric [4] (Noise rate $\in \{0, 0.1, 0.3, 0.5, 0.7, 0.9\}$)

GAN configurations: DCGAN [7], WGAN-GP [8], CT-GAN [9], SN-GAN [10]

Comparison: AC-GAN vs. rAC-GAN, cGAN vs. rcGAN

Evaluation metrics: FID [11], Intra FID [3], GAN-test [12], GAN-train [12]

We tested 336 conditions in total.

[5] Krizhevsky. 2009. [6] van Rooyen et al. NIPS 2015. [4] Patrini et al. CVPR 2017.
[7] Radford et al. ICLR 2016. [8] Gulrajani et al. NIPS 2017. [9] Wei et al. ICLR 2018. [10] Miyato et al. ICLR 2018.
[11] Heusel et al. NIPS 2017. [3] Miyato & Koyama. ICLR 2018. [12] Shmelkov et al. ECCV 2018.

Experiment: Quantitative results

Comparison between the proposed model and the baselines across all the conditions

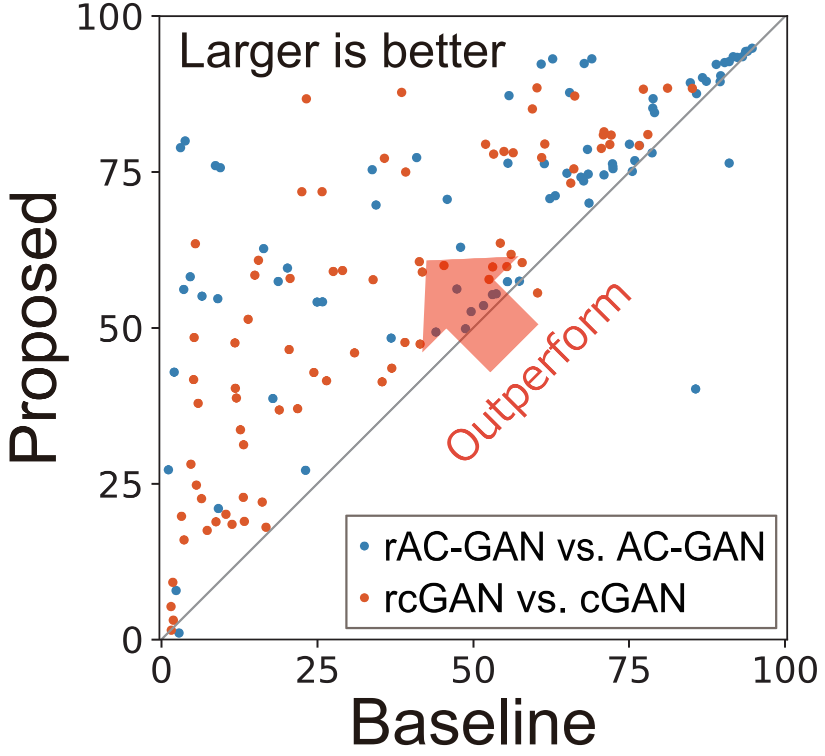

(a) GAN-test

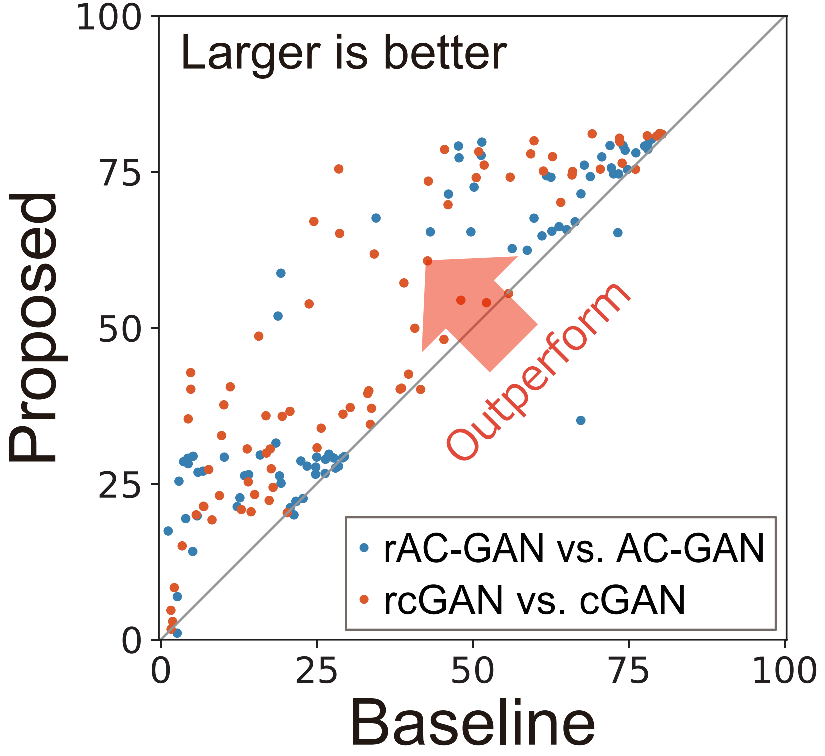

(b) GAN-train

The proposed models outperform the baselines in most cases.

Experiment: Qualitative results I

CIFAR-10 symmetric noise (uniform noise)

Baseline

Proposed

AC-CT-GAN (36.4)

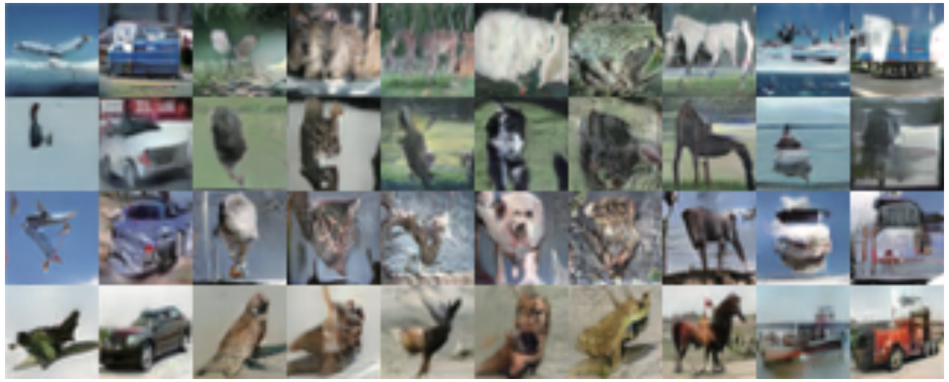

rAC-CT-GAN (**30.2**)

Significantly degraded

cSN-GAN (72.0)

rcSN-GAN (**28.6**)

The number indicates Intra FID. A smaller value is better.

Experiment: Qualitative results II

CIFAR-10 asymmetric noise (class-dependent noise, e.g., cat ⇄ dog)

