How Much Deep Learning does Neural Style Transfer Really Need? An Ablation Study (Supplementary Material)

Len Du
Australian National University
Len.Du@anu.edu.au

This supplementary file is to convince the readers that the results in the main paper is general to different combinations of content/style images. So we simply reproduce the original diagrams in the paper with different content and style images. Images are taken from [2].

References

[1] K. He, Y. Wang, and J. Hopcroft. A powerful generative model using random weights for the deep image representation. In *Advances in Neural Information Processing Systems*, pages 631–639, 2016.

[2] Y. Jing, Y. Yang, Z. Feng, J. Ye, Y. Yu, and M. Song. Neural style transfer: A review. *IEEE Transactions on Visualization and Computer Graphics*, 2019.

(a) Content Image (b) Style Image (c) Baseline (d) $\beta/\alpha = 10^2$ (e) Removing the content loss

Figure 1: Removing the content loss

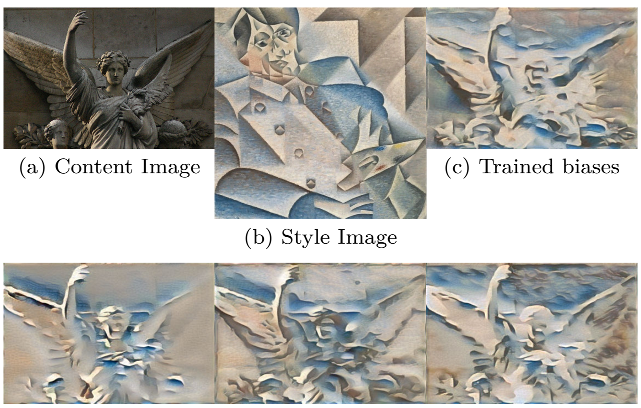

(a) Content Image (b) Style Image (c) Trained biases

(d) All zero biases (e) Fixed biases (0.5) (f) Reordered biases

Figure 2: Varing the biases while keeping trained weights

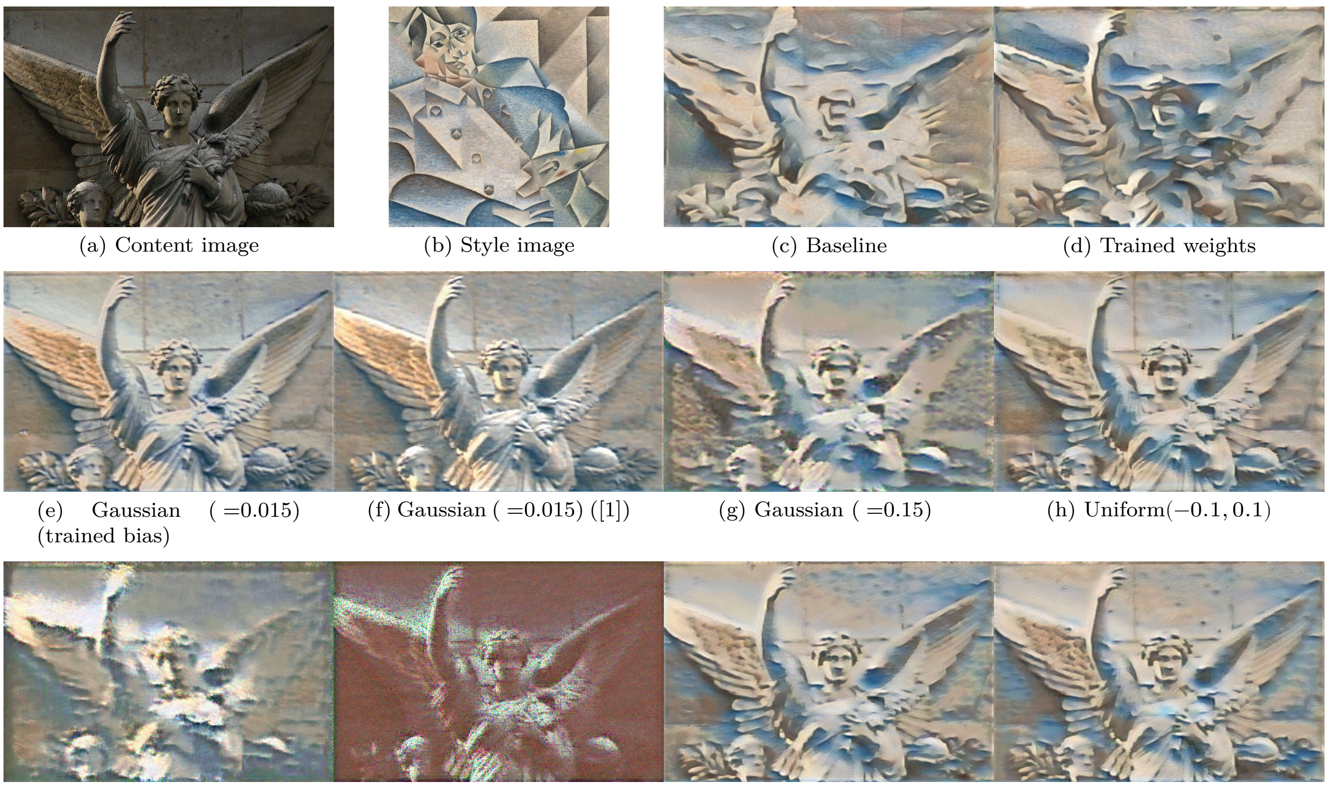

(a) Content image
(b) Style image
(c) Baseline
(d) Trained weights
(e) Gaussian (=0.015) (trained bias)
(f) Gaussian (=0.015) ([1])
(g) Gaussian (=0.15)
(h) Uniform$(-0.1, 0.1)$
(i) Uniform$(-0.5, 0.5)$
(j) Uniform$(-1, 1)$
(k) Shuffled weights
(l) Shuffled weights (trained biases)

Figure 3: Continuously or densely distributed weights

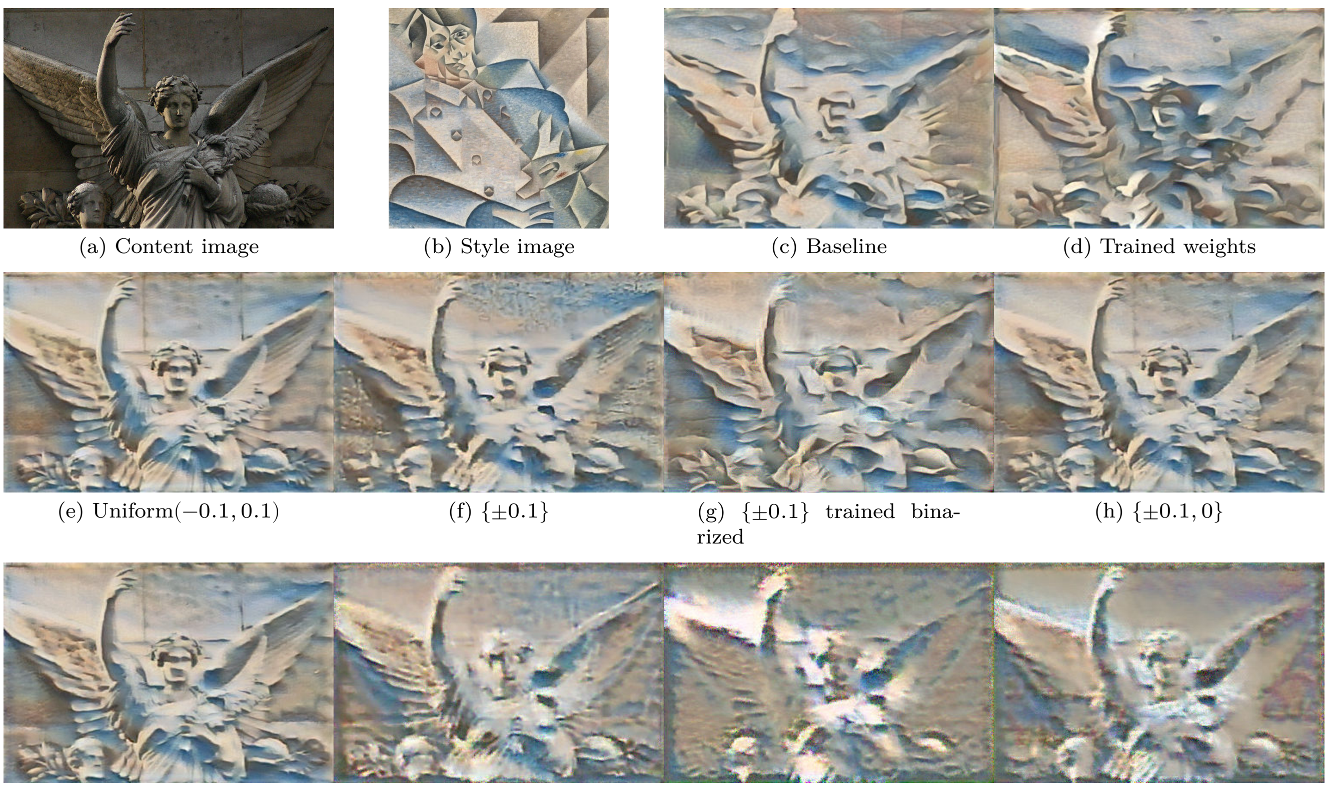

(a) Content image (b) Style image (c) Baseline (d) Trained weights

(e) Uniform$(-0.1, 0.1)$ (f) $\{\pm 0.1\}$ (g) $\{\pm 0.1\}$ trained binarized (h) $\{\pm 0.1, 0\}$

(i) $\{\pm 0.1, 0, 0\}$ (j) $\{\pm 0.2\}$ (k) $\{\pm 0.4\}$ (l) $\{\pm 0.4, \pm 0.2, \pm 0.1, 0\}$

Figure 4: Symmetric discrete weights with fixed biases (0.5)

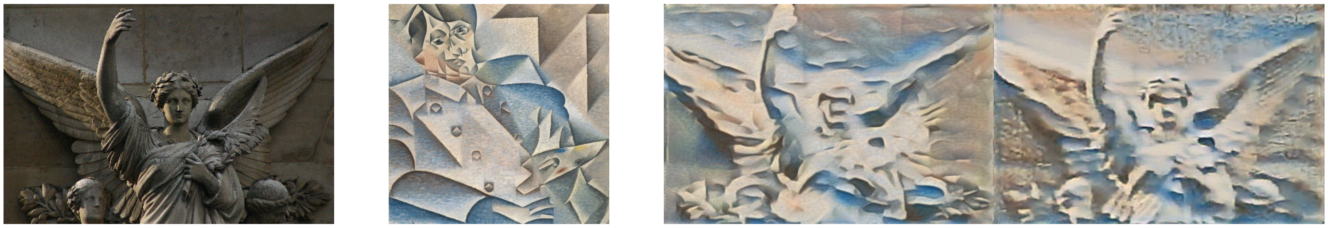

(a) Content image (b) Style image (c) Baseline (d) $\{\pm 0.1\}$

(e) No ReLU no pooling no doubling (f) No ReLU (g) Neither pooling nor doubling (NPND) (h) With doubling but no pooling

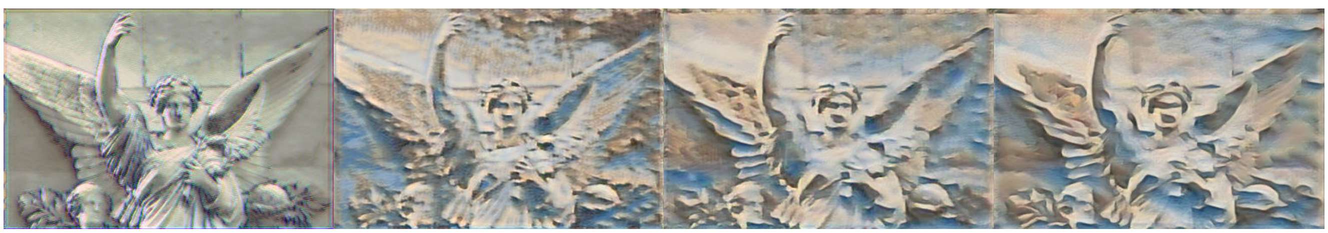

(i) NPND, 4 filters (j) NPND, 16 filters (k) NPND, 71 filters (l) NPND, 512 filters

