

---

# Variational Mixture of HyperGenerators for Learning Distributions over Functions

---


Ignacio Peis

[ipeaz@dtu.dk](mailto:ipeaz@dtu.dk)

Technical University of Denmark




**ICML**  
International Conference  
On Machine Learning



Andaluz.IA

*Honolulu, Hawaii  
July 2023*



# Collaborators



Batuhan Koyuncu<sup>1</sup>



Pablo Sánchez<sup>1,2,3</sup>



Ignacio Peis<sup>4</sup>



Pablo M. Olmos<sup>5</sup>




Isabel Valera<sup>1,6</sup>

- <sup>1</sup> University of Saarland
- <sup>2</sup> Max Planck Institute for Intelligent Systems
- <sup>3</sup> Sony AI
- <sup>4</sup> Technical University of Denmark (DTU)
- <sup>5</sup> Universidad Carlos III de Madrid
- <sup>6</sup> Max Planck Institute for Software Systems

[Paper]



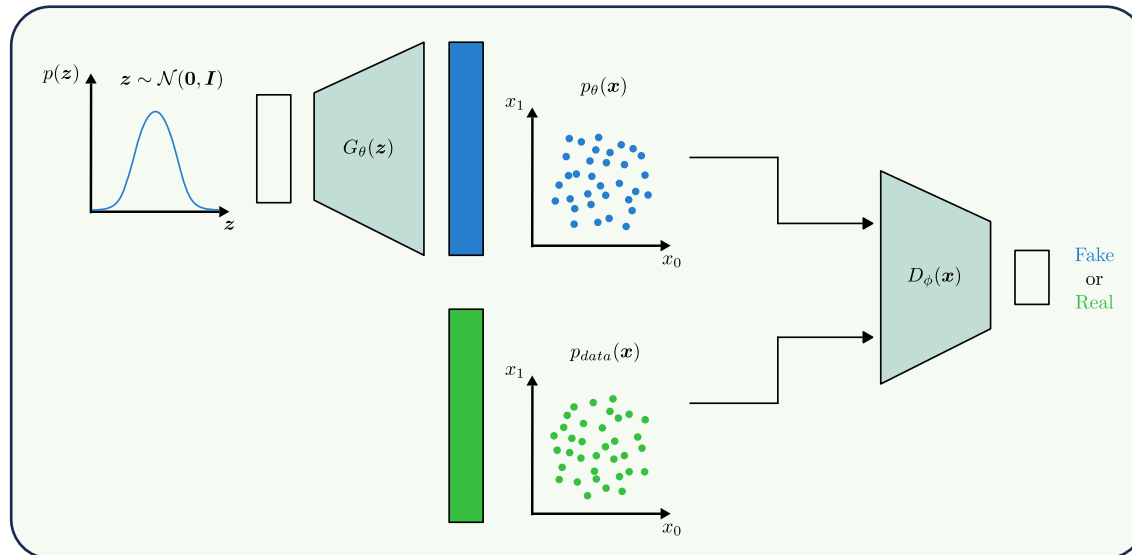
[Code]



# Deep Generative Models

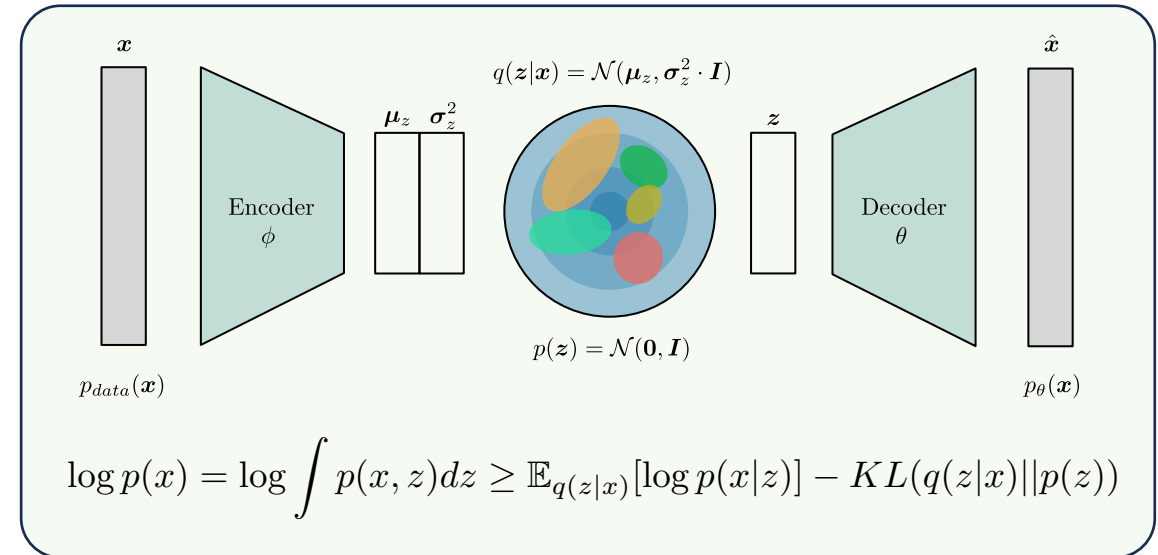
- Learning **probability distributions** on data using Deep Neural Networks.

