Label-Noise Robust Generative Adversarial Networks

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

[1]The University of Tokyo [2]RIKEN

CVPR LONG BEACH CALIFORNIA June 16-20, 2019

Introduction

Objective: Label-noise robust image generation

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

(a) *Clean* labeled data

Unobservable

(b) *Noisy* labeled data

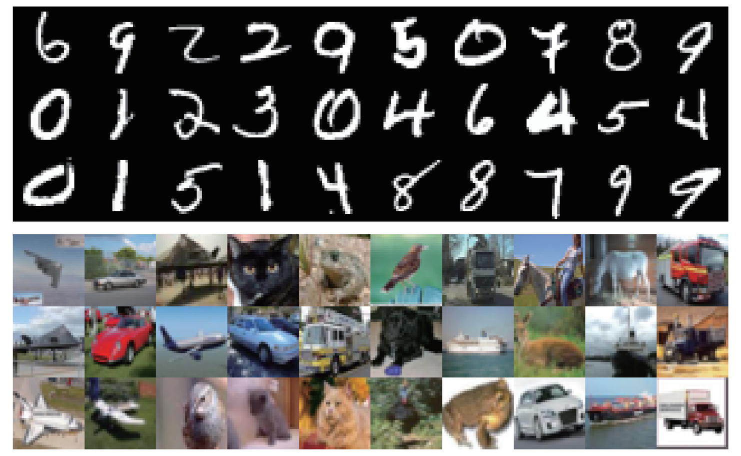

Observable

Challenge and contribution

- **Naïve conditional generative models** construct a generator conditioned on **observable (*noisy*)** labels (c).
- Our proposed **rGANs (label-noise robust GANs)** can construct a generator conditioned on ***clean* labels** (d) even when trained with ***noisy* labeled data** (b).

(c) cGAN trained with (b) [Baseline]

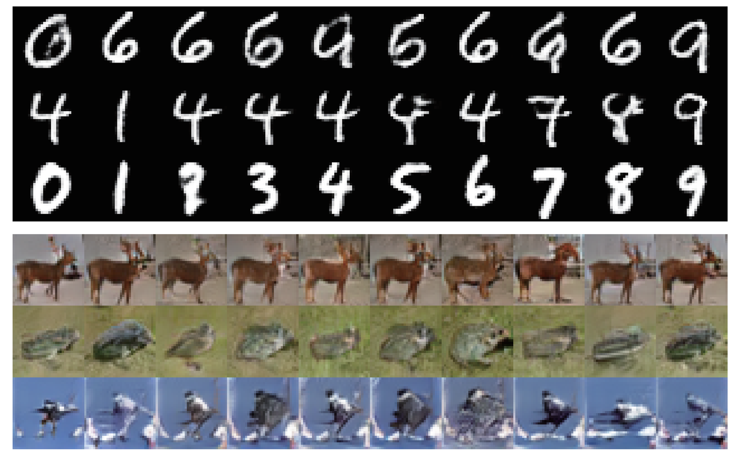

Fits observable (*noisy*) labels

(d) rcGAN trained with (b) [Proposed]

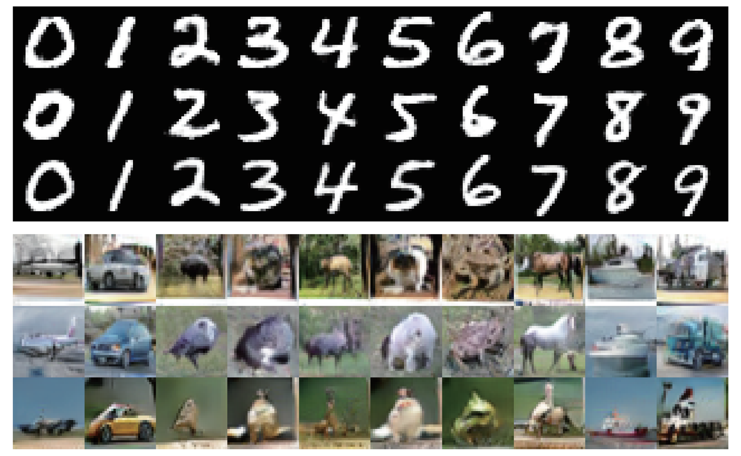

Robust to label noise

Key idea

Incorporation of noise transition model

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

Noisy label *Clean* label

Method

Baseline 1: AC-GAN (auxiliary classifier GAN) [Odena et al. 2017]