Figure 5: Removal of structure

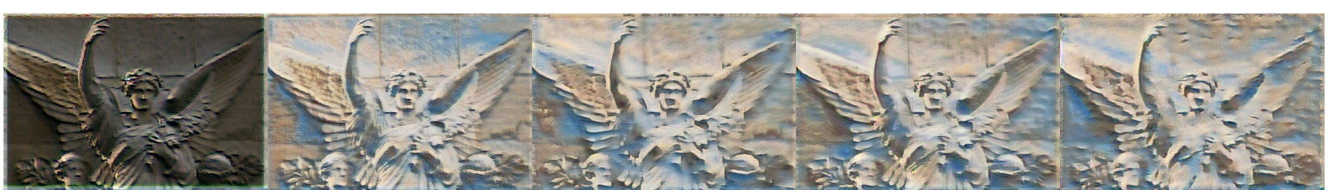

(a) Basic set 1, strategy 4 (32-channel) (b) Basic set 2, strategy 1 (18-channel) (c) Basic set 3, strategy 3 (64-channel) (d) 2×2 kernels, random, 64-channel (e) 3×3 kernels, random, 64-channel

Figure 6: Alternative first-layer convolution kernels

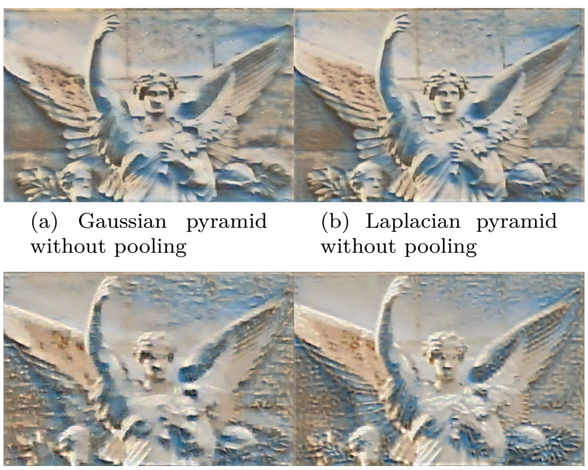

(a) Gaussian pyramid without pooling

(b) Laplacian pyramid without pooling

(c) Gaussian pyramid with pooling

(d) Laplacian pyramid with pooling

Figure 7: Using image pyramids instead of multi-layer nonlinearities

(a) Content Image (b) Style Image (c) Baseline (d) $\beta/\alpha = 10^2$ (e) Removing the content loss

Figure 8: Removing the content loss

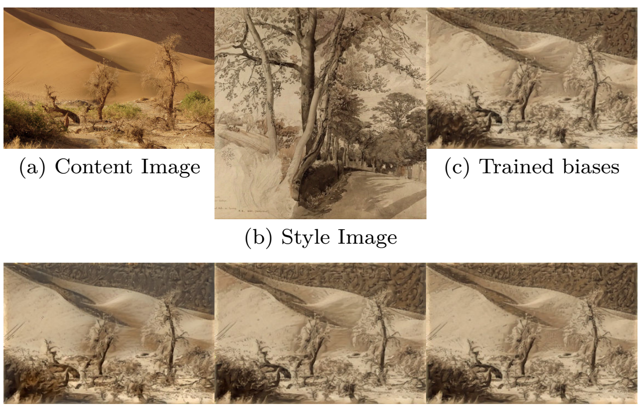

(a) Content Image (b) Style Image (c) Trained biases

(d) All zero biases (e) Fixed biases (0.5) (f) Reordered biases

Figure 9: Varing the biases while keeping trained weights

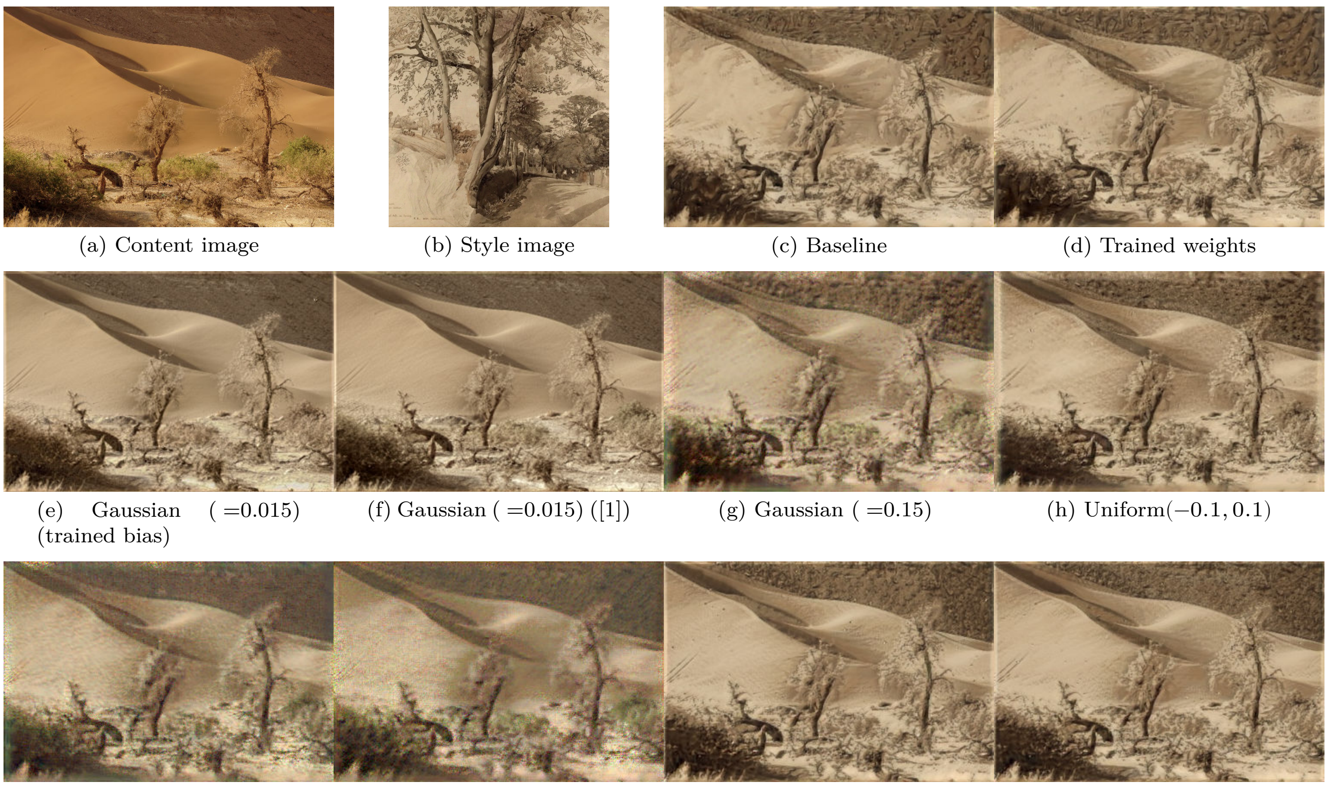

(a) Content image (b) Style image (c) Baseline (d) Trained weights

(e) Gaussian (=0.015) (trained bias) (f) Gaussian (=0.015) ([1]) (g) Gaussian (=0.15) (h) Uniform$(-0.1, 0.1)$

(i) Uniform$(-0.5, 0.5)$ (j) Uniform$(-1, 1)$ (k) Shuffled weights (l) Shuffled weights (trained biases)

Figure 10: Continuously or densely distributed weights

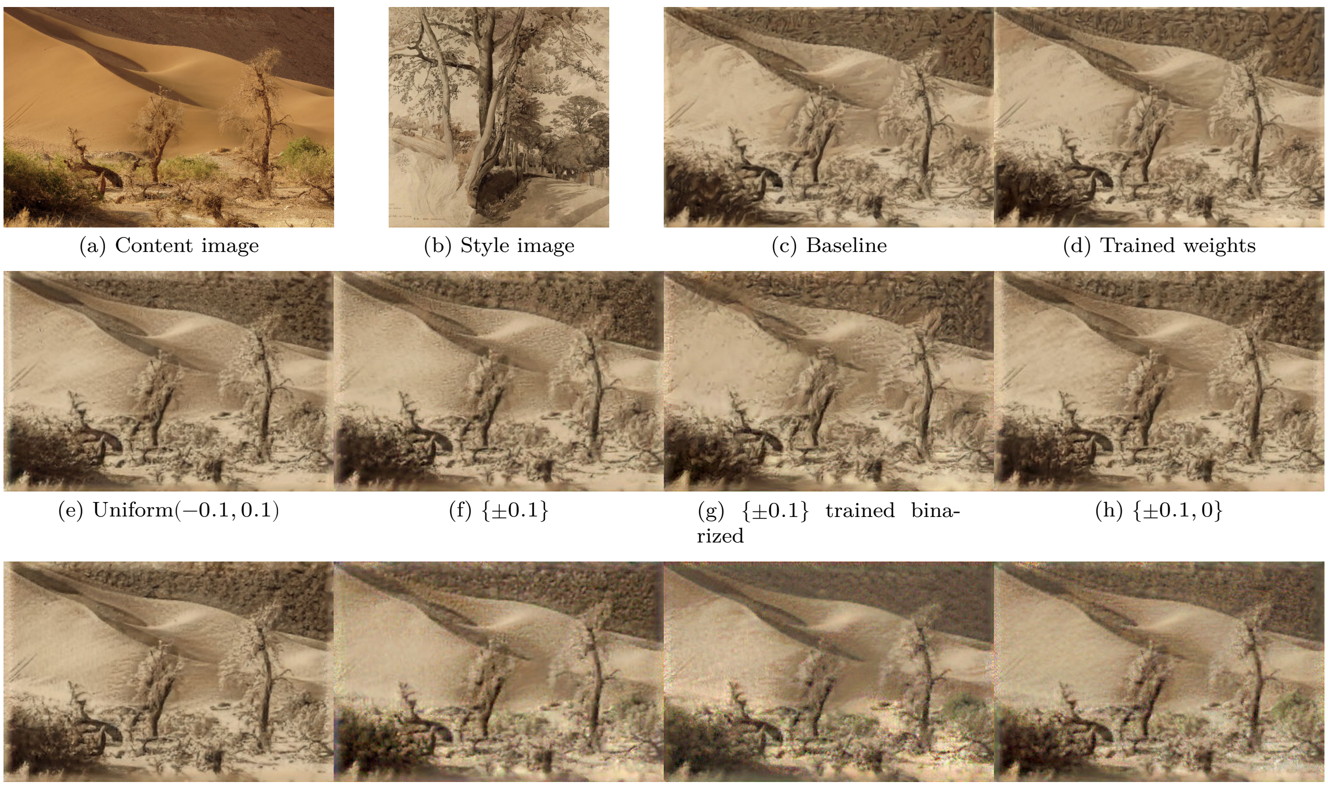

(a) Content image

(b) Style image

(c) Baseline

(d) Trained weights

(e) Uniform$(-0.1, 0.1)$

(f) $\{\pm 0.1\}$

(g) $\{\pm 0.1\}$ trained binarized

(h) $\{\pm 0.1, 0\}$

(i) $\{\pm 0.1, 0, 0\}$

(j) $\{\pm 0.2\}$

(k) $\{\pm 0.4\}$

(l) $\{\pm 0.4, \pm 0.2, \pm 0.1, 0\}$

Figure 11: Symmetric discrete weights with fixed biases (0.5)