Implicitly



Generative Adversarial Networks (GANs [1])


Explicitly

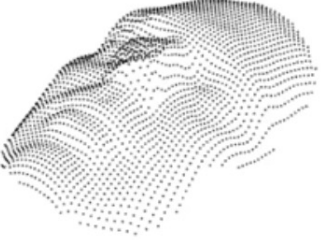


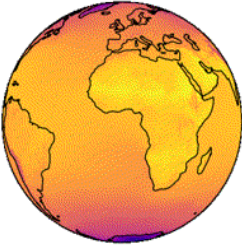
Variational Autoencoders (VAEs [2])  
 Denoising Diffusion Probabilistic Models (DDPMs [3])  
 Score-based models [4]  
 Energy-based models [5]

# Discretization of data

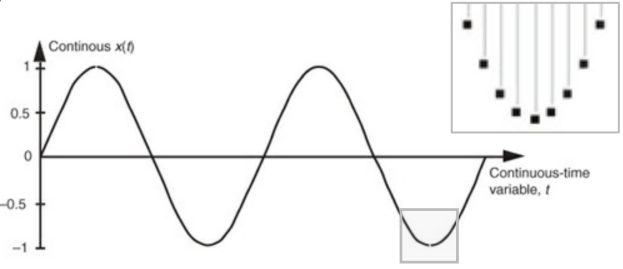
- We typically deal with discretized versions of data that are continuous in nature.


2D Images 

3D shapes 

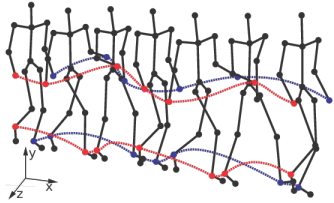
Polar data 

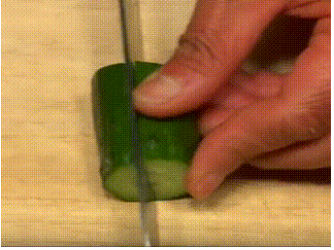
Spatial


Signals 

Audios 

Temporal

Motion sequences 

Videos 

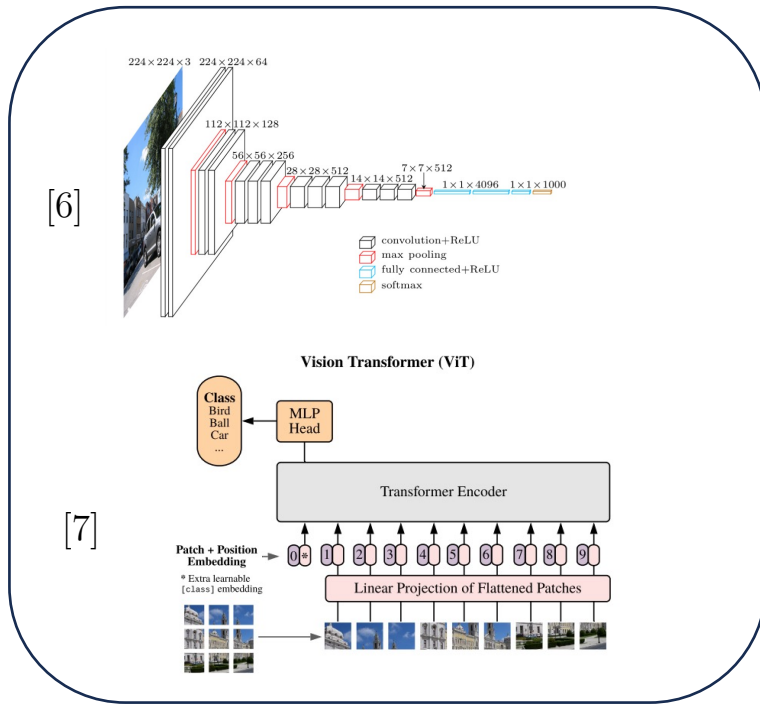


Spatio-temporal

# NNs to exploit discretized data


- DNNs are tailored to the data nature.

