RICE

NEURAL INFORMATION PROCESSING SYSTEMS

STRAINER: Learning Transferable Features for Implicit Neural Representations

Kushal Vyas, Ahmed Imtiaz Humayun, Aniket Dashpute

Richard G Baraniuk, Ashok Veeraraghavan, Guha Balakrishnan

Rice University, Houston. USA.

kushalvyas.github.io/strainer.html

Implicit Neural Representations(INRs)

$$v = f_\phi(x, y, z)$$

Input:
Spatial Coordinates
(x)

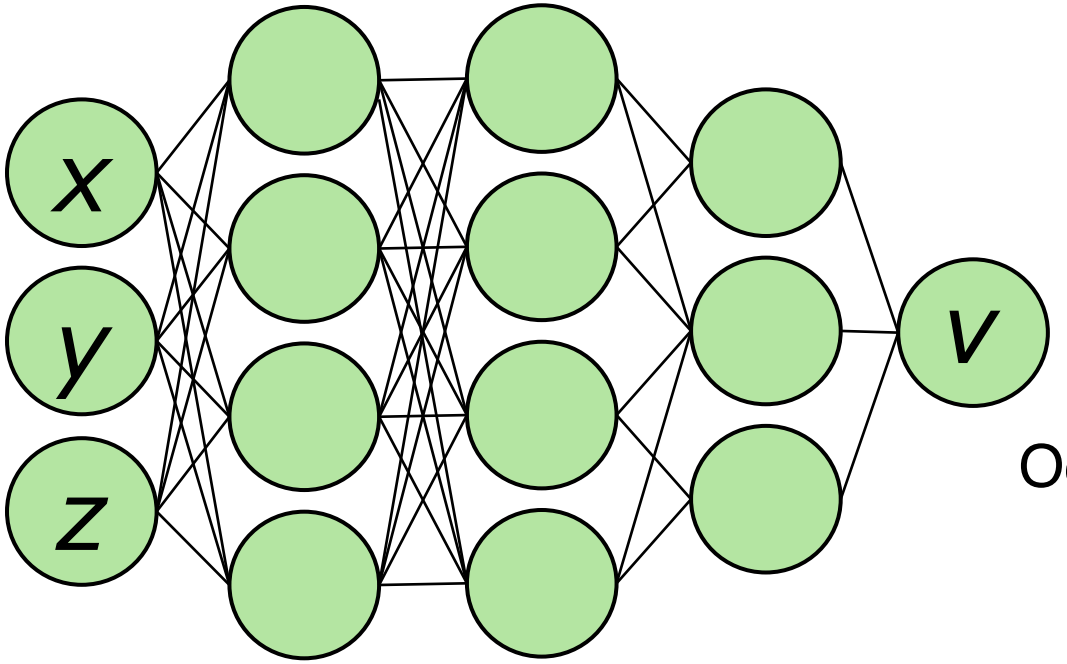

Output:
Pixels
Voxels
Occupancies

- Signal Specific
- Less generalizable

INRs fit images in a coarse-to-fine manner

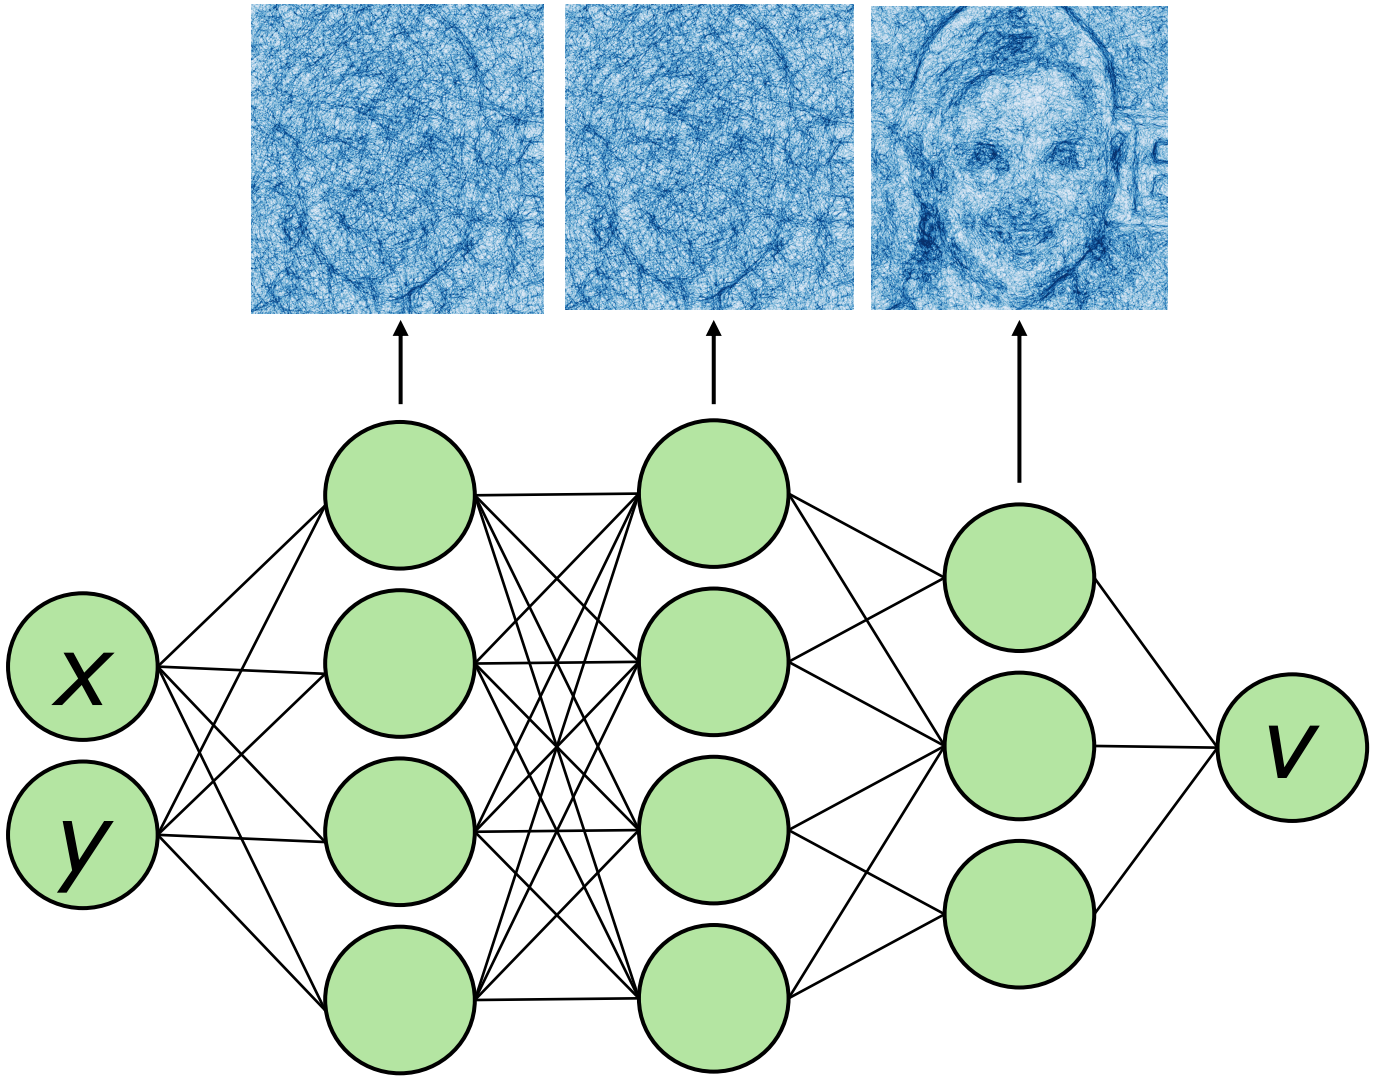

Input:
Spatial Coordinates
(x)