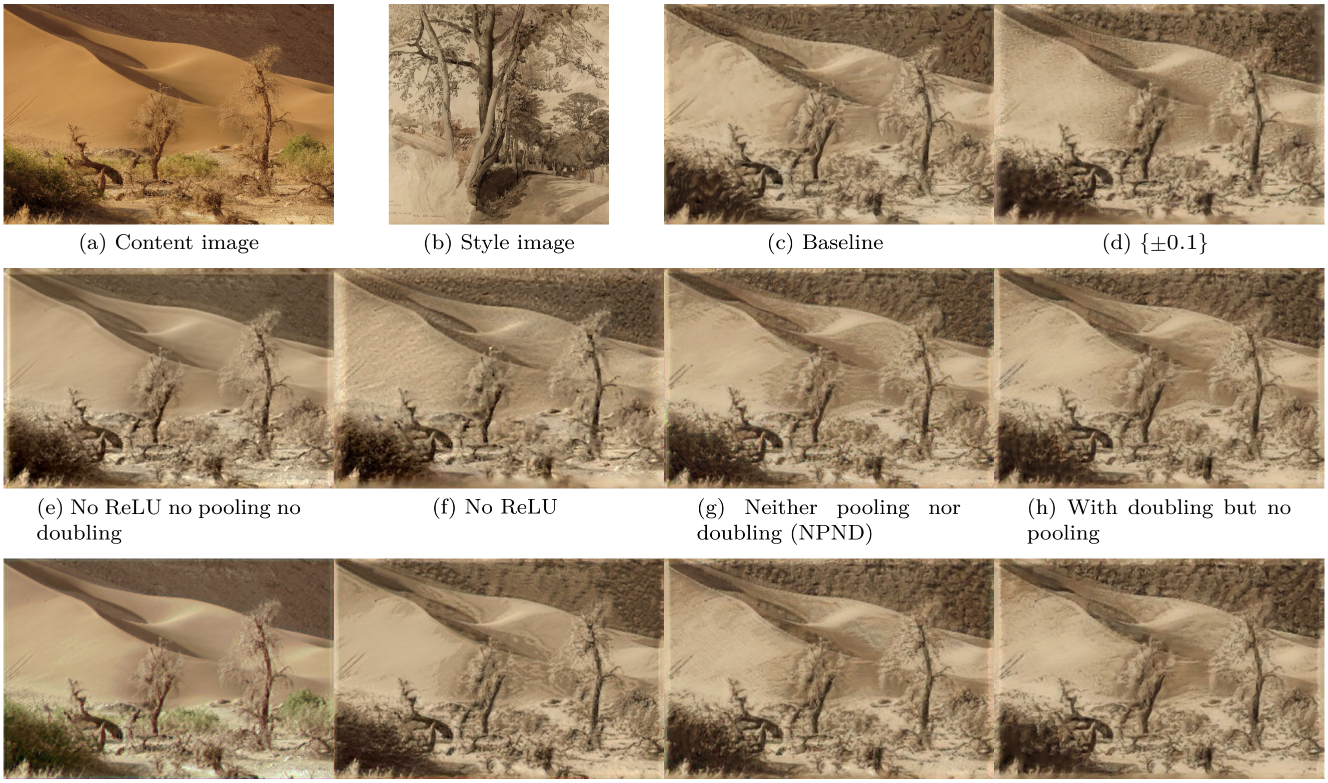

(a) Content image
(b) Style image
(c) Baseline
(d) $\{\pm 0.1\}$
(e) No ReLU no pooling no doubling
(f) No ReLU
(g) Neither pooling nor doubling (NPND)
(h) With doubling but no pooling
(i) NPND, 4 filters
(j) NPND, 16 filters
(k) NPND, 71 filters
(l) NPND, 512 filters

Figure 12: Removal of structure

(a) Basic set 1, strategy 4 (32-channel)
(b) Basic set 2, strategy 1 (18-channel)
(c) Basic set 3, strategy 3 (64-channel)
(d) 2×2 kernels, random, 64-channel
(e) 3×3 kernels, random, 64-channel

Figure 13: Alternative first-layer convolution kernels

(a) Gaussian pyramid without pooling

(b) Laplacian pyramid without pooling

(c) Gaussian pyramid with pooling

(d) Laplacian pyramid with pooling

Figure 14: Using image pyramids instead of multi-layer nonlinearities

(a) Content Image (b) Style Image (c) Baseline (d) $\beta/\alpha = 10^2$ (e) Removing the content loss

Figure 15: Removing the content loss

(a) Content Image

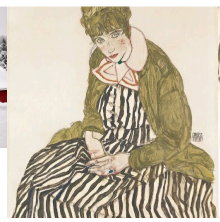

(b) Style Image

(c) Trained biases

(d) All zero biases

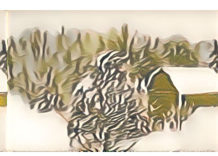

(e) Fixed biases (0.5)

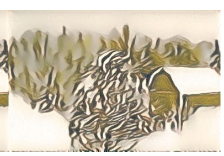

(f) Reordered biases

Figure 16: Varing the biases while keeping trained weights

(a) Content image (b) Style image (c) Baseline (d) Trained weights

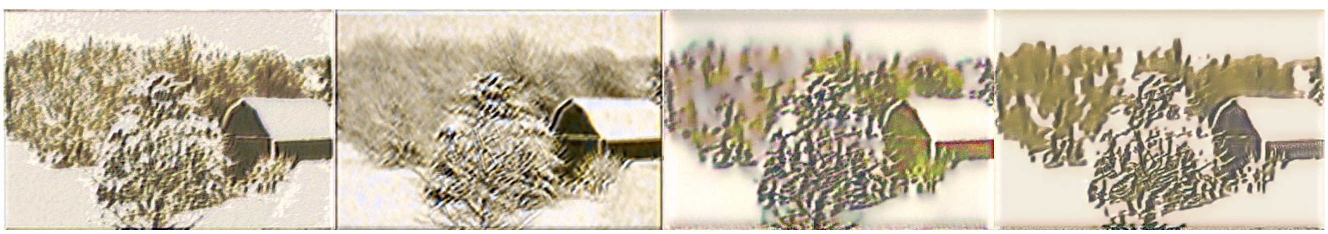

(e) Gaussian (=0.015) (trained bias) (f) Gaussian (=0.015) ([1]) (g) Gaussian (=0.15) (h) Uniform$(-0.1, 0.1)$

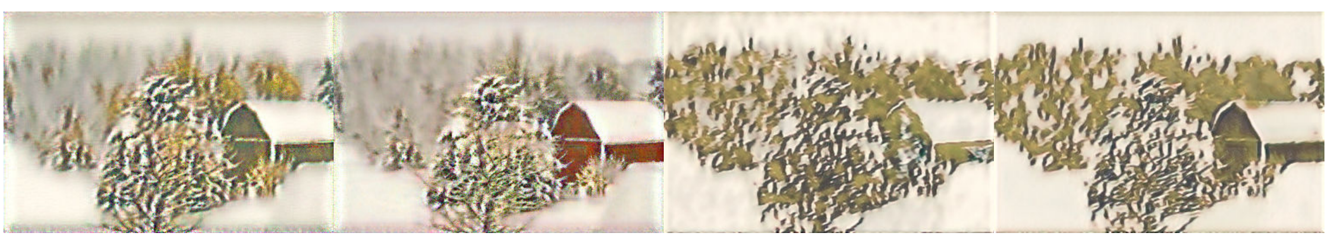

(i) Uniform$(-0.5, 0.5)$ (j) Uniform$(-1, 1)$ (k) Shuffled weights (l) Shuffled weights (trained biases)

Figure 17: Continuously or densely distributed weights

(a) Content image (b) Style image (c) Baseline (d) Trained weights

(e) Uniform$(-0.1, 0.1)$ (f) $\{\pm 0.1\}$ (g) $\{\pm 0.1\}$ trained binarized (h) $\{\pm 0.1, 0\}$

(i) $\{\pm 0.1, 0, 0\}$ (j) $\{\pm 0.2\}$ (k) $\{\pm 0.4\}$ (l) $\{\pm 0.4, \pm 0.2, \pm 0.1, 0\}$

Figure 18: Symmetric discrete weights with fixed biases (0.5)

(a) Content image (b) Style image (c) Baseline (d) $\{\pm 0.1\}$

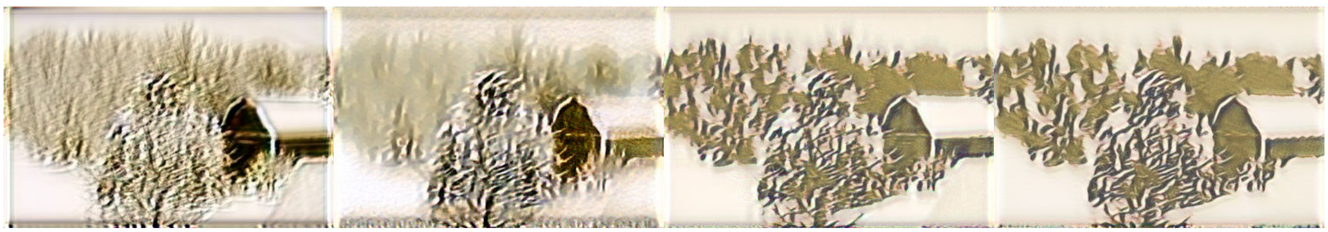

(e) No ReLU no pooling no doubling (f) No ReLU (g) Neither pooling nor doubling (NPND) (h) With doubling but no pooling

(i) NPND, 4 filters (j) NPND, 16 filters (k) NPND, 71 filters (l) NPND, 512 filters

Figure 19: Removal of structure

(a) Basic set 1, strategy 4 (32-channel) (b) Basic set 2, strategy 1 (18-channel) (c) Basic set 3, strategy 3 (64-channel) (d) 2×2 kernels, random, 64-channel (e) 3×3 kernels, random, 64-channel

Figure 20: Alternative first-layer convolution kernels

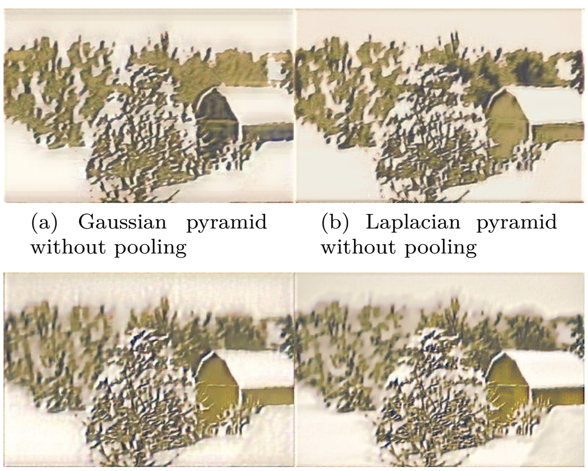

(a) Gaussian pyramid without pooling

(b) Laplacian pyramid without pooling

(c) Gaussian pyramid with pooling

(d) Laplacian pyramid with pooling

Figure 21: Using image pyramids instead of multi-layer nonlinearities

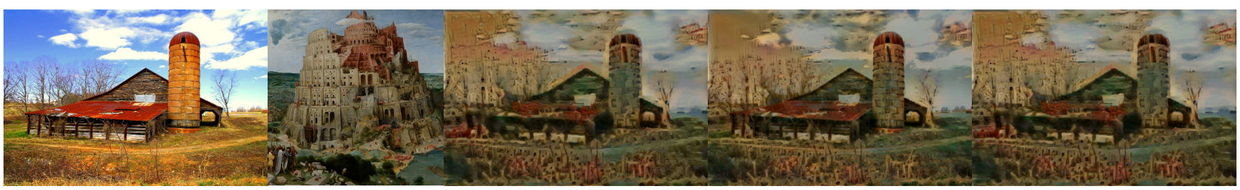

(a) Content Image (b) Style Image (c) Baseline (d) $\beta/\alpha = 10^2$ (e) Removing the content loss