CNNs, Vision Transformers



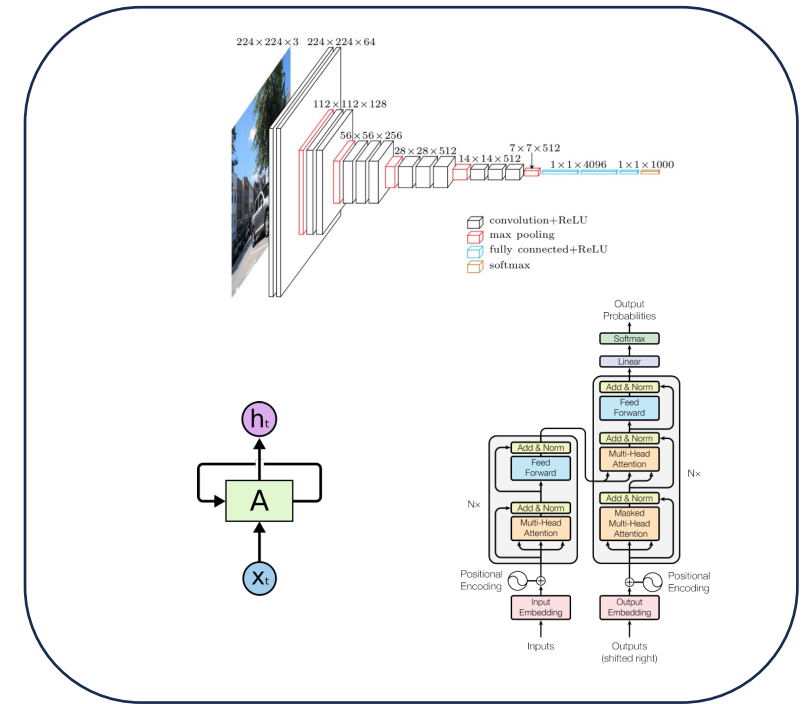
Spatial

RNNs, Transformers



Temporal

CNNs, RNNs, Transformers




Spatio-temporal

[6] (Simonyan et al., 2014) [7] (Dosovitskiy et al., 2020) [8] (Hochreiter et al., 1997) [9] (Vaswani et al., 2017)



# Real data is continuous in nature

- What if we approximate the underlying **continuous functions**?

$$f : \mathbb{R}^3 \rightarrow \{0, 1\}, f(x_1, x_2, x_3) = p$$



$$f : \mathbb{R}^3 \rightarrow \mathbb{R}^3, f(x_1, x_2, t) = (r, g, b)$$




$$f : \mathbb{R}^2 \rightarrow \mathbb{R}^3, f(x_1, x_2) = (r, g, b)$$


$$f : \mathbb{R}^2 \rightarrow \mathbb{R}, f(\varphi, \lambda) = T$$

# Real data is continuous in nature

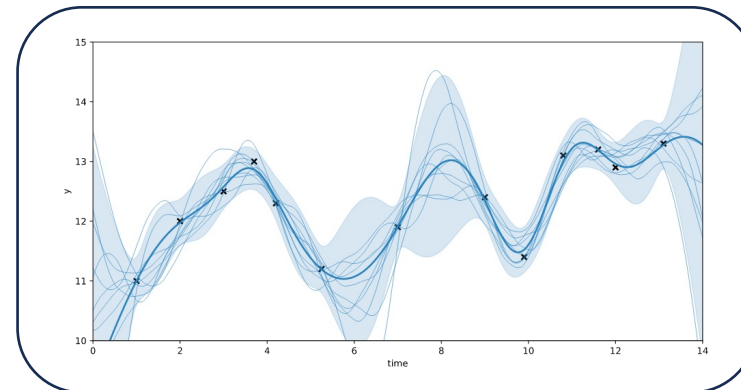
- Learning the distribution of a function allows for **naturally handling**:




Inpainting



Outpainting



Conditional generation

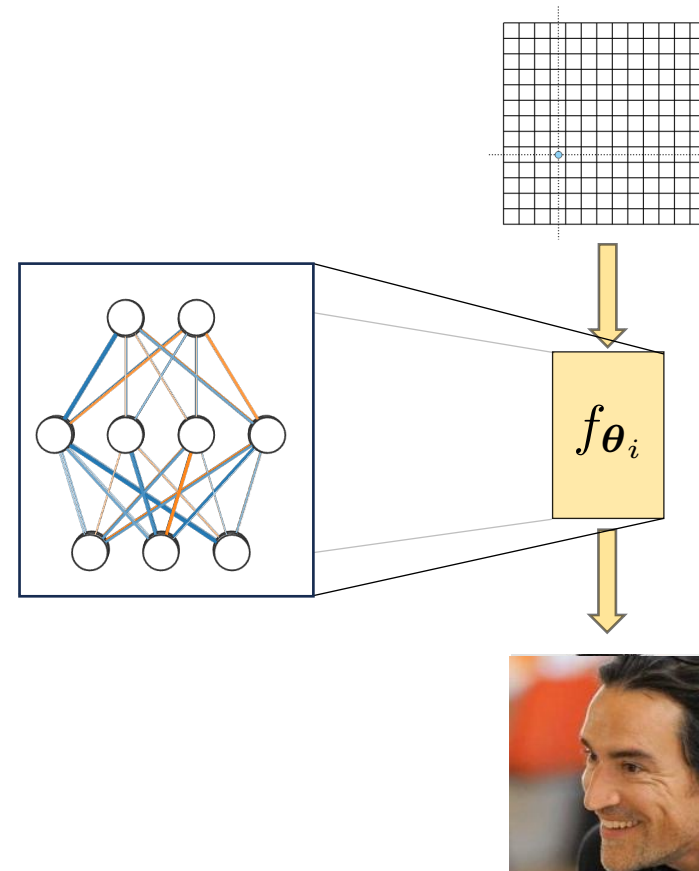
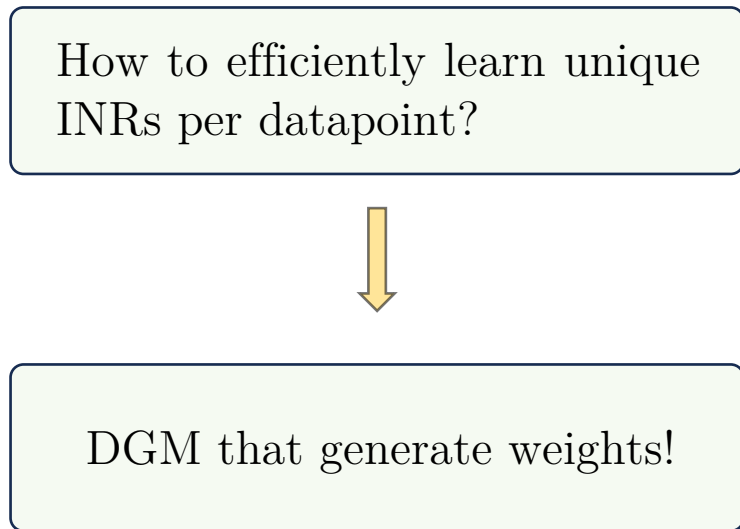


Super-resolution

- We can use the same neural architecture independently of the data nature.
- Information to store will be independent of the data size.

# Implicit Neural Representations

- INRs [10-12] can approximate these functions.