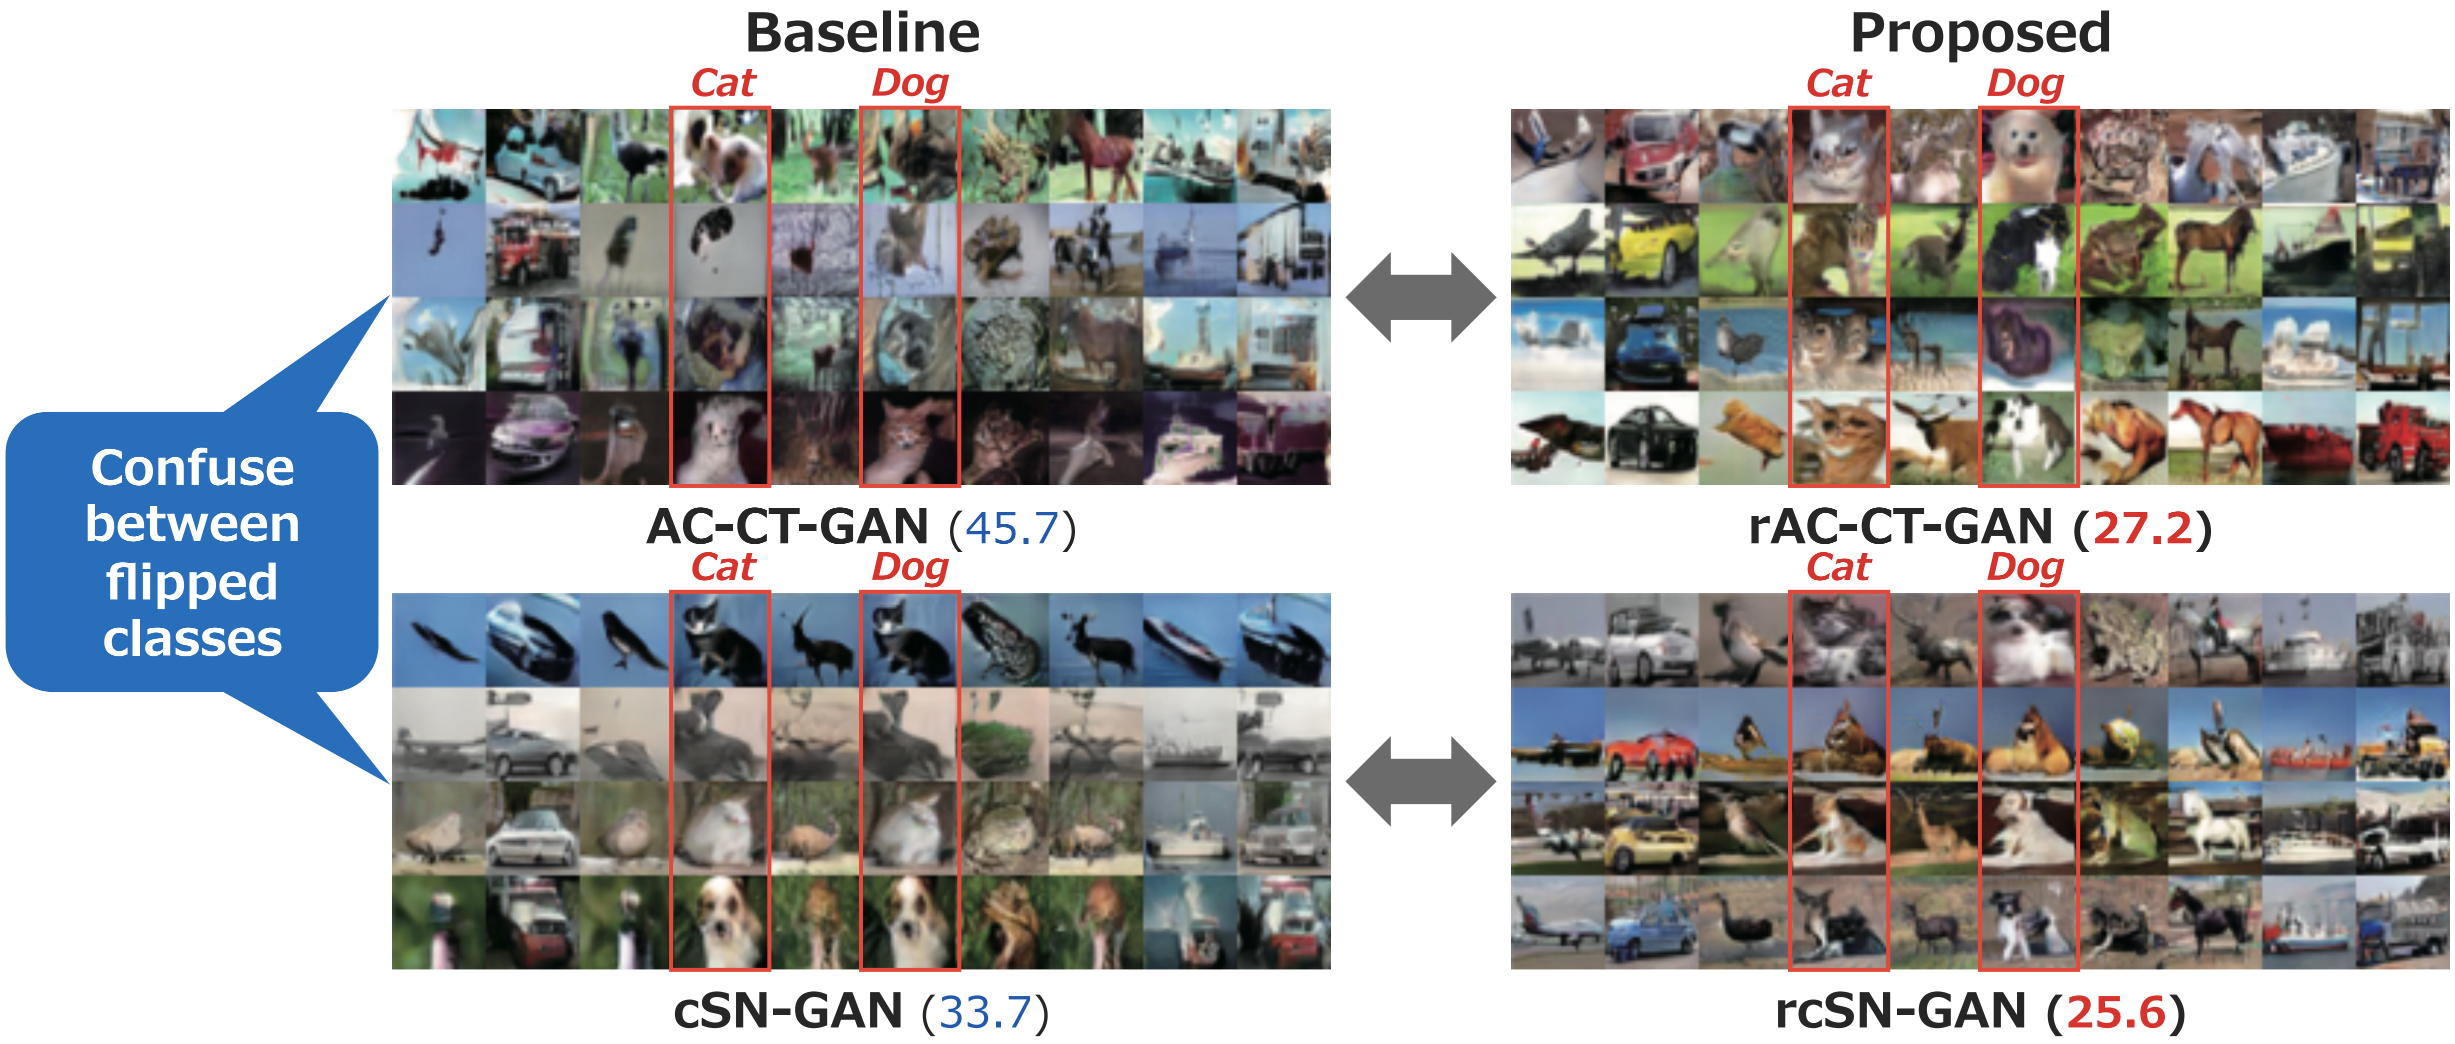

The number indicates Intra FID. A smaller value is better.

Further analyses

Effects of estimated noise transition model

We examined the effect when
the noise transition model is *estimated from data* [4].

$$\bar{\boldsymbol{x}}^i = \underset{\boldsymbol{x} \in \mathcal{X}'}{\operatorname{argmax}}\, C'(\tilde{y} = i | \boldsymbol{x})$$

$$T'_{i,j} = C'(\tilde{y} = j | \bar{\boldsymbol{x}}^i)$$

Robust two-step training algorithm [4]

Evaluation of improved technique

We validated the effect of an *improved technique*,
which we developed to boost the performance
in severely noisy setting.

$$\mathcal{L}_{\mathrm{MI}} = \mathbb{E}_{\boldsymbol{z} \sim p(\boldsymbol{z}), \hat{y}^g \sim p(\hat{y})} [-\log Q(\hat{y} = \hat{y}^g | G(\boldsymbol{z}, \hat{y}^g))]$$

Loss for improving the performance

Evaluation on real-world noise

We tested on Clothing1M [13],
which includes *real-world noisy labeled data*.