Figure 22: Removing the content loss

(a) Content Image

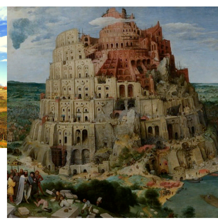

(b) Style Image

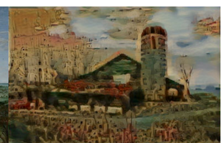

(c) Trained biases

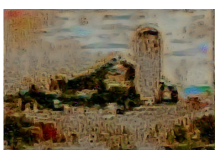

(d) All zero biases

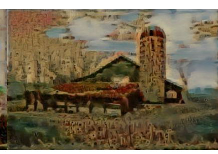

(e) Fixed biases (0.5)

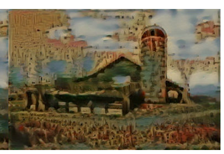

(f) Reordered biases

Figure 23: Varing the biases while keeping trained weights

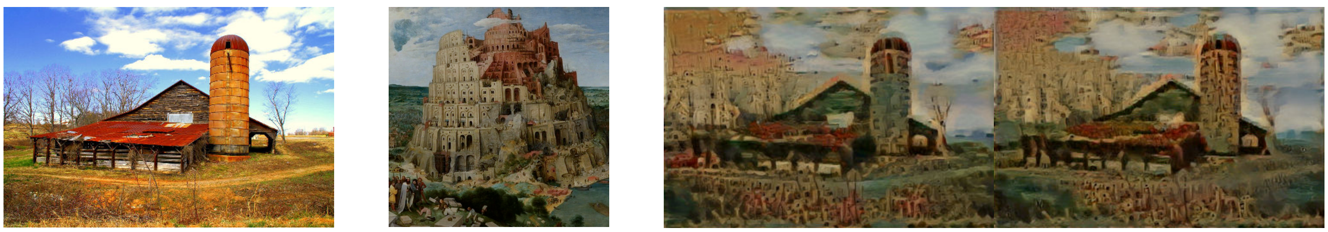

(a) Content image (b) Style image (c) Baseline (d) Trained weights

(e) Gaussian (=0.015) (trained bias) (f) Gaussian (=0.015) ([1]) (g) Gaussian (=0.15) (h) Uniform($-0.1, 0.1$)

(i) Uniform($-0.5, 0.5$) (j) Uniform($-1, 1$) (k) Shuffled weights (l) Shuffled weights (trained biases)

Figure 24: Continuously or densely distributed weights

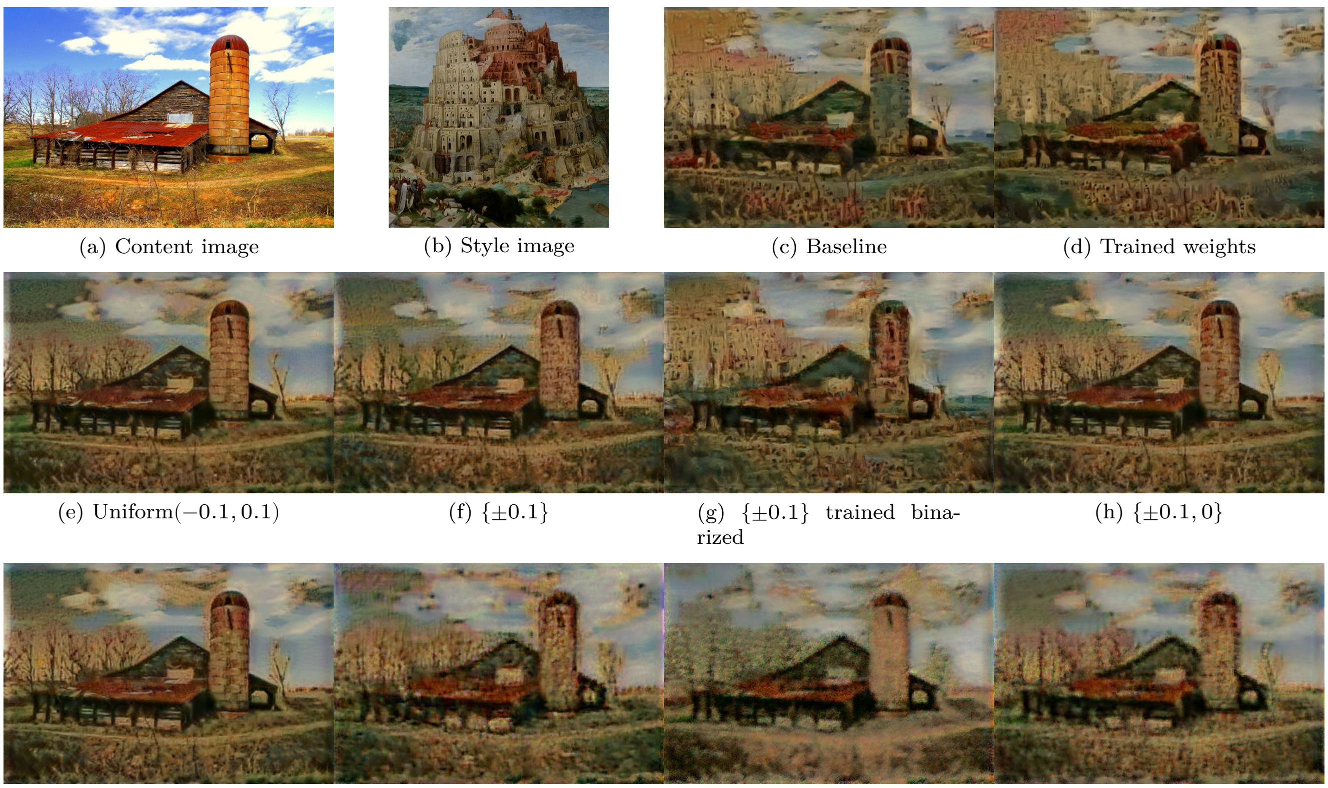

(a) Content image (b) Style image (c) Baseline (d) Trained weights

(e) Uniform($-0.1, 0.1$) (f) $\{\pm 0.1\}$ (g) $\{\pm 0.1\}$ trained binarized (h) $\{\pm 0.1, 0\}$

(i) $\{\pm 0.1, 0, 0\}$ (j) $\{\pm 0.2\}$ (k) $\{\pm 0.4\}$ (l) $\{\pm 0.4, \pm 0.2, \pm 0.1, 0\}$

Figure 25: Symmetric discrete weights with fixed biases (0.5)

(a) Content image (b) Style image (c) Baseline (d) $\{\pm 0.1\}$

(e) No ReLU no pooling no doubling (f) No ReLU (g) Neither pooling nor doubling (NPND) (h) With doubling but no pooling

(i) NPND, 4 filters (j) NPND, 16 filters (k) NPND, 71 filters (l) NPND, 512 filters

Figure 26: Removal of structure

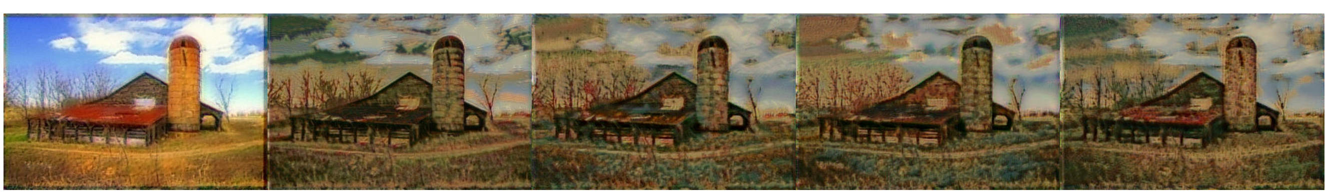

(a) Basic set 1, strategy 4 (32-channel) (b) Basic set 2, strategy 1 (18-channel) (c) Basic set 3, strategy 3 (64-channel) (d) 2×2 kernels, random, 64-channel (e) 3×3 kernels, random, 64-channel

Figure 27: Alternative first-layer convolution kernels

(a) Gaussian pyramid without pooling

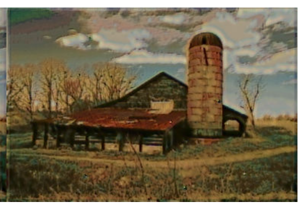

(b) Laplacian pyramid without pooling

(c) Gaussian pyramid with pooling

(d) Laplacian pyramid with pooling

Figure 28: Using image pyramids instead of multi-layer nonlinearities

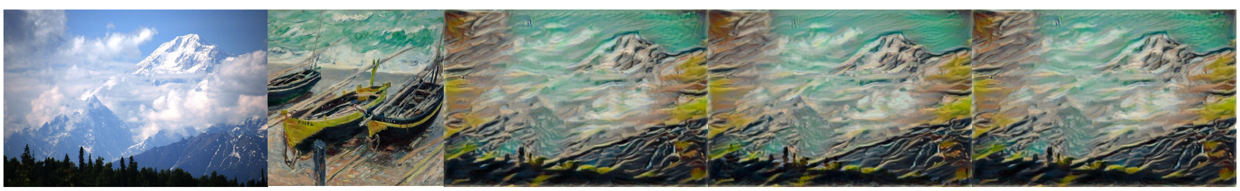

(a) Content Image (b) Style Image (c) Baseline (d) $\beta/\alpha = 10^2$ (e) Removing the content loss

Figure 29: Removing the content loss

(a) Content Image

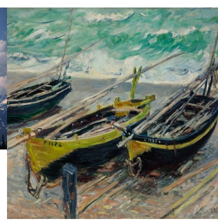

(b) Style Image

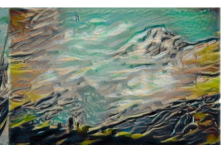

(c) Trained biases

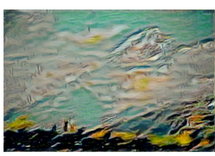

(d) All zero biases

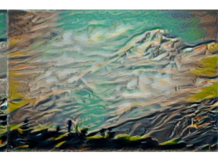

(e) Fixed biases (0.5)

