Density Approximation in Deep Generative Models with Kernel Transfer Operators

Zhichun Huang [1] Rudra Chakraborty [2] Vikas Singh [3]

[1]Carnegie Mellon University [2]University of California, Berkeley [3]University of Wisconsin-Madison

Motivation

Our work is based on an observation of decoder-based generative models, where a deterministic operator is learned to map a simple, known distribution $p_{\mathcal{Z}}$ to the target data distribution $p_{\mathcal{X}}$.

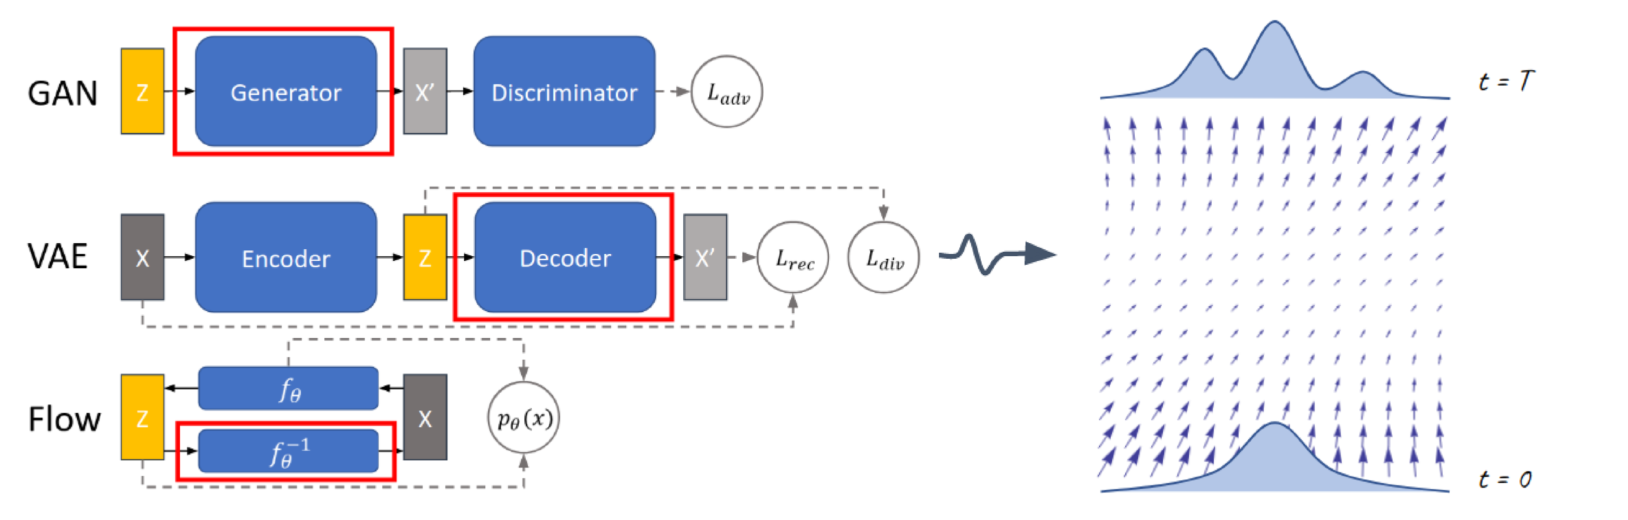

Figure 1:Decoder based generative models can be viewed as discrete-time evaluations of a dynamical system

As observed in [1], decoders → discrete-time dynamical systems, which allows us to link generative models to well studied objects in dynamical systems, including the transfer operator.

Transfer Operator in Dynamical Systems

For a non-singular deterministic mapping f on a measure space $(\mathbb{X}, \mathfrak{B}, \mu)$ such that $f(\mathcal{Z}) \sim p_{\mathcal{X}}$, the transfer operator (or Perron-Frobenius operator) $\mathcal{P} : \mathbb{M}^1_+(\mathbb{X}) \to \mathbb{M}^1_+(\mathbb{X})$ is a linear operator in the space of probability densities defined as

$$\mathcal{P} \in \left\{ \int_\Lambda (\mathcal{P} p_{\mathcal{Z}}) d\mu = \int_{f^{-1}(\Lambda)} p_{\mathcal{Z}} d\mu, \quad \forall \Lambda \in \mathfrak{B} \right\}$$

With this definition, we have $\mathcal{P} p_{\mathcal{Z}} = p_{\mathcal{X}}$. Once we obtain $\mathcal{P}$, we can use it to conveniently transfer $p_{\mathcal{Z}}$ to $p_{\mathcal{X}}$. Q: *Can we learn $\mathcal{P}$ directly for generative models?* A: Can be challenging, but learning in Reproducing Kenrel Hilbert Space (RKHS) helps.