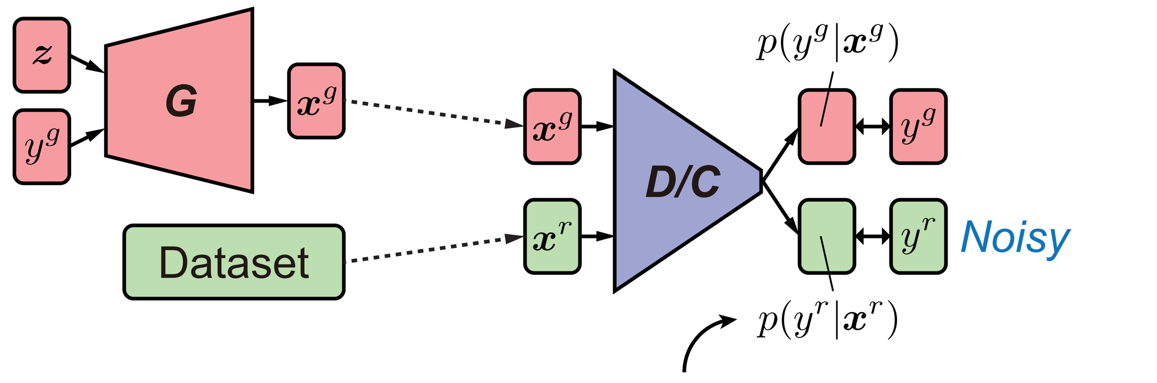

Generator $G(\boldsymbol{z}, y^g)$

Discriminator/Classifier $D(\boldsymbol{x})$ / $C(y|\boldsymbol{x})$

Limitation: ***C*** can fit ***noisy*** **labels** when trained with ***noisy*** **labels**.

Proposal 1: rAC-GAN (label-noise robust AC-GAN)

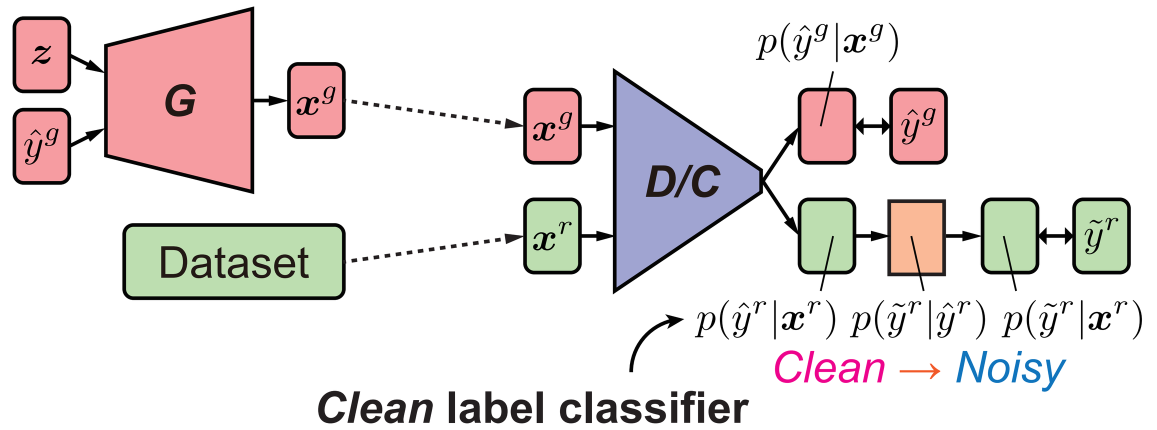

Generator $G(\boldsymbol{z}, \hat{y}^g)$

Discriminator/Classifier $D(\boldsymbol{x})$ / $\sum p(\tilde{y}|\hat{y})\hat{C}(\hat{y}|\boldsymbol{x})$

Solution: We correct ***C*** prediction using the **noise transition model**.

Baseline 2: cGAN (conditional GAN) [Mirza & Osindero. 2014, Miyato & Koyama. 2018]