$$\mathbf{X}^{(i)} = \left\{ \mathbf{x}_d^{(i)} \right\}_{d=1}^D$$

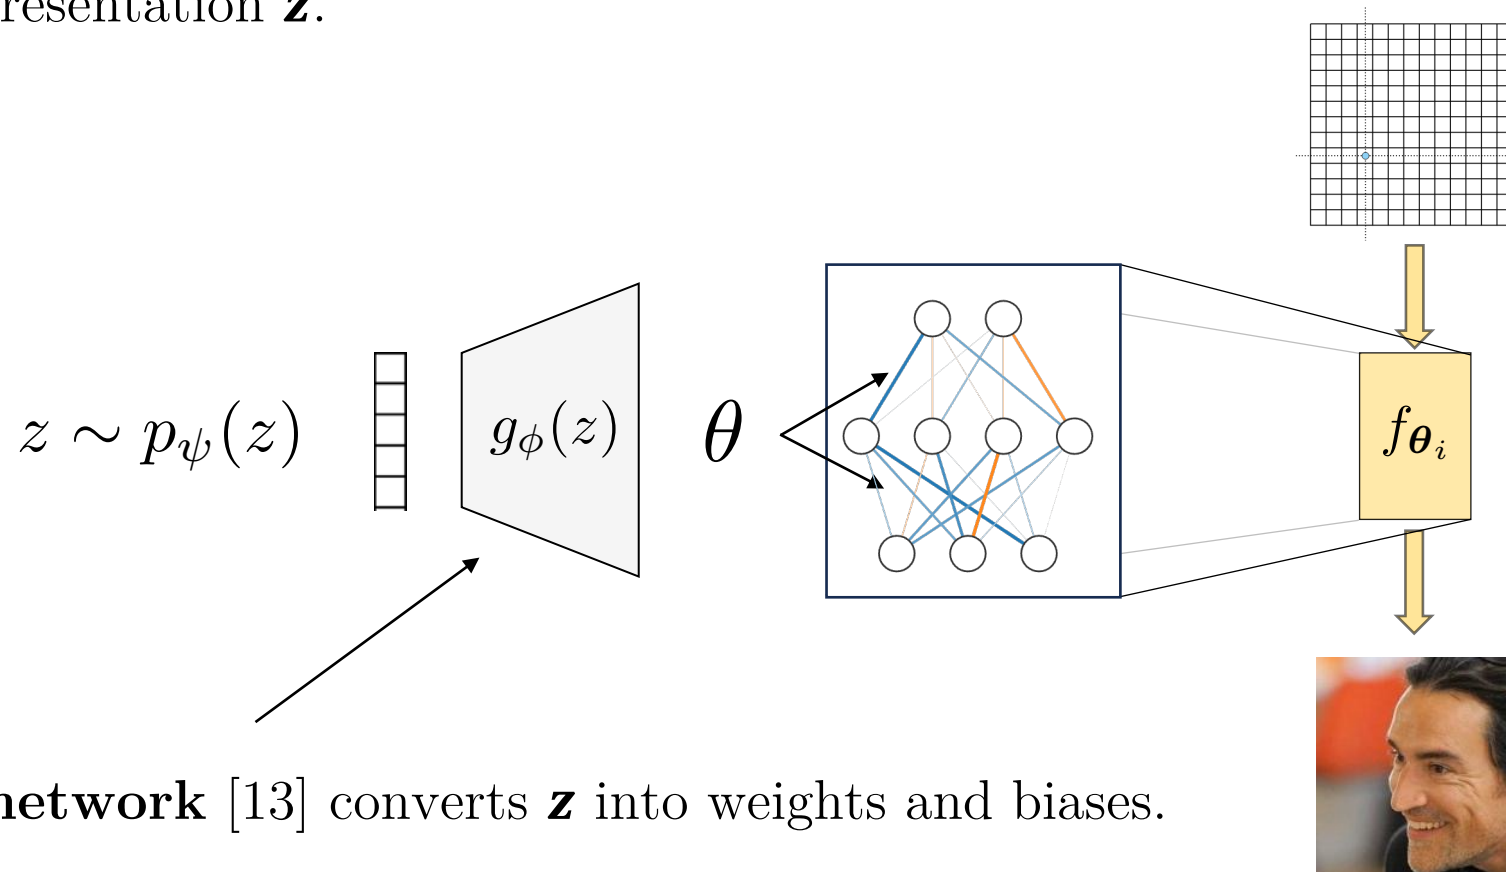
$$\mathbf{Y}^{(i)} = \left\{ \mathbf{y}_d^{(i)} \right\}_{d=1}^D$$

[10] (Sitzmann et al., 2020)   [11] (Mescheder et al., 2019)   [12] (Stanley et al., 2007)



# Our proposed method

- Every set of weights and biases,  $\theta_i$ , comes from a reduced latent representation  $\mathbf{z}$ .



$$\mathbf{X}^{(i)} = \left\{ \mathbf{x}_d^{(i)} \right\}_{d=1}^D$$

- A hypernetwork [13] converts  $\mathbf{z}$  into weights and biases.


$$\mathbf{Y}^{(i)} = \left\{ \mathbf{y}_d^{(i)} \right\}_{d=1}^D$$

[13] (Ha et al., 2017)

# Our proposed method

- To learn the parameters of our model, we opt by using **Amortized Variational Inference**, and optimize the following ELBO.

$$\max_{\phi, \psi, \gamma} \mathcal{L}(\phi, \psi, \gamma; \mathbf{Y}, \mathbf{X}) = \max_{\phi, \gamma} \mathbb{E}_{q_{\gamma}(z|\mathbf{Y}, \mathbf{X})} [\log p_{\theta}(\mathbf{Y}|\mathbf{X}, z)] - D_{KL}(q_{\gamma}(z|\mathbf{Y}, \mathbf{X}) || p_{\psi}(z))$$




[14] (Wu et al., 2019)

# Our proposed method

- We incorporate a **Mixture of HyperGenerators** for increased flexibility.

