Non-Stationary Texture Synthesis by Adversarial Expansion

Yang Zhou, Zhen Zhu, Dani Lischinski, Xiang Bai, Daniel Cohen-Or, Hui Huang

SIGGRAPH 2018

Presented by Wenjing Wang
2018.09.12

Classical Example-Based Texture Synthesis Tasks

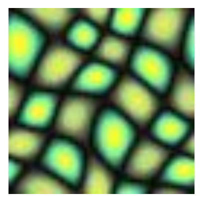

Input

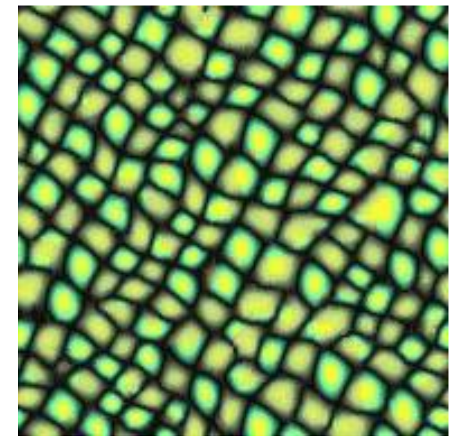

Output

- Stationary

- Regular

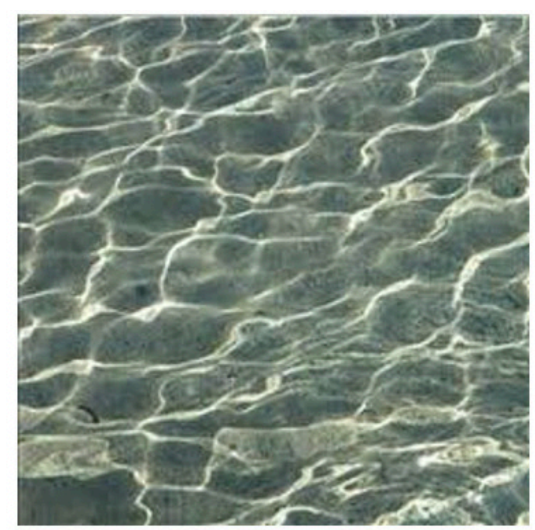

- Near-regular

- Stochastic

Classical Example-Based Texture Synthesis Methods

- Pixel-based synthesis

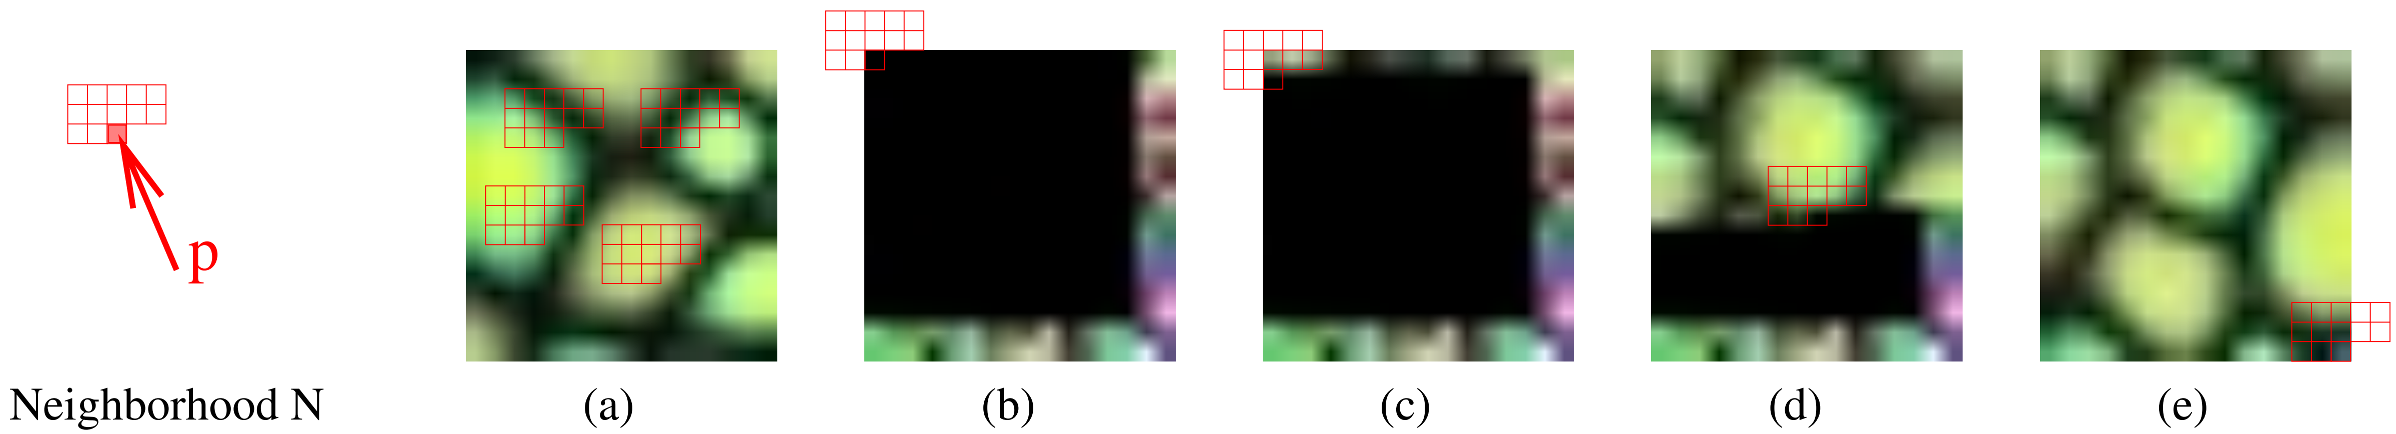

- Patch-based synthesis

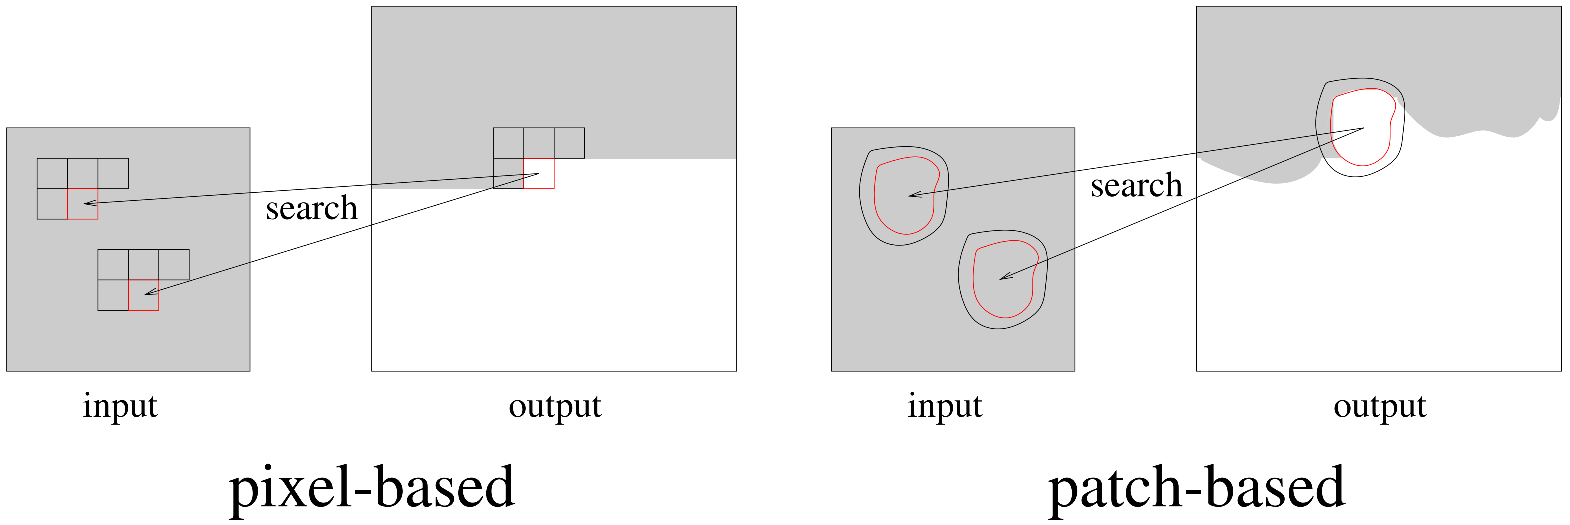

Classical Example-Based Texture Synthesis Methods

➤ Image Analogy

➤ Layered Shape Synthesis

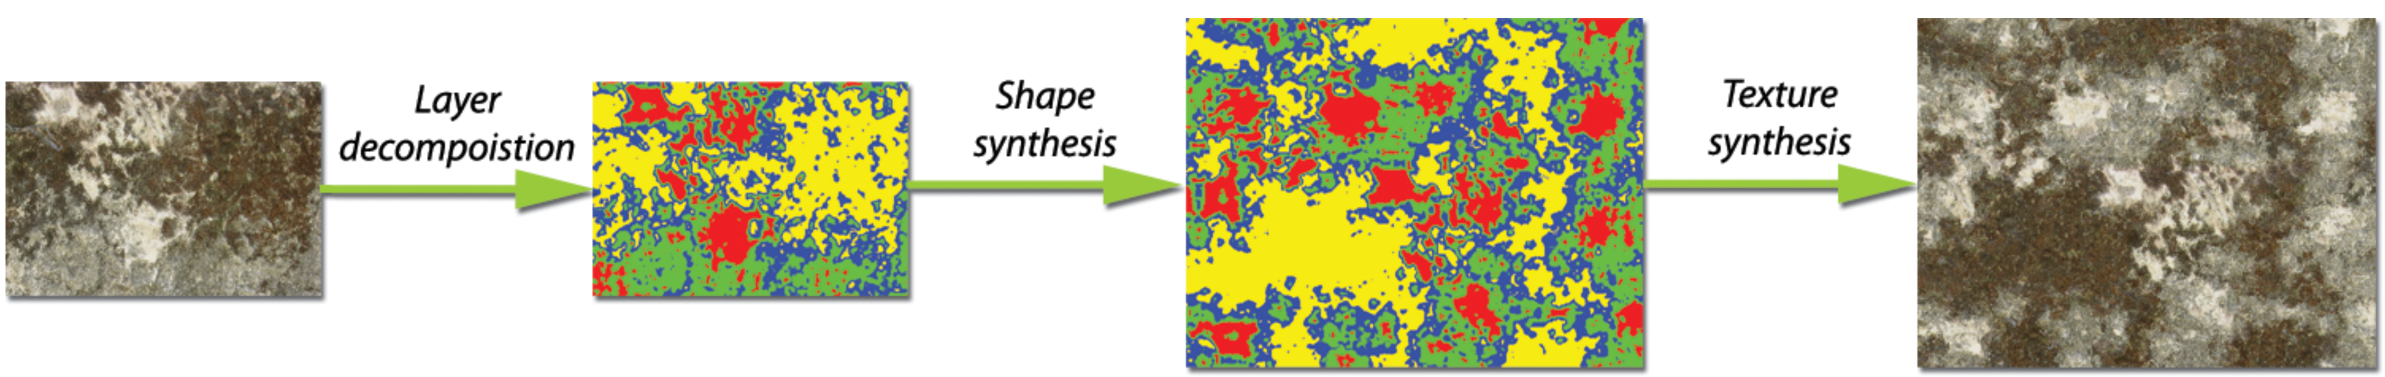

Classical Example-Based Texture Synthesis Methods

➤ With Guidance

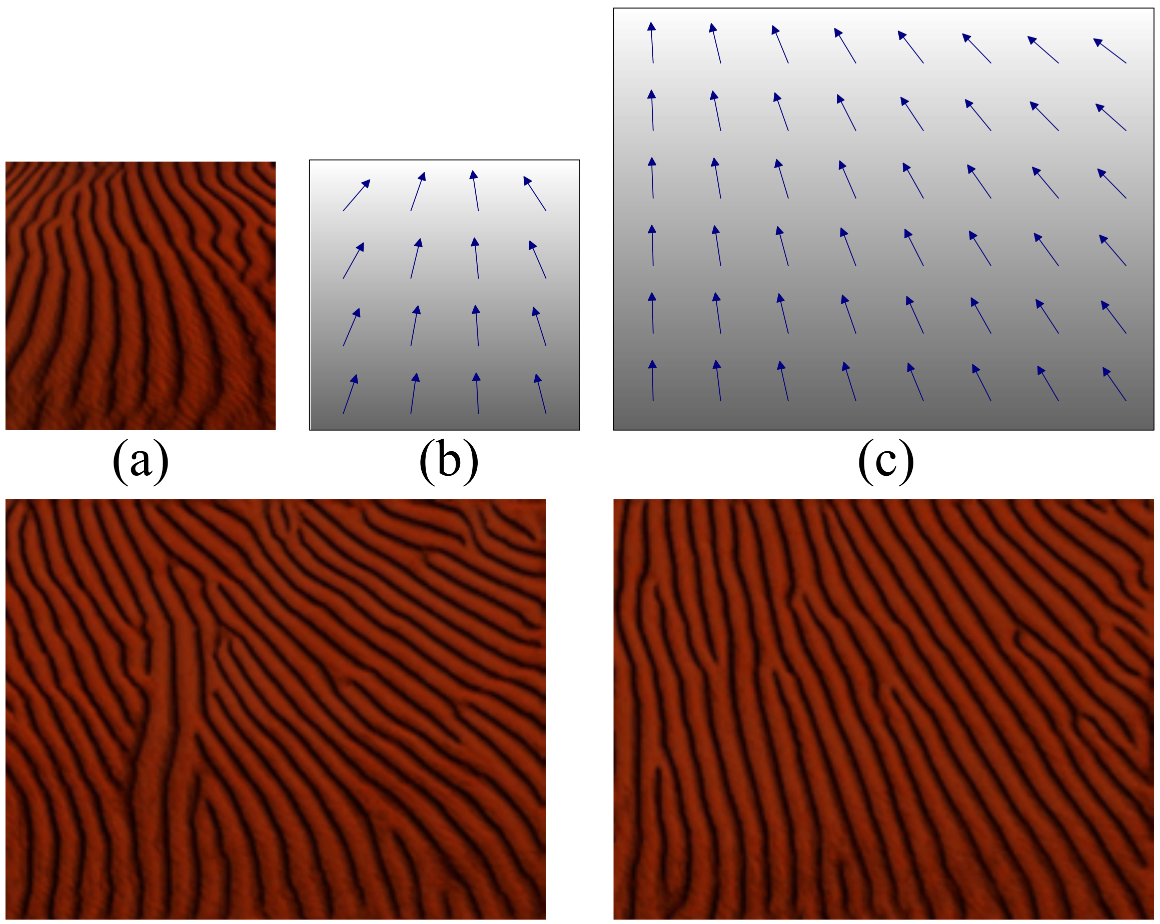

(a) (b) (c)

(d) (e)

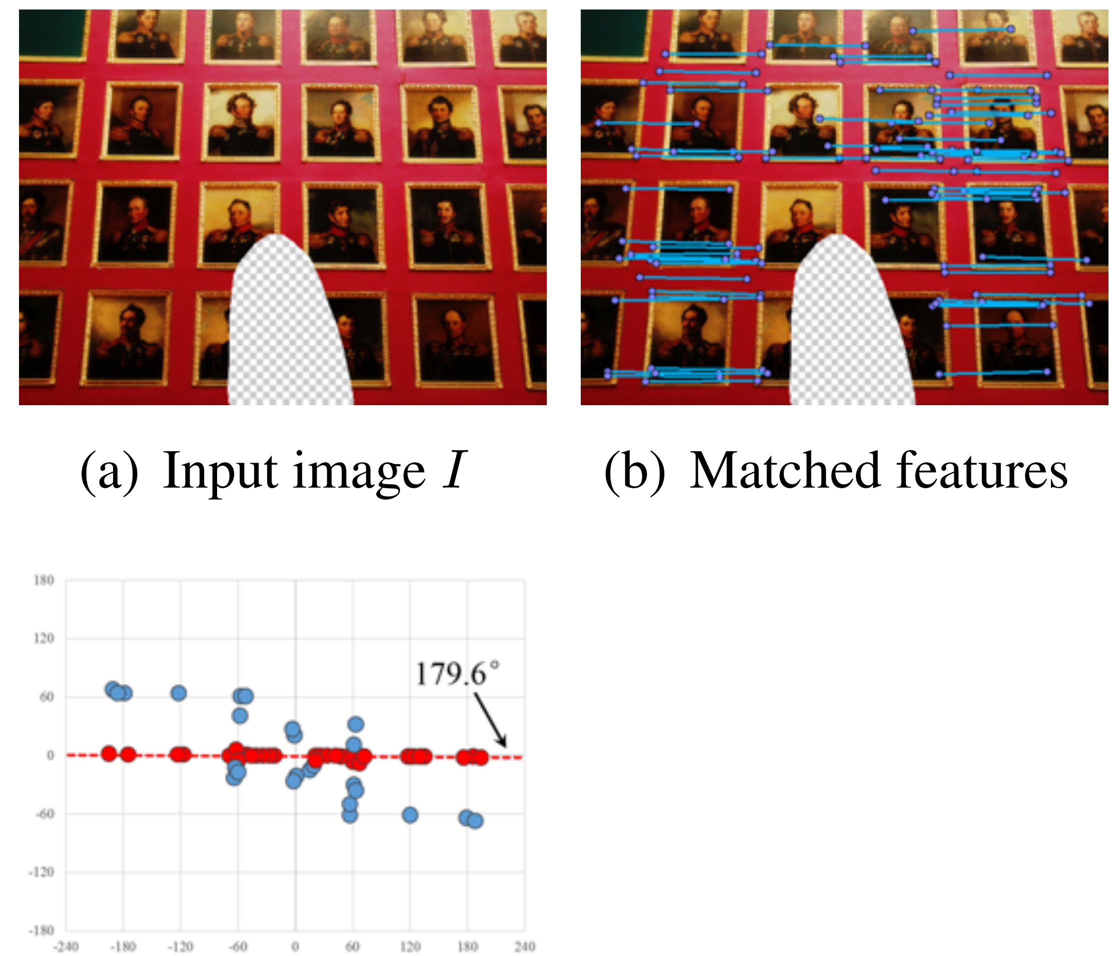

(a) Input image I (b) Matched features

(c) Structure lines

What Is Non-Stationary Textures?

