Spider GANs: Leveraging Friendly Neighbors to Accelerate GAN Training

Siddarth Asokan[1] and Chandra Sekhar Seelamantula[2]
[1]Robert Bosch Centre for Cyber-Physical Systems,
[2]Department of Electrical Engineering
Indian Institute of Science, Bangalore, India

JUNE 18-22, 2023
CVPR
VANCOUVER, CANADA

1. Generative Adversarial Networks

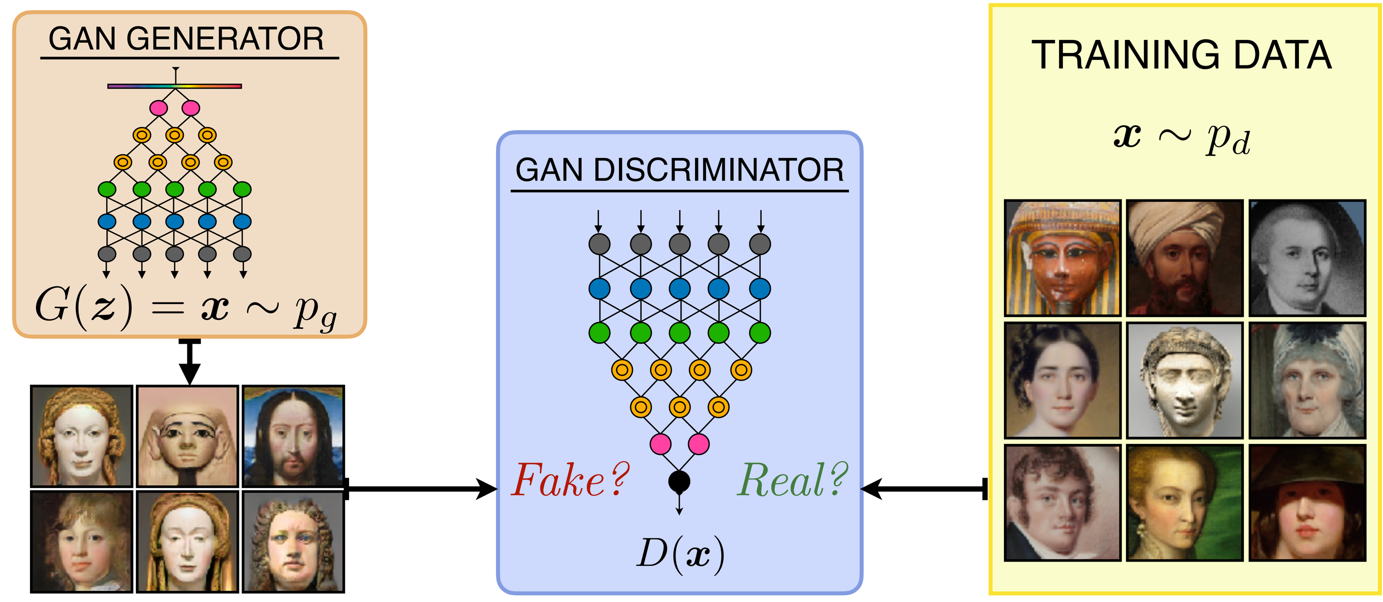

- The generator G transforms noise $\boldsymbol{z} \sim \mathcal{N}(\mathbf{0}, \mathbb{I})$ into fake samples $\boldsymbol{x} = G(\boldsymbol{z})$.
- The GAN[1] discriminator D classifies $\boldsymbol{x}$ as real or fake.
- The optimal generator learns to confuses the discriminator!

