Robustness of Conditional GANs to Noisy Labels

Spotlight presentation, NeurIPS 2018

Kiran K. Thekumparampil[1] Ashish Khetan[1]
Zinan Lin[2] Sewoong Oh[1]

[1]University of Illinois at Urbana-Champaign

[2]Carnegie Mellon University

Poster #5, Tue, Dec 4 2018

Conditional GAN (cGAN) is vital for achieving high quality

- **Input:** Labeled real samples (X, Y)
- **Output:** Fake samples for label Y

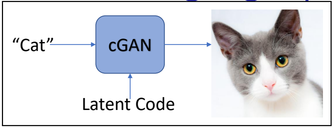

Conditional GAN (cGAN) is vital for achieving high quality

- **Input:** Labeled real samples (X, Y)
- **Output:** Fake samples for label Y

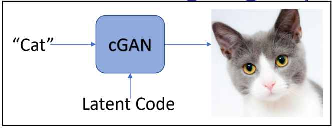

[Brock et al. 2018]

Conditional GAN (cGAN) is vital for achieving high quality

- **Input:** Labeled real samples (X, Y)
- **Output:** Fake samples for label Y

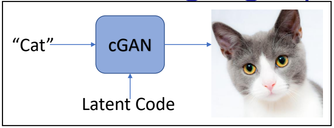

[Brock et al. 2018]

- **Visual quality:** cGAN $>>$ GAN

[https://github.com/tensorflow/models/tree/master/research/gan]

Conditional GAN is sensitive to noise in labels

cGAN trained with noisy labels produces samples

- that are **biased**, generating examples from wrong classes, and,
- of **lower quality** (red boxes).

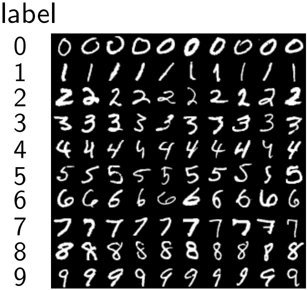

real data

Conditional GAN is sensitive to noise in labels

cGAN trained with noisy labels produces samples

- that are **biased**, generating examples from wrong classes, and,
- of **lower quality** (red boxes).

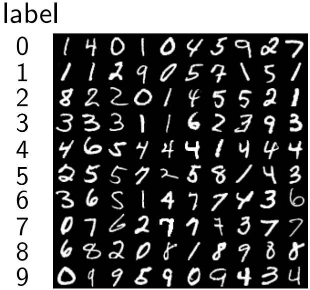

noisy real data

Conditional GAN is sensitive to noise in labels

cGAN trained with noisy labels produces samples

- that are **biased**, generating examples from wrong classes, and,
- of **lower quality** (red boxes).

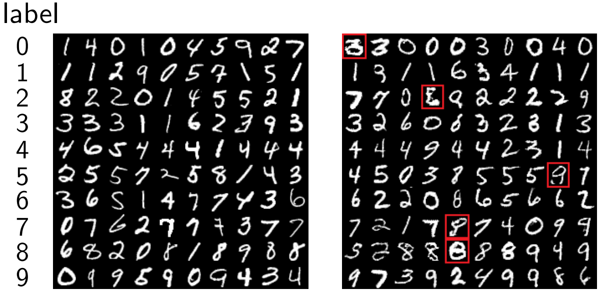

noisy real data standard cGAN

Conditional GAN is sensitive to noise in labels

cGAN trained with noisy labels produces samples

- that are **biased**, generating examples from wrong classes, and,
- of **lower quality** (red boxes).

noisy real data | standard cGAN | our RCGAN

Conditional GAN (cGAN)