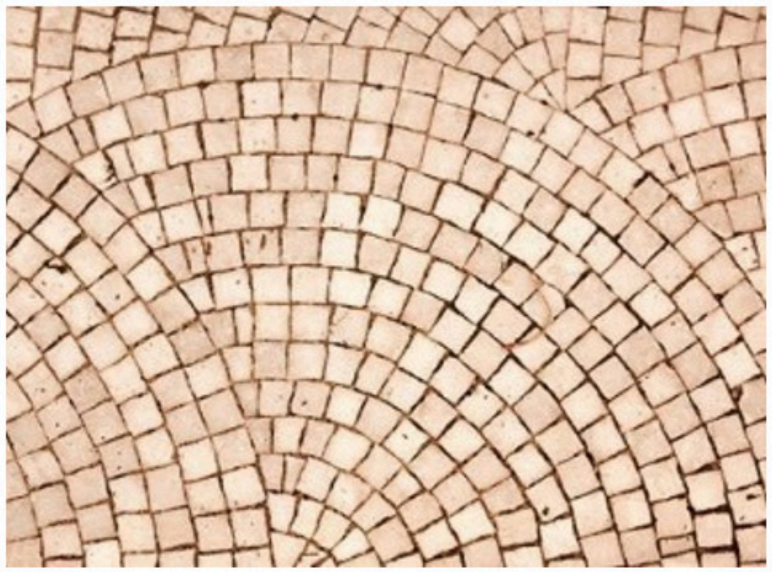

- Textures with large scale structures

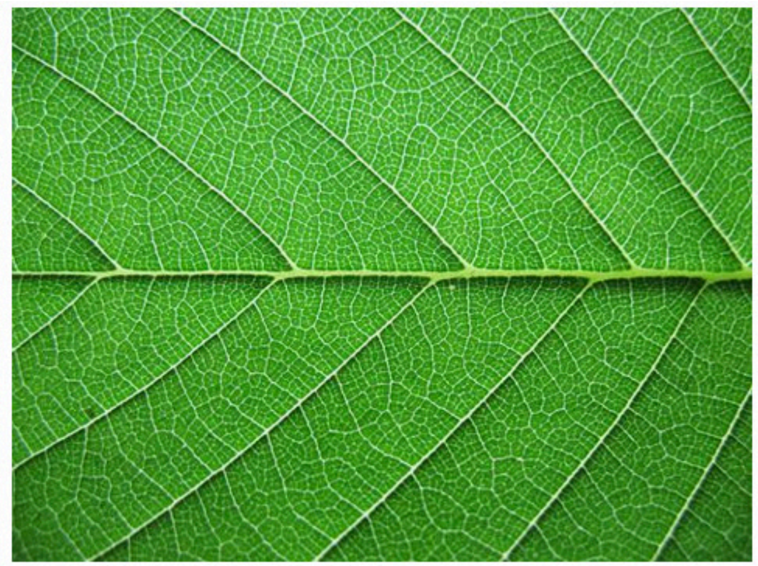

- Spatially variant

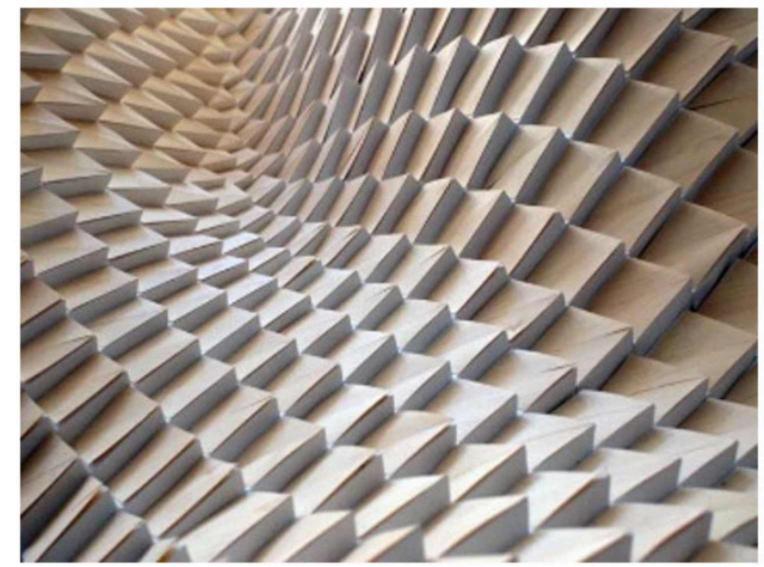

- Inhomogeneous textures

And This Paper Is…

- The first Method than can handle these challenging textures

Output *Input* *Output*

Contents

Method

➤ Train a GAN for each input exemplar.

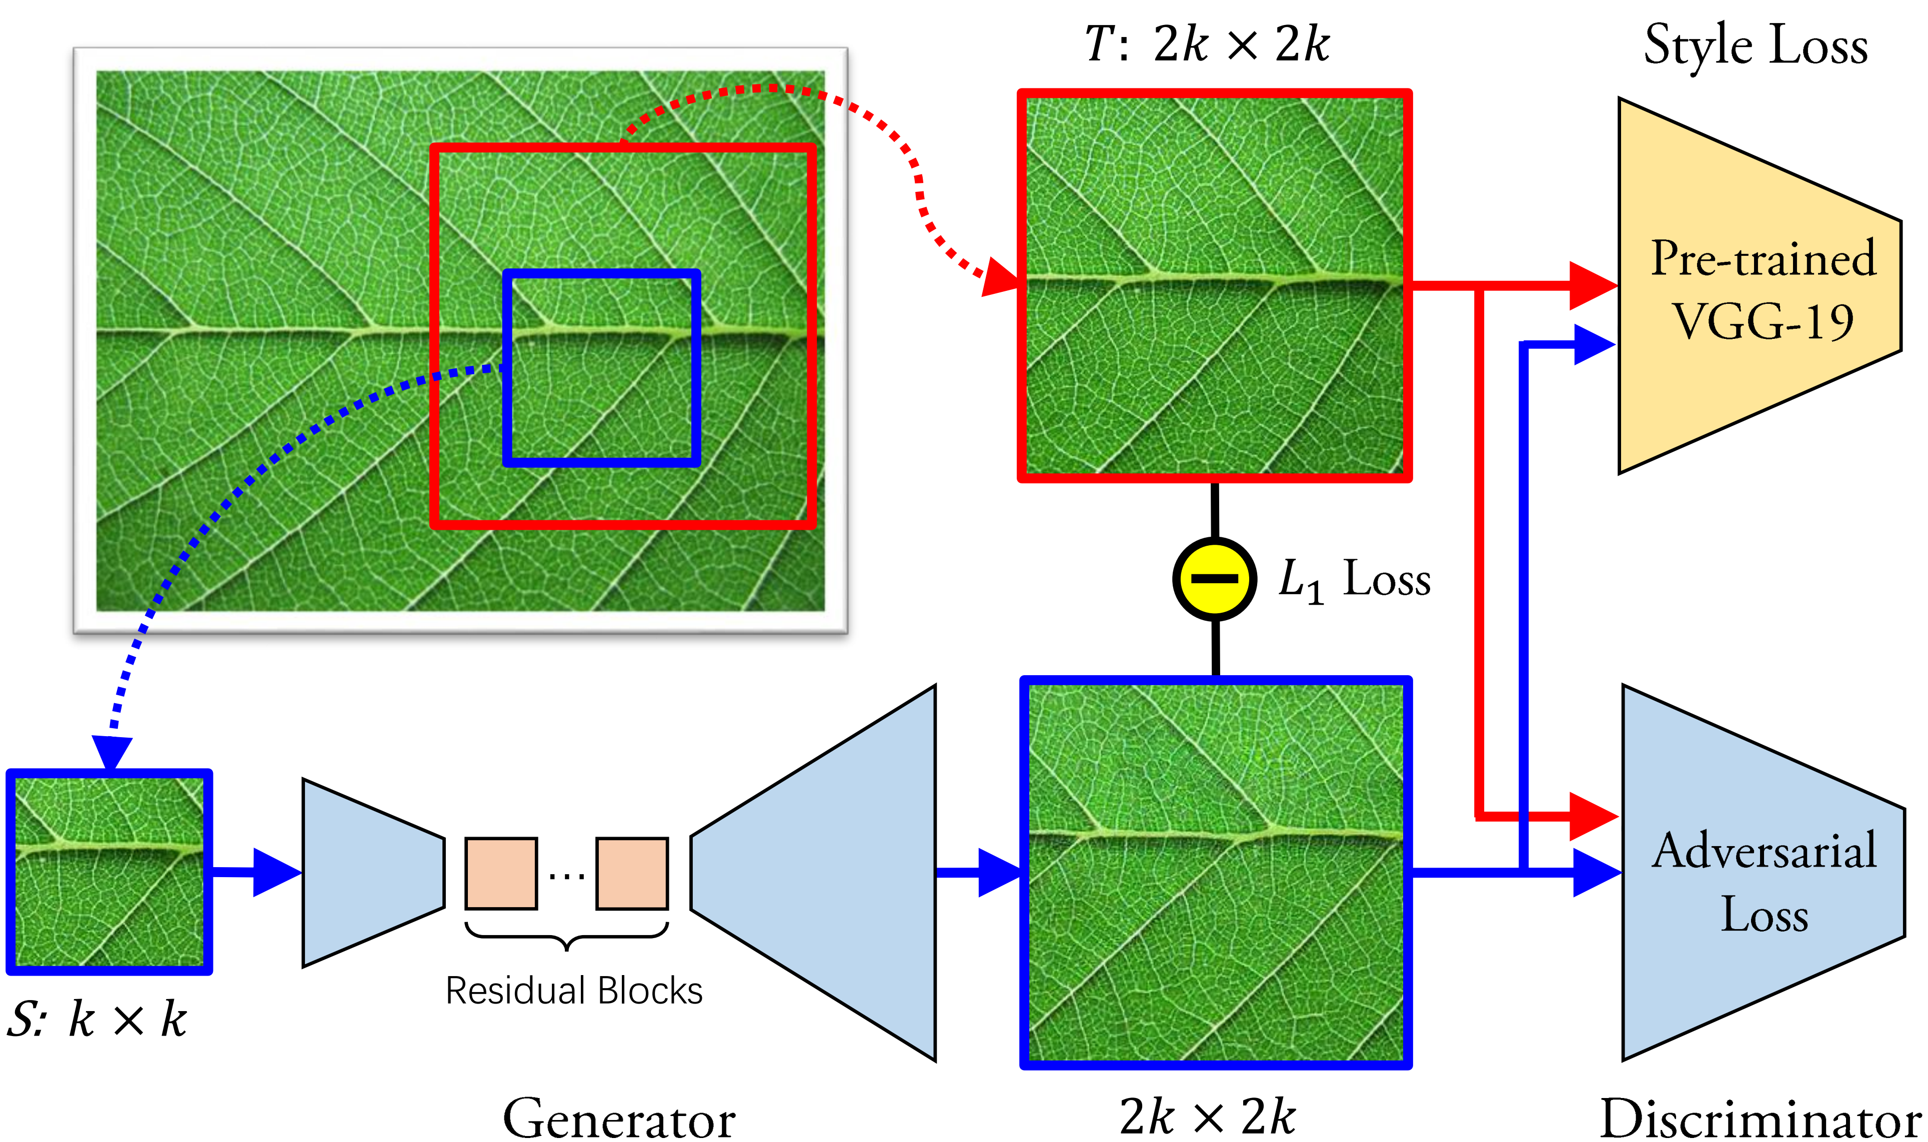

Typically $k = 128$, exemplars are of size 600×400.

Method

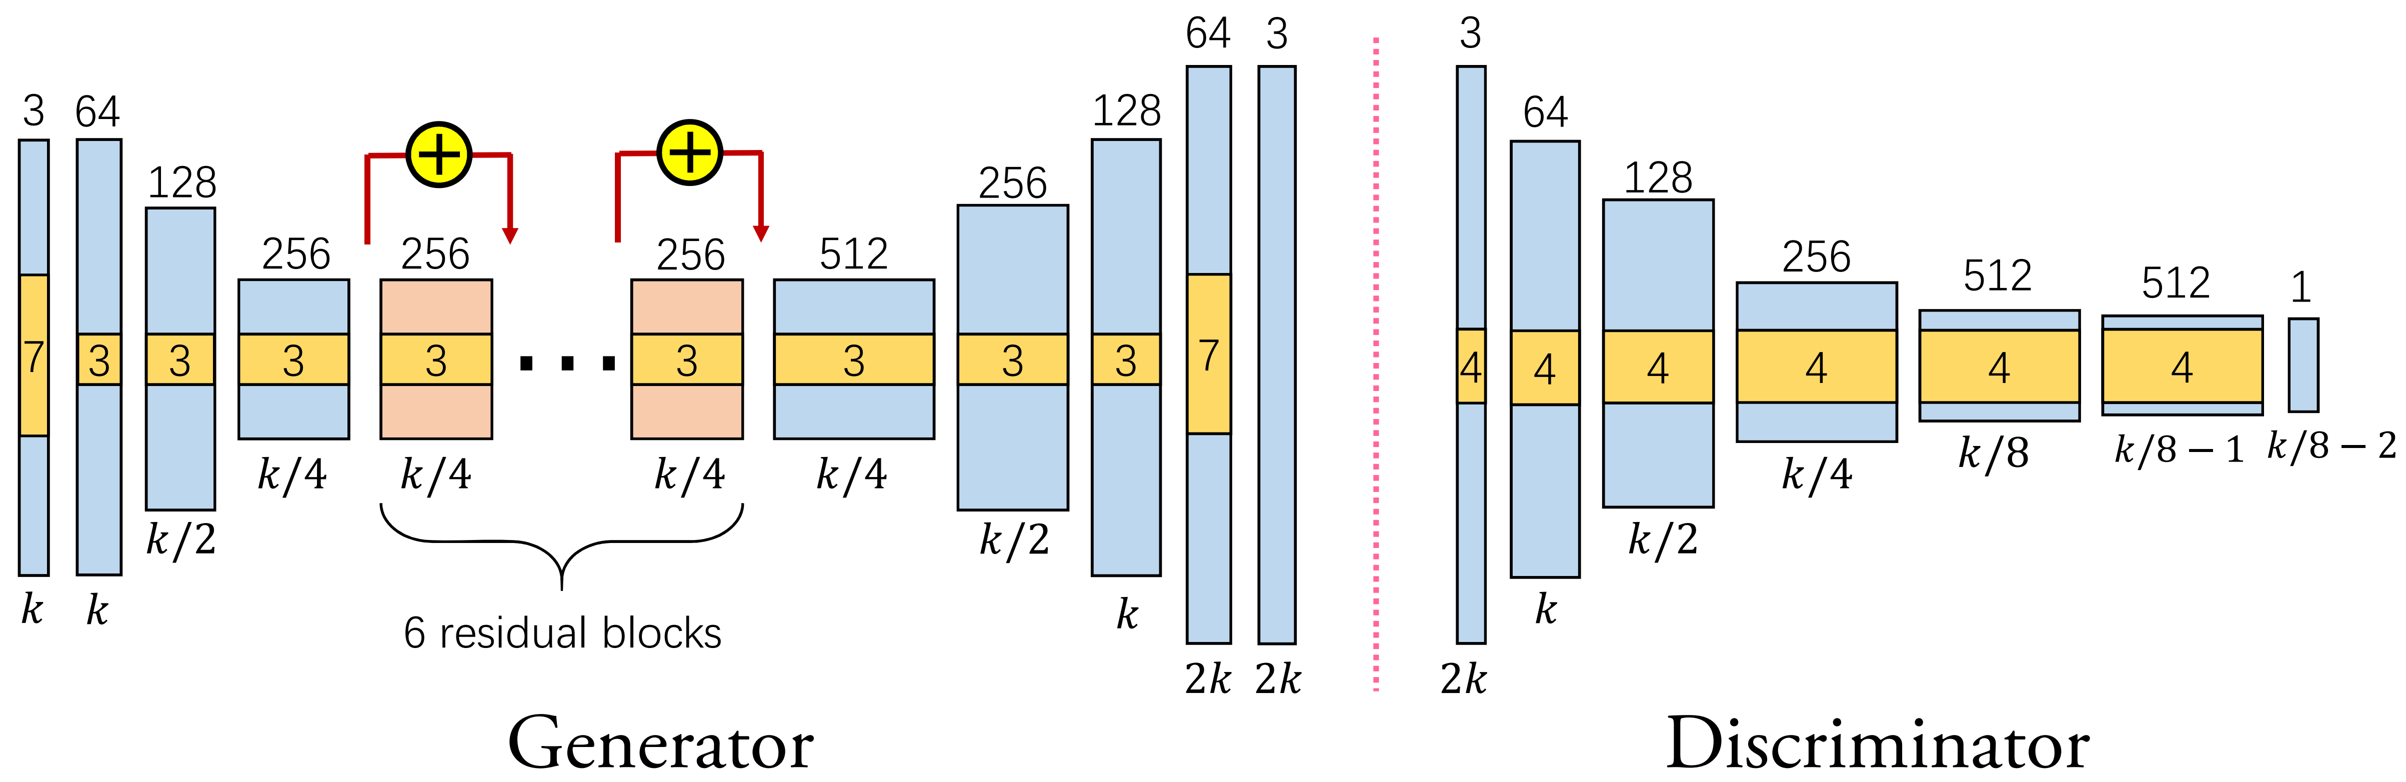

- The receptive field should approach the size of the source block.
- Fully convolutional architecture → arbitrary-sized inputs at test time
- PatchGAN discriminator