STRAINER: Learning transferable features for INRs

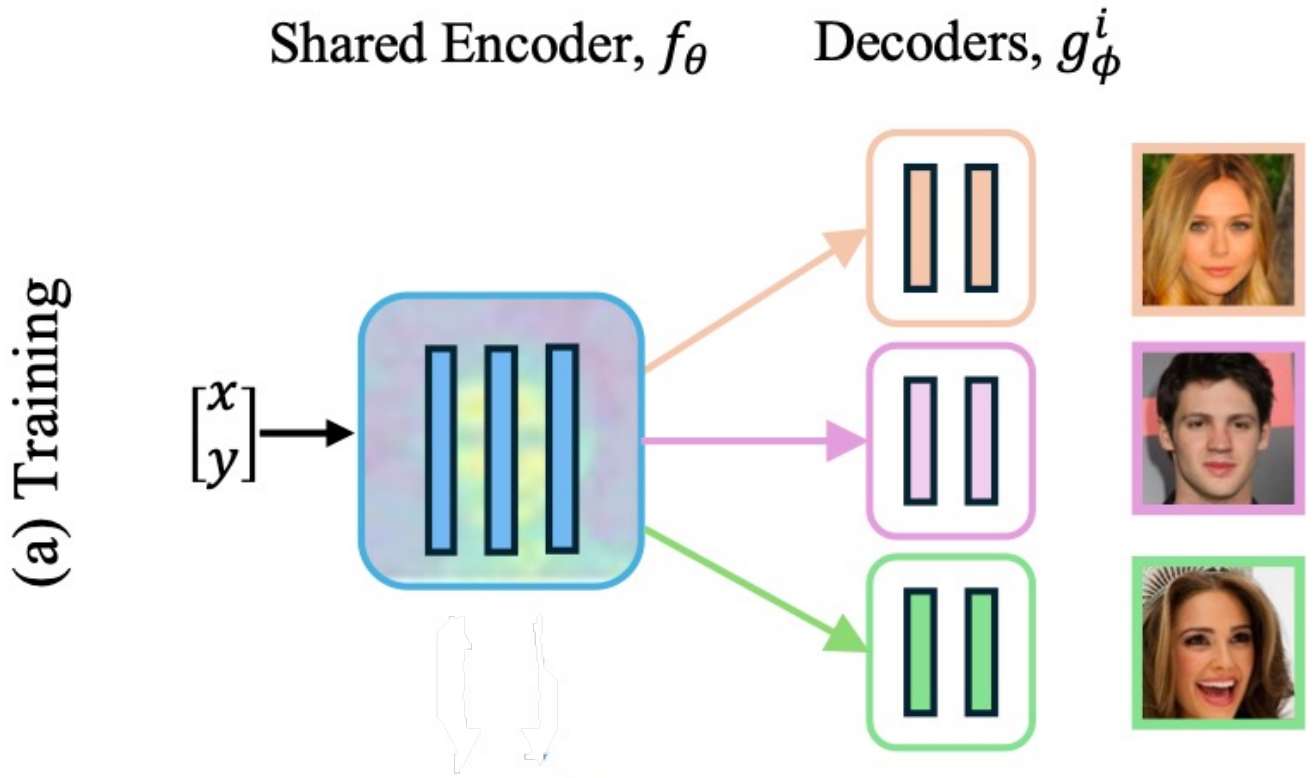

STRAINER: Learning transferable