A L2 Pairwise Distance

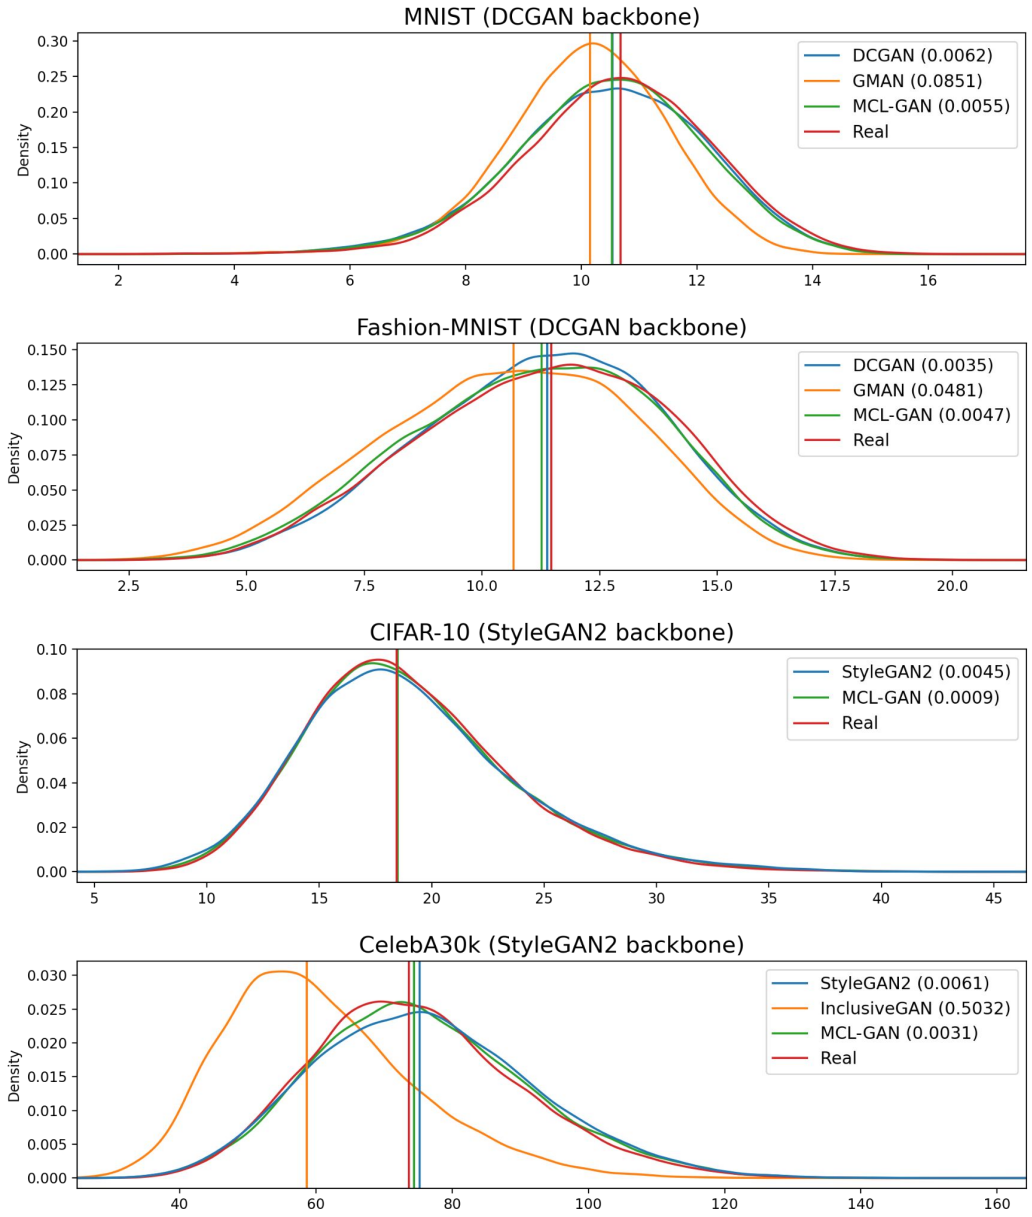

Figure A: Distribution of L2 pairwise distance. We randomly sample a pair of images from model samples (or training dataset in case of real) then plot the distribution from the statistics of pairwise distances following [1]. The vertical lines of each plot indicate the median values of the corresponding statistics and we also provide approximate KL divergences with respect to real data in the legends of the plots.

B Fashion-MNIST and partial MNIST

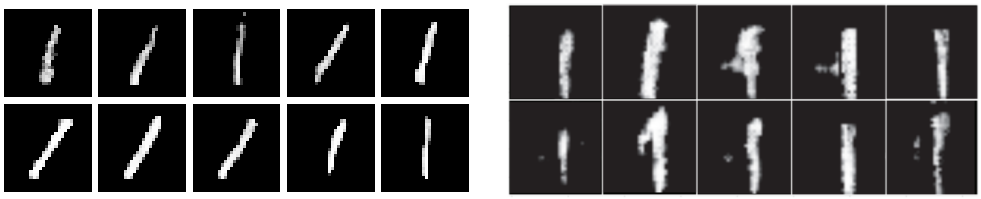

MCL-GAN (M=10) [2]

Figure B: Random "1" samples. We also show the result of [2] which is captured from Figure 3 of [2] for comparison. Note that [2] shows the most confident samples while our results are randomly sampled but have superior image quality.

C Nearest neighbors from train dataset

Figure C: Nearest neighbors from training data. We visualize nearest neighbors from training data of some images generated by MCL-GAN with StyleGAN2 [3] backbone. For each generated sample, we provide 3 nearest neighbors in terms of L2 distance on the pixel space. The nearest samples are sorted in increasing order in terms of the distance to the corresponding generated sample from top to bottom.

References

[1] Metz, L., Poole, B., Pfau, D., Sohl-Dickstein, J.: Unrolled generative adversarial networks. In ICLR. (2017)

[2] Zhong, P., Mo, Y., Xiao, C., Chen, P., Zheng, C.: Rethinking generative mode coverage: A pointwise guaranteed approach. In NeurIPS. Volume 32. (2019)

[3] Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of stylegan. In CVPR. (2020) 8110–8119