Method

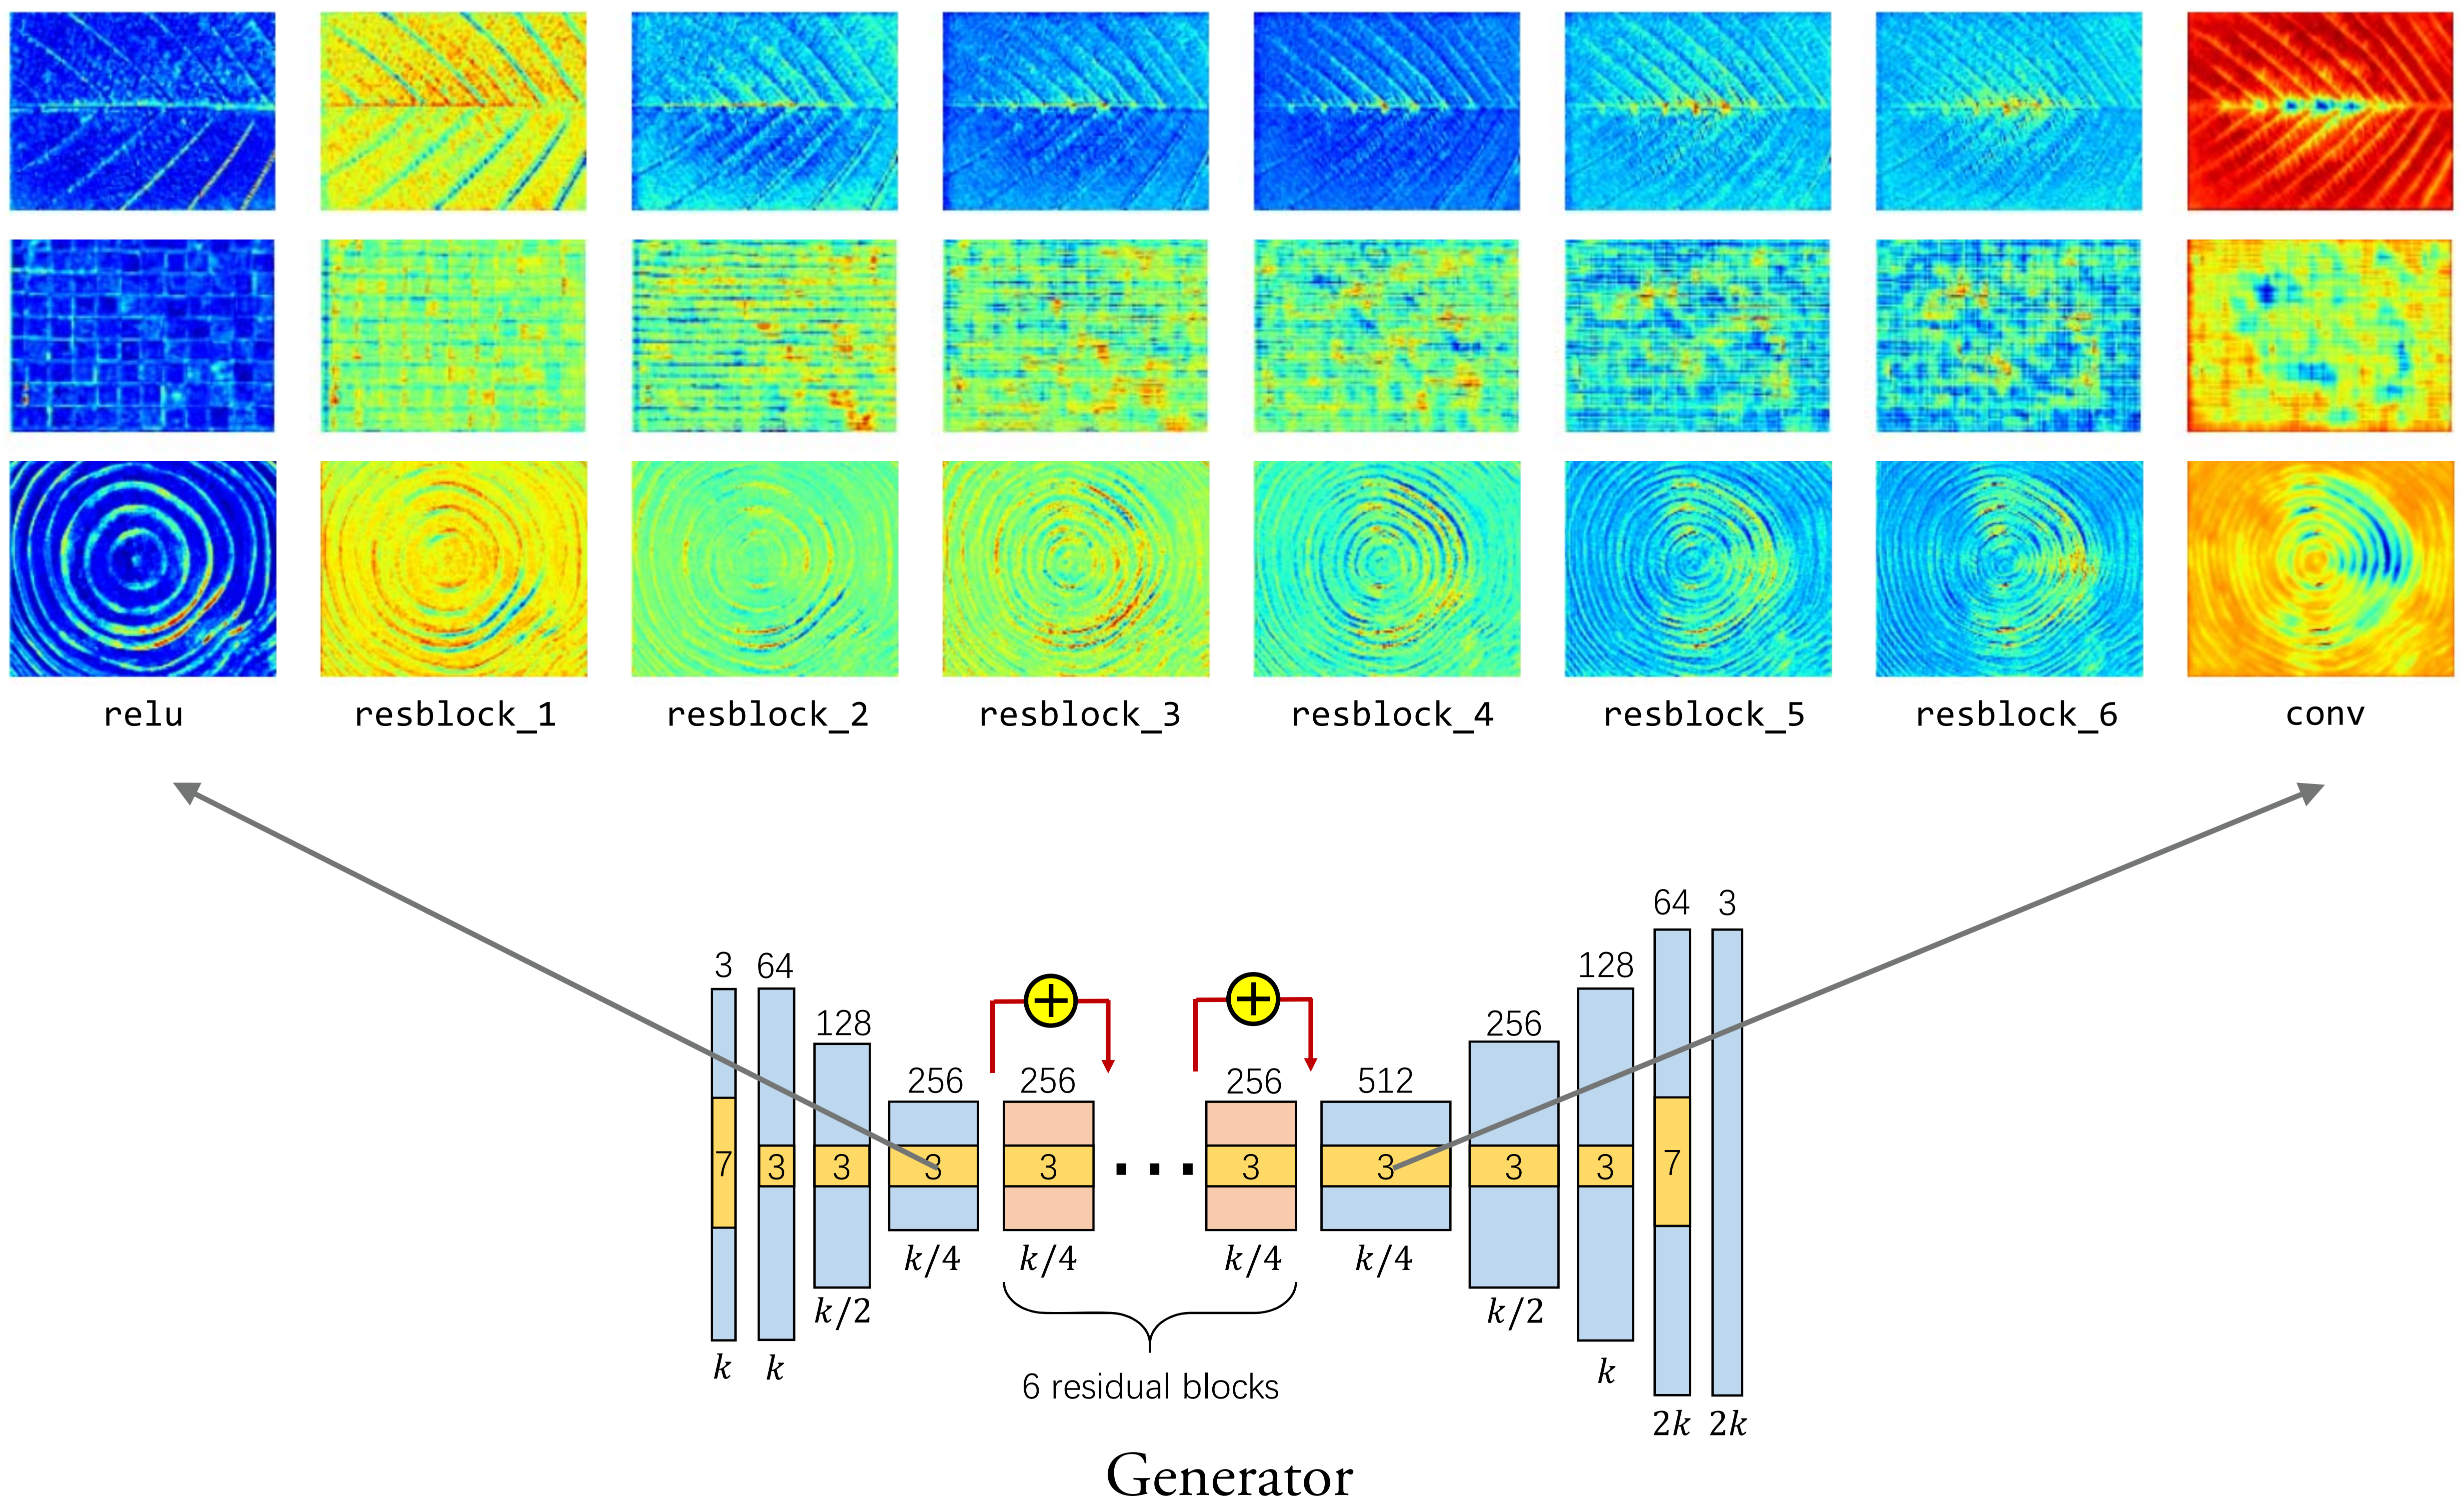

relu
resblock_1
resblock_2
resblock_3
resblock_4
resblock_5
resblock_6
conv
3
64
128
256
256
256
512
256
128
64
3
7
3
3
3
3
3
3
3
3
7
k
k
k/2
k/4
k/4
k/4
k/4
k/2
k
2k
2k
6 residual blocks
Generator

Method

- Loss: $\mathcal{L}_{total} = \mathcal{L}_{adv} + \lambda_1 \mathcal{L}_{L_1} + \lambda_2 \mathcal{L}_{style},$
- adv: standard adversarial loss
- L1: L1 loss
- style: perceptual style loss

- Training: NVIDIA Titan Xp GPU (12GB)
- 5 hours for 100,000 iterations (usually <2hour for 36,000 iterations)