[1]Goodfellow et al., NeurIPS 14

2. The Spider GAN Philosophy

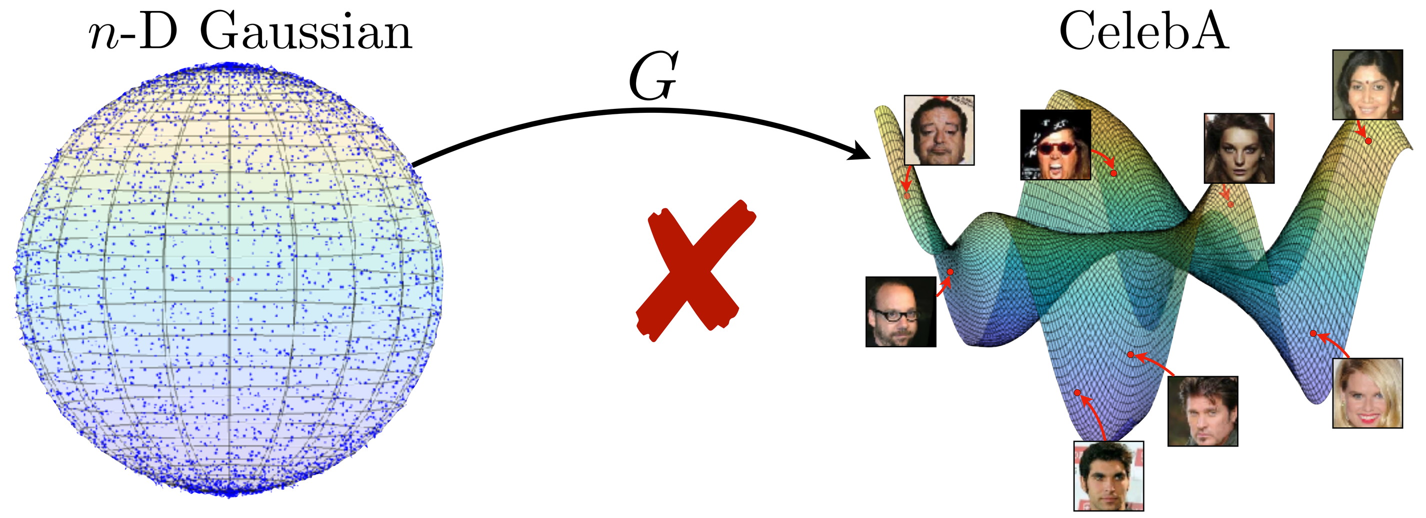

- **Standard GANs** provide high-dimensional Gaussian noise as input to the generator.
- In **Spider GANs** we propose to provide images drawn form closely-related *friendly neighborhood* datasets as input!