Qualitative results on Clothing1M

[4] Patrini et al. CVPR 2017. [13] Xiao et al. CVPR 2015.

Thank you!

Our code is publicly available at https://github.com/takuhirok/rGAN/

Poster

Session: Synthesis

#: 133

Time: Tuesday 15:20-18:00

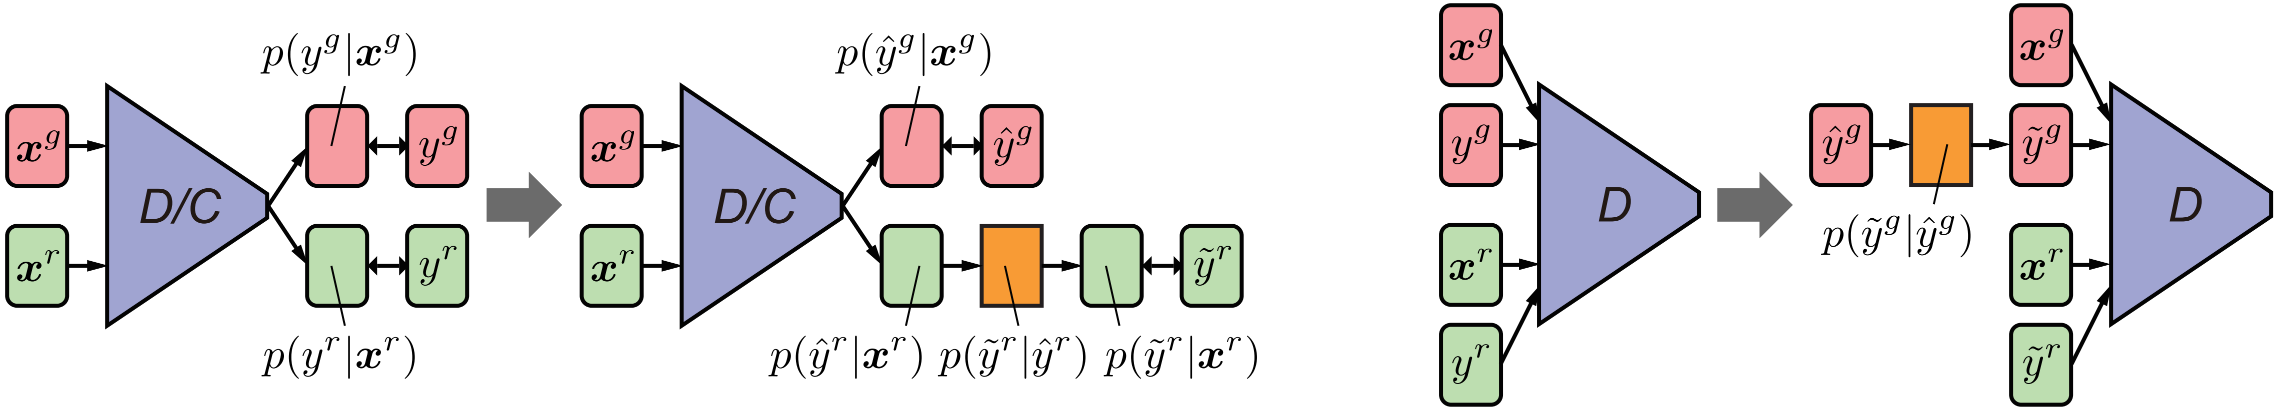

AC-GAN [1] **rAC-GAN** **cGAN** [2, 3] **rcGAN**

[1] Odena et al. ICML 2017. [2] Mirza & Osindero. arXiv 2014. [3] Miyato & Koyama. ICLR 2018.