Soft Curriculum for Learning Conditional GANs with Noisy-Labeled and Uncurated Unlabeled Data: Supplementary Material

Kai Katsumata[1] Duc Minh Vo[1] Tatsuya Harada[1,2] Hideki Nakayama[1]
[1]The University of Tokyo, [2]RIKEN
{katsumata,vmduc,nakayama}@nlab.ci.i.u-tokyo.ac.jp, harada@mi.t.u-tokyo.ac.jp

A. Details of compared methods.

We provide the training procedure for compared methods. All methods build upon the DiffAug CR-GAN [7] baseline except the experiments of Tab. 6. DiffAug CR-GAN only use labeled samples for training. Other baseline methods have label distribute function C for unlabeled samples. S^3GAN [2], OSSGAN [1], and CurriculumGAN have a classifier trained simultaneously. RandomGAN has $C(x) = \boldsymbol{e}^y$ where y are sampled from $\mathcal{Y}$ uniformly. Here, $\boldsymbol{e}^y$ is the y-th standard basis vector of $\mathbb{R}^K$. SingleGAN has $C(x) = [1/K, 1/K, \dots, 1/K]^\intercal$. Semi-supervised baselines assign temporal labels to unlabeled samples using C. Namely, the labels the discriminator takes are defined as

$$\boldsymbol{y}' = \begin{cases} y & x \text{ is labeled} \\ C(x) & \text{otherwise,} \end{cases} \tag{1}$$

where x is a sample, and y is a corresponding label. For extended baselines with re-labeling strategy and CurriculumGAN, the labels for the discriminator are defined as

$$\boldsymbol{y}' = \begin{cases} (y + C(x))/2 & x \text{ is labeled} \\ C(x) & \text{otherwise.} \end{cases} \tag{2}$$

CurriculumGAN assigns the confidence c to real unlabeled samples in the curriculum learning manner of open-set semi-supervised learning [6] with Otsu-threshold [3]. To achieve the best performance, CurriculumGAN corrects labels of labeled data and applies the curriculum learning to only unlabeled data. In other words, it assigns the high confidence $c = 1$ to all labeled data.

B. Additional results

We run additional experiments on TinyImageNet [5] to see the behavior of our method on different numbers of unlabeled samples. We conduct experiments on five configurations: 10%, 30%, 50%, 70%, 100% usage ratios. Other parameters are the same: 150 closed-set classes, 30% label noise, and 20% labeled samples. Our method achieves a consistent performance gain on different numbers of unlabeled samples, as shown in Fig. A.

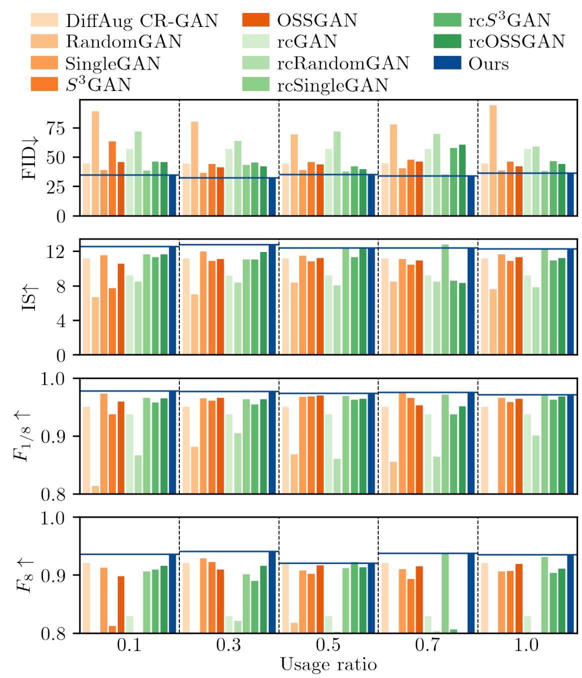

Figure A. Quantitative comparison over different numbers of unlabeled samples. We report the results of the experiments on the TinyImageNet dataset with 150 classes, 20% labeled data, 30% label noise ratio, and usage ratio for unlabeled data of $\{10\%, 30\%, 50\%, 70\%, 100\%\}$. Our method yields better performance stably (blue line).

Figure B provide qualitative comparison in the experiments on ImageNet [4] with 200 closed-set classes and 4% labeled data. We can observe the performance gains of our method in terms of both image quality and diversity.

References

[1] Kai Katsumata, Duc Minh Vo, and Hideki Nakayama. OSSGAN: open-set semi-supervised image generation. In *CVPR*, pages 11185–11193, 2022. 1

Figure B. Visual comparisons of class-conditional image synthesis results on ImageNet. Our method produces diverse images while respecting the given condition.

[2] Mario Lučić, Michael Tschannen, Marvin Ritter, Xiaohua Zhai, Olivier Bachem, and Sylvain Gelly. High-fidelity image generation with fewer labels. In *ICML*, pages 4183–4192, 2019. 1

[3] Nobuyuki Otsu. A threshold selection method from gray-level histograms. *IEEE Transactions on Systems, Man, and Cybernetics*, 9(1):62–66, 1979. 1

[4] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet large scale visual recognition challenge. 1

[5] Jiayu Wu, Qixiang Zhang, and Guoxi Xu. Tiny ImageNet challenge. 1

[6] Qing Yu, Daiki Ikami, Go Irie, and Kiyoharu Aizawa. Multi-task curriculum framework for open-set semi-supervised learning. In *ECCV*, pages 438–454, 2020. 1

[7] Shengyu Zhao, Zhijian Liu, Ji Lin, Jun-Yan Zhu, and Song Han. Differentiable augmentation for data-efficient GAN training. In *NeurIPS*, 2020. 1