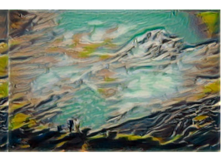

(f) Reordered biases

Figure 30: Varing the biases while keeping trained weights

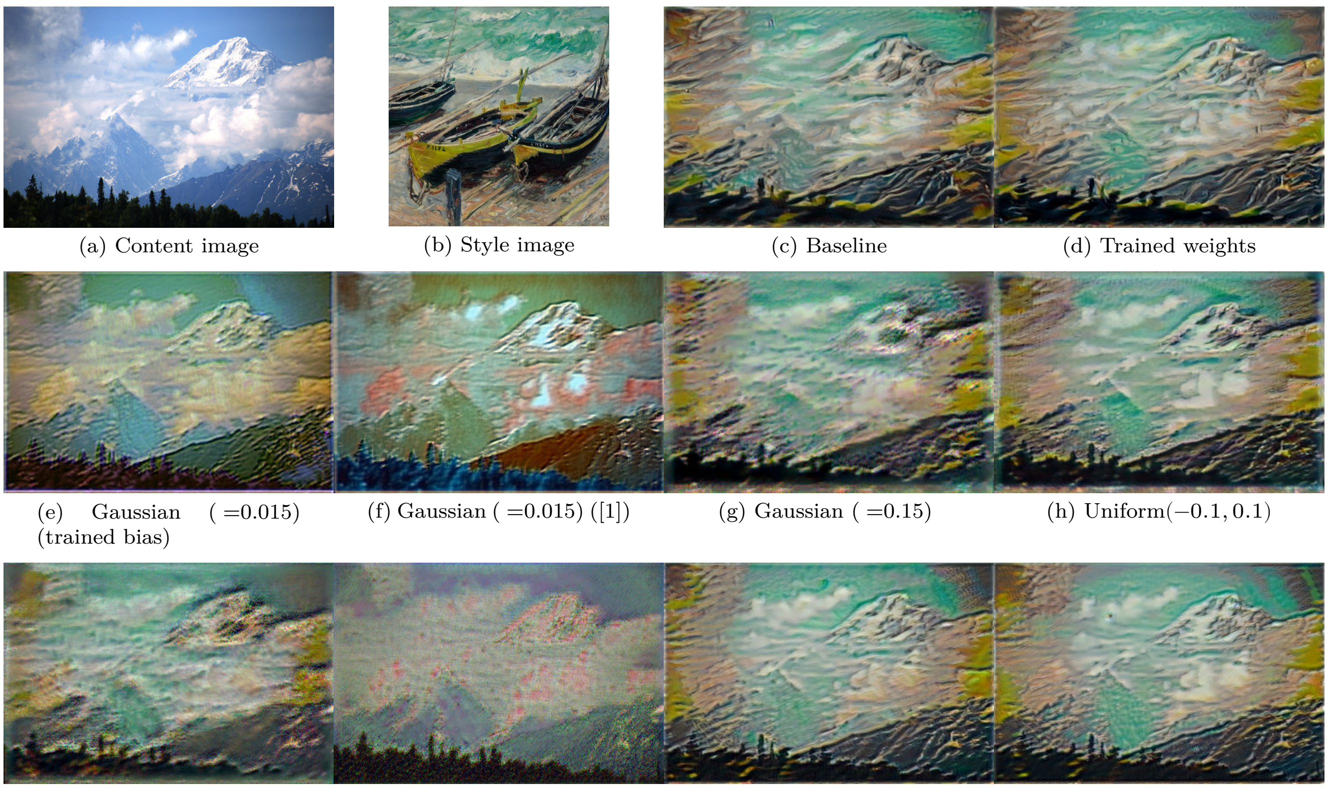

(a) Content image

(b) Style image

(c) Baseline

(d) Trained weights

(e) Gaussian (=0.015) (trained bias)

(f) Gaussian (=0.015) ([1])

(g) Gaussian (=0.15)

(h) Uniform$(-0.1, 0.1)$

(i) Uniform$(-0.5, 0.5)$

(j) Uniform$(-1, 1)$

(k) Shuffled weights

(l) Shuffled weights (trained biases)

Figure 31: Continuously or densely distributed weights

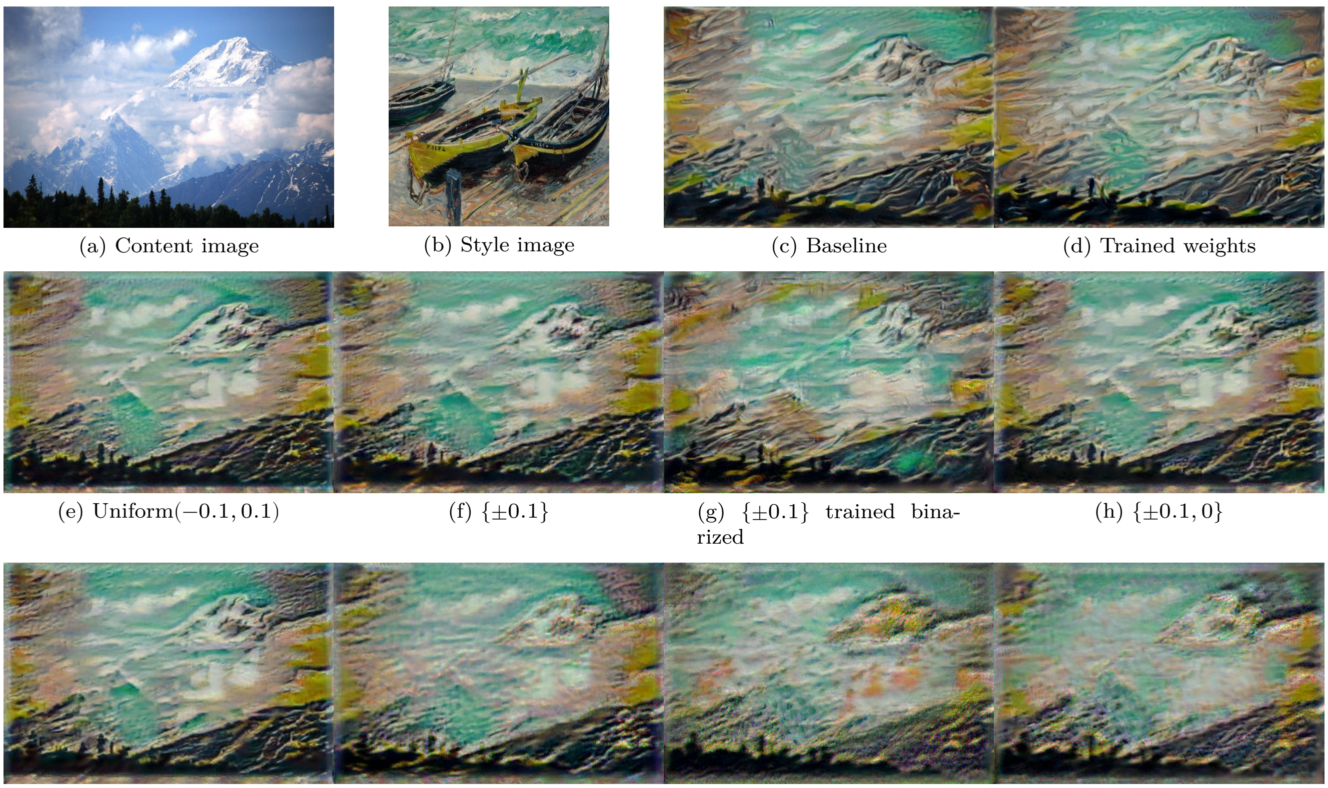

(a) Content image

(b) Style image

(c) Baseline

(d) Trained weights

(e) Uniform$(-0.1, 0.1)$

(f) $\{\pm 0.1\}$

(g) $\{\pm 0.1\}$ trained binarized

(h) $\{\pm 0.1, 0\}$

(i) $\{\pm 0.1, 0, 0\}$

(j) $\{\pm 0.2\}$

(k) $\{\pm 0.4\}$

(l) $\{\pm 0.4, \pm 0.2, \pm 0.1, 0\}$

Figure 32: Symmetric discrete weights with fixed biases (0.5)

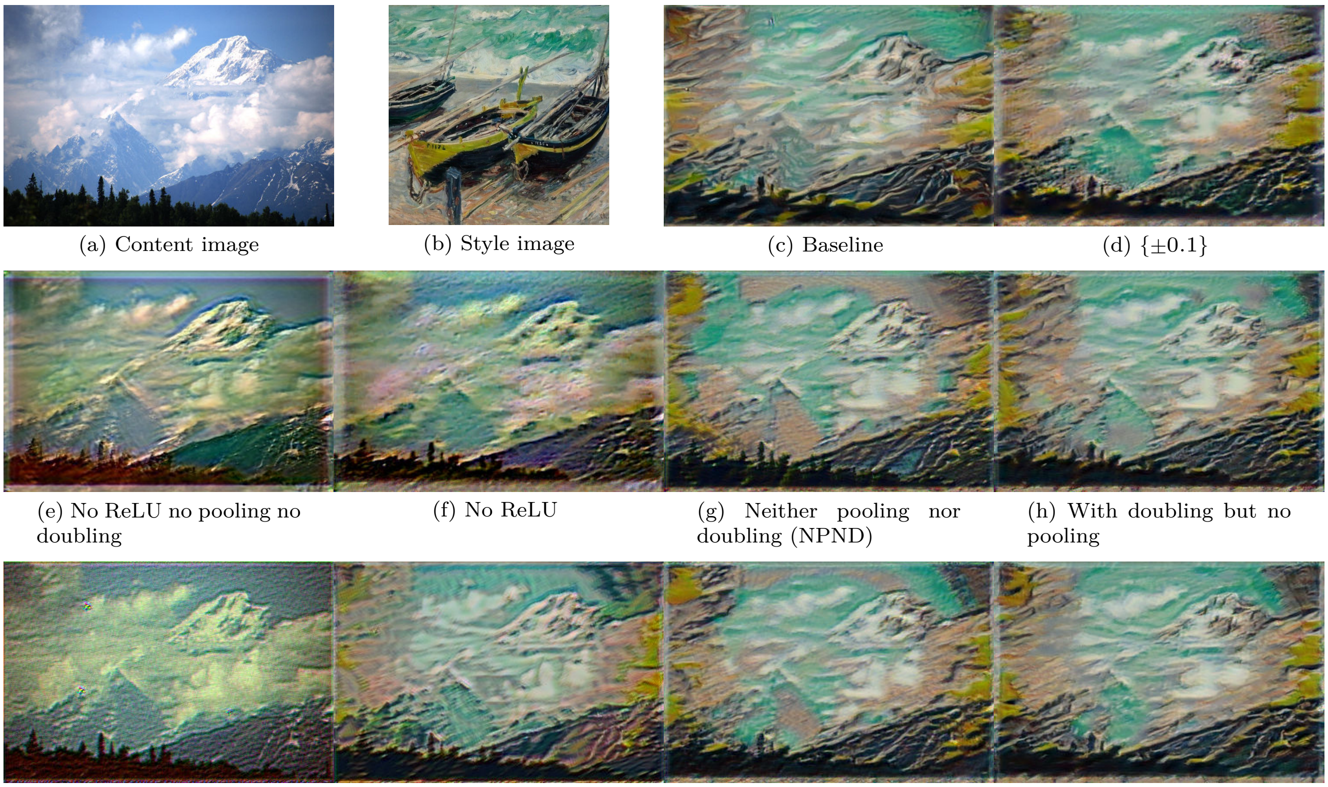

(a) Content image (b) Style image (c) Baseline (d) $\{\pm 0.1\}$

(e) No ReLU no pooling no doubling (f) No ReLU (g) Neither pooling nor doubling (NPND) (h) With doubling but no pooling

(i) NPND, 4 filters (j) NPND, 16 filters (k) NPND, 71 filters (l) NPND, 512 filters

Figure 33: Removal of structure

(a) Basic set 1, strategy 4 (32-channel) (b) Basic set 2, strategy 1 (18-channel) (c) Basic set 3, strategy 3 (64-channel) (d) 2×2 kernels, random, 64-channel (e) 3×3 kernels, random, 64-channel

Figure 34: Alternative first-layer convolution kernels

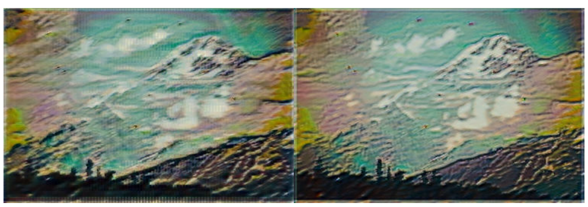

(a) Gaussian pyramid without pooling

(b) Laplacian pyramid without pooling

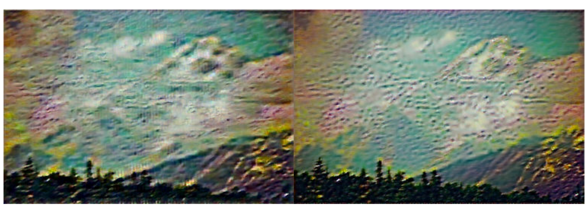

(c) Gaussian pyramid with pooling

(d) Laplacian pyramid with pooling

Figure 35: Using image pyramids instead of multi-layer nonlinearities

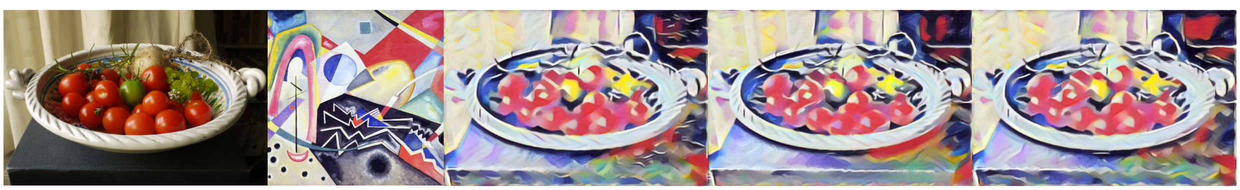

(a) Content Image (b) Style Image (c) Baseline (d) $\beta/\alpha = 10^2$ (e) Removing the content loss