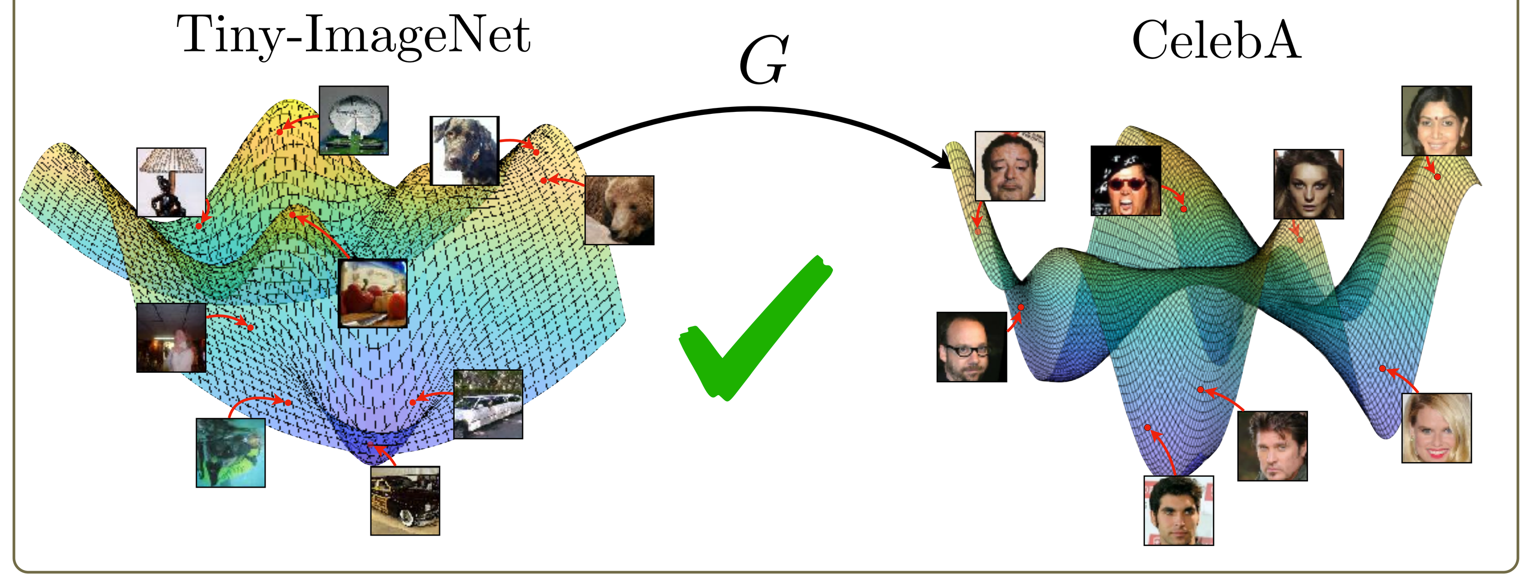

3. The Signed Inception Distance (SID)

- The polyharmonic spline interpolator or order m in n-D[2]:

$$\tilde{D}_p^*(\boldsymbol{x}) = \frac{\kappa_{m,n}}{N}\left(\sum_{\boldsymbol{g}_j \sim p_g} \psi_{m,n}(\boldsymbol{x} - \boldsymbol{g}_j) - \sum_{\boldsymbol{d}_i \sim p_d} \psi_{m,n}(\boldsymbol{x} - \boldsymbol{d}_i)\right),$$

$$\text{where } \psi_{m,n}(\boldsymbol{x}) = \begin{cases} \|\boldsymbol{x}\|^{2m-n} & \text{if } 2m-n<0 \text{ or } n \text{ is odd}, \\ \|\boldsymbol{x}\|^{2m-n}\ln(\|\boldsymbol{x}\|) & \text{if } 2m-n\geq 0 \text{ and } n \text{ is even}. \end{cases}$$

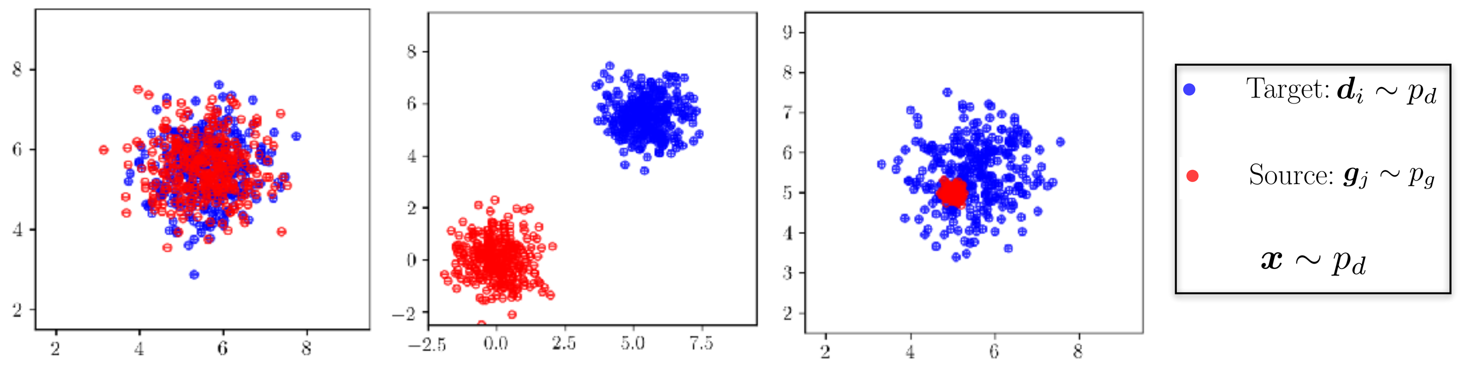

(a) $D_p^*(\boldsymbol{x}) = 0$. (b) $D_p^*(\boldsymbol{x}) > 0$. (c) $D_p^*(\boldsymbol{x}) < 0$.