- Requires rich basis functions → RKHS spanned by infinite bases
- Samples rather than densities → easy to compute empirical kernel mean embedding
- Cannot apply directly on samples → reproducing property of RKHS

Kernel Perron-Frobenius Operator (kPF)

[3] proposes an operator in RKHS that transfers the densities in the mean embedded form. Define $\phi(z) = k(\cdot, z), \psi(x) = l(\cdot, x)$ as the feature mappings of RKHS $\mathcal{H}$ and $\mathcal{G}$. The kernel mean embeddings (KME) of $p_{\mathcal{Z}}$ and $p_{\mathcal{X}}$ are given by

$$\mu_{\mathcal{Z}} = \mathbb{E}_{z \sim p_{\mathcal{Z}}}[\phi(z)], \quad \mu_{\mathcal{X}} = \mathbb{E}_{x \sim p_{\mathcal{X}}}[\psi(x)]$$

Note that KME is injective for characteristic kernels. The kernel Perron-Frobenius operator $\mathcal{P}_{\mathcal{E}}$ is defined using the (uncentered) covariance/cross-covariance operators.

$$\mathcal{P}_{\mathcal{E}} = \mathcal{C}_{\mathcal{Z}\mathcal{X}} \mathcal{C}^{-1}_{\mathcal{Z}\mathcal{Z}}, \quad \text{where } \mathcal{C}_{\mathcal{X}\mathcal{Z}} = \mathbb{E}_{(x,z) \sim p_{\mathcal{Z},\mathcal{X}}}[\psi(x) \otimes \phi(z)], \mathcal{C}_{\mathcal{Z}\mathcal{Z}} = \mathbb{E}_{z \sim p_{\mathcal{Z}}}[\phi(z) \otimes \phi(z)]$$

We have $\mu_{\mathcal{X}} = \mathcal{P}_{\mathcal{E}} \mu_{\mathcal{Z}}$ under certain conditions.

Main Result

In the context of generative modeling, we propose to use the empirical form of kPF to transfer $p_{\mathcal{Z}}$ to $p_{\mathcal{X}}$. Let $\Phi = [\phi(z_i)]_{i \in [n]}$ and $\Psi = [\psi(x_i)]_{i \in [n]}$. The empirical kPF is given by

$$\hat{\mathcal{P}}_{\mathcal{E}} = \hat{\mathcal{C}}_{\mathcal{Z}\mathcal{X}} (\hat{\mathcal{C}}_{\mathcal{Z}\mathcal{Z}} + \lambda I)^{-1} = \Psi(\Phi^\top \Phi + \lambda I)^{-1} \Phi$$

Suppose we have an exact preimage map ψ^{-1}, a generated sample x^* can be constructed as $x^* = \psi^{-1}(\Psi^*) = \psi^{-1}(\hat{\mathcal{P}}_{\mathcal{E}} k(\cdot, z^*))$, where $z^* \sim p_{\mathcal{Z}}$ and $\Psi^* = \hat{\mathcal{P}}_{\mathcal{E}} k(\cdot, z^*)$ is called a *transferred sample* in RKHS. We can show that $\mu_{x^*} = \mathbb{E}[\psi(x^*)] = \mu_{\mathcal{X}}$, indicating a match in distribution

Image Generation

It can be hard to generative images due to the inaccurate preimages in high dimensional space. However, data often lie on low-dimensional manifolds. Following [2], we generate high-dimensional data by:

1. Train a regularized autoencoder (E, D) to learn a mapping to a low-dimensional latent space. The regularization is used to encourage *smoothness* of the latent space
2. Construct kPF is using *i.i.d.* samples of the known distribution $Z = \{z_i\}_{i \in [n]} \sim p^n_{\mathcal{Z}}$ and the latent representations of the training data $H = \{E(x_i)\}_{i \in [n]}$
3. Compute the approximate preimages of the transferred samples $h^* \sim \psi^{-1}(\Psi^*)$ and output the decoded image $x^* = D(h^*)$