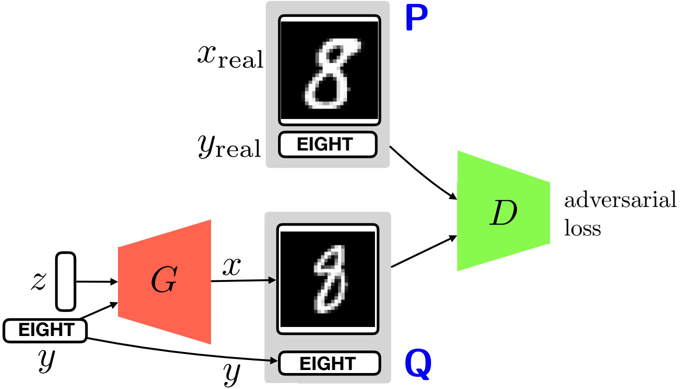

$$\min_{\mathbf{Q}} \mathrm{JS}(\mathbf{P} \,||\, \mathbf{Q})$$

Conditional GAN under noisy labeled data

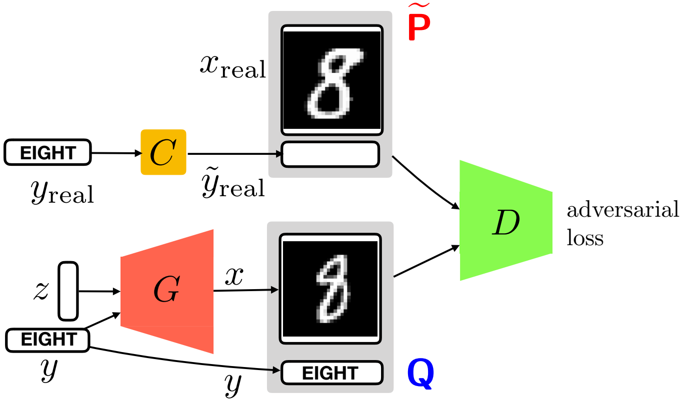

$$\min_{\mathbf{Q}} \mathrm{JS}(\widetilde{\mathbf{P}} \,||\, \mathbf{Q})$$

Robust Conditional GAN (RCGAN) Architecture

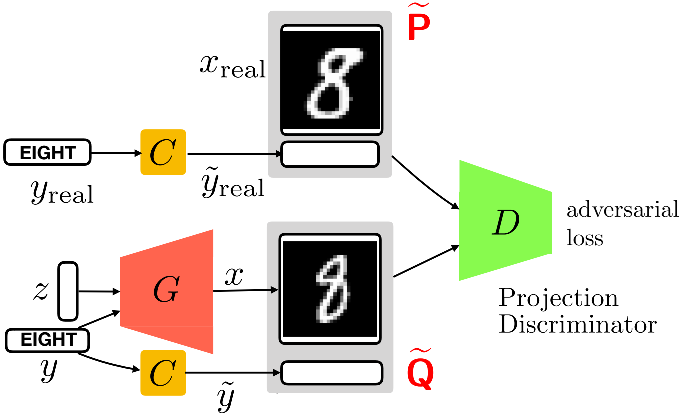

$$\min_{\mathbf{Q}} \mathrm{JS}(\widetilde{\mathbf{P}} \,||\, \widetilde{\mathbf{Q}})$$

[Bora et al. 2018, Miyato et al. 2018, Sukhbaatar et al. 2015]

Minimizing noisy divergence minimizes true divergence

Let $\widetilde{P}$ & $\widetilde{Q}$ be the noisy labeled versions of P & Q.

Theorem 1 (Population-level Analysis)

$$\left.\begin{array}{l} \mathrm{TV}\left(\widetilde{P},\widetilde{Q}\right) \;\le \mathrm{TV}\left(P,Q\right) \;\le\; M_C\,\mathrm{TV}\left(\widetilde{P},\widetilde{Q}\right) \\ \mathrm{JS}\left(\widetilde{P}\,\middle\|\,\widetilde{Q}\right) \;\le \mathrm{JS}(P \,\|\, Q) \;\le\; M_C\sqrt{8\,\mathrm{JS}\left(\widetilde{P}\,\middle\|\,\widetilde{Q}\right)} \end{array}\right\} \Longrightarrow \widetilde{\mathbf{Q}} = \widetilde{\mathbf{P}} \Rightarrow \mathbf{Q} = \mathbf{P}$$

where TV: Total Variation, JS: Jensen-Shannon divergence and $M_C \triangleq \max_i \sum_j \left|(C^{-1})_{ij}\right|$.

Minimizing noisy divergence minimizes true divergence

Let $\widetilde{P}$ & $\widetilde{Q}$ be the noisy labeled versions of P & Q.

Theorem 1 (Population-level Analysis)

$$\left.\begin{array}{l} \mathrm{TV}\left(\widetilde{P},\widetilde{Q}\right) \le \mathrm{TV}(P,Q) \le M_C\,\mathrm{TV}\left(\widetilde{P},\widetilde{Q}\right) \\ \mathrm{JS}\left(\widetilde{P}\,\middle\|\,\widetilde{Q}\right) \le \mathrm{JS}(P\,\|\,Q) \le M_C\sqrt{8\,\mathrm{JS}\left(\widetilde{P}\,\middle\|\,\widetilde{Q}\right)} \end{array}\right\} \Longrightarrow \widetilde{\mathbf{Q}} = \widetilde{\mathbf{P}} \Rightarrow \mathbf{Q} = \mathbf{P}$$

where TV*: Total Variation, JS: Jensen-Shannon divergence and* $M_C \triangleq \max_i \sum_j \left|(C^{-1})_{ij}\right|$.