- Testing: 4–5 ms for a 600×400 input

Contents

- Authorship
- Methods
- Experiments

Experiments

➤ both non-stationary and stationary textures

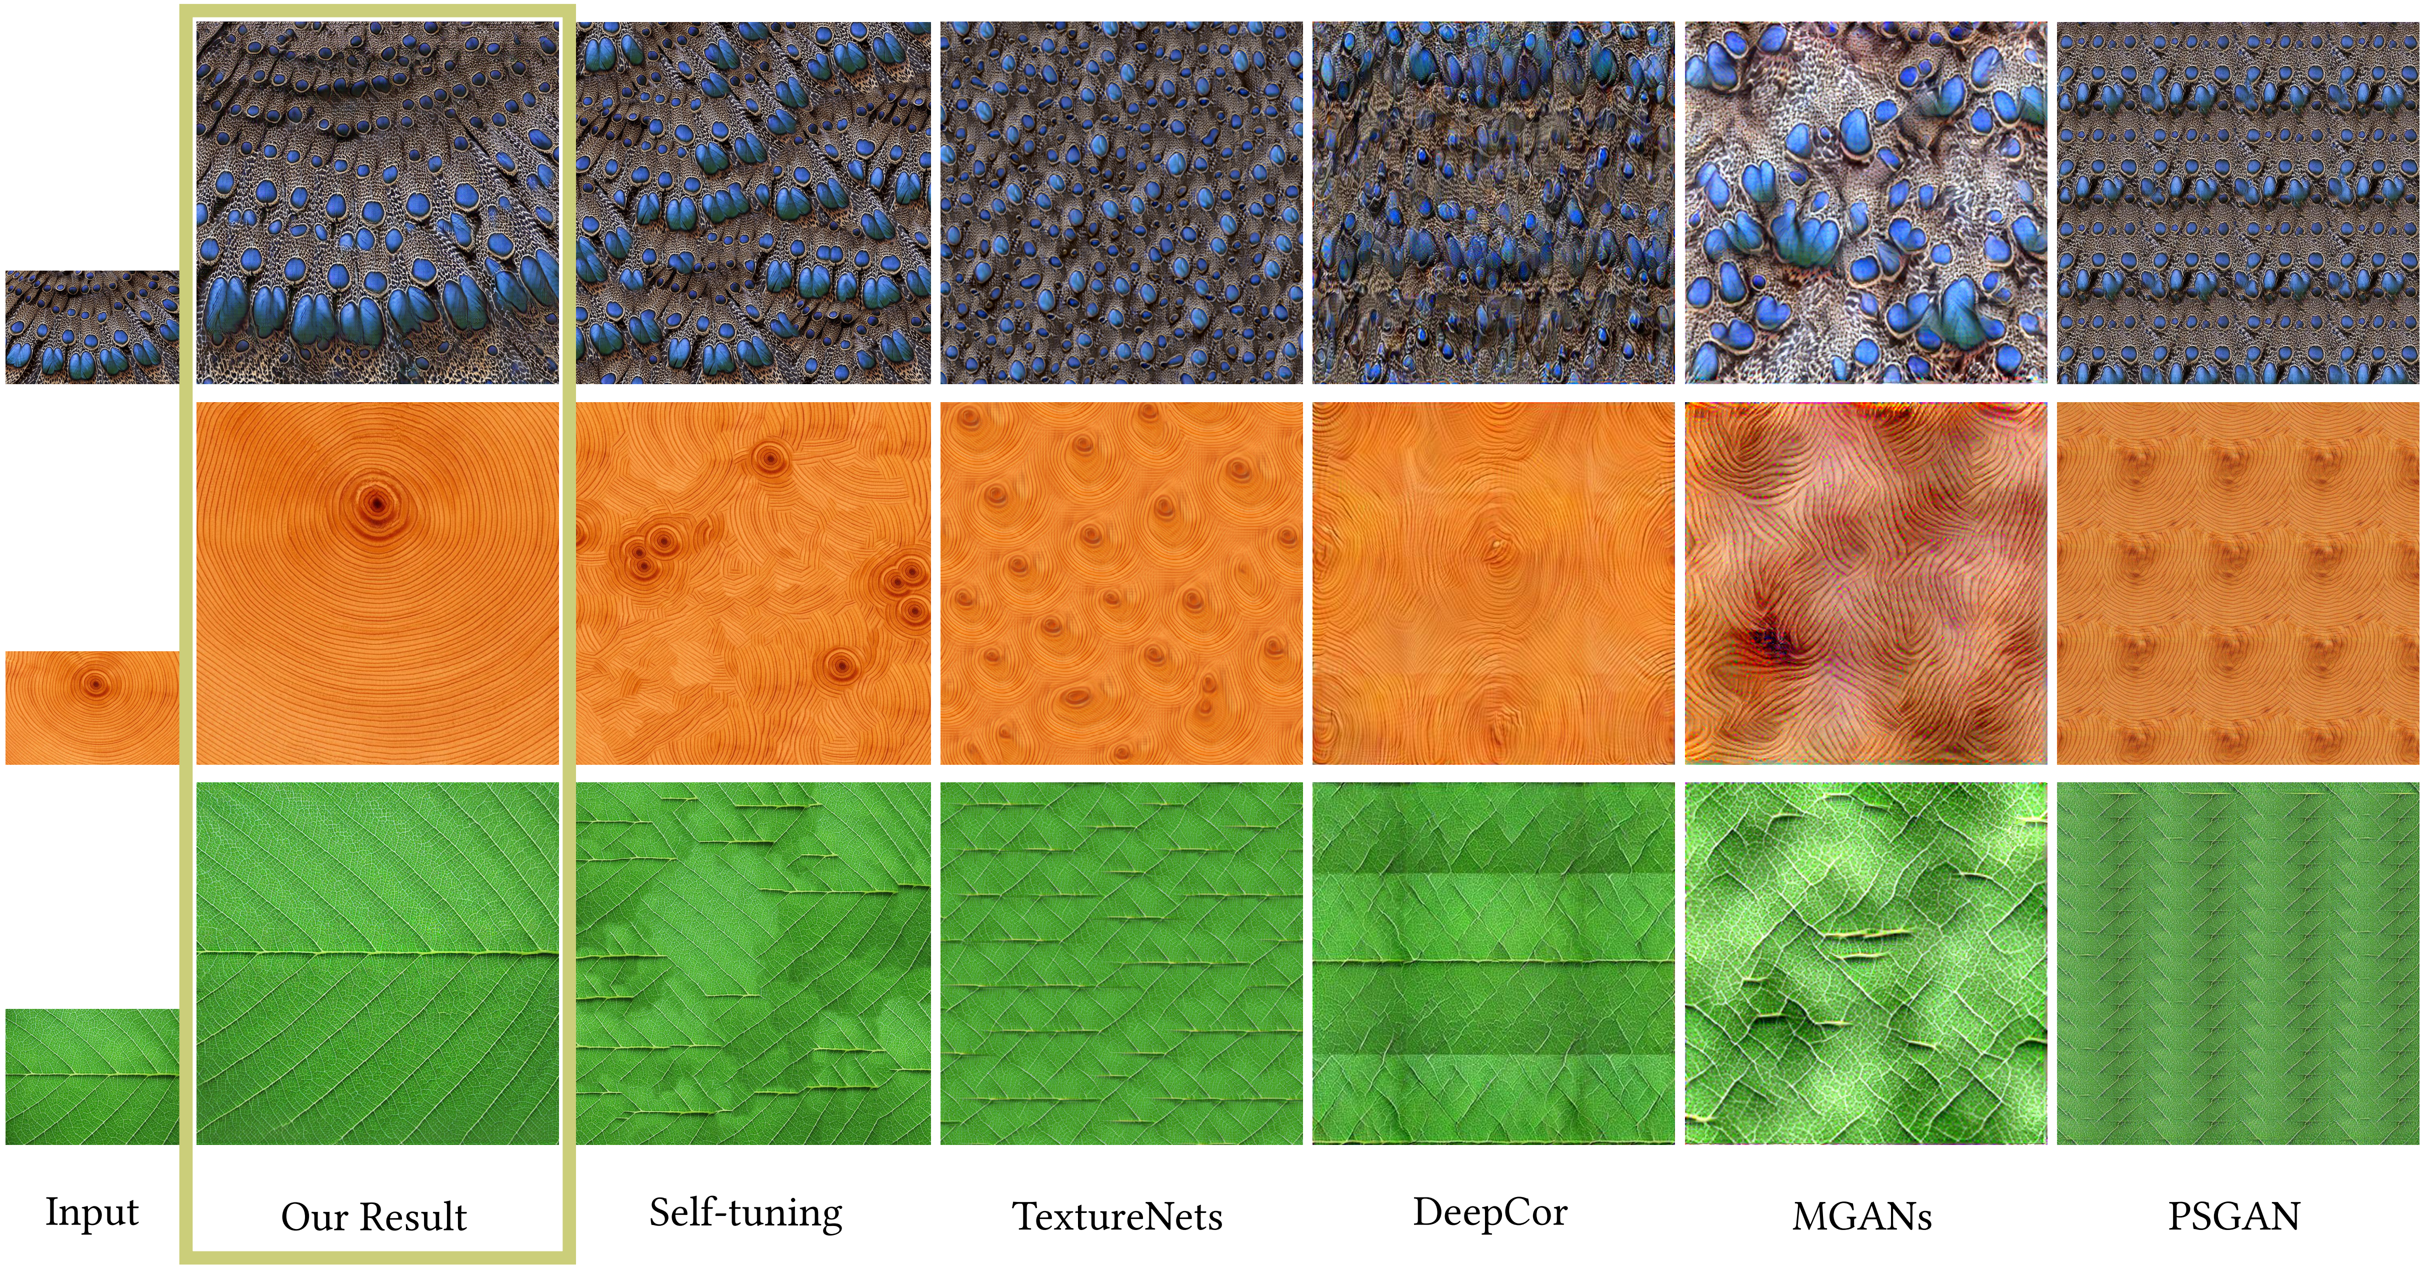

Experiments

- both non-stationary and stationary textures

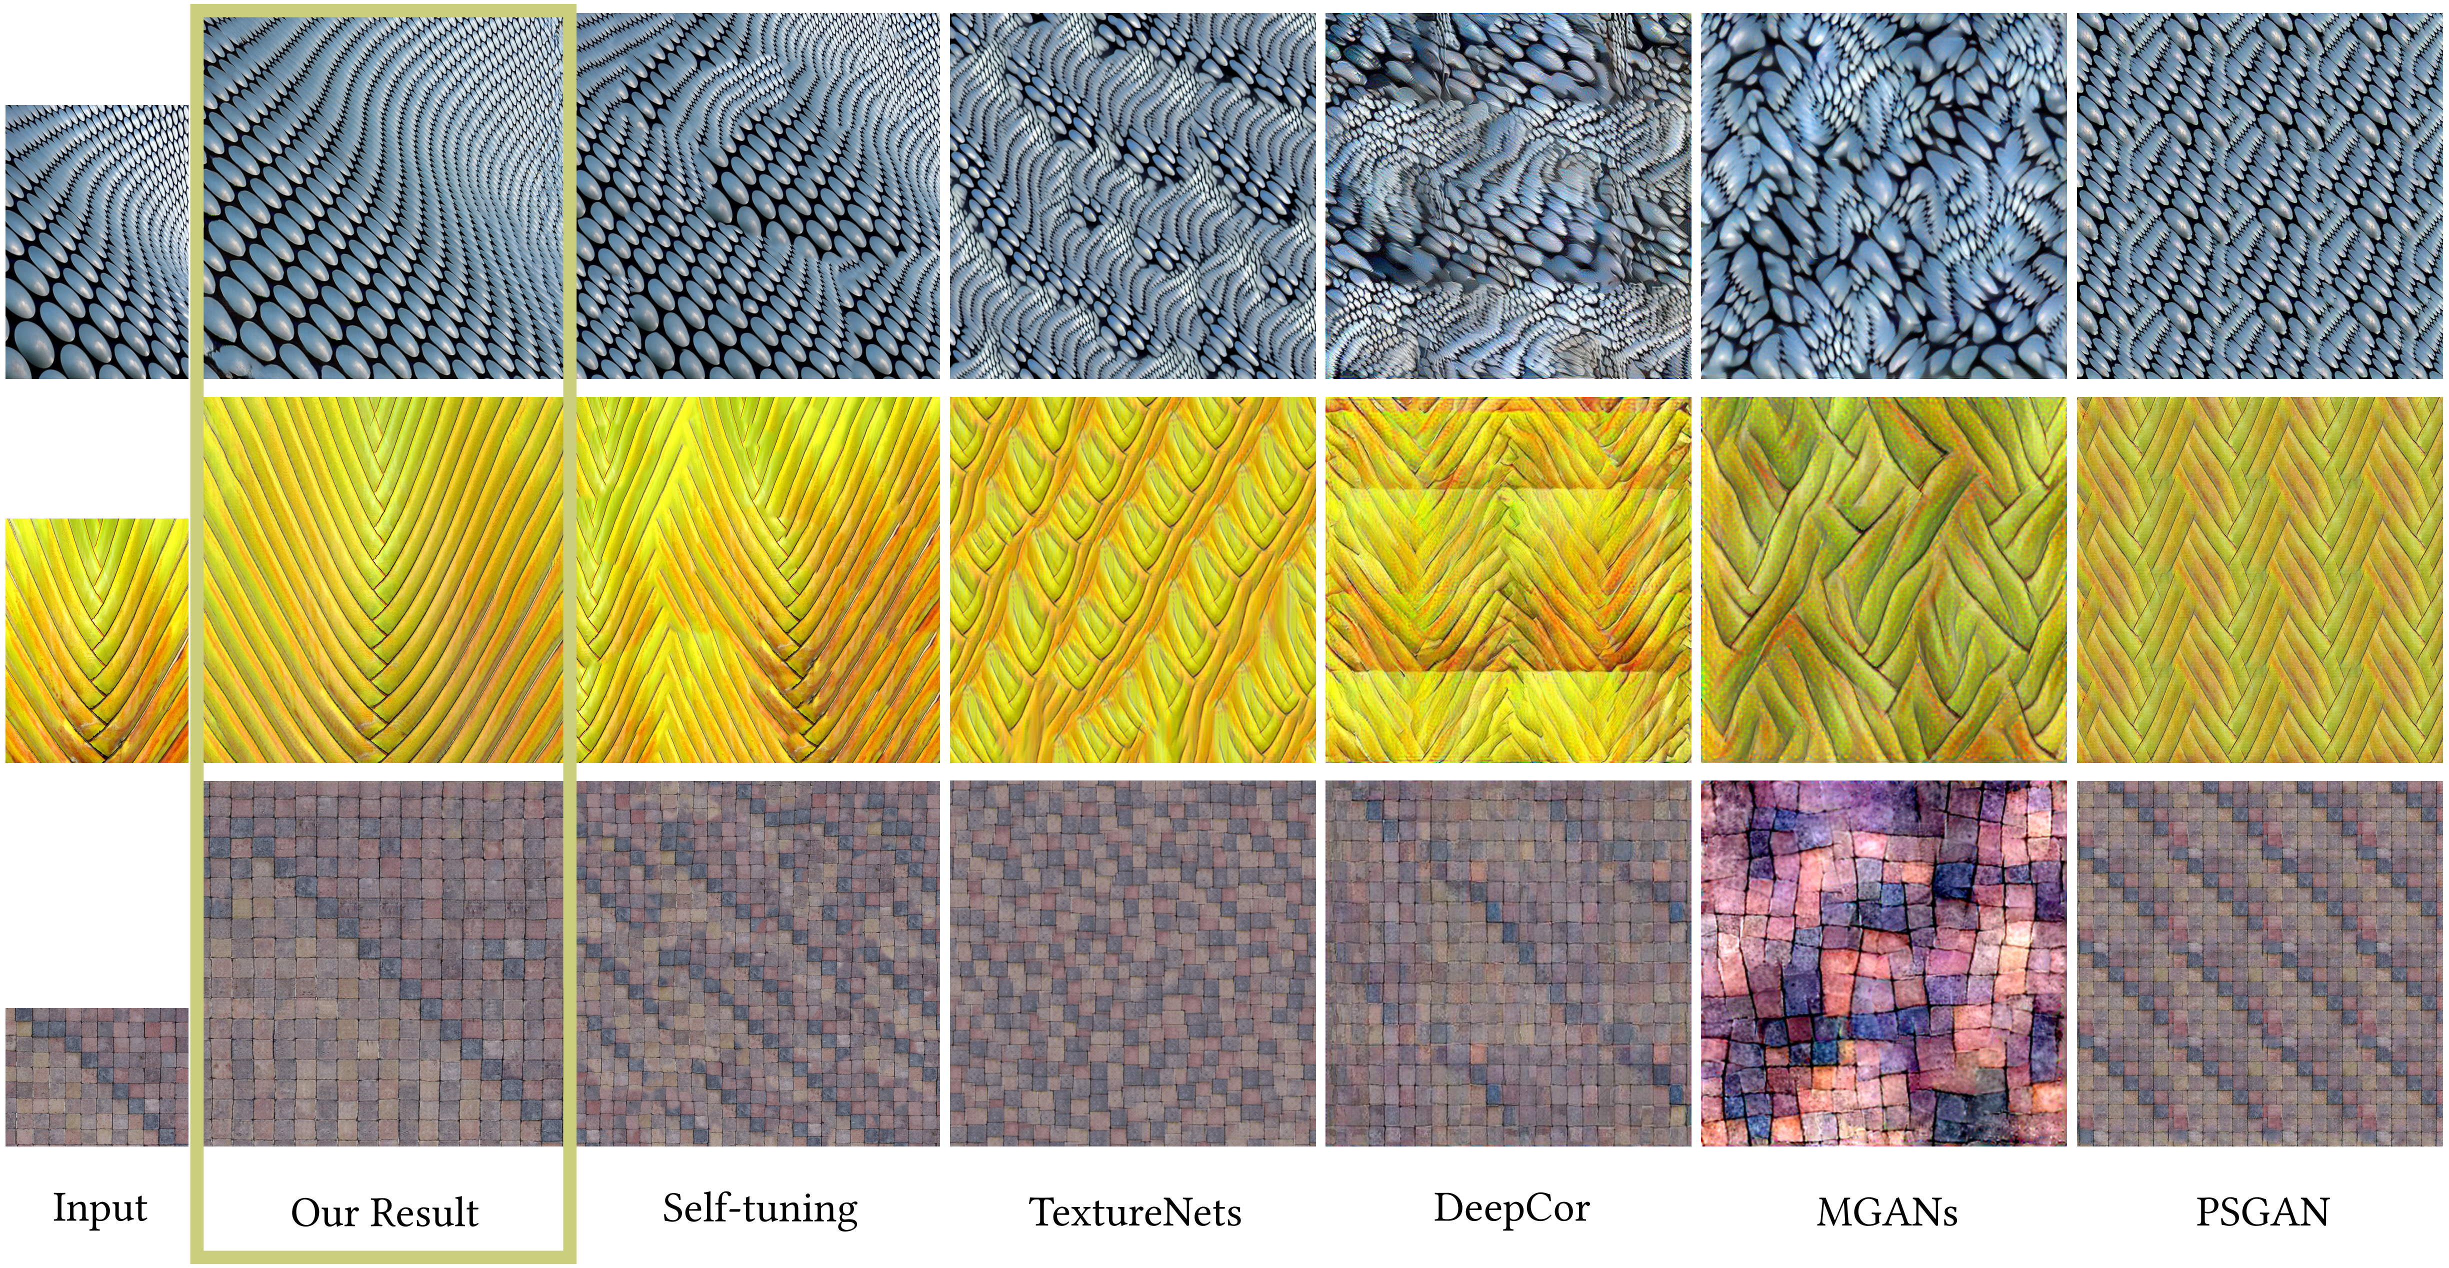

Experiments

➤ both non-stationary and stationary textures

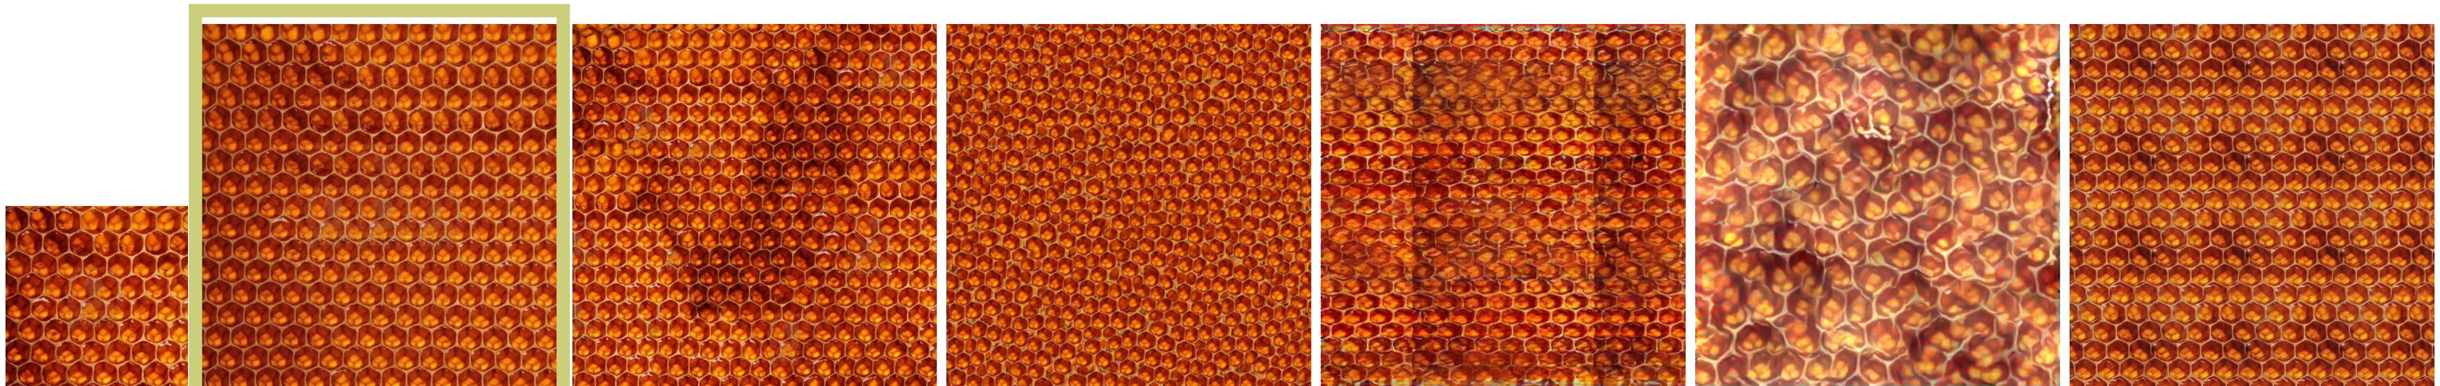

Input	Our Result	Self-tuning	TextureNets	DeepCor	MGANs	PSGAN
	5 hour	*20 min*	*1 hour*	*8 hour*		*12 hour*

Experiments

➤ Diversification

➤ Taking two random crops as input.