Figure 36: Removing the content loss

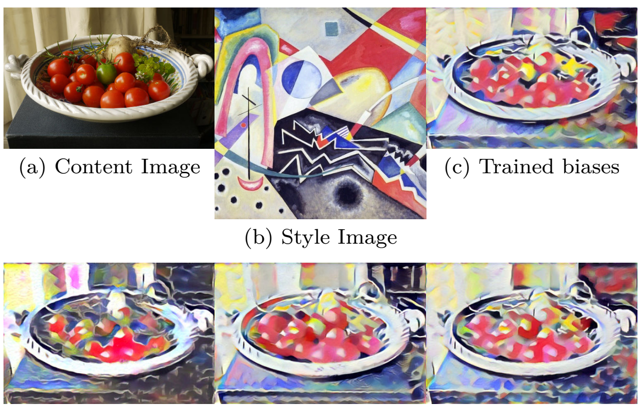

(a) Content Image (b) Style Image (c) Trained biases

(d) All zero biases (e) Fixed biases (0.5) (f) Reordered biases

Figure 37: Varing the biases while keeping trained weights

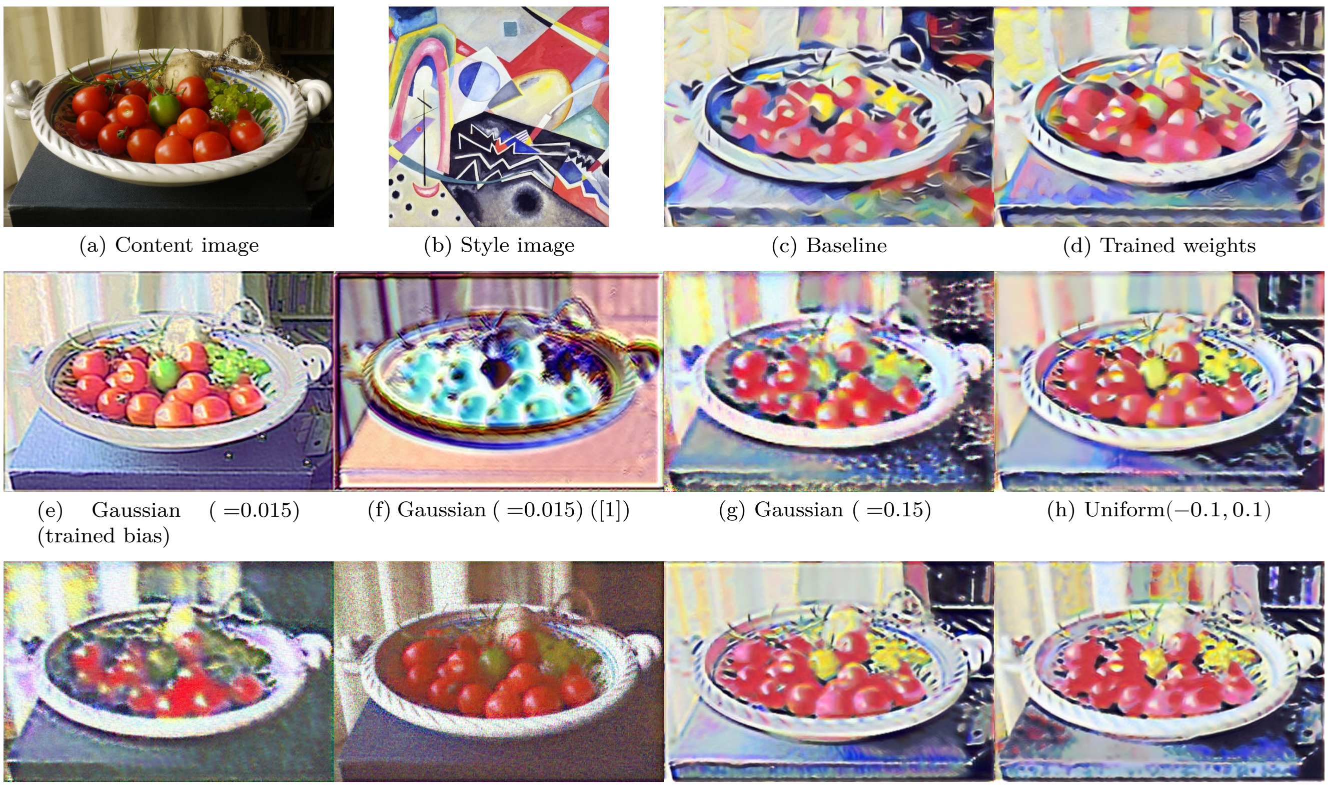

(a) Content image
(b) Style image
(c) Baseline
(d) Trained weights
(e) Gaussian (=0.015) (trained bias)
(f) Gaussian (=0.015) ([1])
(g) Gaussian (=0.15)
(h) Uniform$(-0.1, 0.1)$
(i) Uniform$(-0.5, 0.5)$
(j) Uniform$(-1, 1)$
(k) Shuffled weights
(l) Shuffled weights (trained biases)

Figure 38: Continuously or densely distributed weights

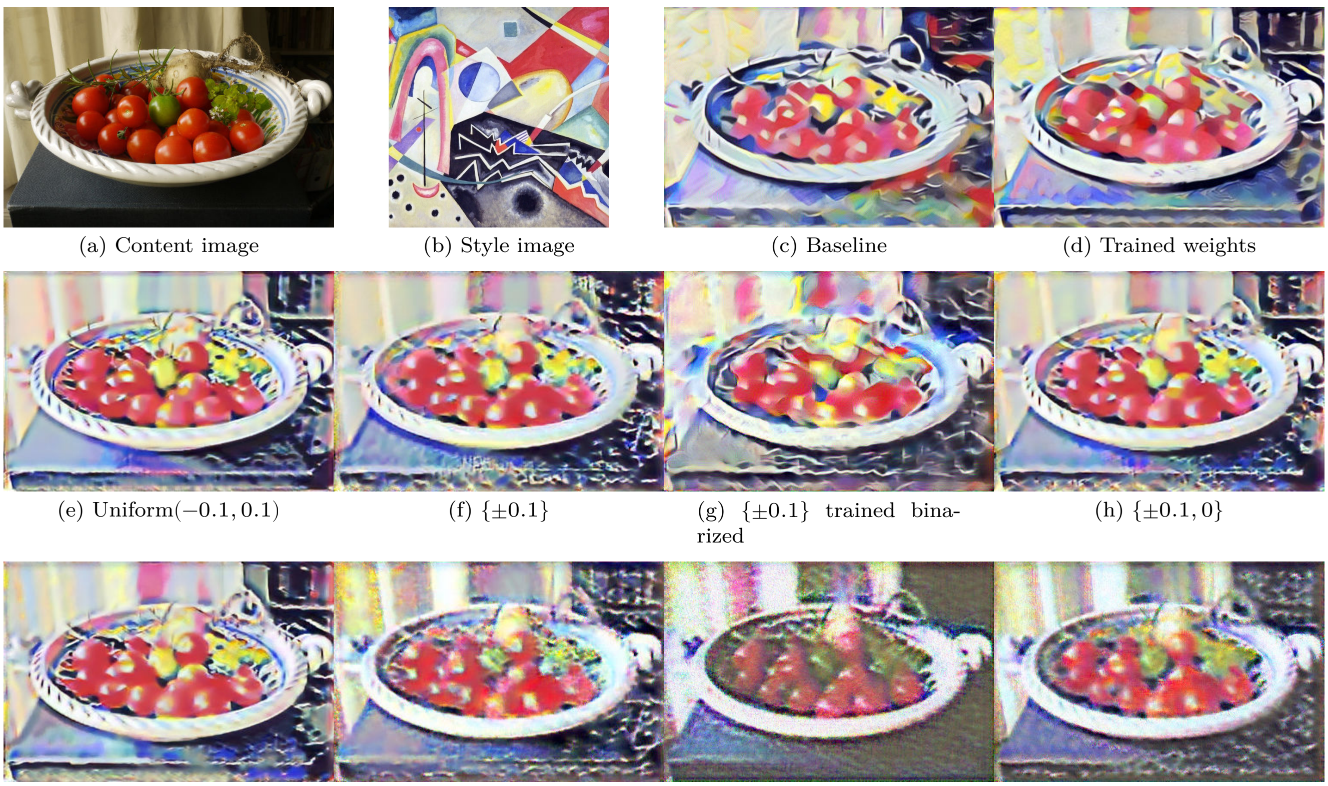

(a) Content image (b) Style image (c) Baseline (d) Trained weights

(e) Uniform$(-0.1, 0.1)$ (f) $\{\pm 0.1\}$ (g) $\{\pm 0.1\}$ trained binarized (h) $\{\pm 0.1, 0\}$

(i) $\{\pm 0.1, 0, 0\}$ (j) $\{\pm 0.2\}$ (k) $\{\pm 0.4\}$ (l) $\{\pm 0.4, \pm 0.2, \pm 0.1, 0\}$

Figure 39: Symmetric discrete weights with fixed biases (0.5)

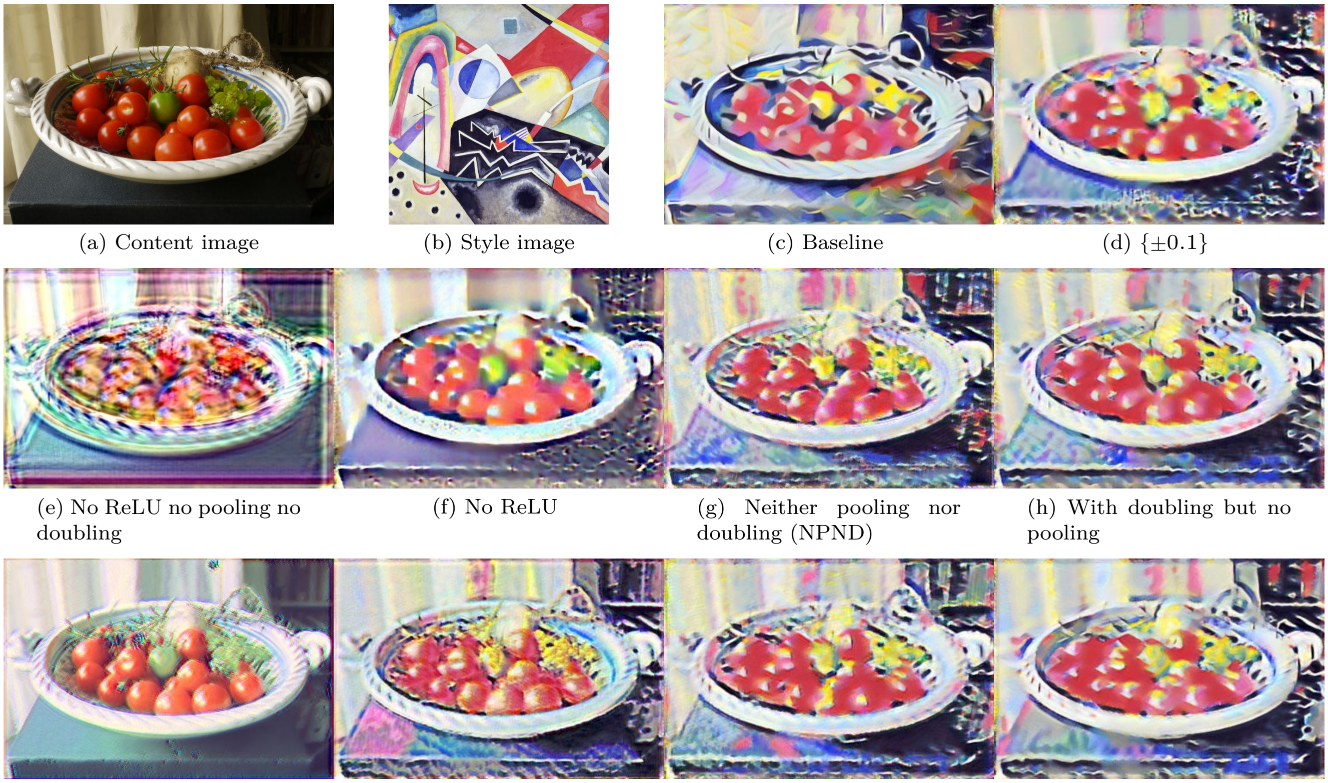

(a) Content image (b) Style image (c) Baseline (d) $\{\pm 0.1\}$

(e) No ReLU no pooling no doubling (f) No ReLU (g) Neither pooling nor doubling (NPND) (h) With doubling but no pooling

