Non-Parametric Priors For Generative Adversarial Networks

Rajhans Singh[1], Pavan Turaga[1,2], Suren Jayasuriya[1,2],
Ravi Garg[3], Martin W. Braun[3]

[1]School of Electrical, Computer, and Energy Engineering, Arizona State University
[2]School of Arts, Media and Engineering, Arizona State University
[3]Intel Corporation

Arizona State University

Generative Adversarial Network (GAN)

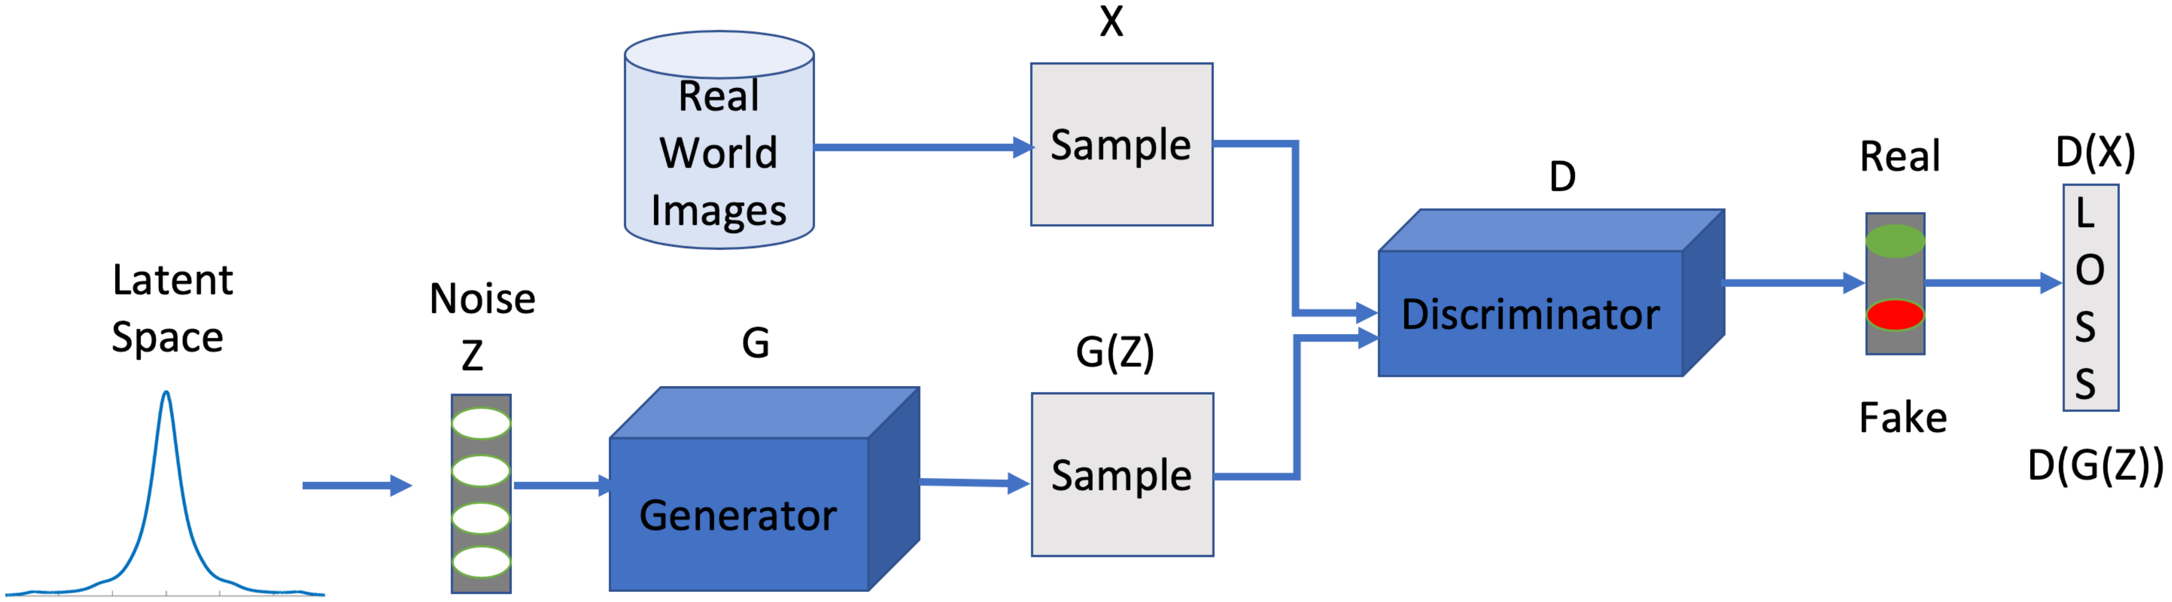

Distribution Mismatch Issue

- Simple parametric latent space distributions like Normal and Uniform suffer from distribution mismatch issue: **the prior distribution does not match the interpolated point's distribution.**
- The GANs trained on the prior points to generate realistic images, when used to interpolate between points, **often loose fidelity in image quality** due to distribution mismatch.