$$\mathcal{L}(\mathbf{Y}, \mathbf{X}; \psi, \phi, \gamma) = \sum_{d=1}^D \mathbb{E}_{q_{\gamma_z}(z|\mathbf{Y}, \mathbf{X})} \left[ \sum_{k=1}^K \log p_{\theta_k}(\mathbf{y}_d | \mathbf{x}_d) \cdot \pi_{dk} \right] - D_{KL}(q_{\gamma_z}(z | \mathbf{X}, \mathbf{Y}) \| p_{\psi_z}(z)) - D_{KL}(q_{\gamma_c}(\mathbf{C} | z, \mathbf{X}, \mathbf{Y}) \| p_{\psi_c}(\mathbf{C} | z, \mathbf{X}))$$

Reconstruction  
 KL of the continuous latent variable  
 KL of the discrete latent variable



# Our proposed method



- We incorporate a **Mixture of HyperGenerators** for increased flexibility.



Image Reconstruction with Mixture of HyperGenerators

# Our proposed method


- To alleviate the **holes** problem, similarly like Latent Diffusion models [15], we learn the prior as a planar Flow [16].



$$z^{(i)} \sim p_\psi(z)$$

# Results

- We achieve **comparable sampling quality** and diversity wrt baselines.




Model	CELEBA HQ			SHAPES3D		
	↓ FID	↑ Precision	↑ Recall	↓ FID	↑ Precision	↑ Recall
[17] GASP (Dupont et al., 2022b)	<b>14.01 ± 0.18</b>	<b>0.81 ± 0.0</b>	<b>0.43 ± 0.01</b>	118.66 ± 0.64	0.01 ± 0.0	0.16 ± 0.01
[18] Funcata (Dupont et al., 2022a)	40.40	-	-	57.81 ± 0.15	0.06 ± 0.0	0.13 ± 0.0
VaMoH	66.27 ± 0.18	0.65 ± 0.0	0.0 ± 0.0	<b>56.25 ± 0.57</b>	<b>0.08 ± 0.0</b>	<b>0.64 ± 0.01</b>


[17] (Dupont et al., 2022a) [18] (Dupont et al., 2022b)

# Results

- We naturally generate samples with **any desired resolution**.