(i) NPND, 4 filters (j) NPND, 16 filters (k) NPND, 71 filters (l) NPND, 512 filters

Figure 40: Removal of structure

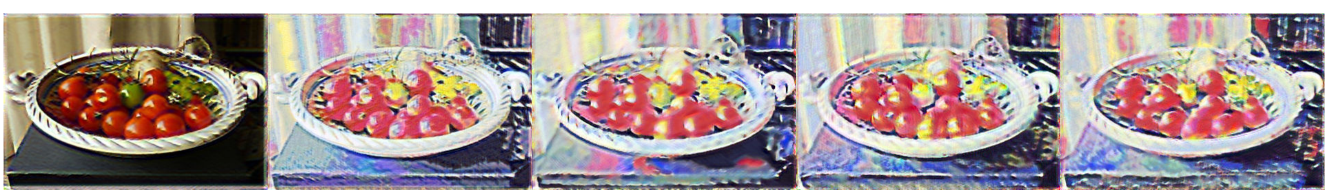

(a) Basic set 1, strategy 4 (32-channel) (b) Basic set 2, strategy 1 (18-channel) (c) Basic set 3, strategy 3 (64-channel) (d) 2×2 kernels, random, 64-channel (e) 3×3 kernels, random, 64-channel

Figure 41: Alternative first-layer convolution kernels

(a) Gaussian pyramid without pooling

(b) Laplacian pyramid without pooling

(c) Gaussian pyramid with pooling

(d) Laplacian pyramid with pooling

Figure 42: Using image pyramids instead of multi-layer nonlinearities

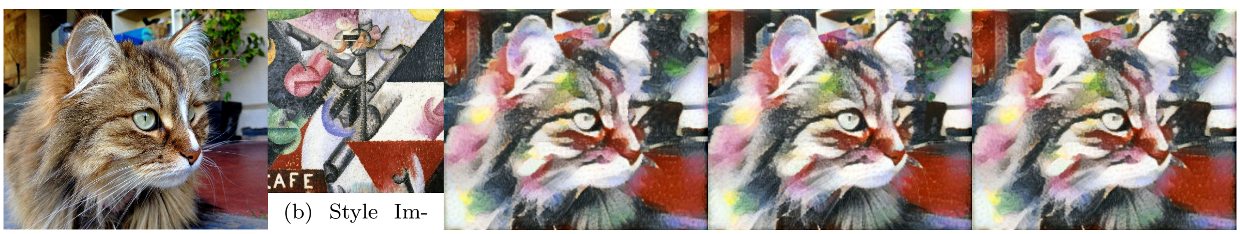

(a) Content Image (b) Style Image (c) Baseline (d) $\beta/\alpha = 10^2$ (e) Removing the content loss

Figure 43: Removing the content loss

(a) Content Image

(b) Style Image

(c) Trained biases

(d) All zero biases

(e) Fixed biases (0.5)

(f) Reordered biases

Figure 44: Varing the biases while keeping trained weights

(a) Content image

(b) Style image

(c) Baseline

(d) Trained weights

(e) Gaussian (=0.015) (trained bias)

(f) Gaussian (=0.015) ([1])

(g) Gaussian (=0.15)

(h) Uniform($-0.1, 0.1$)

(i) Uniform($-0.5, 0.5$)

(j) Uniform($-1, 1$)

(k) Shuffled weights

(l) Shuffled weights (trained biases)

Figure 45: Continuously or densely distributed weights

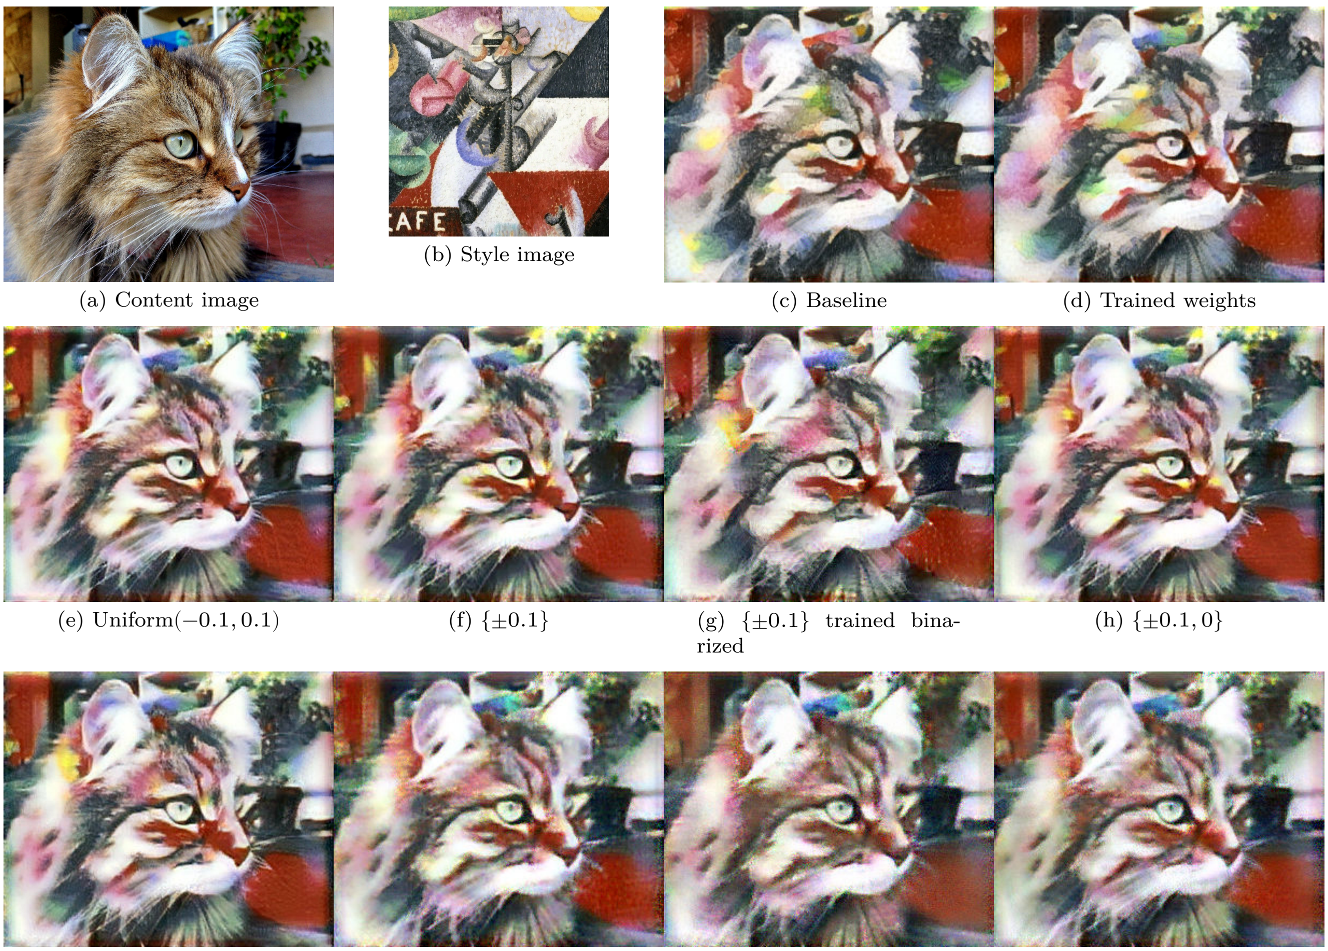

(a) Content image
(b) Style image
(c) Baseline
(d) Trained weights
(e) Uniform$(-0.1, 0.1)$
(f) $\{\pm 0.1\}$
(g) $\{\pm 0.1\}$ trained binarized
(h) $\{\pm 0.1, 0\}$
(i) $\{\pm 0.1, 0, 0\}$
(j) $\{\pm 0.2\}$
(k) $\{\pm 0.4\}$
(l) $\{\pm 0.4, \pm 0.2, \pm 0.1, 0\}$

Figure 46: Symmetric discrete weights with fixed biases (0.5)

(a) Content image

(b) Style image

(c) Baseline

(d) $\{\pm 0.1\}$

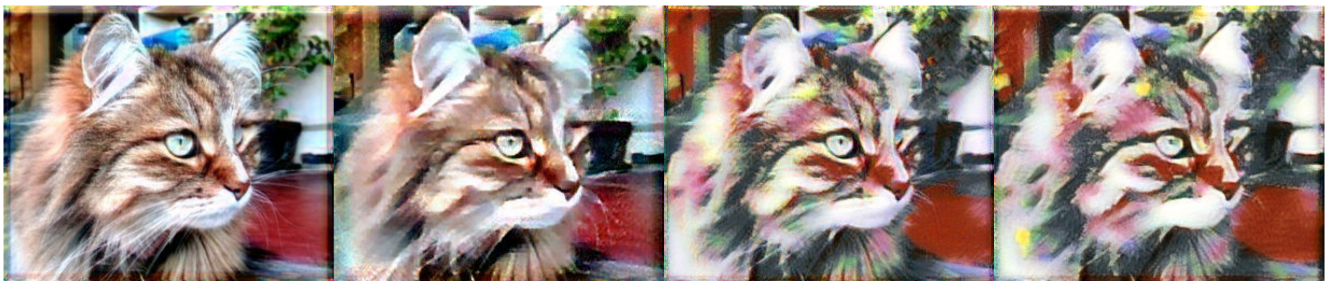

(e) No ReLU no pooling no doubling

(f) No ReLU

(g) Neither pooling nor doubling (NPND)

(h) With doubling but no pooling

(i) NPND, 4 filters

(j) NPND, 16 filters

(k) NPND, 71 filters

(l) NPND, 512 filters

Figure 47: Removal of structure

(a) Basic set 1, strategy 4 (32-channel)

(b) Basic set 2, strategy 1 (18-channel)

(c) Basic set 3, strategy 3 (64-channel)