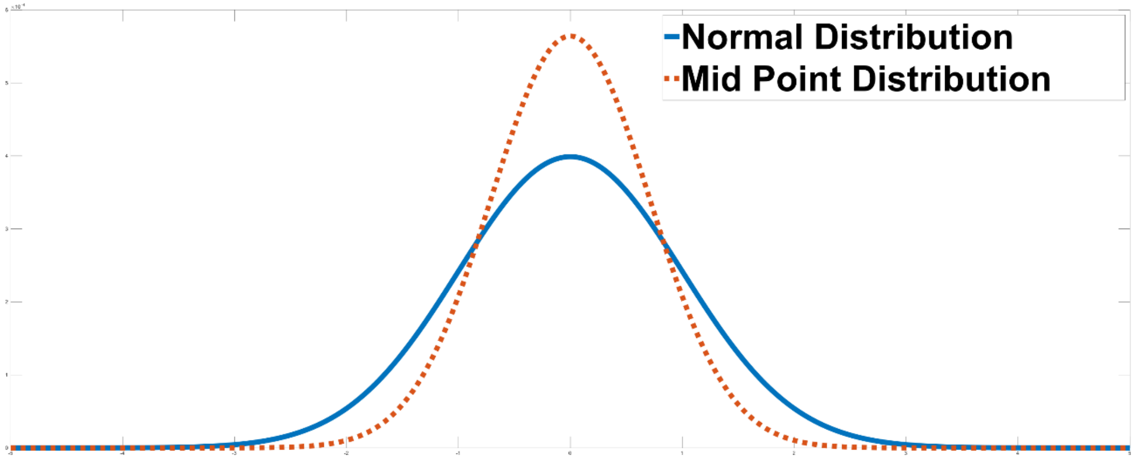

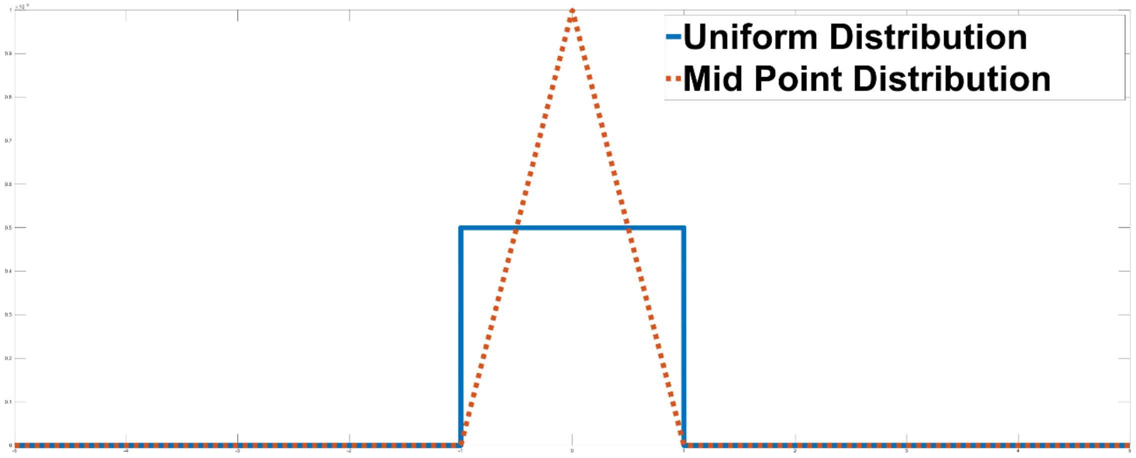

Non-Parametric Distribution

- If $x_1, x_2 \sim f_X(x)$, then the pdf of the interpolated point: $(1-\lambda)x_1 + \lambda x_2$, for some $\lambda \,\epsilon [0,1]$, is given by

$$Q(x;\lambda) = \frac{1}{|\lambda(1-\lambda)|} f_X\left(\frac{x}{\lambda}\right) * f_X\left(\frac{x}{1-\lambda}\right)$$

- The goal is to minimize the KL divergence between $P(x) = f_X(x)$ and $Q(x;\lambda)$.
- $P(x)$ is restricted a compact domain, i.e. [0,1] and discretize into 2^{10} bins to obtain a tractable solution.
- Variance constraint is added to avoid delta function.

$$\min_P f(P||Q)$$

$$s.t. \sum_{i=1}^{n} p_i = 1,\ \frac{1}{n}\left(\sum_{i=1}^{n} i^2 p_i - \left(\sum_{i=1}^{n} i p_i\right)^2\right) \geq \varepsilon,\ \ p_i \geq 0$$