Neural Network Distance ($d_{\mathcal{F}}$) w.r.t a class of parametric discriminator functions $\mathcal{F}$ is known to generalize [Arora et al. 2017]

Minimizing noisy divergence minimizes true divergence

Let $\widetilde{P}_n$ & $\widetilde{Q}_n$ be the empirical noisy real and generated distributions.

Theorem 2 (Finite Sample Analysis)

If $\mathcal{F}$ satisfies **inclusion condition***, then* $\exists c > 0$ *such that*

$$d_{\mathcal{F}}(\widetilde{P}_n, \widetilde{Q}_n) - \epsilon \;\; \leq \;\; d_{\mathcal{F}}(P, Q) \;\; \leq \;\; M_C \left(d_{\mathcal{F}}(\widetilde{P}_n, \widetilde{Q}_n) + \epsilon \right)$$

with probability at least $1 - e^{-p}$ *for any* $\varepsilon > 0$ *and* $n \geq cp \log\left(pL/\epsilon\right)/\epsilon^2$ *when* $\mathcal{F}$ *is L-Lipschitz in p parameters*

Projection Discriminator satisfies inclusion condition

RCGAN generates correct class (MNIST)

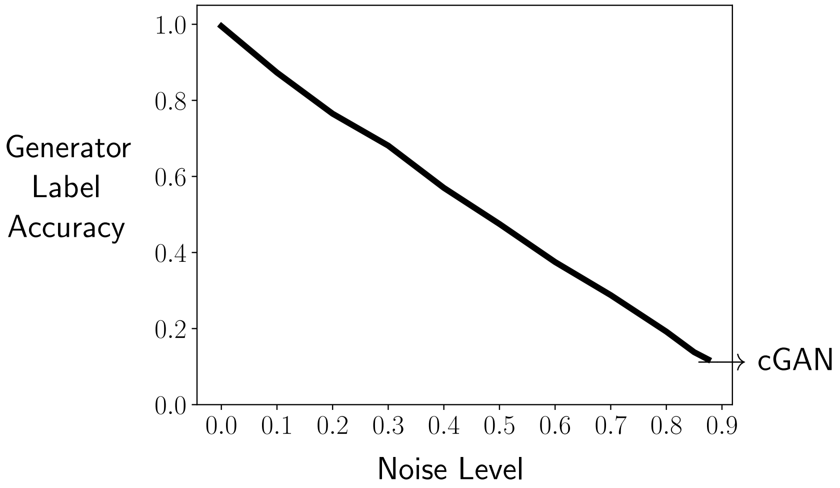

RCGAN generates correct class (MNIST)

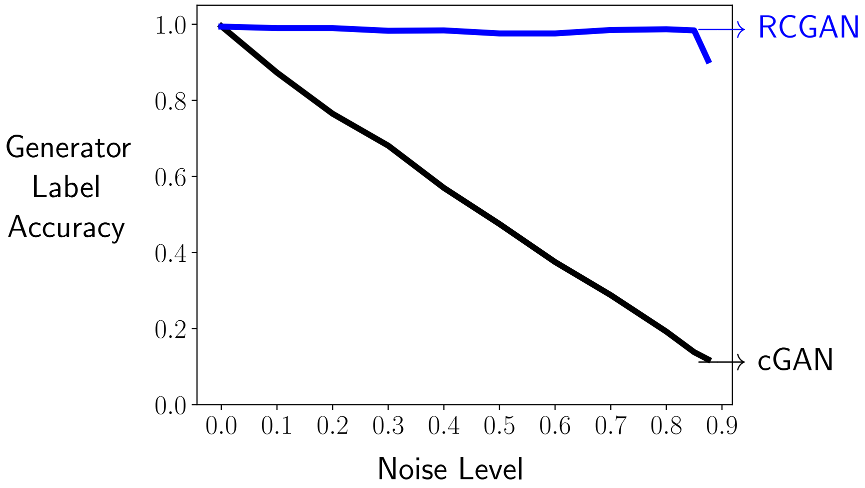

RCGAN generates correct class (MNIST)

