	IDBM	DBFS (Ours)
MNIST → EMNIST	**10.9**	13.7
EMNIST → MNIST	**8.2**	9.7

Table 1: Comparison with finite-dimensional baseline on unpaired image transfer (MNIST ↔ EMNIST). We compute the FID score. Lower is better.

	NDP	SP-SGM	DBFS (Ours)
Quadratic	≥ 99.0	5.4 ± 0.7	$\mathbf{5.1 \pm 0.4}$
Melbourne	12.8 ± 0.4	$\mathbf{5.3 \pm 0.7}$	9.67 ± 0.45
Gridwatch	16.3 ± 1.8	4.7 ± 0.5	$\mathbf{3.9 \pm 0.4}$

Table 2: Comparison with infinite-dimensional baselines on 1D function generation. We compute a Power(%) of a kernel two-sample test. Lower is better.

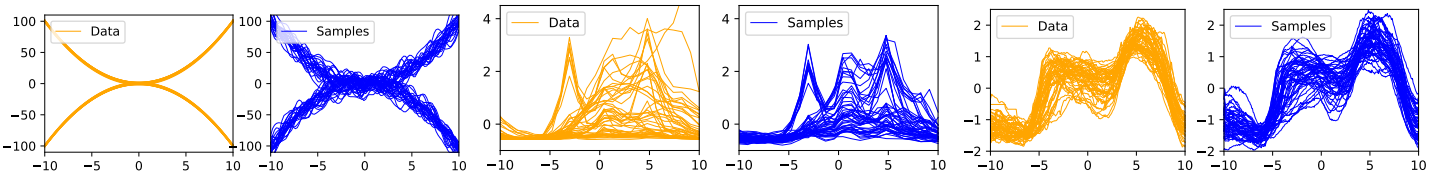

(a) Quadratic (b) Melbourne (c) Gridwatch

Figure 1: (Left) Real data and (Right) generated samples from our model.

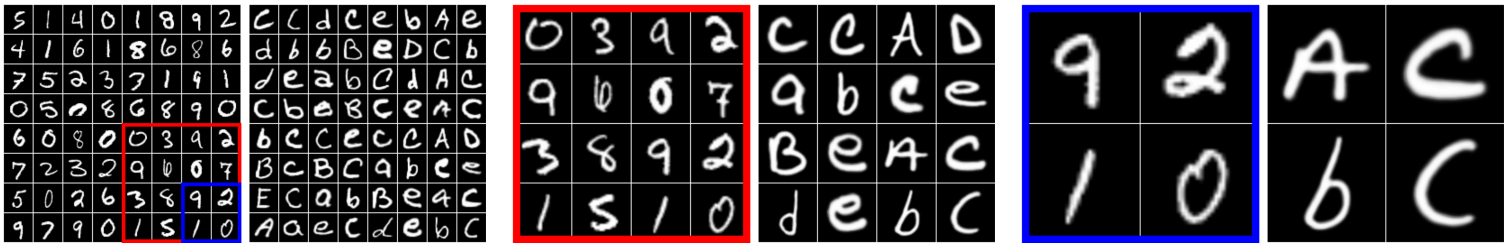

(a) 32^2 (observed resolution) (b) 64^2 (unseen resolution) (c) 128^2 (unseen resolution)

Figure 2: (Left) Real data and (Right) generated samples from our model. For generation at unseen resolutions (64^2 and 128^2), the images within the red and blue boxes were upsampled (using bi-linear transformation) from the observed resolution (32^2).

(a) 64^2 (observed resolution) (b) 128^2 (unseen resolution)

Figure 3: (Left) Real data and (Right) generated samples from our model. For generation at unseen resolution (128^2), the images within the red box were upsampled (using bi-linear transformation) from the observed resolution (64^2).

(a) 64^2 (observed resolution) (b) 128^2 (unseen resolution)

Figure 4: (Left) Real data and (Right) generated samples from our model. For generation at unseen resolution (128^2), the images within the red box were upsampled (using bi-linear transformation) from the observed resolution (64^2).