Non-Parametric Prior

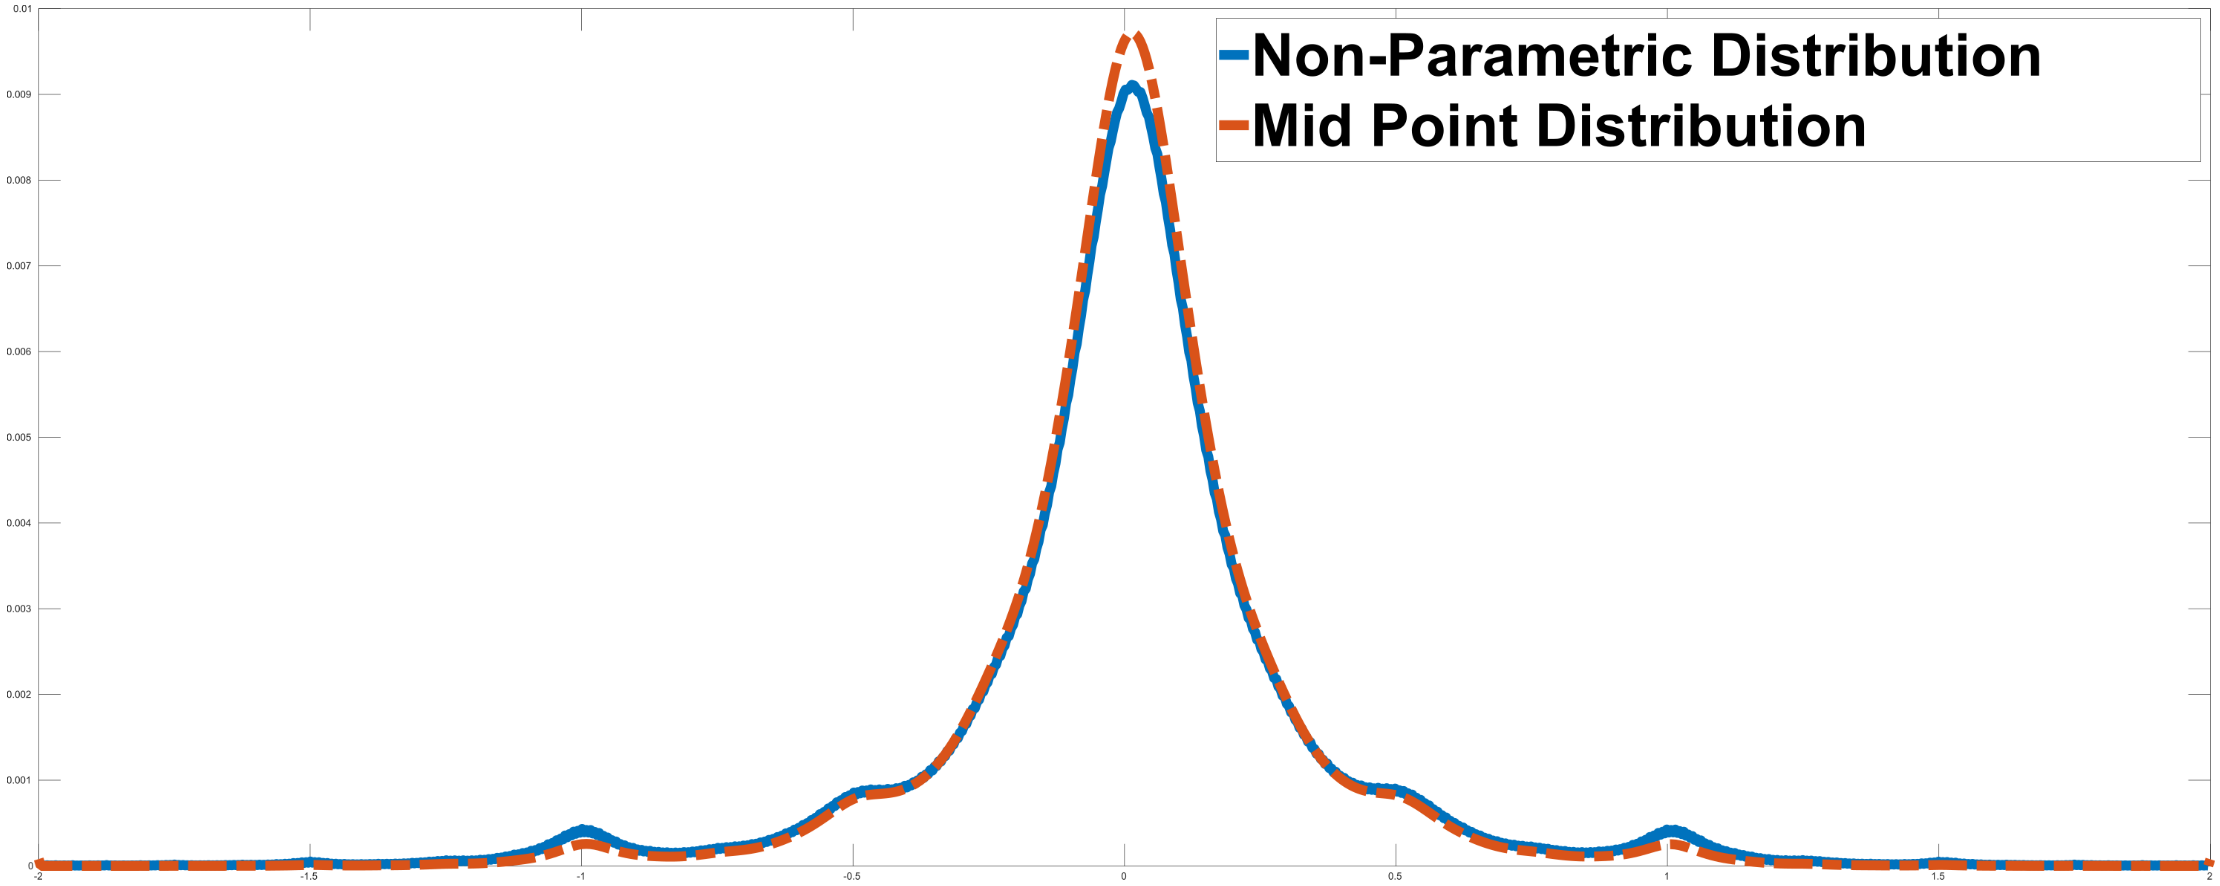

Results

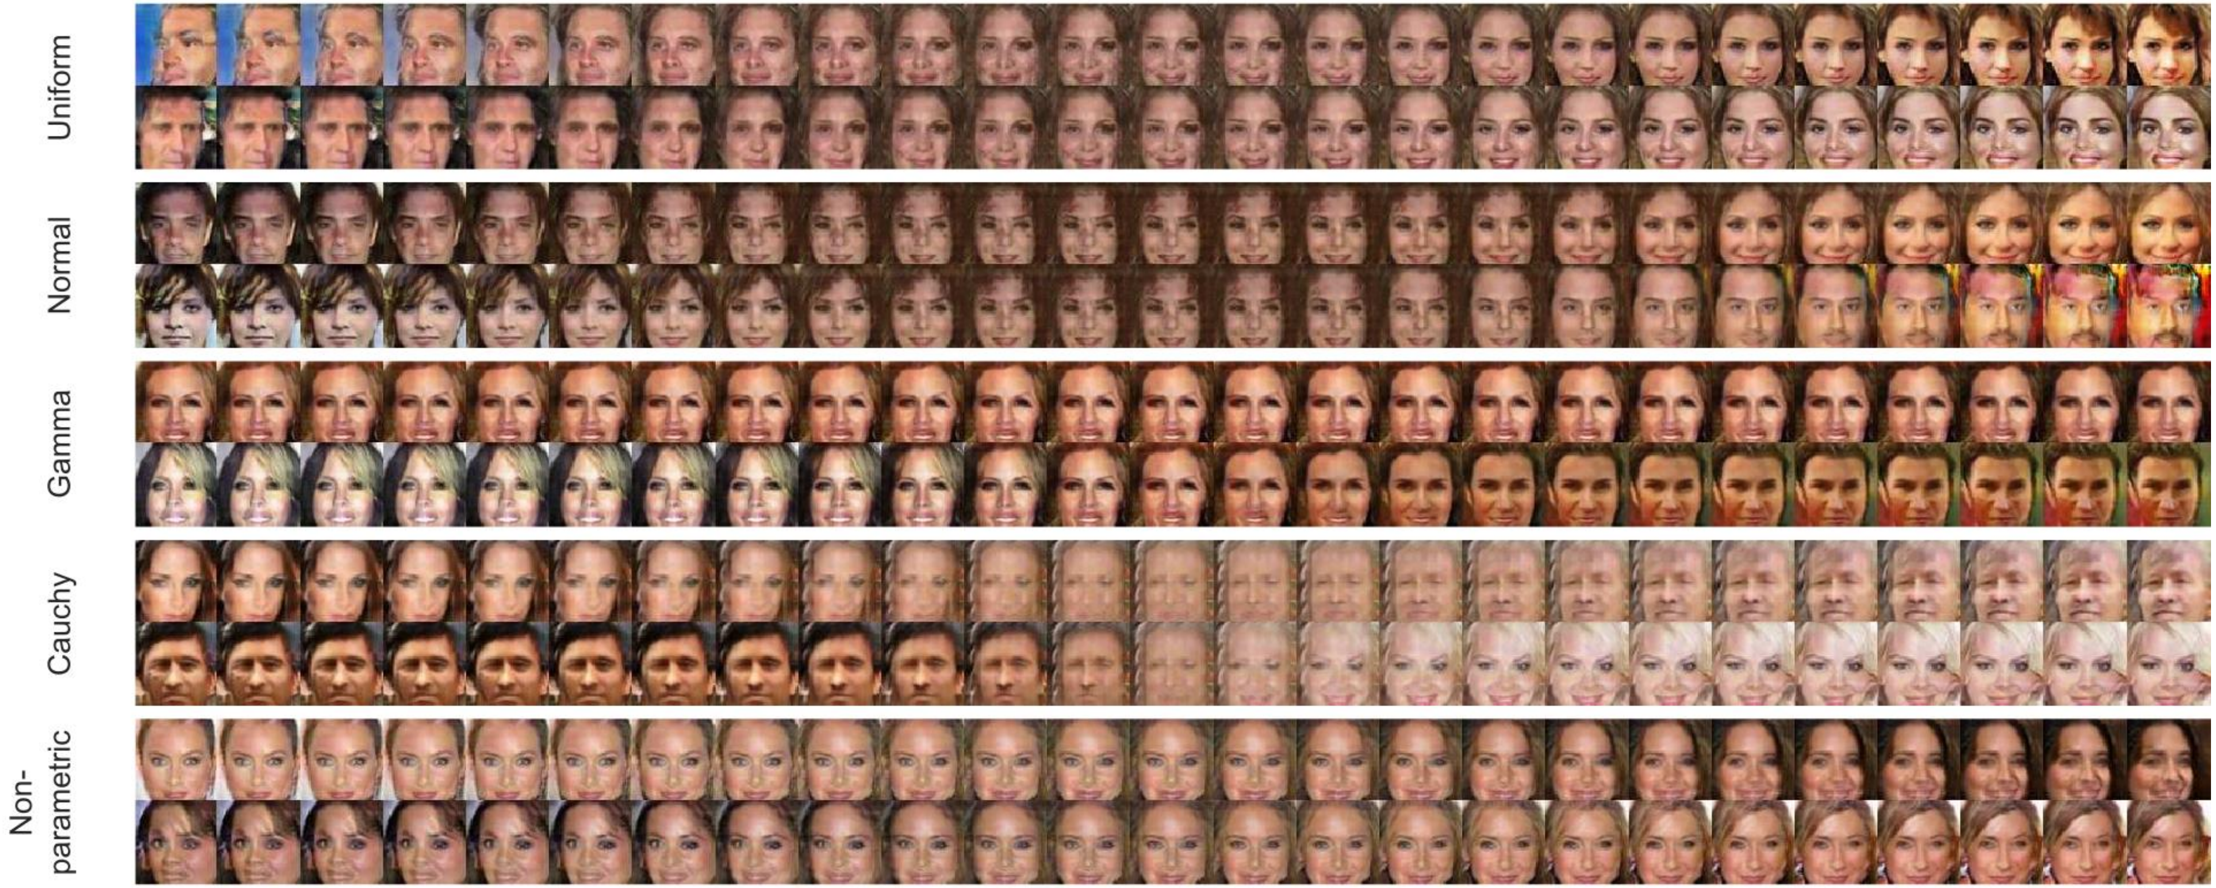

Interpolation (left to right) through the origin on CelebA dataset

Gamma Prior: [Kilcher et al., 2018]
Cauchy Prior: [Lesniak ´ et al., 2019]

Quantitative Results

Distribution	CelebA				LSUN Bedroom			
	Inception Score		FID Score		Inception Score		FID Score	
	Prior	Mid-Point	Prior	Mid-Point	Prior	Mid-Point	Prior	Mid-Point
Uniform	1.843	1.369	24.055	40.371	2.969	2.649	42.998	76.412
Normal	1.805	1.371	26.173	42.136	2.812	2.591	64.682	108.49
Gamma	1.776	1.618	29.912	28.608	2.930	2.808	162.44	161.37
Cauchy	1.625	1.628	59.601	60.128	**3.148**	**3.149**	97.057	97.109
Non-parametric	**1.933**	**1.681**	**17.735**	**19.115**	3.028	2.769	**27.857**	**31.472**

Conclusion

- We derive a non-parametric approach to search for a prior which can address the distribution mismatch problem.
- The proposed prior distribution **provides better qualitative and quantitative results** as compared to the standard priors such as Normal and Uniform distributions.
- The FID score and the Inception score (IS) using the proposed prior are either **the best, or very close to the best** on all the four tested datasets.
- For future work it would be interesting to extend this approach to extrapolation problems, or impose other interesting statistical or physically-motivated constraints over latent spaces.