features for INRs

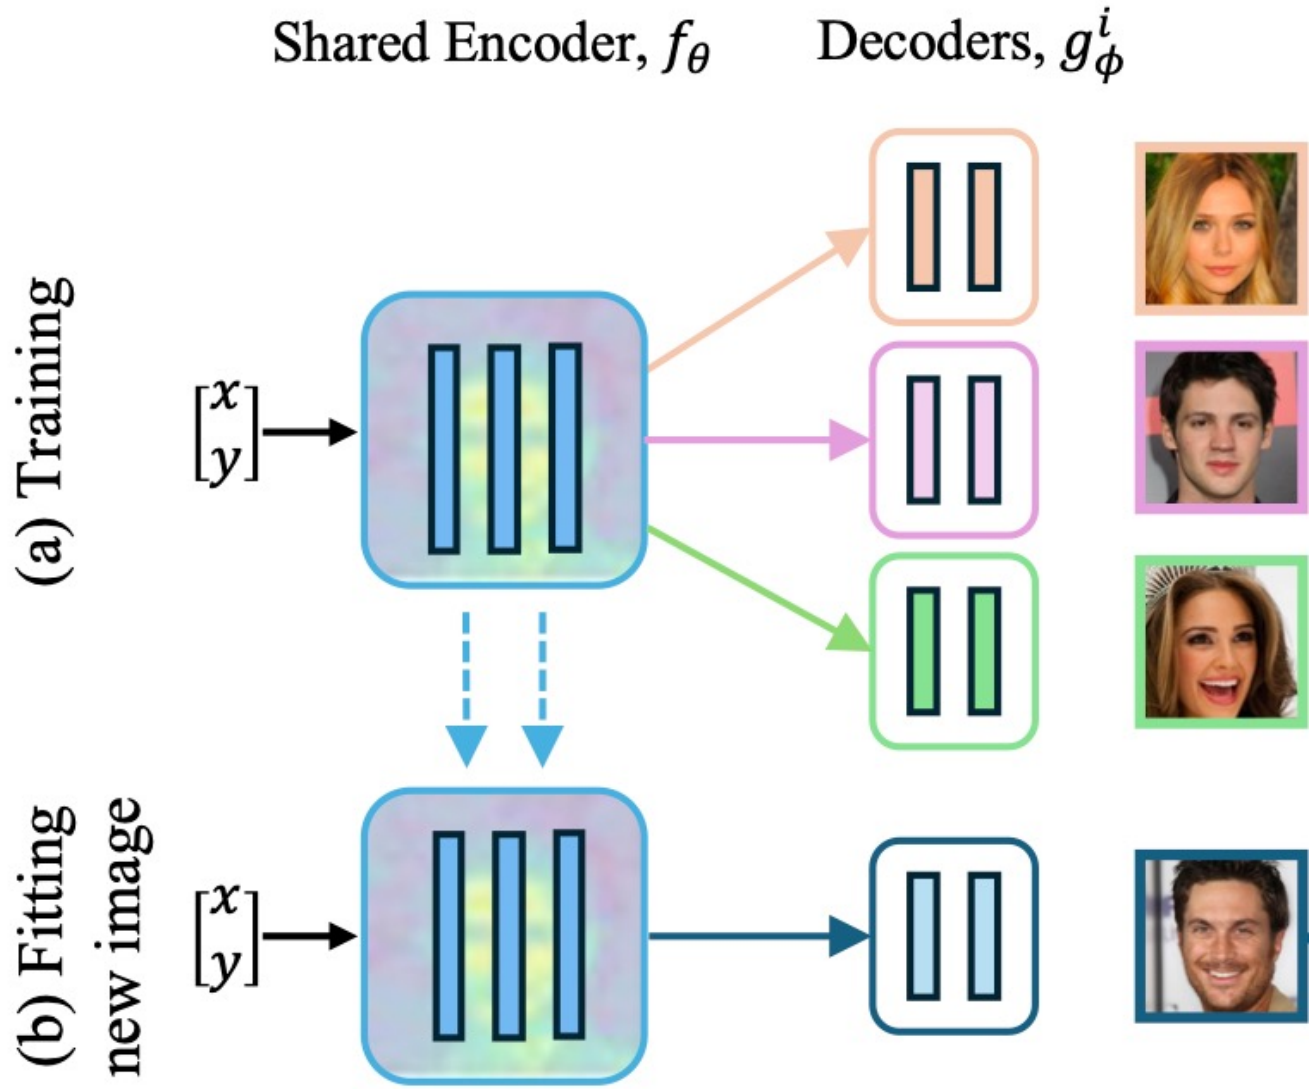

STRAINER: Learning transferable features for INRs