- The **signed distance** ($\mathrm{SD}(p_s\|p_t)$): Average $D_p^*(\boldsymbol{x})$ over $\boldsymbol{x}$.

$$\mathrm{SD}(p_s\|p_t) = \sum_{\boldsymbol{x}\in\mathcal{U}[\mathcal{C}_r]}\left(\sum_{d_i \sim p_s} \psi_{m,n}(\boldsymbol{x} - \boldsymbol{d}_i) - \sum_{g_j \sim p_t} \psi_{m,n}(\boldsymbol{x} - \boldsymbol{g}_j)\right)$$

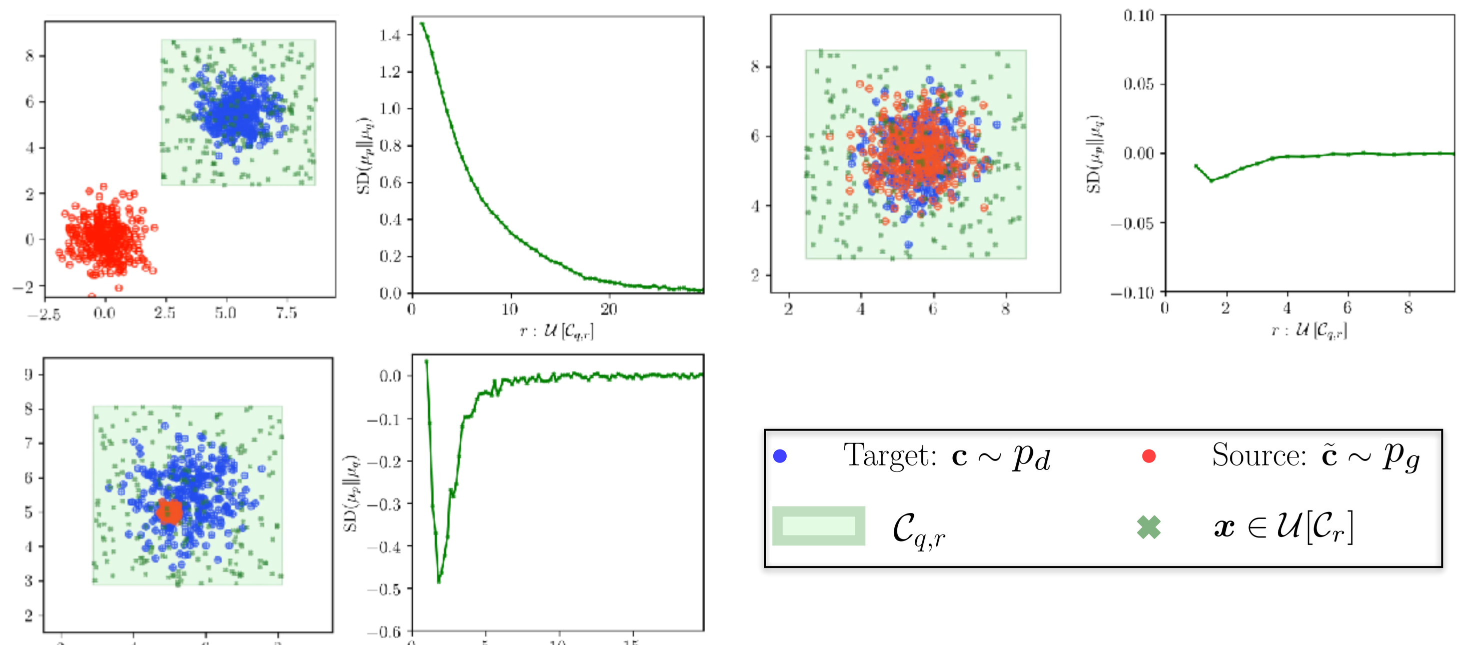

- The **signed Inception distance**, $\mathrm{SID}(p_s\|p_t)$: Compute $\mathrm{SD}(p_s\|p_t)$ over Inception embeddings.

Table: Comparing **FID** and **area under the SID curve (CSID)** in identifying friendly neighbors of target datasets in terms of the first, second and third "friendliest neighbors."

Source \ Target	FID (Source , Target) MNIST	CIFAR-10	TinyImageNet	Ukiyo-E	CSID_m (Source ‖ Target) MNIST	CIFAR-10	TinyImageNet	Ukiyo-E
MNIST	1.2491	258.246	264.250	398.280	0.1863	29.298	9.436	201.550
F-MNIST	**176.813**	188.367	197.057	387.049	**162.962**	19.051	-2.5571	191.010
SVHN	**236.707**	**168.615**	**189.133**	372.444	212.473	**34.534**	**21.668**	214.507
CIFAR-10	**259.045**	5.0724	**64.3941**	303.694	221.337	-0.1487	-7.109	198.991
TinyImageNet	264.309	**64.0312**	6.4854	**257.078**	230.916	**12.892**	0.6743	**197.447**
CelebA	360.773	303.490	**250.735**	**301.108**	**204.794**	**23.685**	**8.829**	**184.170**
Ukiyo-E	396.791	300.511	254.102	5.9137	250.226	39.793	**18.727**	0.5494
Church	350.708	**294.982**	254.991	**267.638**	**212.452**	-4.655	-23.115	**198.750**

[2]Asokan and Seelamantula, INTERPOLATE @ NeurIPS 22

4. Experiments on Spider GANs

- Visual correspondence is not necessary! Spider GANs leverage underlying structural similarity between datasets.