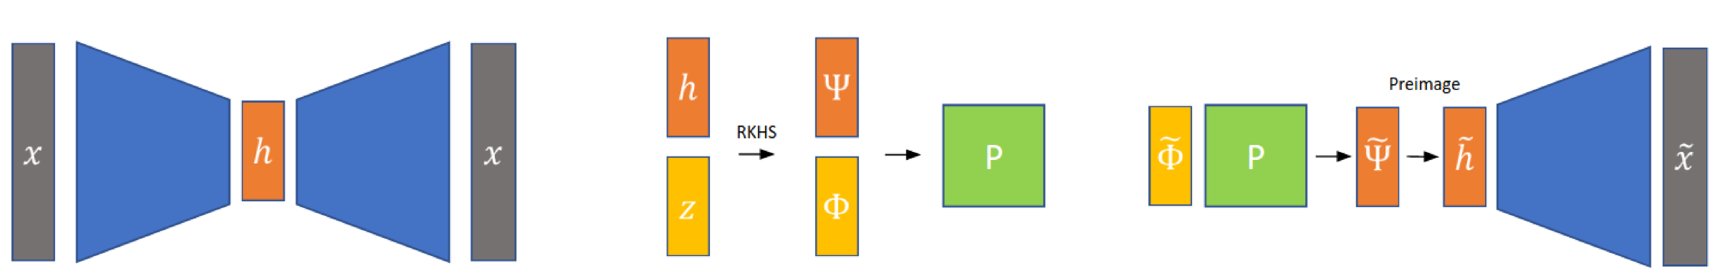

Figure 2:Image generation procedure for kPF

Experimental Results

We evaluated kPF on density approximation with toy distributions and unconditional generation with popular CV datasets. **Result: Better quality & more efficent than deep methods**