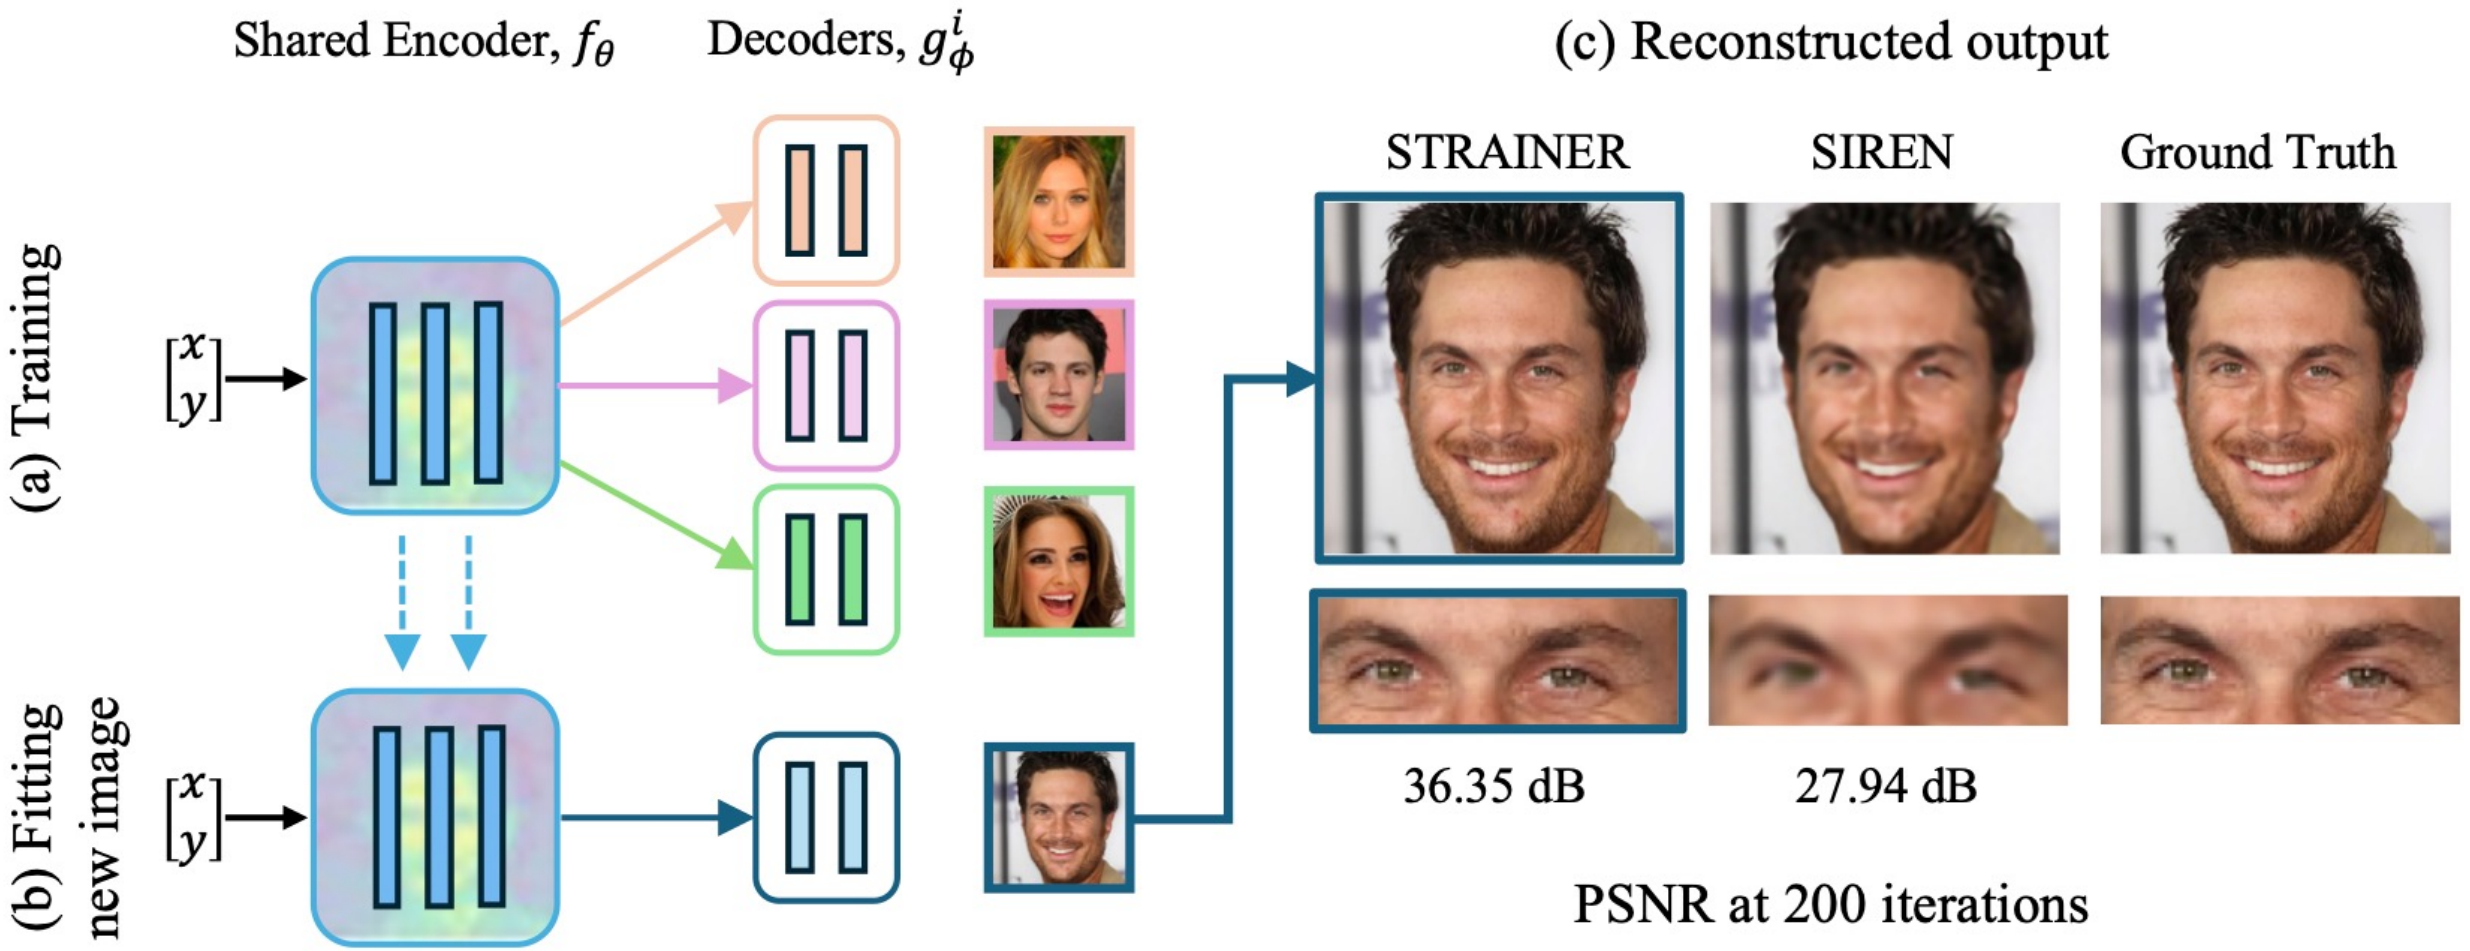

Image Fitting: STRAINER converges fast and with high quality with just 10 training images.