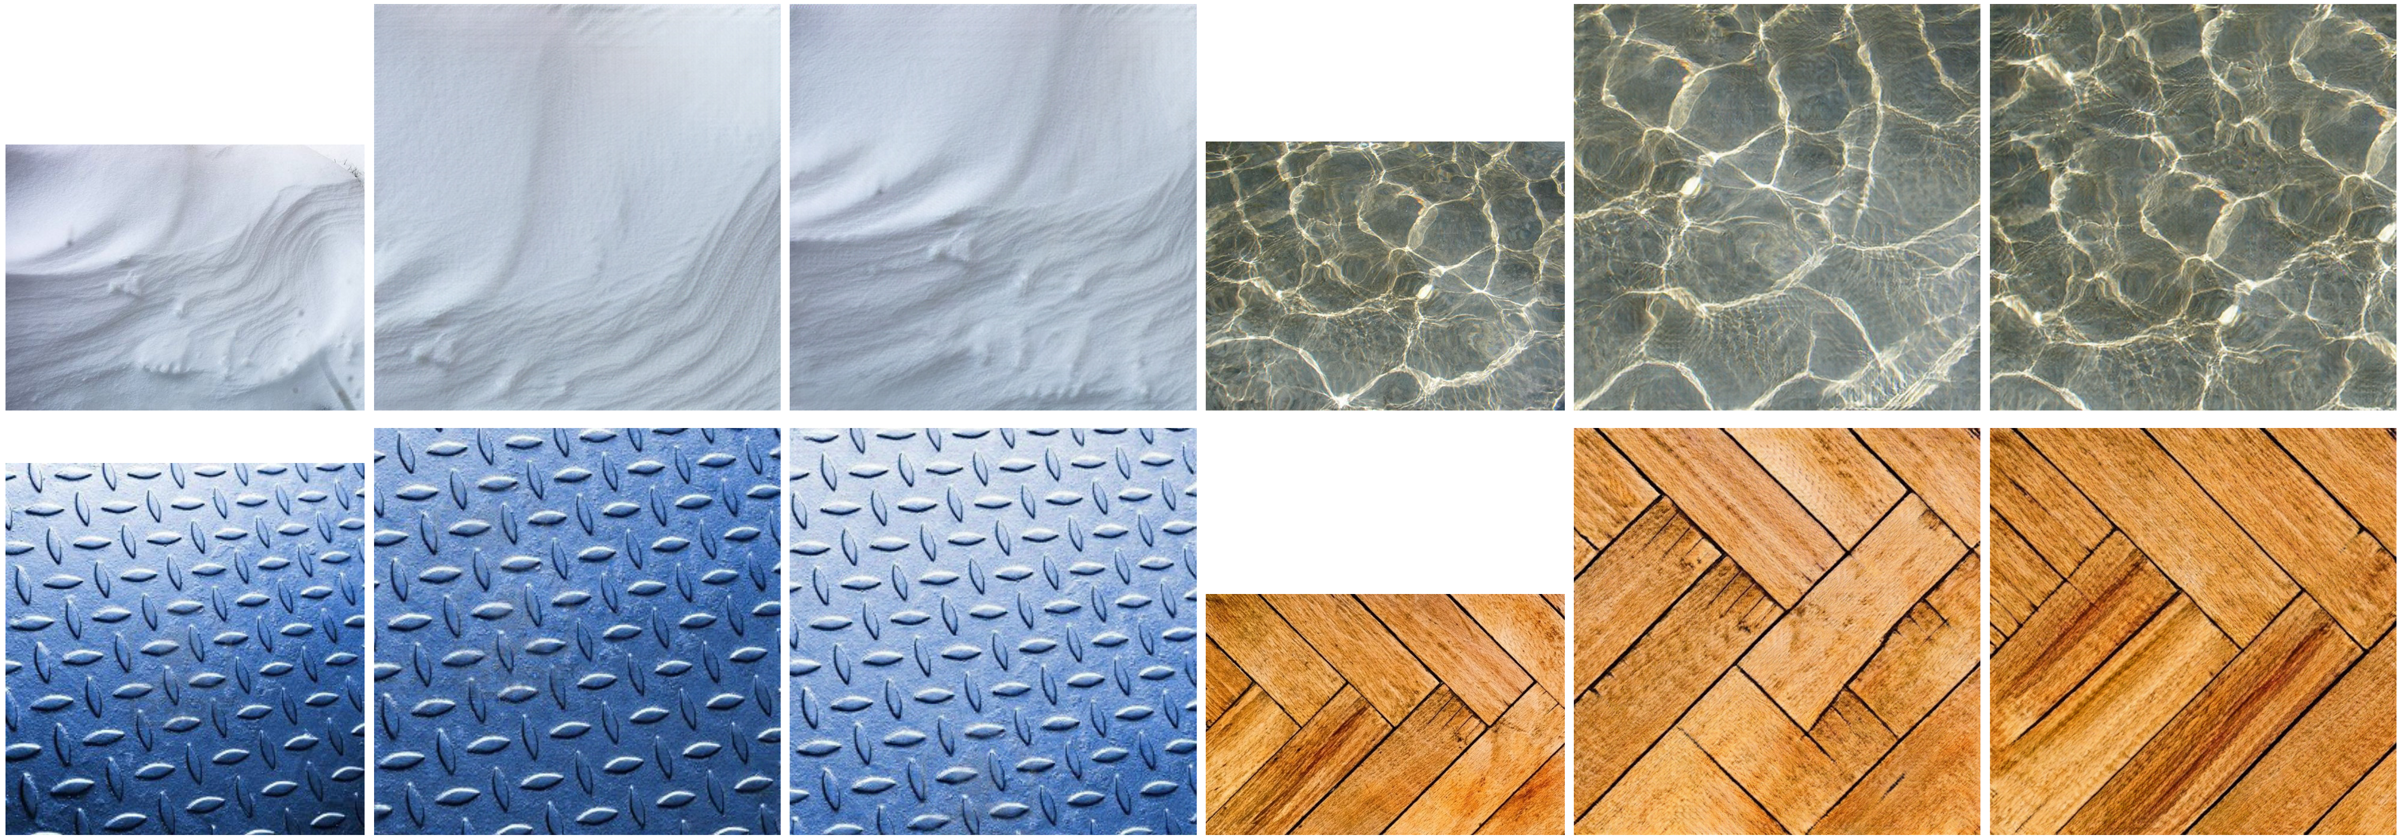

Input *Output1* *Output2* *Input* *Output1* *Output2*

Experiments

- Diversification
- reshuffling or perturbing the source texture

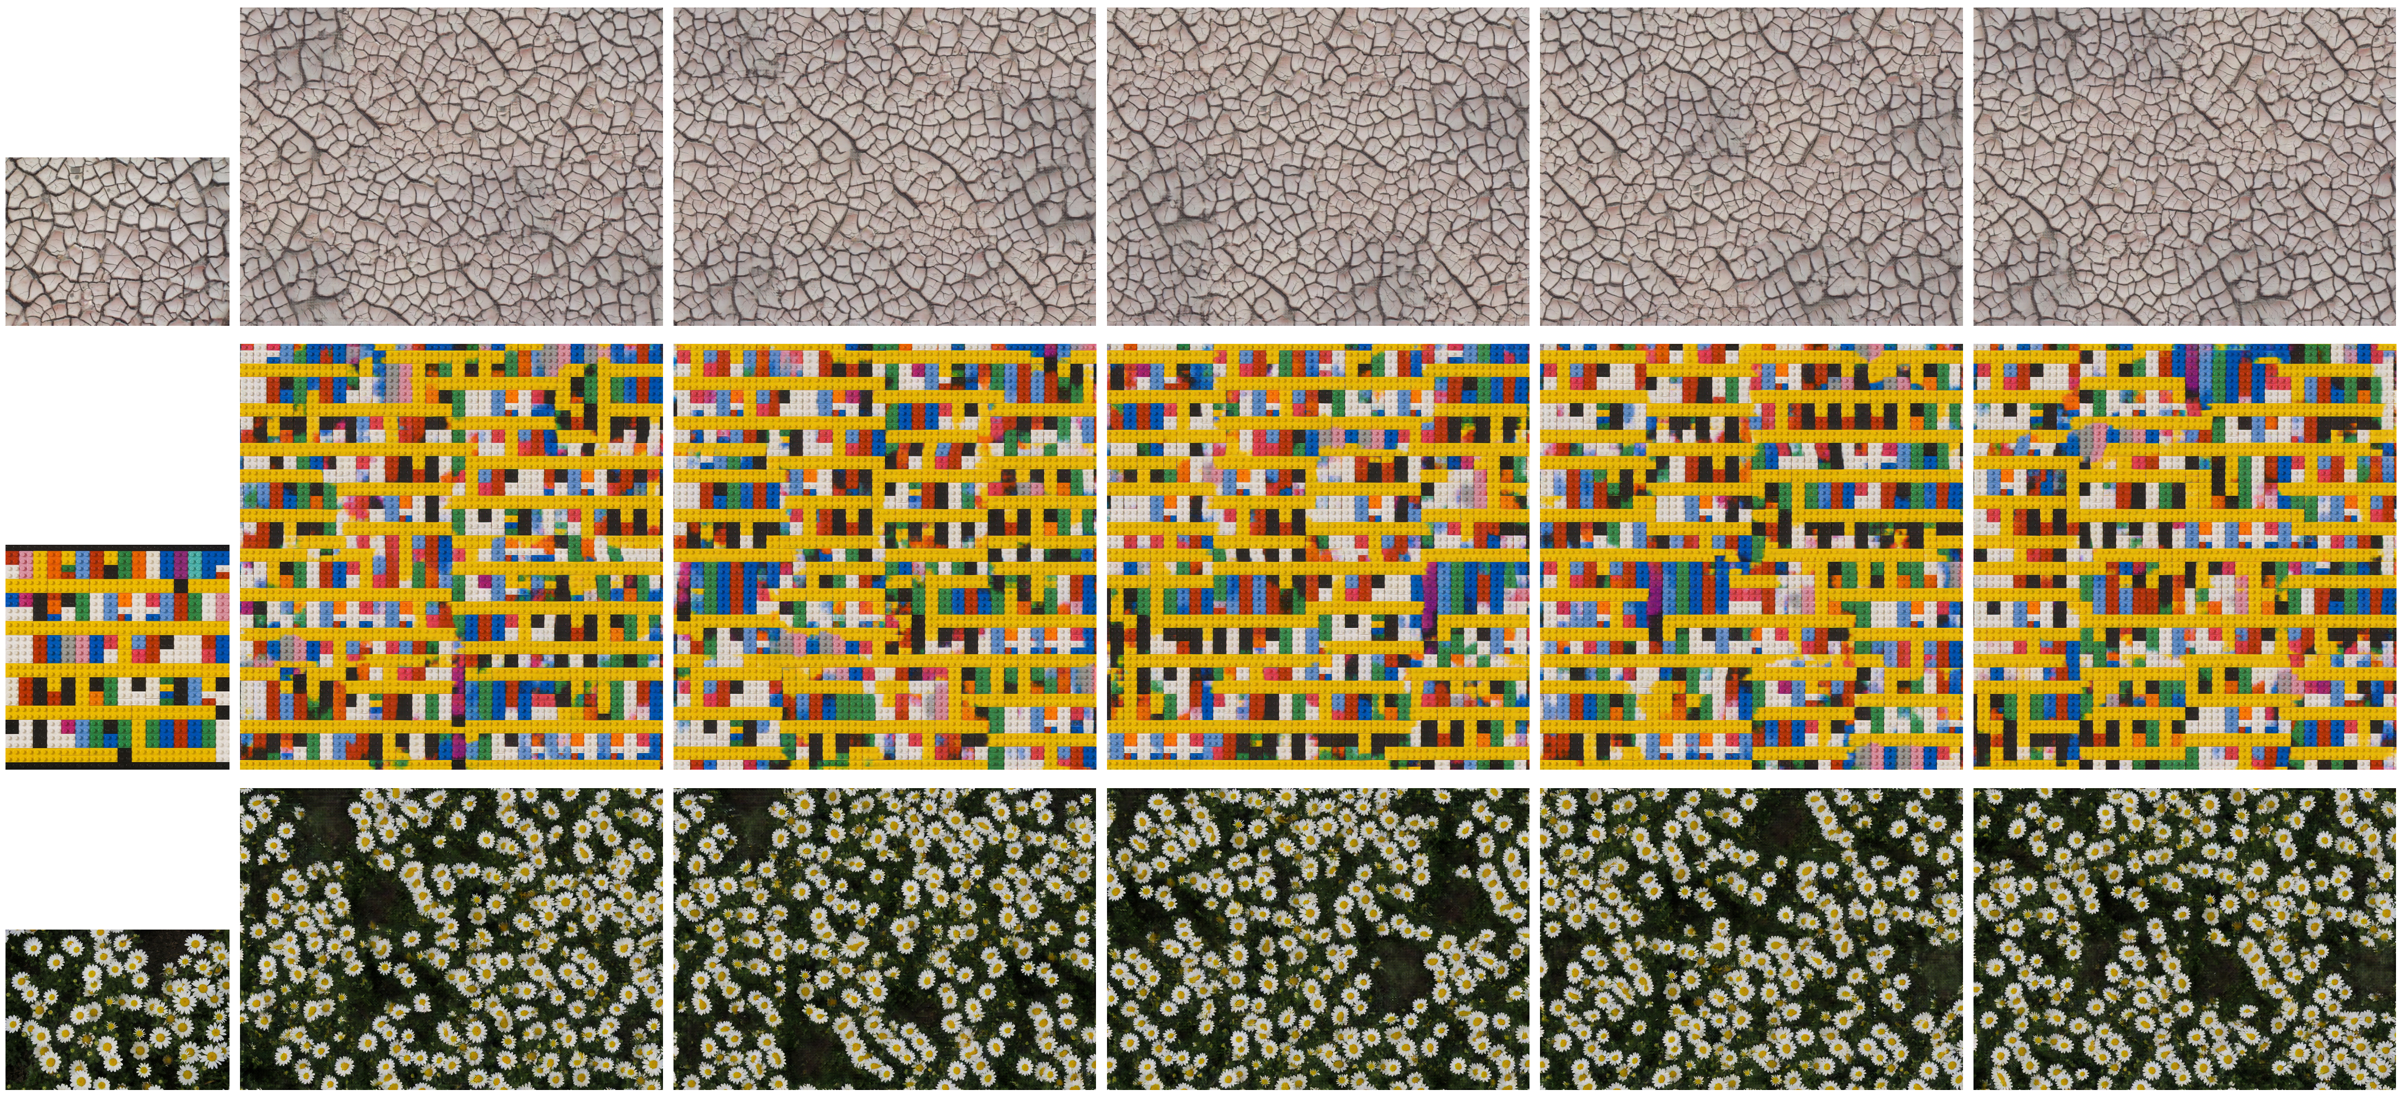

Input *Output1* *Output2* *Output3* *Output4* *Output5*

Experiments

- Diversification
- Perlin noise.