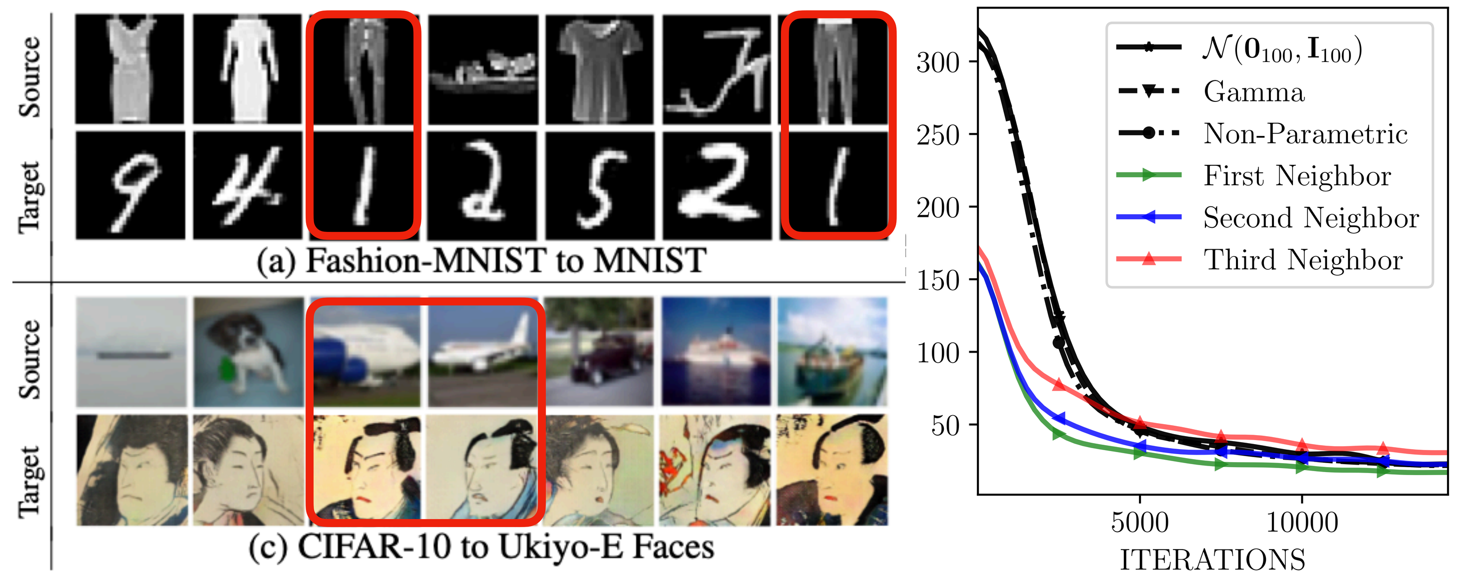

5. Novel Transfer Learning — Cascaded Spider GANs

- The output of a **Stage-one pre-trained GAN** can be used