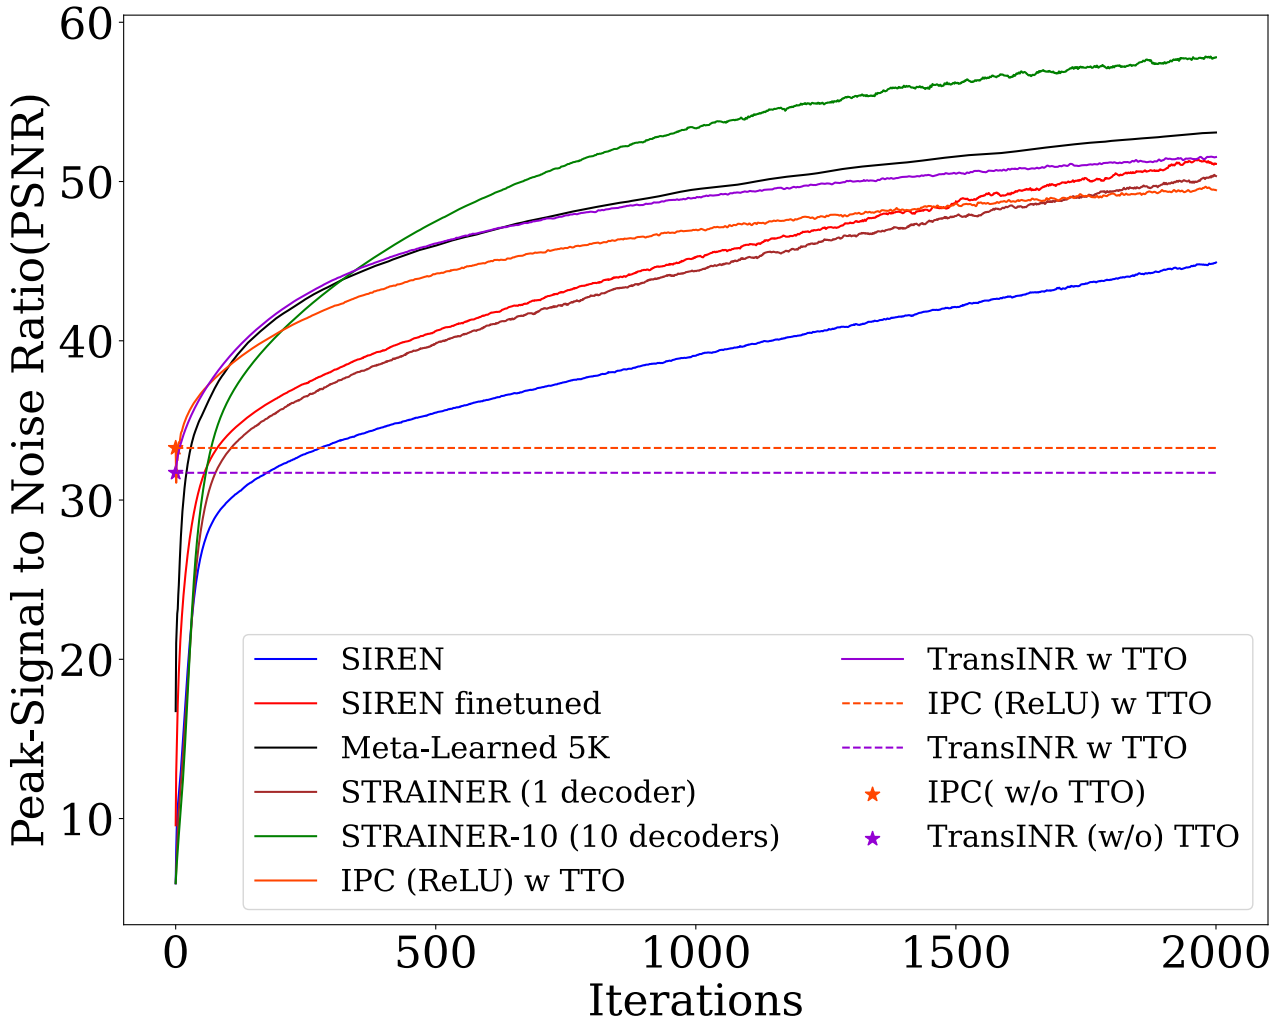

STRAINER also generalizes extremely well for out–of-domain image fitting

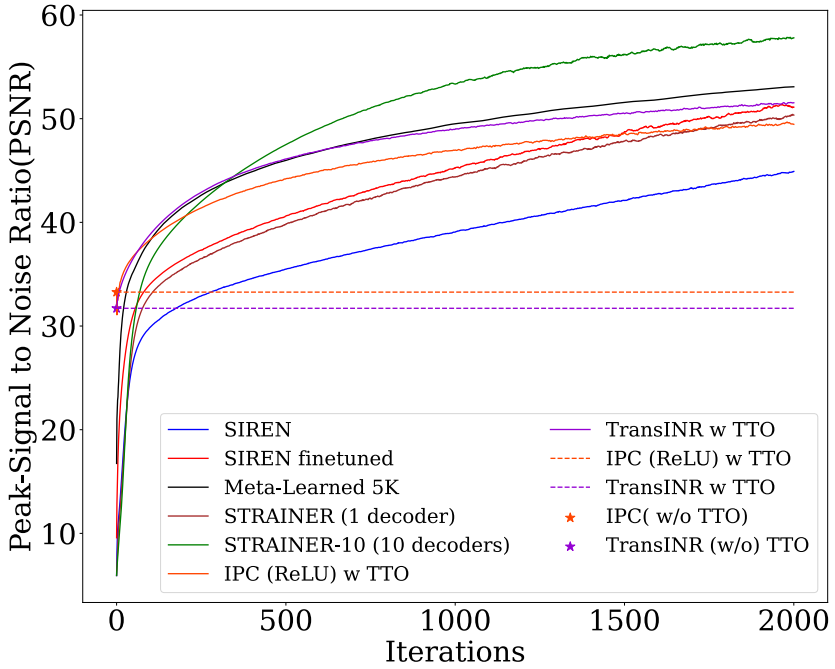

Method	CelebA-HQ (ID) PSNR↑	AFHQ (OOD) PSNR↑	OASIS MRI (OOD) PSNR↑
Meta-learned 5K	53.08	52.40	55.86
Trans INR w/o TTO	31.59	28.63	31.97
Trans INR w TTO	51.86	49.01	55.45
IPC(ReLU + Pos Enc.) w/o TTO	33.27	29.96	33.96
IPC(ReLU + Pos Enc.) w TTO	49.72	47.19	51.35
STRAINER-10	**57.80**	**57.46**	59.50
STRAINER-10 (trained on Flowers[1])	-	56.98	58.52
STRAINER-10 (trained on StanfordCars[2])	-	56.88	**59.66**

STRAINER converges rapidly for inverse problems

	Super Resolution (Fast)		Denoising	
Method	PSNR	# iterations	PSNR	# iterations
SIREN	32.10	3329	26.75 ± 1.67	203 ± 66
STRAINER -10	31.56	1102 ($\approx 3 \times faster$)	26.41 ± 1.39	76 ± 27

Speedup means interesting priors have been encoded!