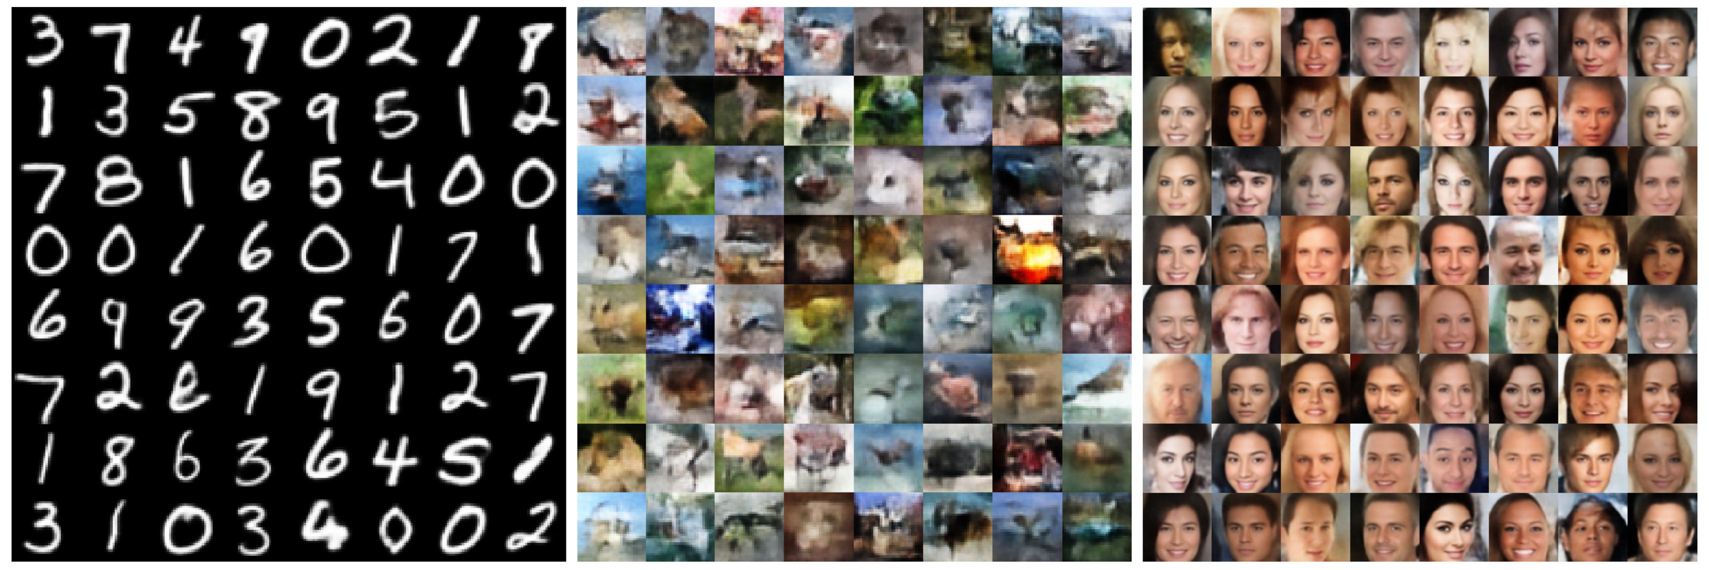

Figure 3:Unconditional image generations

	Glow‡	CAGlow‡	Vanilla VAE	WAE†	2-stage VAE	SRAE$_{Glow}$	SRAE$_{GMM}$	SRAE$_{RBF-kPF}$ (ours)	SRAE$_{NTK-kPF}$ (ours)
MNIST	25.8	26.3	**13.7**	20.4	18.3	23.7	16.7	21.7	21.5
CIFAR-10	-	-	111.0	117.4	110.3	110.7	79.2	**77.9**	**77.5**
CelebA	103.7	104.9	52.1	53.7	44.7	59.8	42.0	**41.9**	**41.0**

Table 1:Comparative FID values. SRAE indicates an autoencoder with spherical latent space and spectral regularization following [2]. Results reported from ‡: [4]. †: [2].