Input

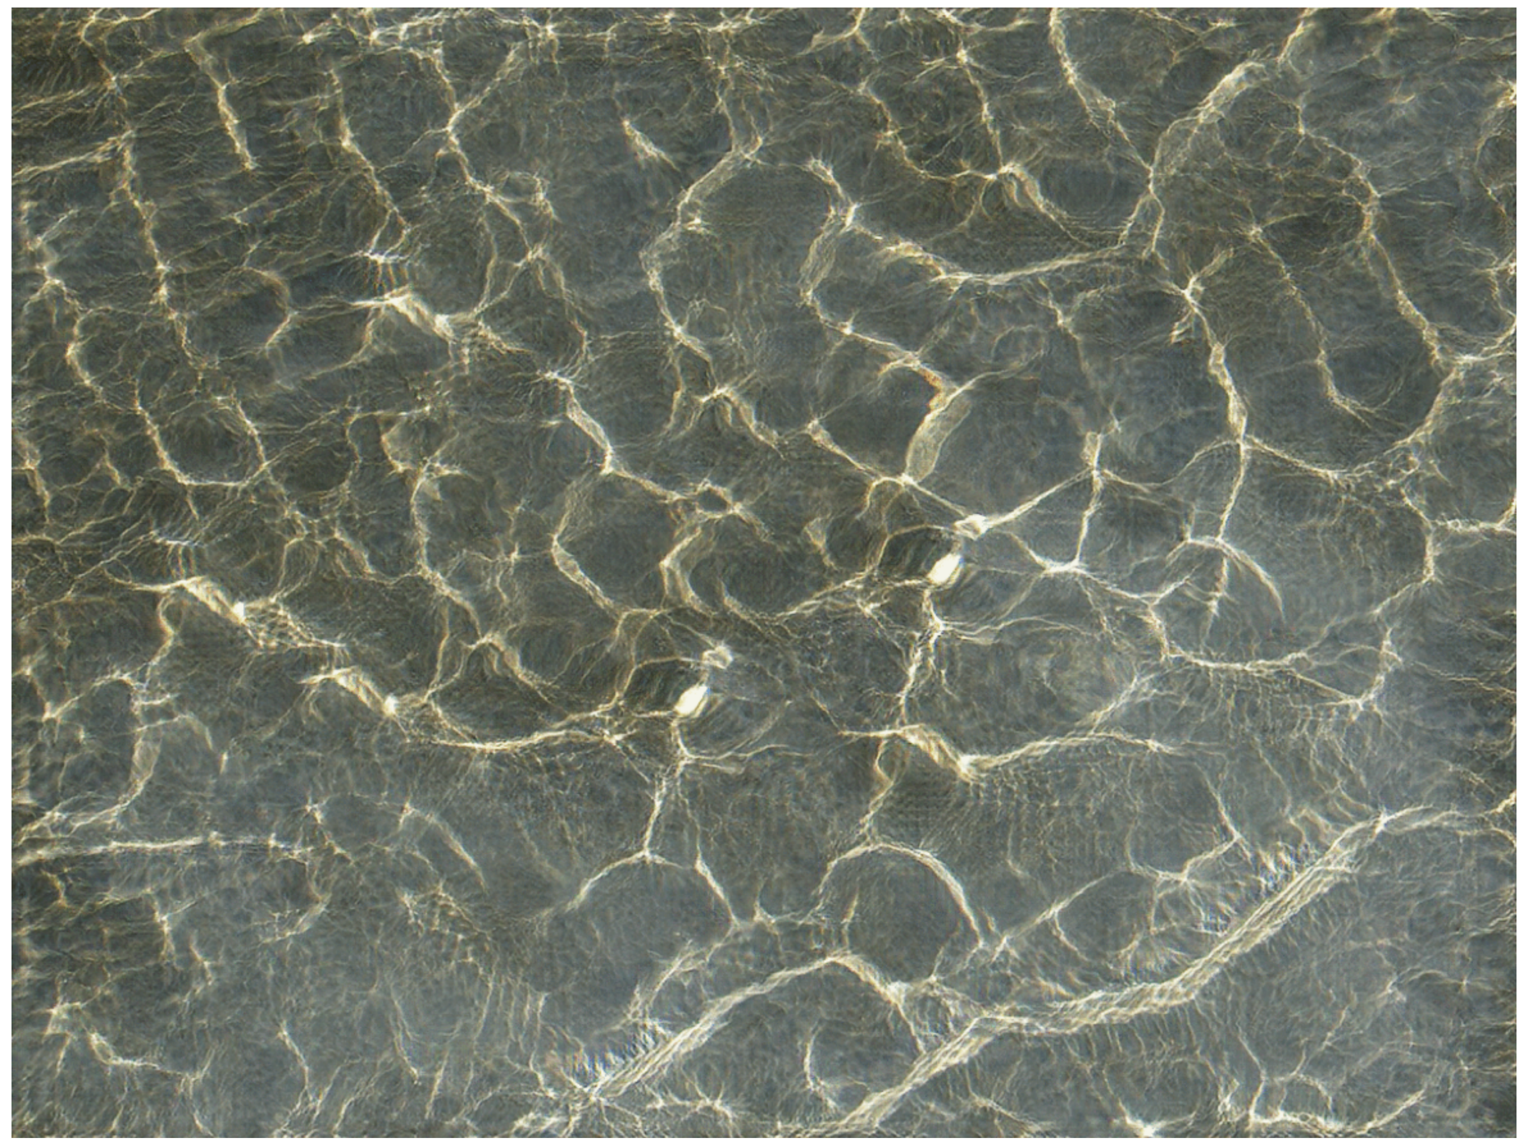

Output

Experiments

- ➤ Ablation studies
- ➤ loss function

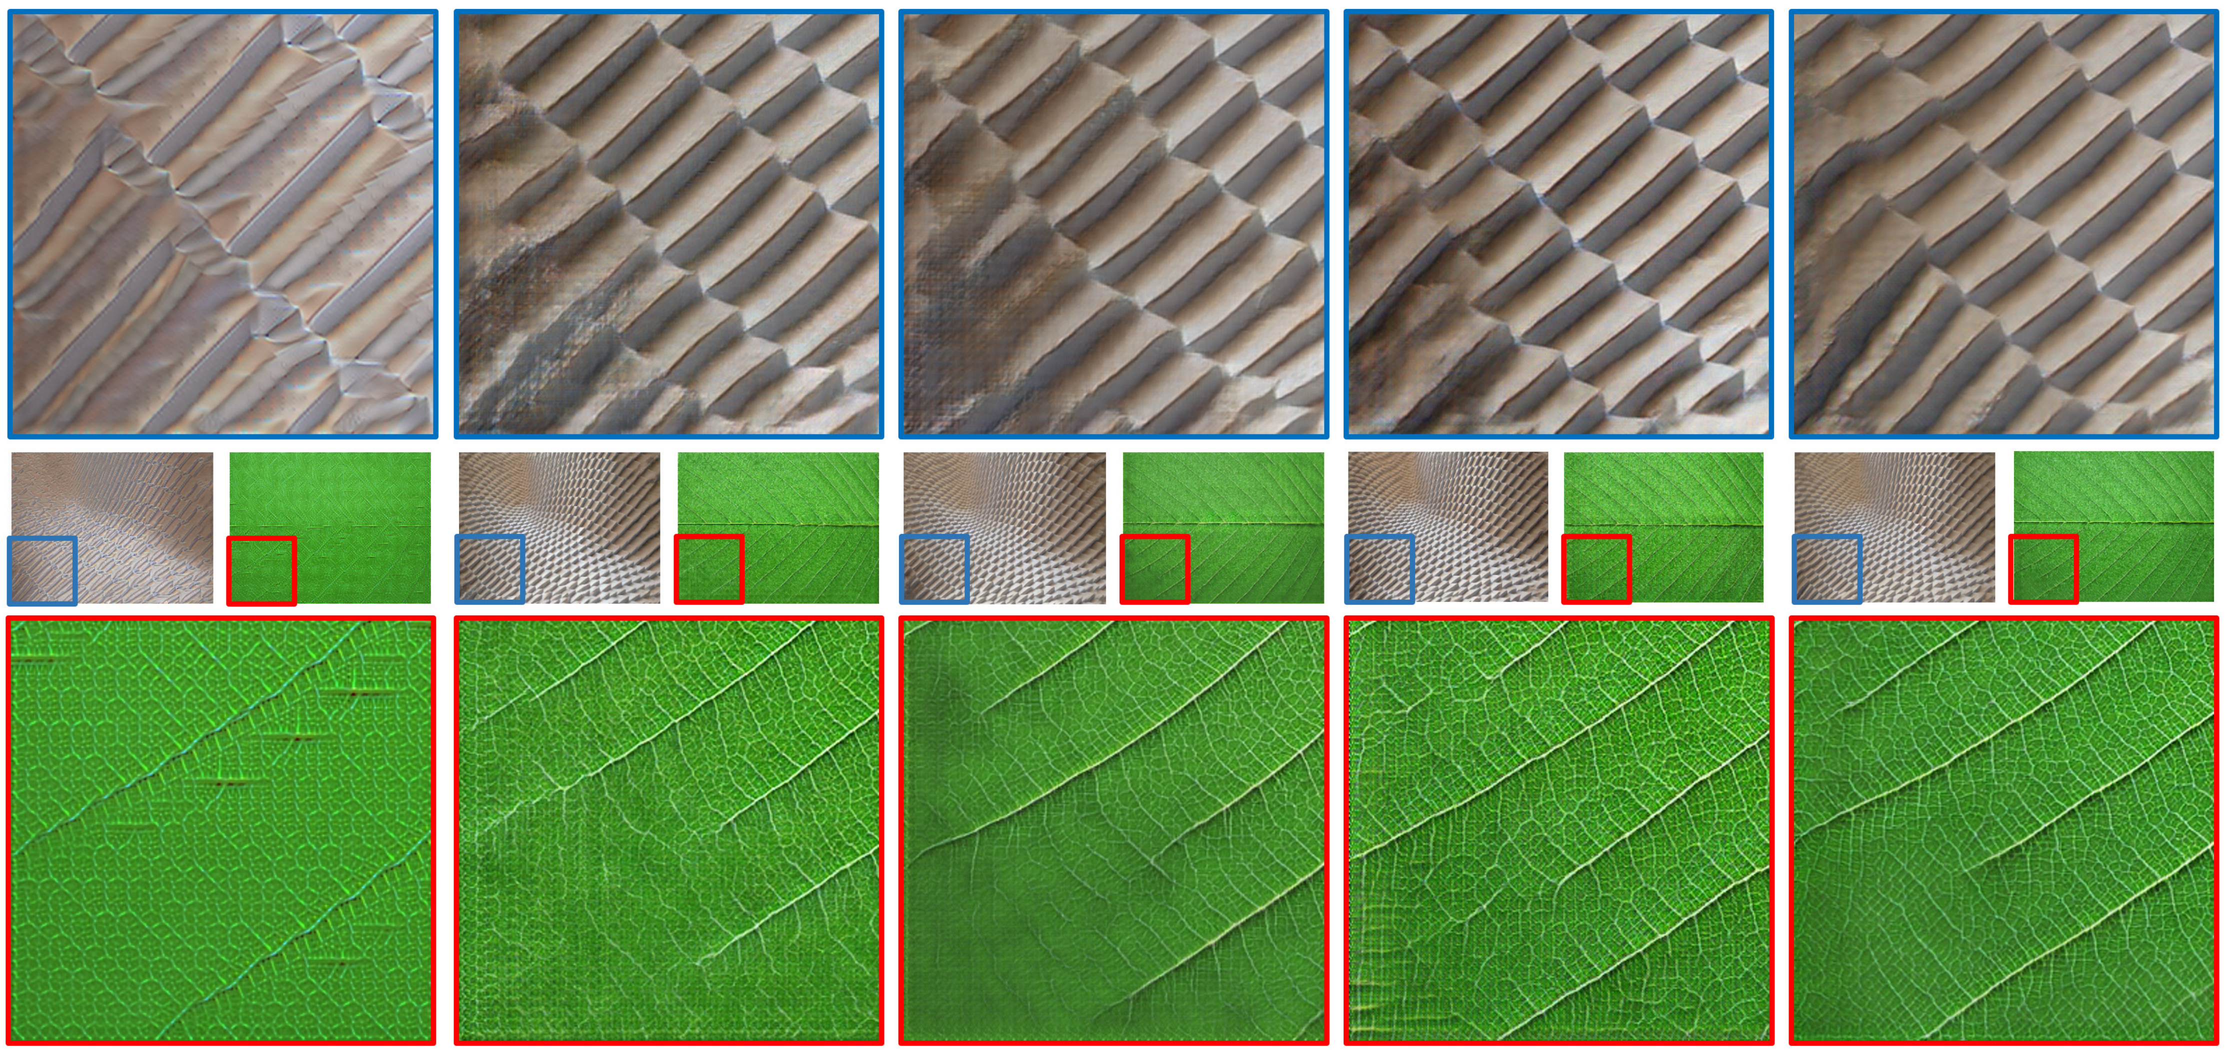

$\mathcal{L}_{L_1} + \lambda_2/\lambda_1\mathcal{L}_{style}$ | $\mathcal{L}_{adv}$ only | $\mathcal{L}_{adv} + \lambda_1\mathcal{L}_{L_1}$ | $\mathcal{L}_{adv} + \lambda_2\mathcal{L}_{style}$ | $\mathcal{L}_{adv} + \lambda_1\mathcal{L}_{L_1} + \lambda_2\mathcal{L}_{style}$

Experiments

- Ablation studies
  - loss function

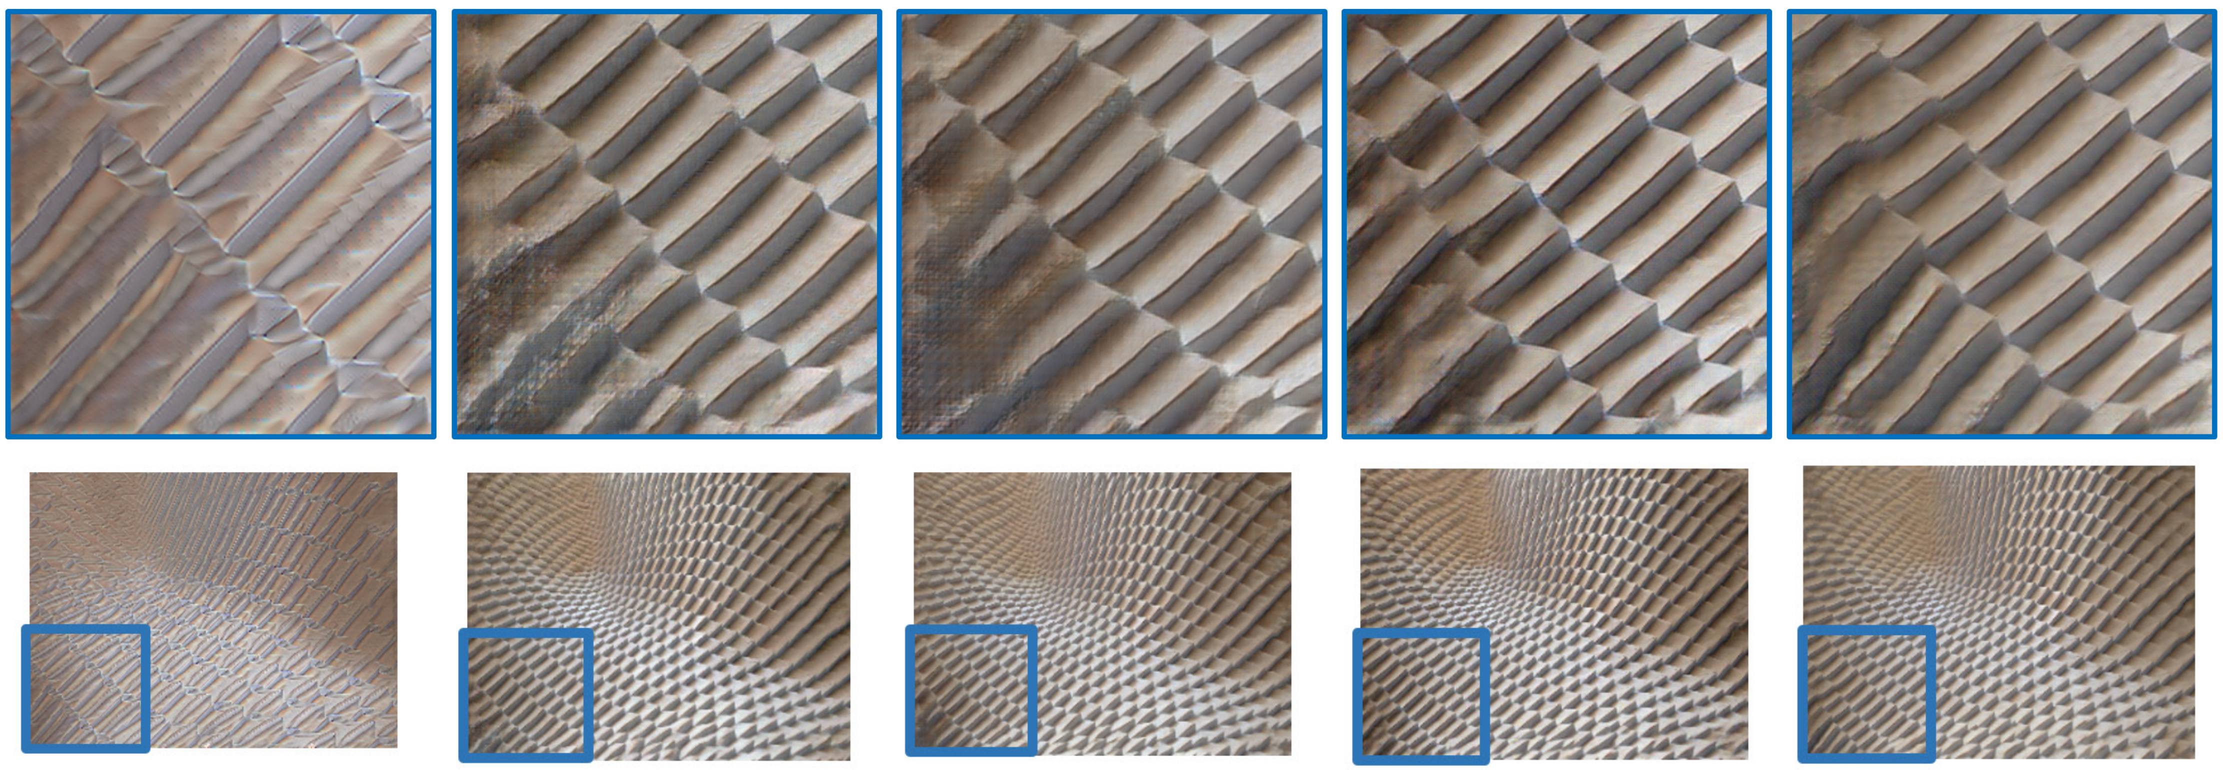

$\mathcal{L}_{L_1} + \lambda_2/\lambda_1 \mathcal{L}_{style}$ | $\mathcal{L}_{adv}$ only | $\mathcal{L}_{adv} + \lambda_1 \mathcal{L}_{L_1}$ | $\mathcal{L}_{adv} + \lambda_2 \mathcal{L}_{style}$ | $\mathcal{L}_{adv} + \lambda_1 \mathcal{L}_{L_1} + \lambda_2 \mathcal{L}_{style}$