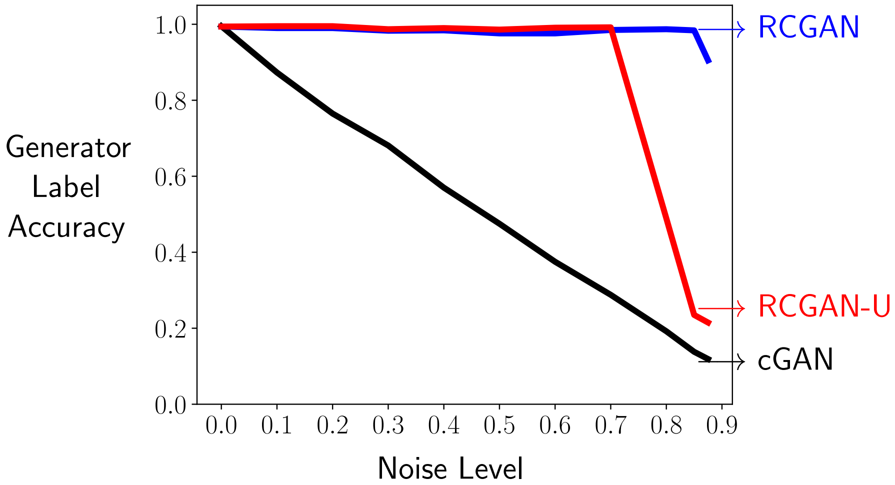

RCGAN improves quality of samples (CIFAR-10)

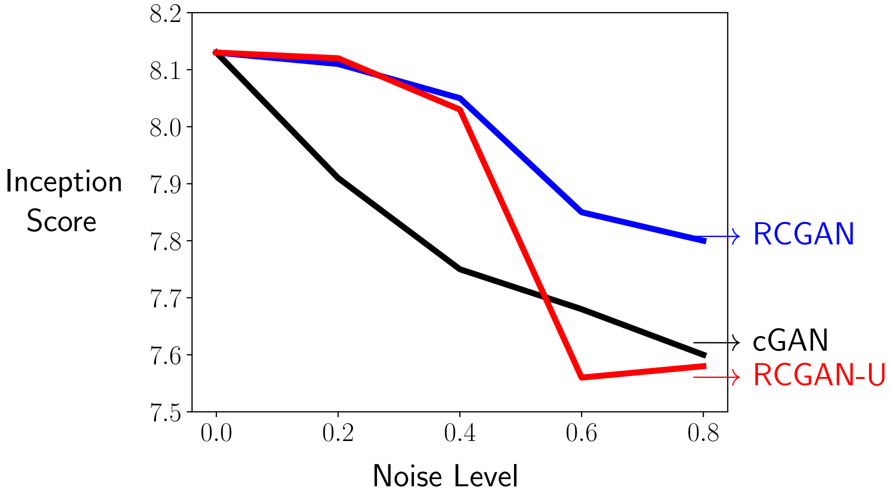

RCGAN can correct noisy training labels (MNIST)

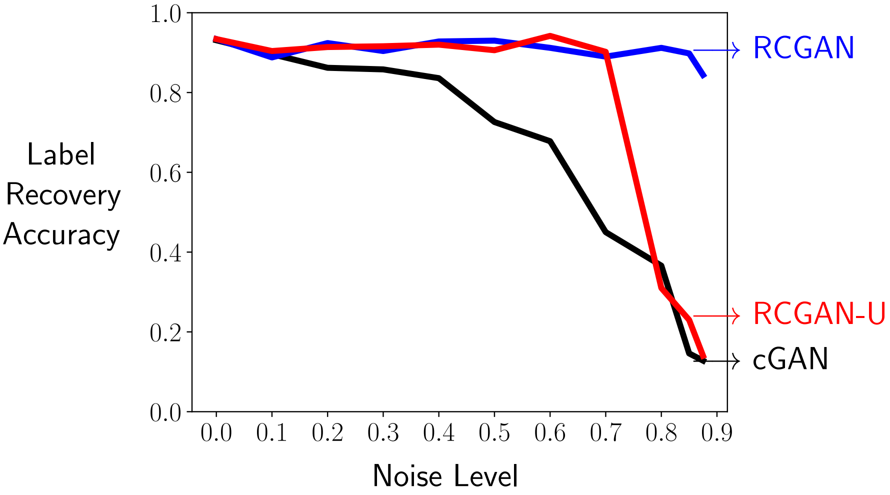

Thank you

Poster #5, Tue, Dec 04

https://github.com/POLane16/Robust-Conditional-GAN

[Arora 2015] S. Arora, R. Ge, Y. Liang, T. Ma, and Y. Zhang. Generalization and equilibrium in generative adversarial nets (GANs), *ICML 2018*.
[Bora 2018] A. Bora, E. Price, and A. G. Dimakis. AmbientGAN: Generative models from lossy measurements, *ICLR, 2018.*
[Brock 2018] A. Brock, J. Donahue, and K. Simonyan. Large scale gan training for high fidelity natural image synthesis, *arXiv preprint arXiv:1809.11096*.
[Miyato 2018] T. Miyato, and M. Koyama. cGANs with projection discriminator. *ICLR, 2018.*
[Sukhbaatar 2015] S. Sukhbaatar, J. Bruna, M. Paluri, L. Bourdev, and R. Fergus. Training convolutional networks with noisy labels. *In ICLR, Workshop, 2015.*