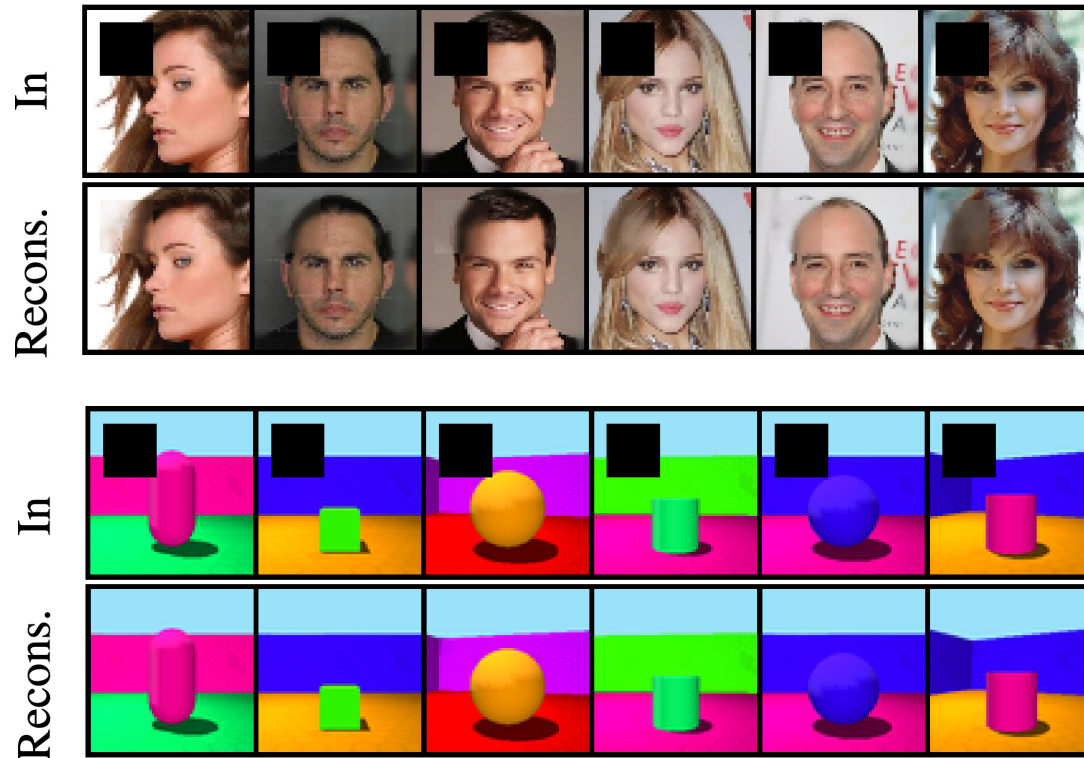
(b) POLYMNIST



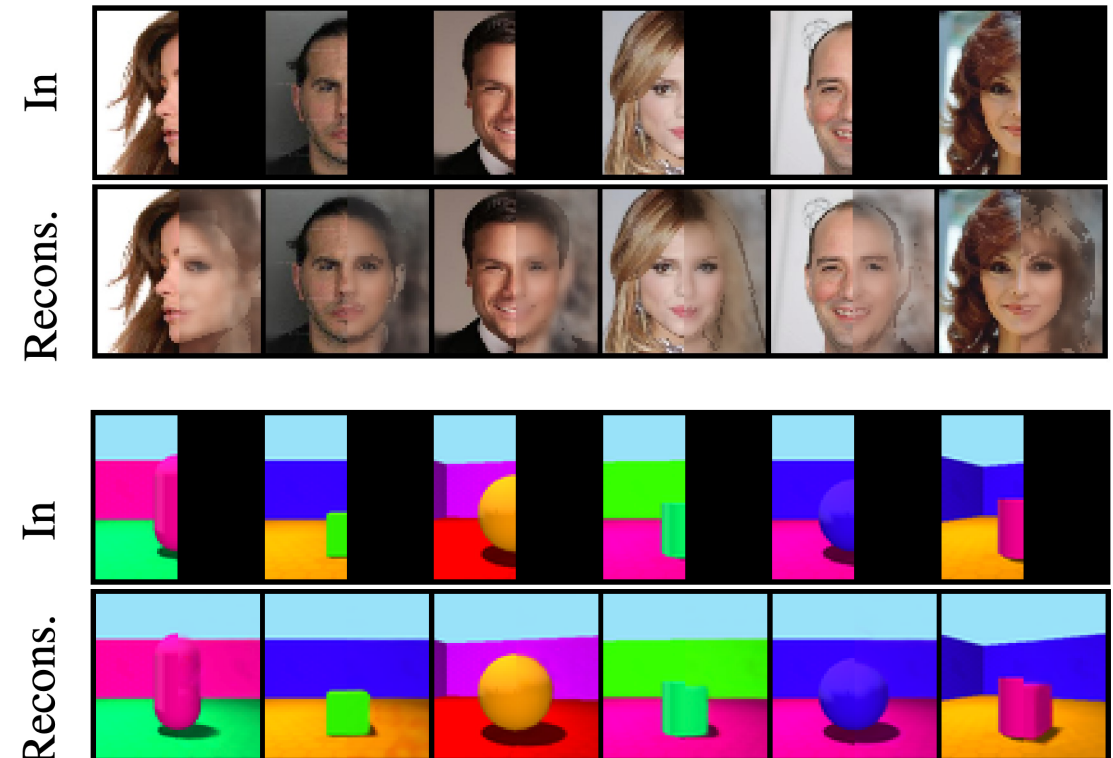
(b) POLYMNIST

# Results

- Our method allows for **efficient conditional generation via inference.**



(a) Missing a patch

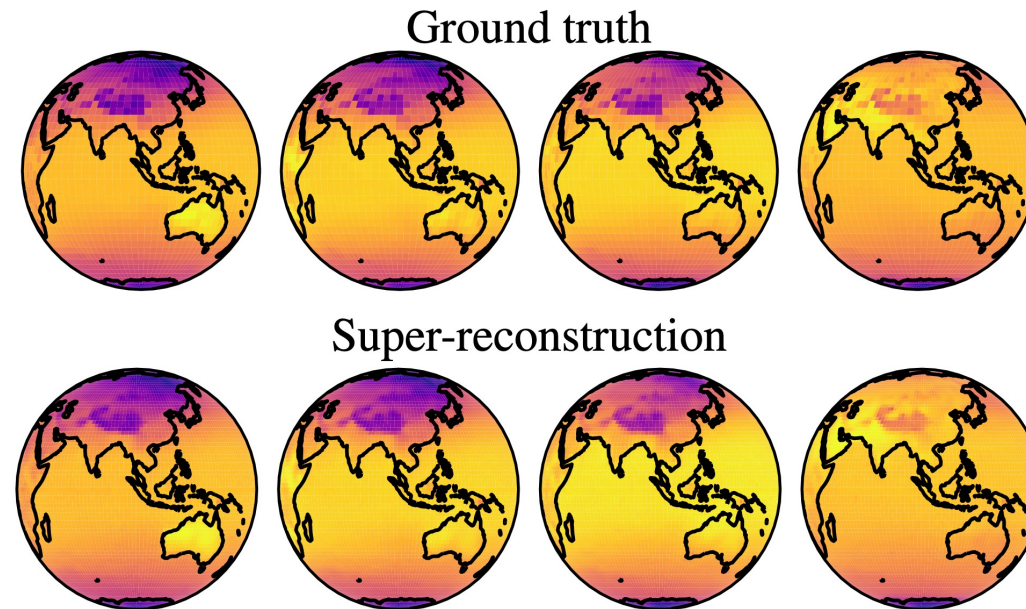


(b) Missing half of the image



# Results

- Our method allows for **efficient conditional generation via inference**.



# Results

- We achieve a **7-11 times faster inference than the alternative!**

Dataset	Model Inference Time (secs)			Speed Improvement	
	VaMoH	Functa (3)	Functa (10)	vs. Functa (3)	vs. Functa (10)
POLYMNIST	<b>0.00453</b>	0.01648	0.05108	<b>x 3.64</b>	<b>x 11.28</b>
SHAPES3D	<b>0.00536</b>	0.01759	0.05480	<b>x 3.28</b>	<b>x 10.22</b>
CELEBA HQ	<b>0.00757</b>	0.01733	0.05381	<b>x 2.29</b>	<b>x 7.11</b>
ERA5	<b>0.00745</b>	0.01899	0.05932	<b>x 2.55</b>	<b>x 7.96</b>
SHAPENET	<b>0.00689</b>	0.02095	0.06576	<b>x 3.04</b>	<b>x 9.54</b>

# Conclusion


- Thanks to **learning distributions of functions**, our proposed **VAMoH** can:
  - “Sample” neural networks for generating new data.
  - **Infer** the **latent representation** of a neural network for conditionally generating data.
  - Use the same neural architecture independently of the nature of the data.
  - Easily perform the **conditional generation at any desired resolution**, while being:
    - ✓ **Robust** to partially observed data.
    - ✓ **Expressive** for generating high-quality data.
    - ✓ **Efficient** in terms of inference.

# References

- [1] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... & Bengio, Y. (2014). Generative adversarial nets. *Advances in neural information processing systems*, 27.
- [2] Kingma, D. P., & Welling, M. (2013). Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- [3] Ho, J., Jain, A., & Abbeel, P. (2020). Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33, 6840-6851.
- [4] Song, Y., & Ermon, S. (2019). Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32.
- [5] LeCun, Y., Chopra, S., Hadsell, R., Ranzato, M., & Huang, F. (2006). A tutorial on energy-based learning. *Predicting structured data*, 1(0).