Visualizing STRAINER partition geometry in input space

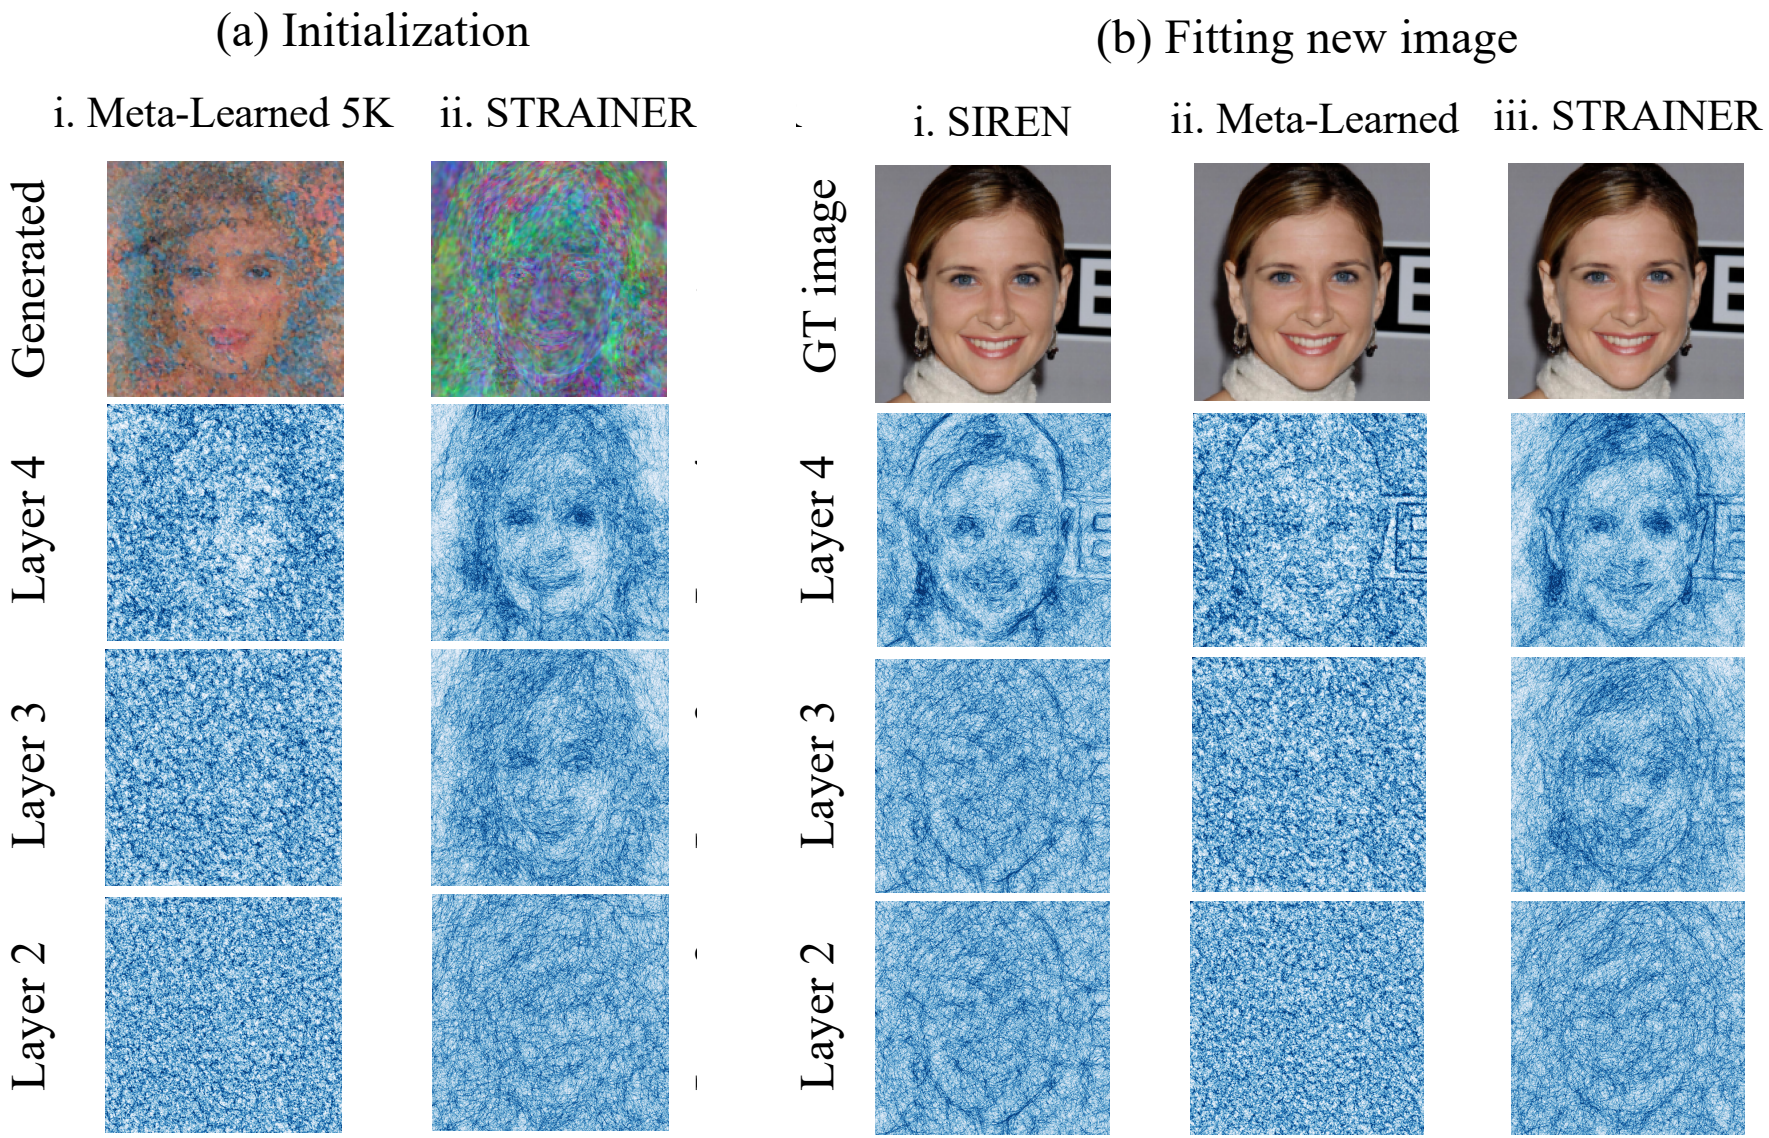

STRAINER learns blazing fast!

Method	# training images	Gradient updates / iteration	Time (Nvidia A100)
SIREN	N/A	N/A	N/A
STRAINER (1 decoder)	1	264,707	11.84s
STRAINER-10 (10 decoders)	10	271,646	24.54s
Meta-learned 5K	10	794,121 ($\approx 3\times$more)	1432.3s = 23.8 min
TransINR[9]	14,000	$\approx 40M$	$\approx$ 1 day
IPC[23] w TTO	14,000	$\approx 40M$	$\approx$ 1 day

RICE

NEURAL INFORMATION PROCESSING SYSTEMS

Kushal Vyas

Ahmed Imtiaz Humayun

Aniket Dashpute

Richard G Baraniuk

Ashok Veeraraghavan

Guha Balakrishnan

For more details:

kushalvyas.github.io/strainer.html