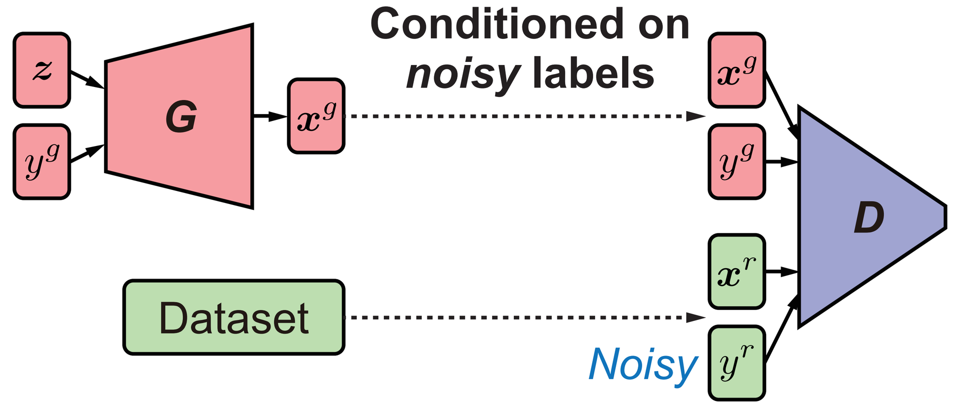

Generator $G(\boldsymbol{z}, y^g)$

Discriminator $D(\boldsymbol{x}, y)$

Limitation: ***G*** is optimized conditioned on ***noisy*** **labels** when trained with ***noisy*** **labels**.

Proposal 2: rcGAN (label-noise robust cGAN)

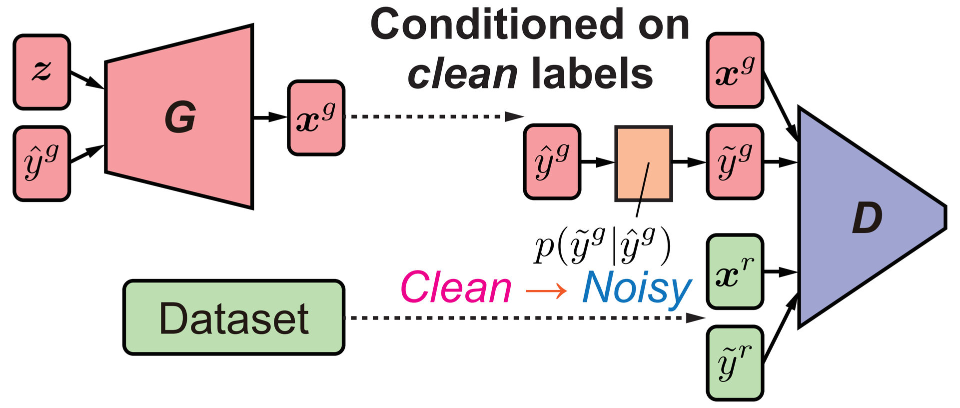

Generator $G(\boldsymbol{z}, \hat{y}^g)$

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

Solution: We correct ***D*** input using the **noise transition model**.

Experiments

Comprehensive study (336 conditions were tested in total)

Dataset: CIFAR-10 and CIFAR-100

Noise: Symmetric noise and asymmetric noise (noise rate μ ∈ {0, 0.1, 0.3, 0.5, 0.7, 0.9})

GAN configuration: DCGAN, WGAN-GP, CT-GAN, and SN-GAN

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

Evaluation metrics: FID, Intra FID, GAN-test, and GAN-train

Quantitative results

Comparison across all conditions

1. GAN-test

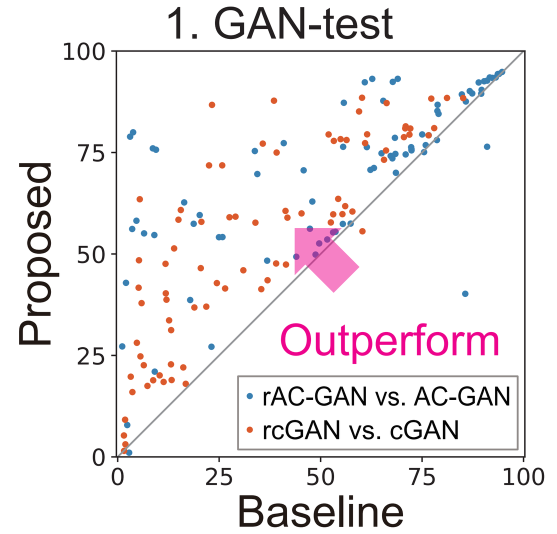

2. GAN-train

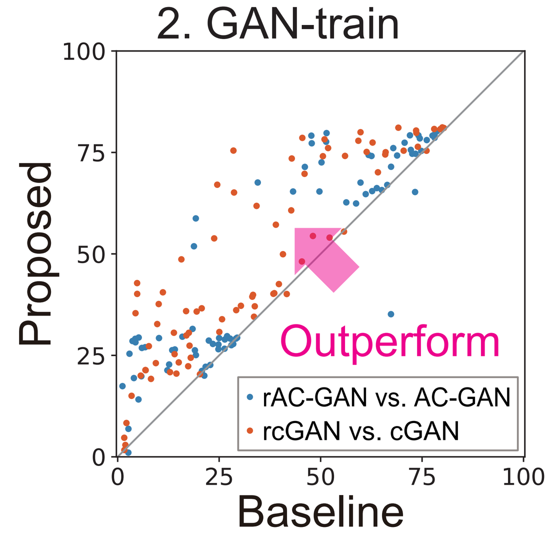

Qualitative results

[Baseline] [Proposed]