Experiments

- Ablation studies
- Discriminator patch size

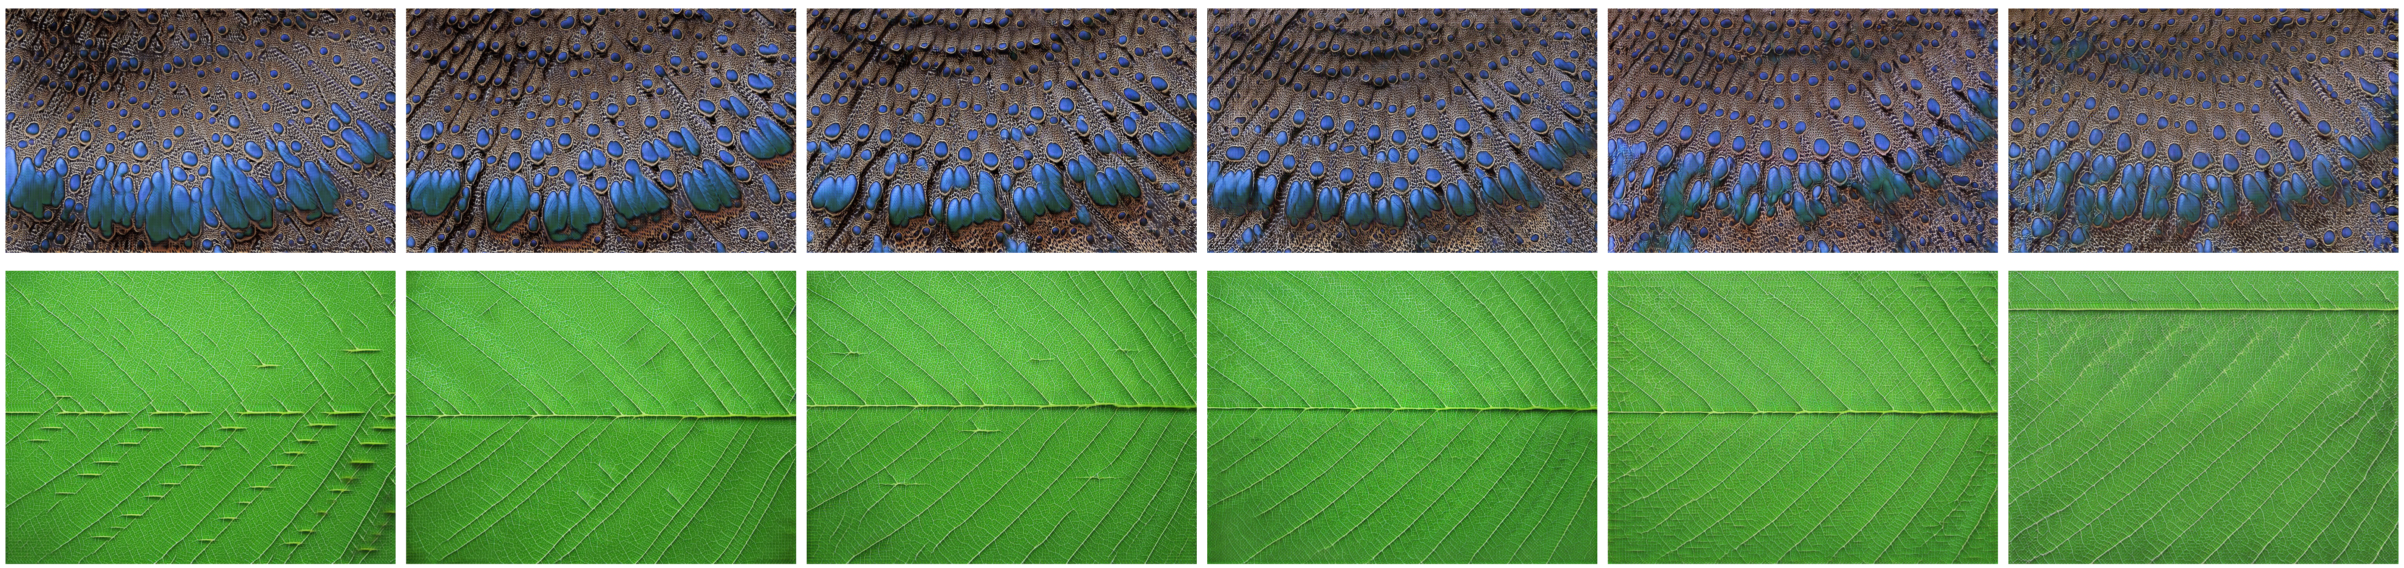

16×16 34×34 70×70 142×142 286×286 574×574

Experiments

- Synthesis stability
- Randomly crop the result and feed it as input again, repeated 4 times, without any re-training or fine-tuning the generator.

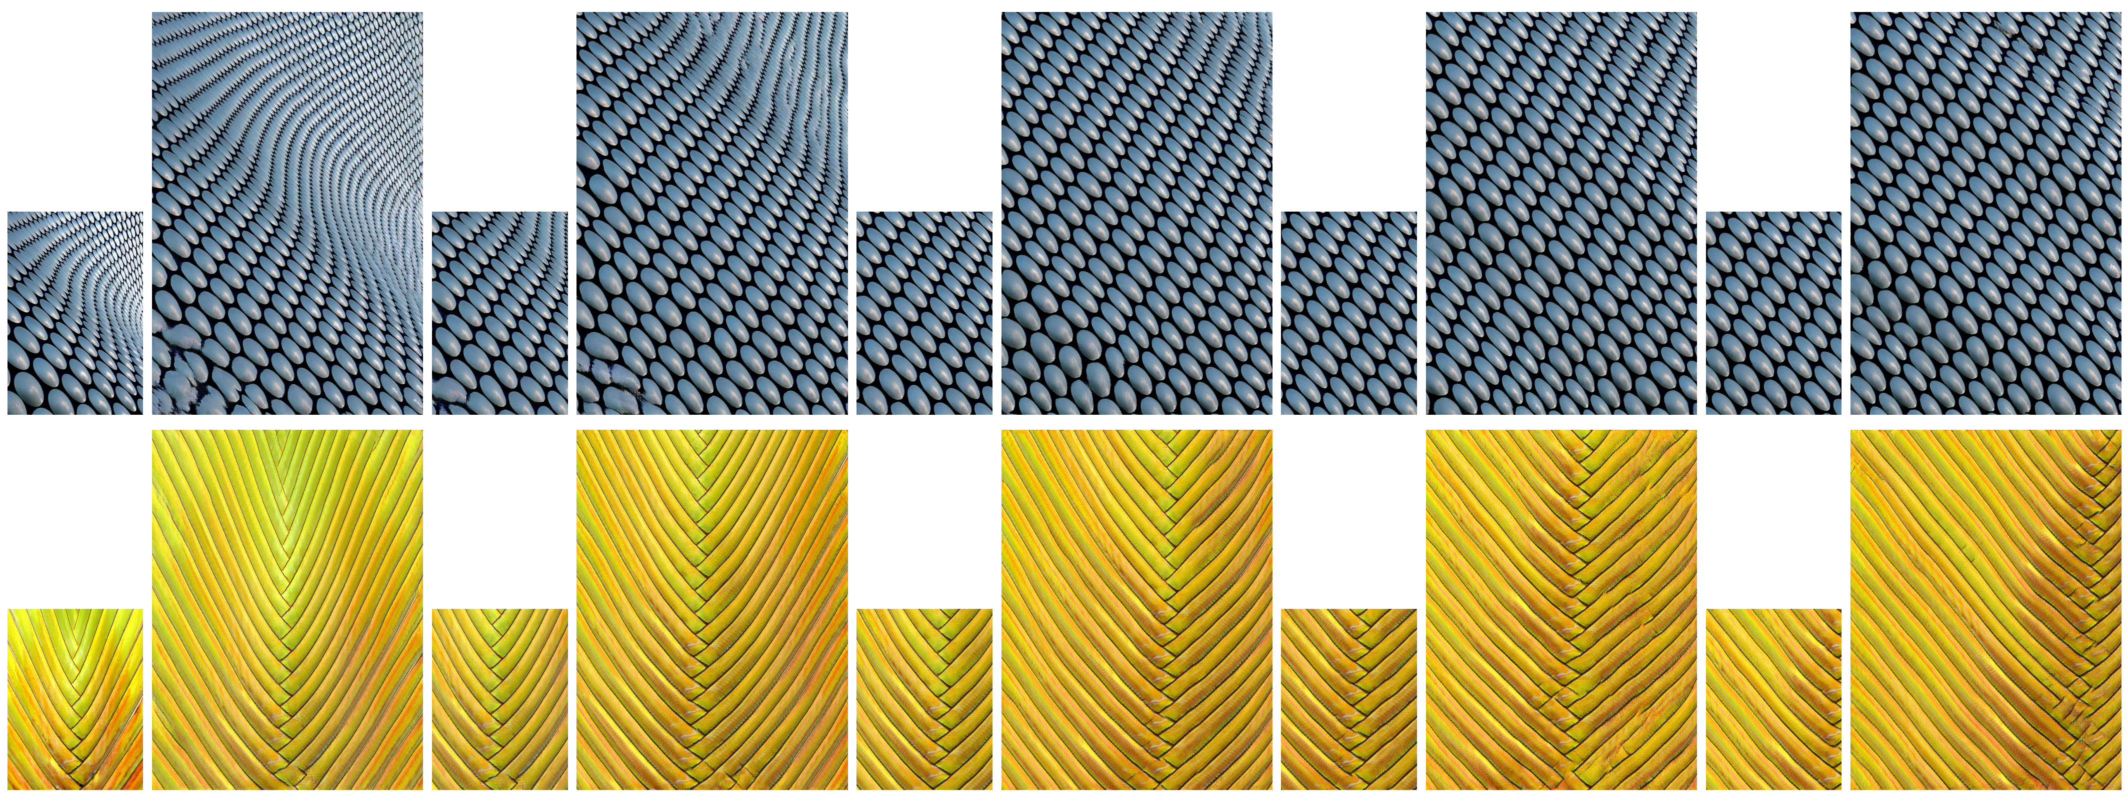

- The final result (rightmost column) is still very sharp and natural looking.

Experiments

➤ Extreme expansion, 32x (expanding 5 times)

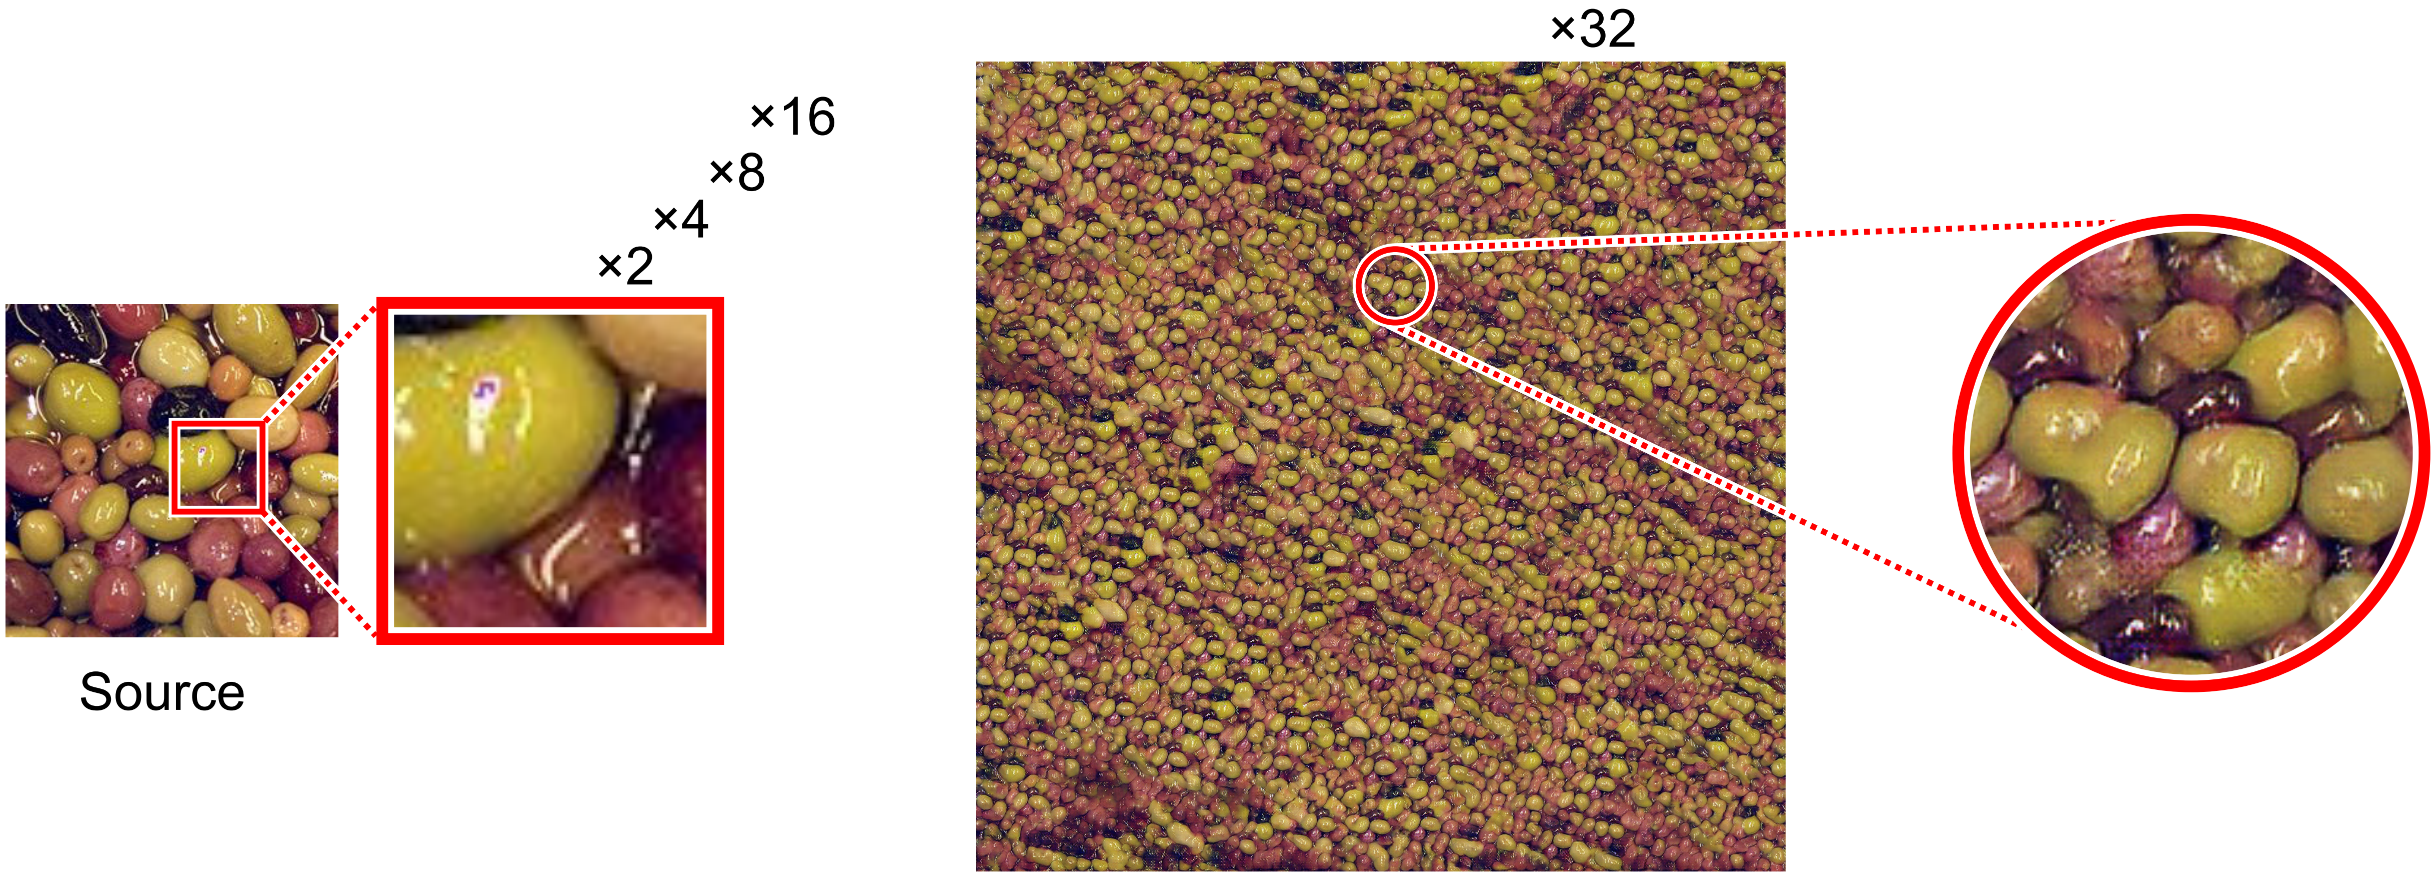

Experiments

- Not suited for artistic style transfer
- How to transfer: feeding the guiding image as input to a trained generator