- [6] Simonyan, K., & Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- [7] Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., ... & Houlsby, N. (2020). An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- [8] Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9(8), 1735-1780.
- [9] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.
- [10] Sitzmann, V., Martel, J., Bergman, A., Lindell, D., & Wetzstein, G. (2020). Implicit neural representations with periodic activation functions. *Advances in neural information processing systems*, 33, 7462-7473.

# References

- [11] Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., & Geiger, A. (2019). Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 4460-4470).
- [12] Stanley, K. O. (2007). Compositional pattern producing networks: A novel abstraction of development. *Genetic programming and evolvable machines*, 8, 131-162.
- [13] Ha, D., Dai, A. M., and Le, Q. V. Hypernetworks. In *International Conference on Learning Representations*, 2017.
- [14] Wu, W., Qi, Z., and Fuxin, L. Pointconv: Deep convolutional networks on 3d point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9621–9630, 2019.
- [15] Rombach, R., Blattmann, A., Lorenz, D., Esser, P., & Ommer, B. (2022). High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 10684-10695).
- [16] Rezende, D. and Mohamed, S. Variational inference with normalizing flows. In *International Conference on Machine Learning*, 2015.
- [17] Dupont, E., Teh, Y. W., and Doucet, A. Generative models as distributions of functions. In *International Conference on Artificial Intelligence and Statistics*, pp. 2989–3015. PMLR, 2022a.
- [18] Dupont, E., Kim, H., Eslami, S. A., Rezende, D. J., and Rosenbaum, D. From data to functa: Your data point is a function and you can treat it like one. In *International Conference on Machine Learning*, pp. 5694–5725. PMLR, 2022



Andaluz.IA

Universidad  
**Carlos III**  
de Madrid



[Paper]



[Code]



# Thank you!