as the input to a **Stage-two Spider GAN**.

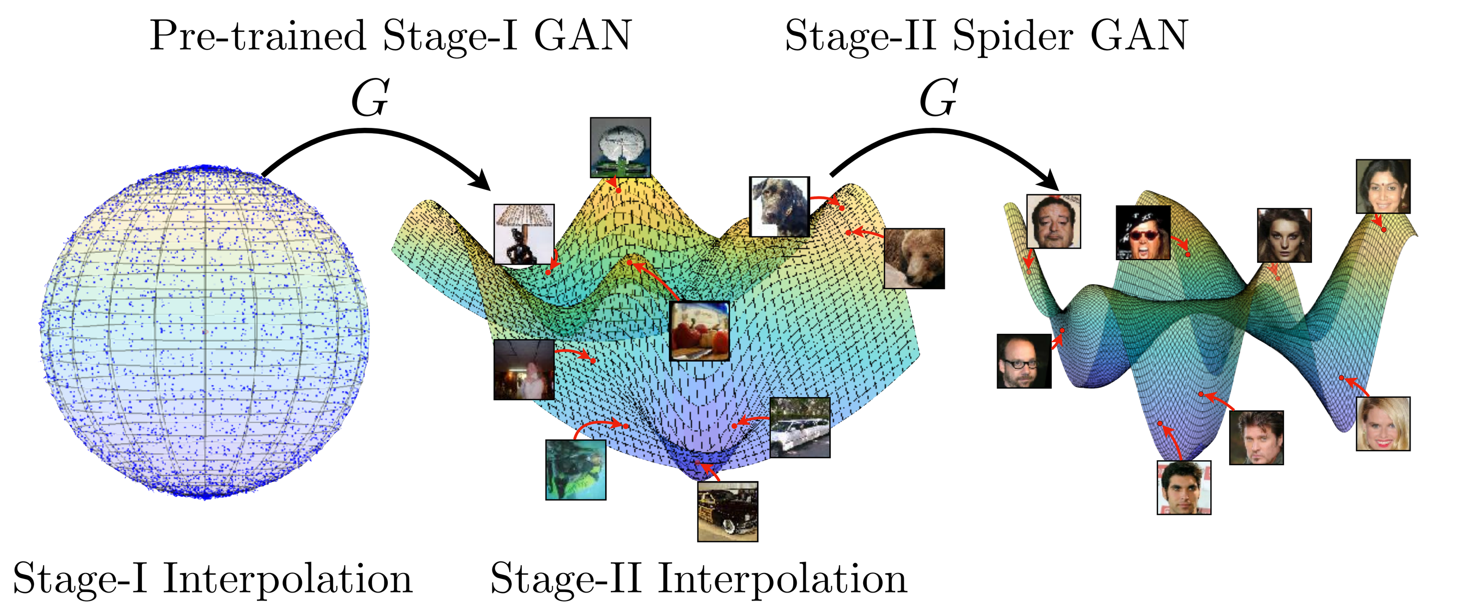

- State-of-the-art performance in one-fifth of the training time!