1. CIFAR-10 symmetric (uniform) noise with μ = 0.7

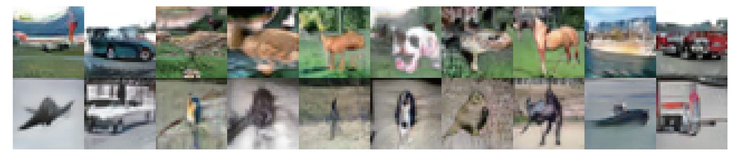

AC-CT-GAN (36.4)

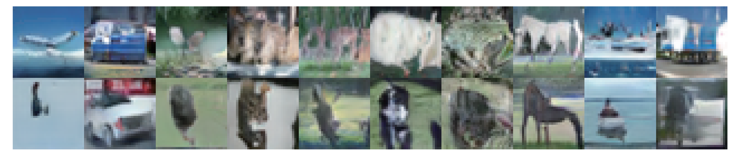

rAC-CT-GAN (30.2)

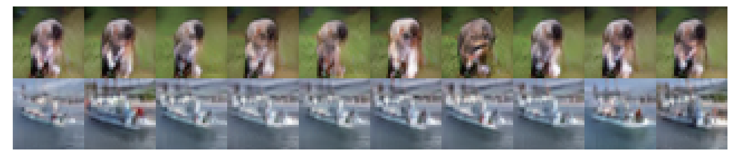

cSN-GAN (72.0)

rcSN-GAN (28.6)

Failed to learn disentangled representations

2. CIFAR-10 asymmetric (class-dependent; e.g., cat ⇄ dog) noise with μ = 0.7

AC-CT-GAN (45.7)

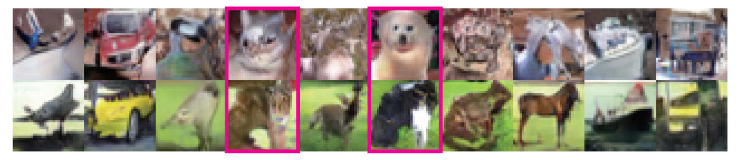

rAC-CT-GAN (27.2)

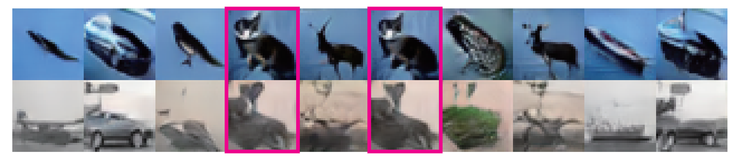

cSN-GAN (33.7)

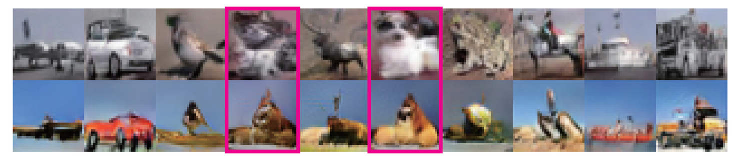

rcSN-GAN (25.6)

Confused between flipped classes

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

Further analyses (see paper for details)

Effects of estimated noise transition model

We examined the effect when *the noise transition model is estimated from data* [Patrini et al. 2017].

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

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

Robust two-step training algorithm

Evaluation of improved technique

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

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

Mutual information regularization

Evaluation on real-world noise

We tested rGANs in *a real-world noise setting* using Clothing1M [Xiao et al. 2015], which includes real-world noisy labeled data.

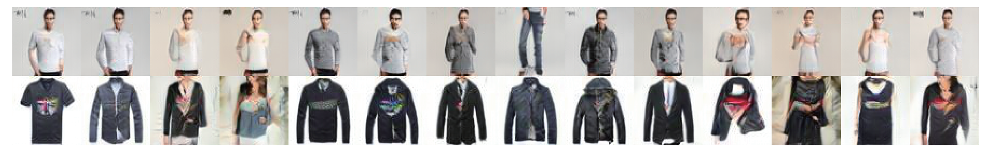

Qualitative results on Clothing1M