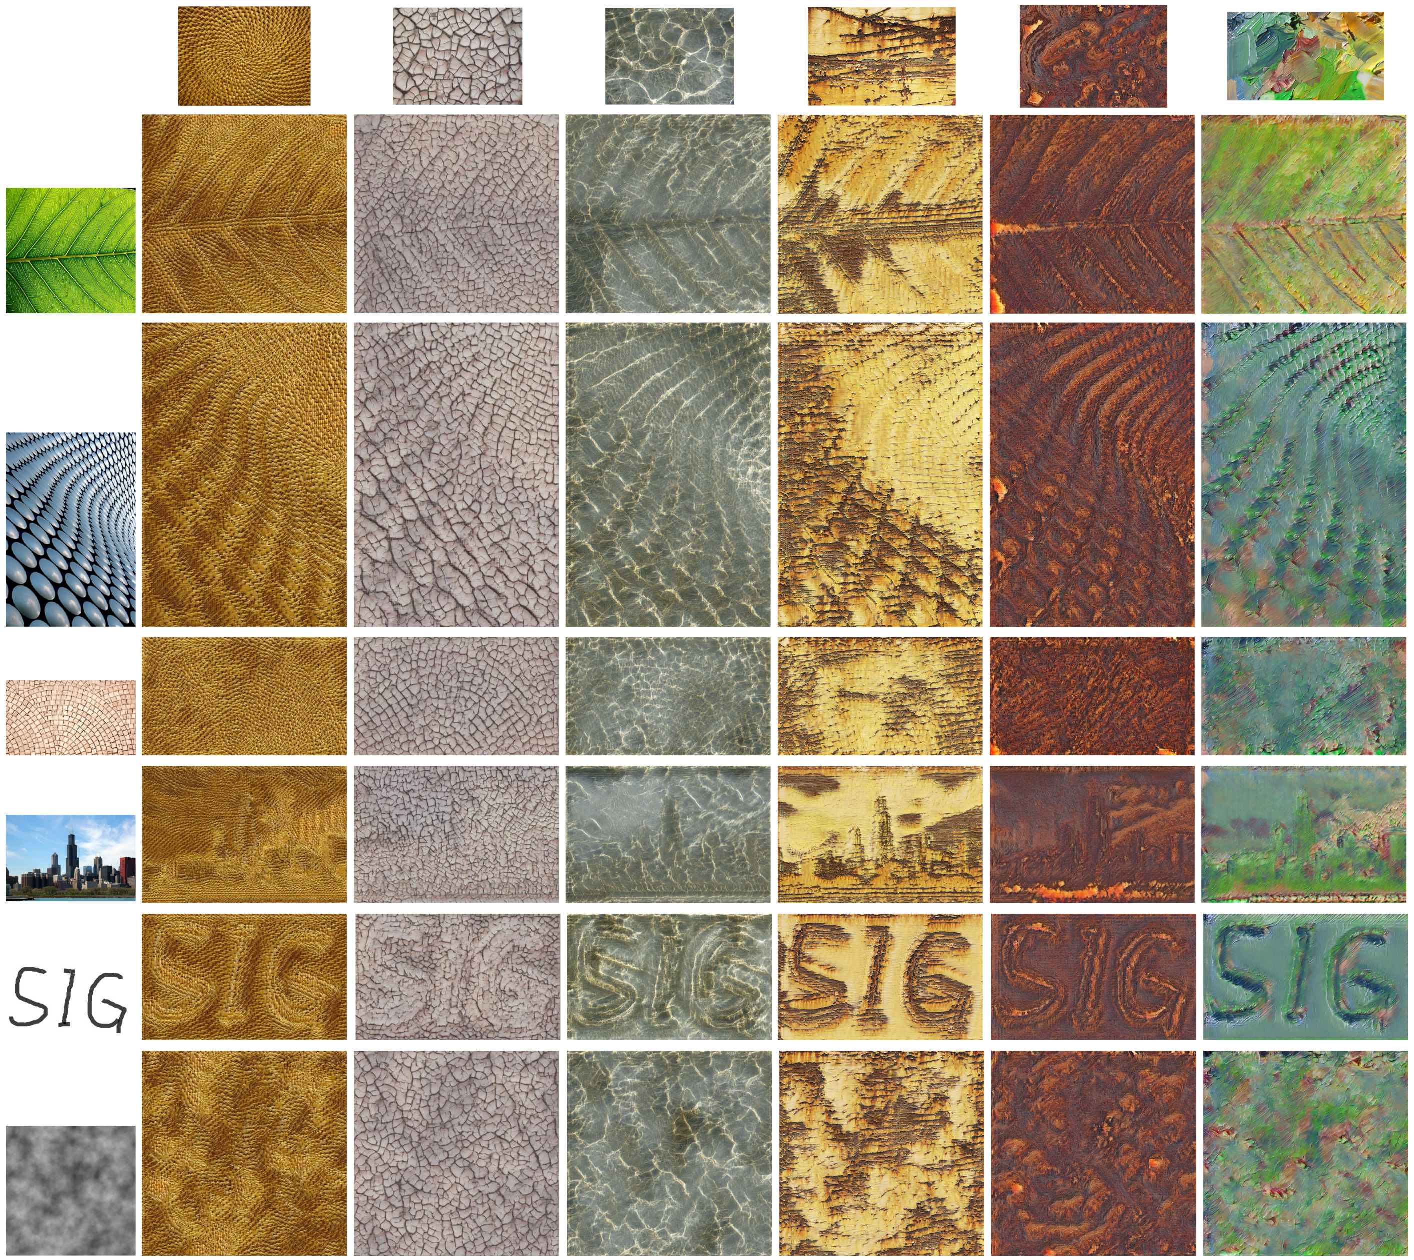

Experiments

- Limitation

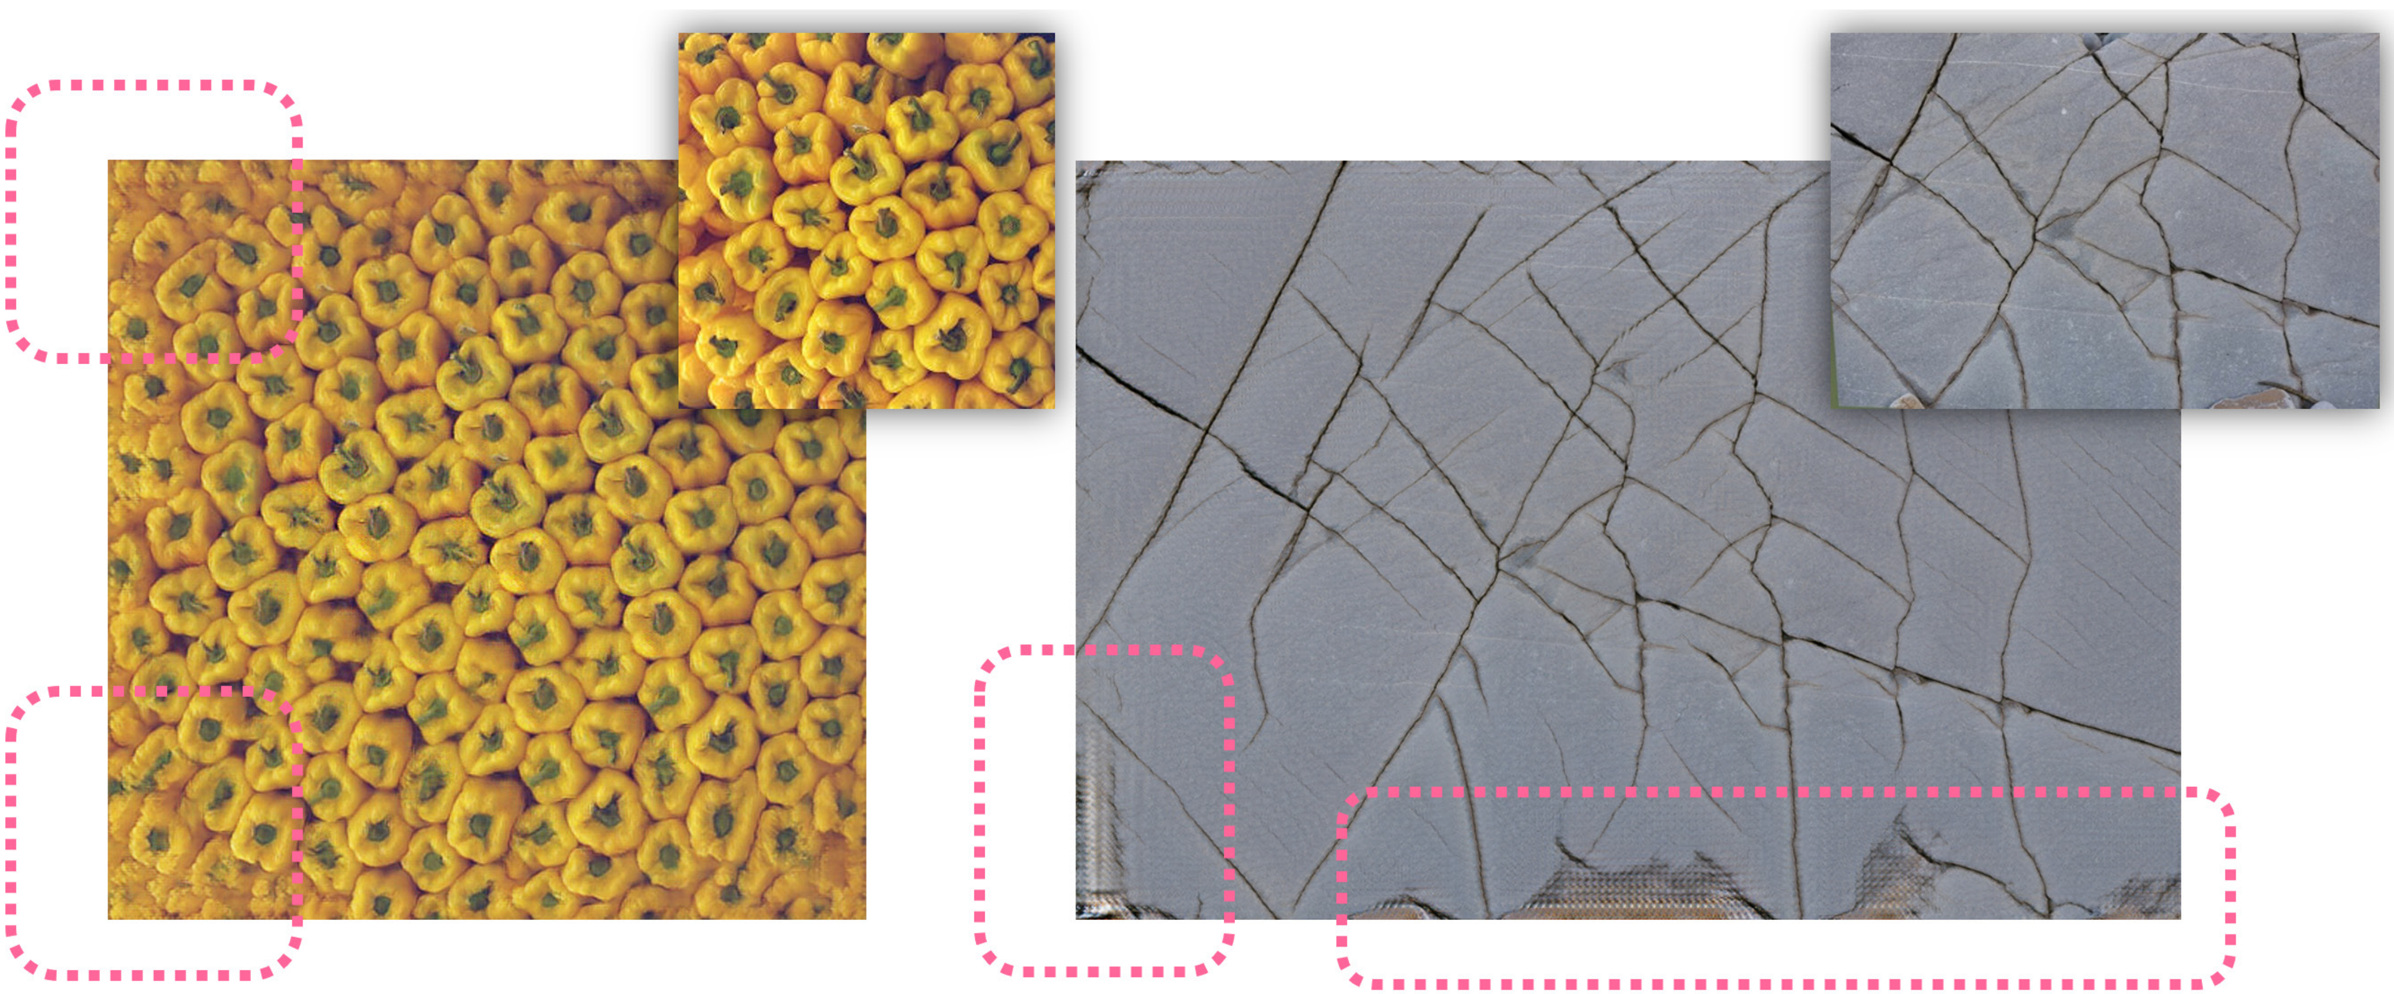

Fig. 16. Artifacts in the border and corner regions.

Experiments

➤ Limitation

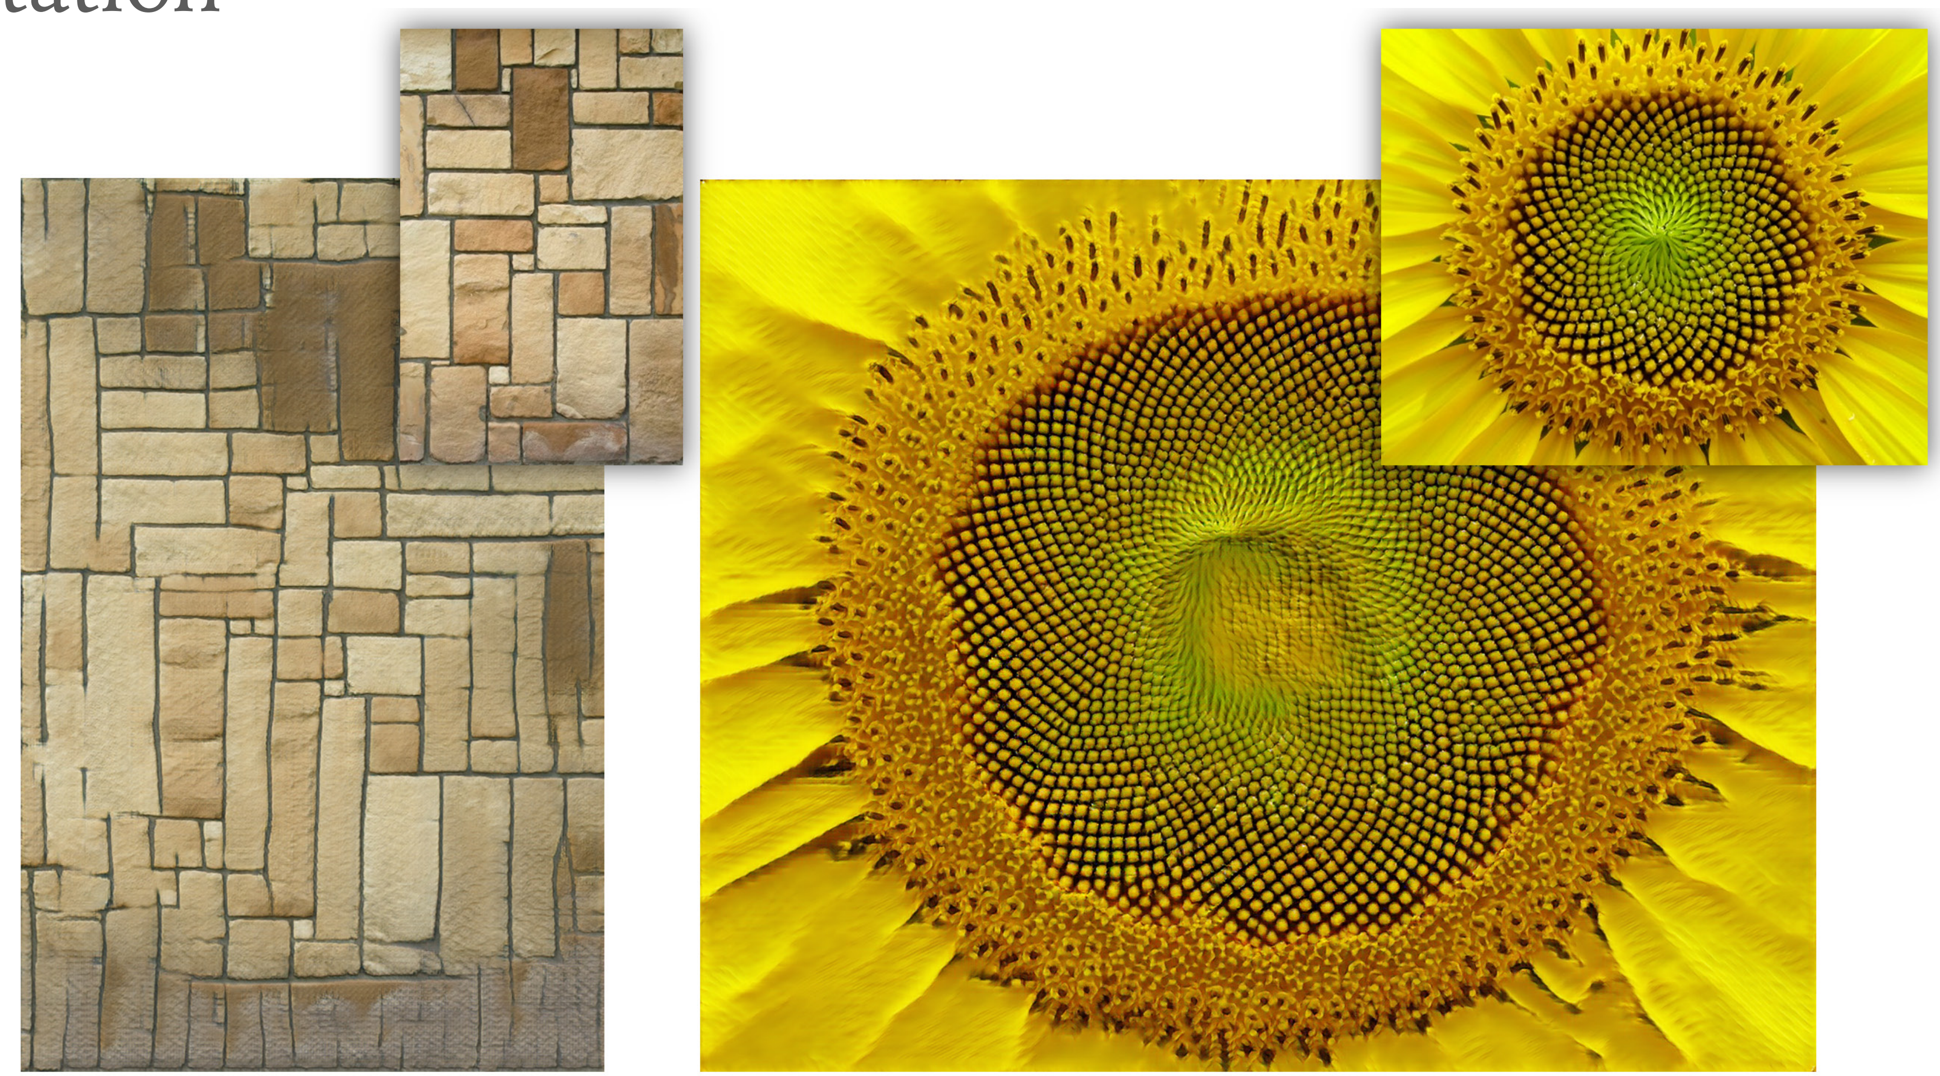

Fig. 17. Failure cases of our method. For the stone tiles texture, our method failed to learn its large scale structure (left). While for the sunflower, our method failed to reproduce the singularity at the center (right).