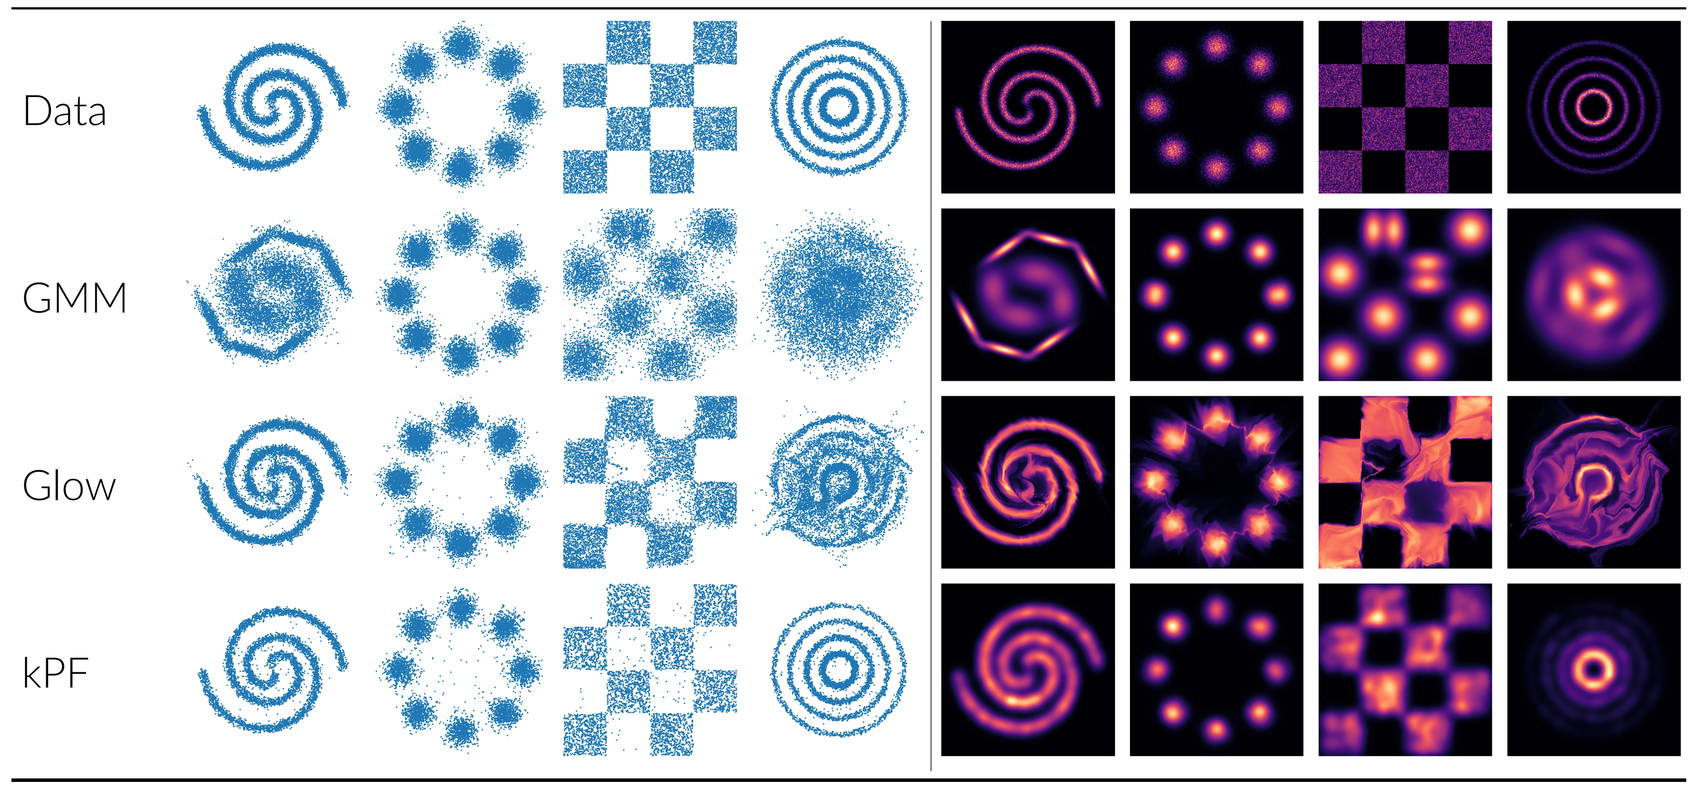

Figure 4:Density approximation on toy distributions

Small data setting: kPF is best suited for small datasets due to the super-quadratic cost to compute the kernel inverse. We evaluated kPF on generating with few examples (<1% of the CelebA dataset), and observed that kPF outperforms other deep methods by a large margin. Additionally, we compared kPF to VAE on generating high-res brain MR images with <500 examples, and kPF yields sharper and statistically consistent samples.

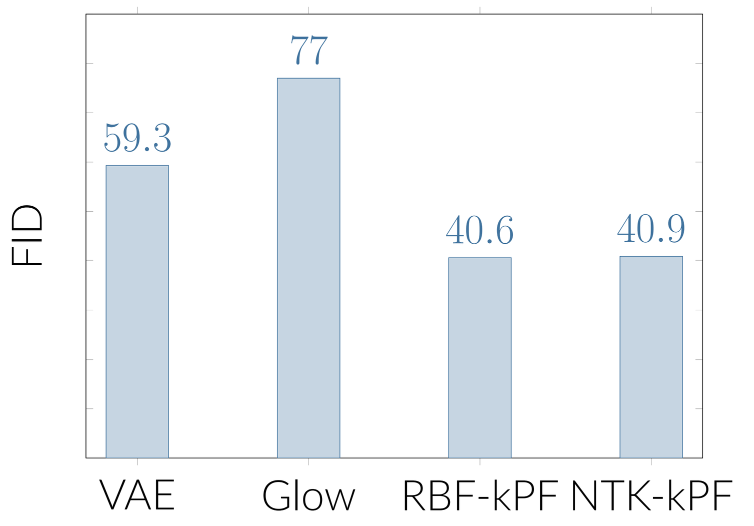

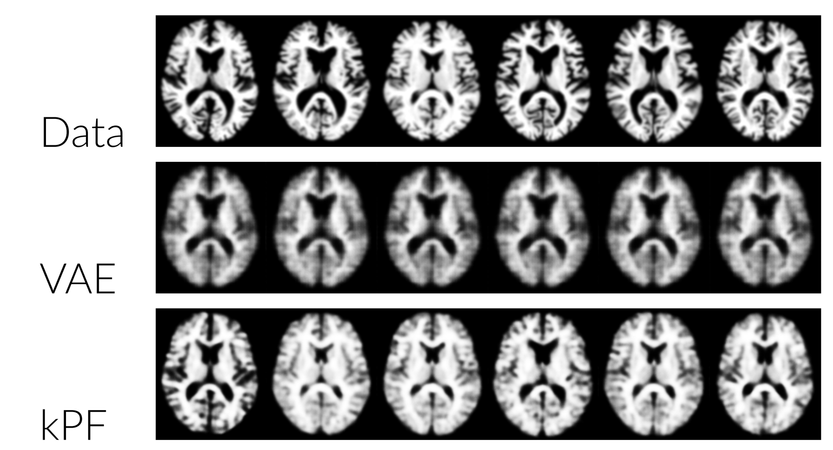

Figure 5:High-res brain MR image generations

References

[1] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 6571--6583. Curran Associates, Inc., 2018.

[2] Partha Ghosh, Mehdi S. M. Sajjadi, Antonio Vergari, Michael Black, and Bernhard Scholkopf. From variational to deterministic autoencoders. In *International Conference on Learning Representations*, 2020.

[3] Stefan Klus, Ingmar Schuster, and Krikamol Muandet. Eigendecompositions of transfer operators in reproducing kernel hilbert spaces. *Journal of Nonlinear Science*, 30(1):283--315, 2020.

[4] Rui Liu, Yu Liu, Xinyu Gong, Xiaogang Wang, and Hongsheng Li. Conditional adversarial generative flow for controllable image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.