6. Experimental Validation on Cascaded Spider StyleGANs

Fig. Images generated by Spider StyleGAN3-T on FFHQ, with inputs drawn from StyleGAN2 pre-trained on Tiny-ImageNet.

StyleGANs: Karras et al., NeurIPS 20, CVPR 21, NeurIPS 22

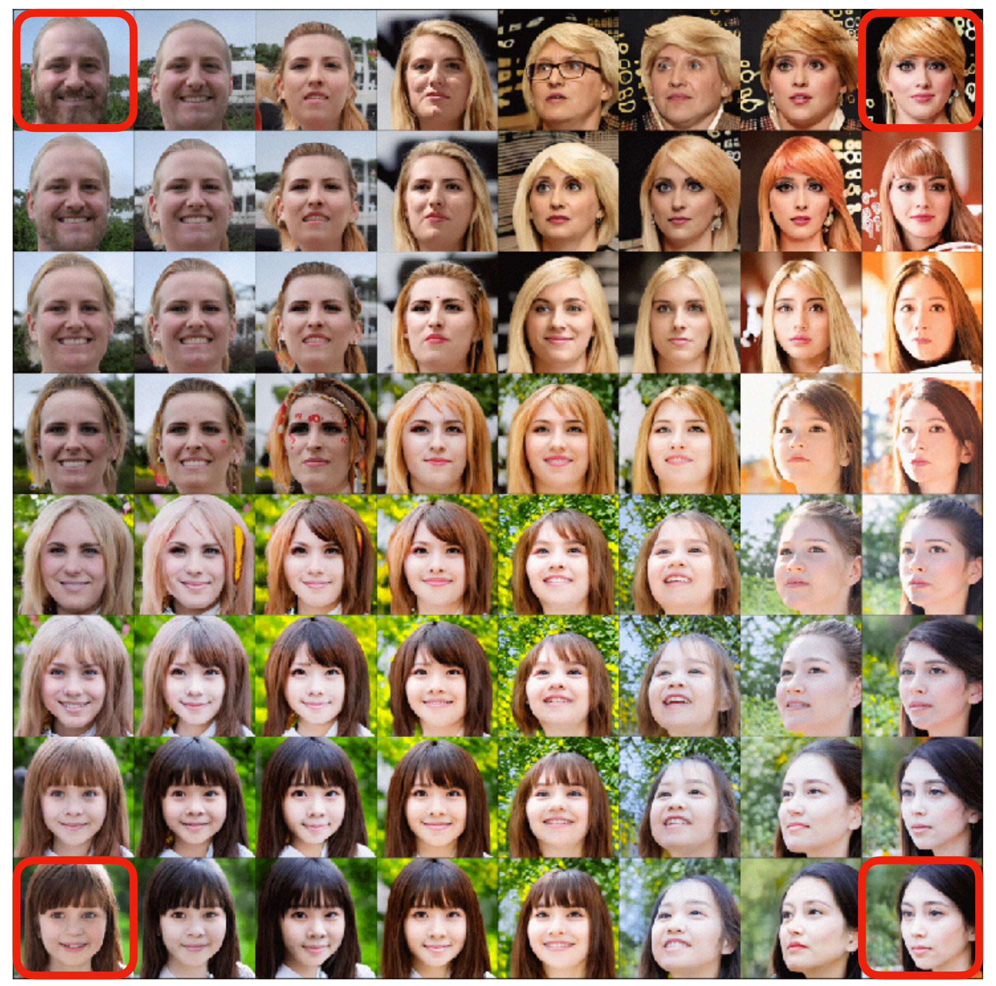

Fig. Interpolations on Stage-I Spider StyleGAN2-ADA

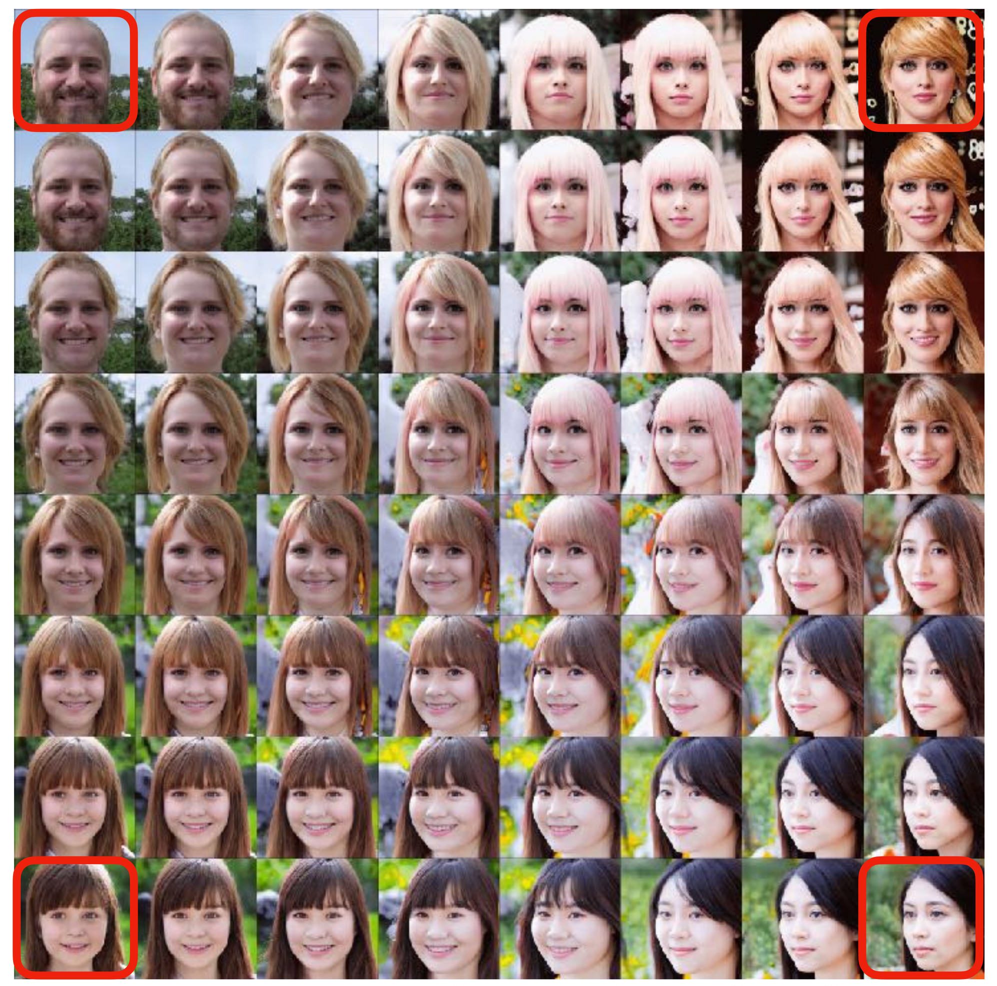

Fig. Interpolations on Stage-II Spider StyleGAN2-ADA

7. Take-home Message

- Trained the generator with inputs from closely related datasets.
- SID can identify friendly neighbors and quantify diversity.
- Spider StyleGANs achieved state-of-the-art performance in a fraction of the training time.

Acknowledgements

This work is supported by the Microsoft Research Ph.D. Fellowship, Qualcomm Innovation Fellowship and Robert Bosch Center for Cyber-Physical Systems Ph.D. Fellowship

Microsoft Qualcomm Robert Bosch Centre for CYBER PHYSICAL SYSTEMS