(d) 2×2 kernels, random, 64-channel

(e) 3×3 kernels, random, 64-channel

Figure 48: Alternative first-layer convolution kernels

(a) Gaussian pyramid without pooling

(b) Laplacian pyramid without pooling

(c) Gaussian pyramid with pooling

(d) Laplacian pyramid with pooling

Figure 49: Using image pyramids instead of multi-layer nonlinearities

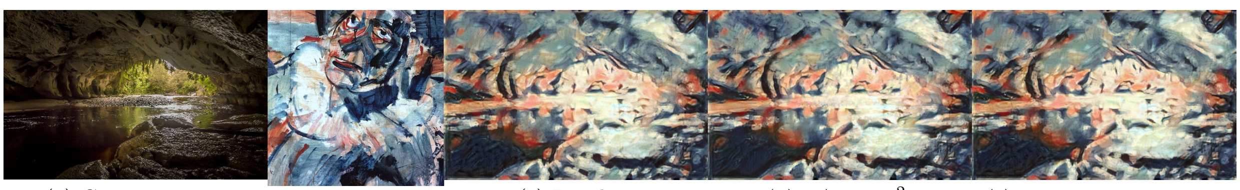

(a) Content Image

(b) Style Image

(c) Baseline

(d) $\beta/\alpha = 10^2$

(e) Removing the content loss

Figure 50: Removing the content loss

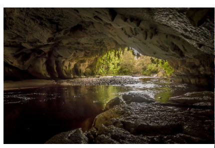

(a) Content Image

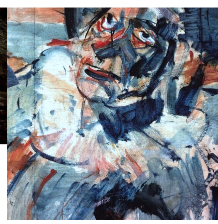

(b) Style Image

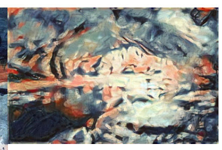

(c) Trained biases

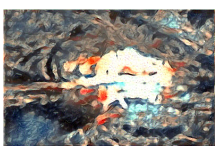

(d) All zero biases

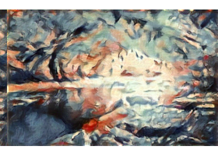

(e) Fixed biases (0.5)

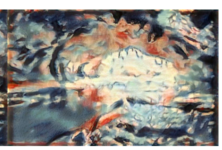

(f) Reordered biases

Figure 51: Varing the biases while keeping trained weights

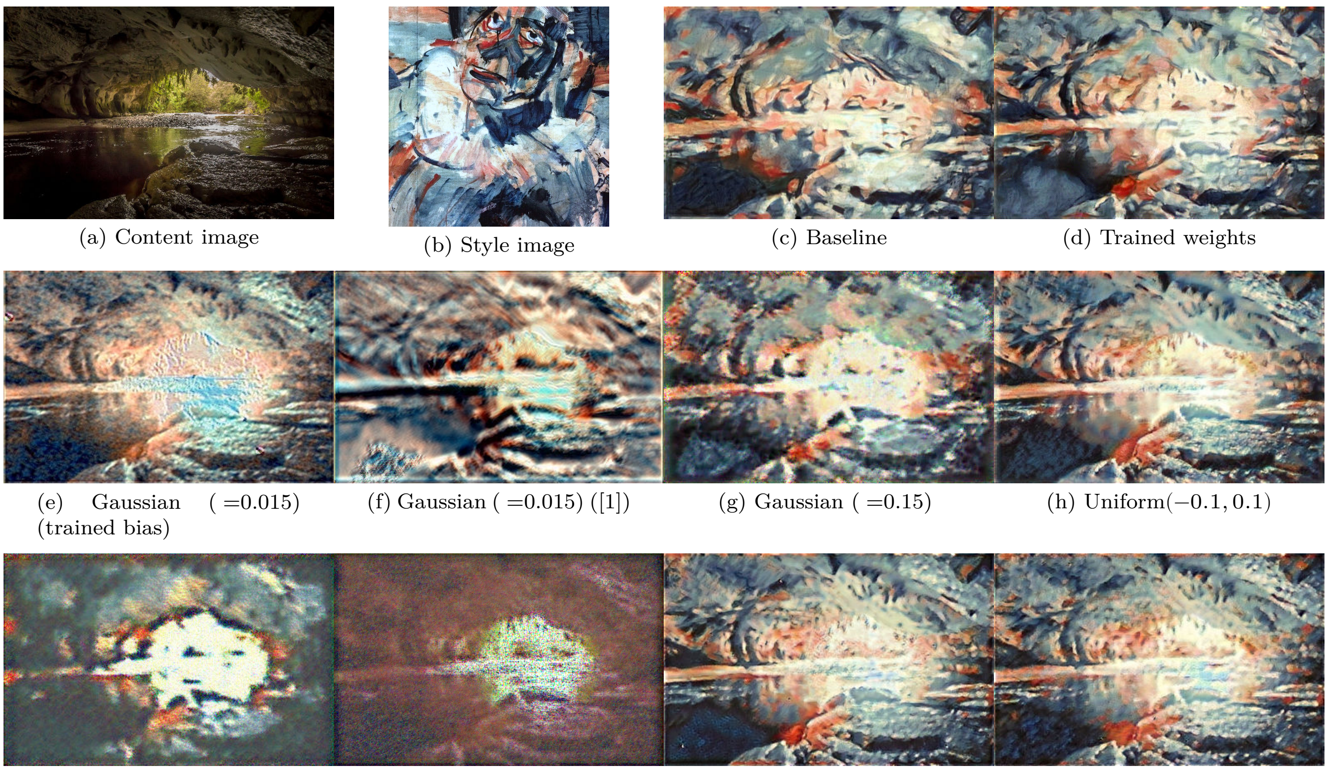

(a) Content image
(b) Style image
(c) Baseline
(d) Trained weights
(e) Gaussian (=0.015) (trained bias)
(f) Gaussian (=0.015) ([1])
(g) Gaussian (=0.15)
(h) Uniform$(-0.1, 0.1)$
(i) Uniform$(-0.5, 0.5)$
(j) Uniform$(-1, 1)$
(k) Shuffled weights
(l) Shuffled weights (trained biases)

Figure 52: Continuously or densely distributed weights

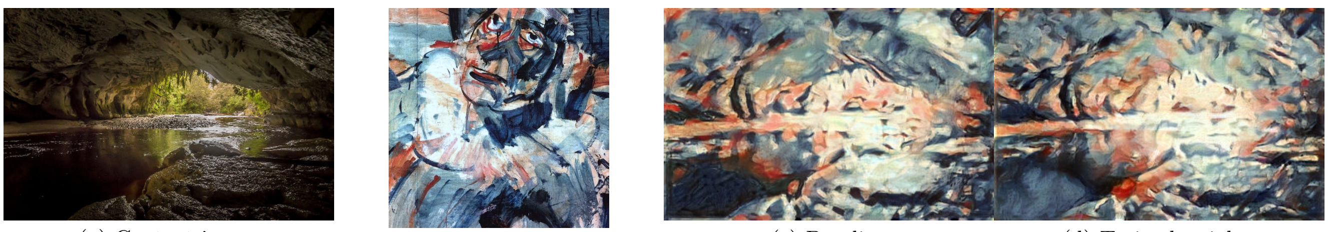

(a) Content image (b) Style image (c) Baseline (d) Trained weights

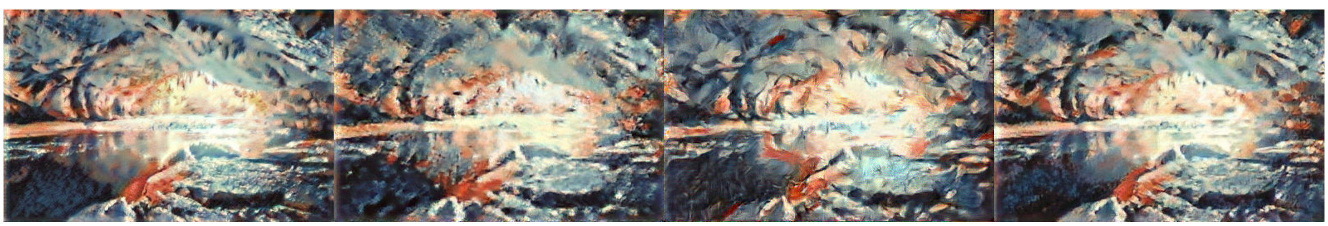

(e) Uniform$(-0.1, 0.1)$ (f) $\{\pm 0.1\}$ (g) $\{\pm 0.1\}$ trained binarized (h) $\{\pm 0.1, 0\}$

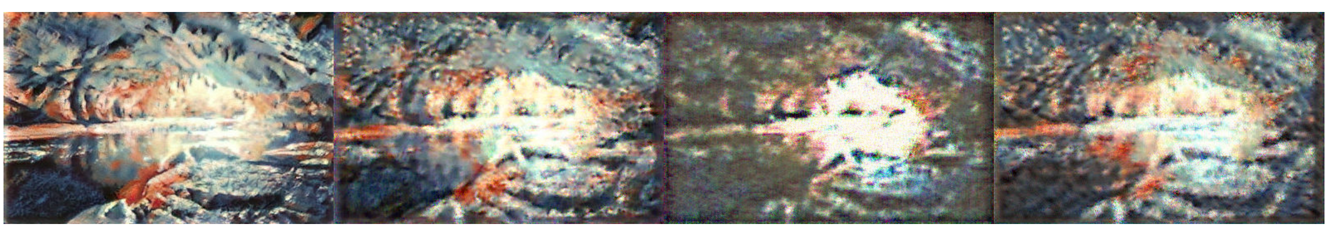

(i) $\{\pm 0.1, 0, 0\}$ (j) $\{\pm 0.2\}$ (k) $\{\pm 0.4\}$ (l) $\{\pm 0.4, \pm 0.2, \pm 0.1, 0\}$

Figure 53: Symmetric discrete weights with fixed biases (0.5)

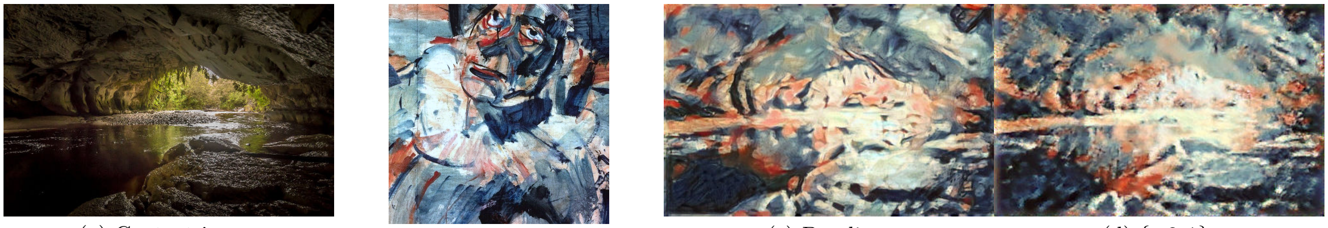

(a) Content image | (b) Style image | (c) Baseline | (d) $\{\pm 0.1\}$

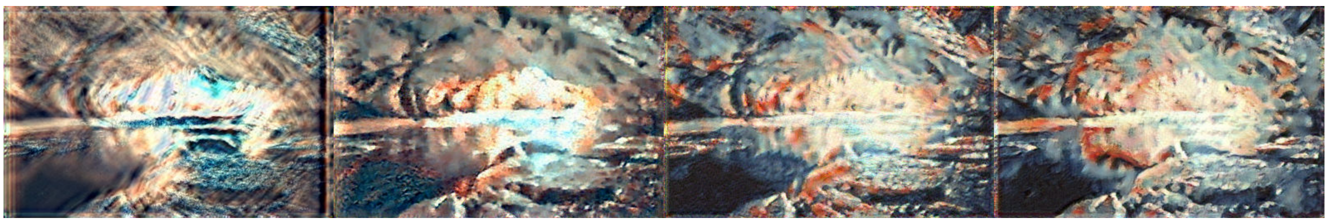

(e) No ReLU no pooling no doubling | (f) No ReLU | (g) Neither pooling nor doubling (NPND) | (h) With doubling but no pooling

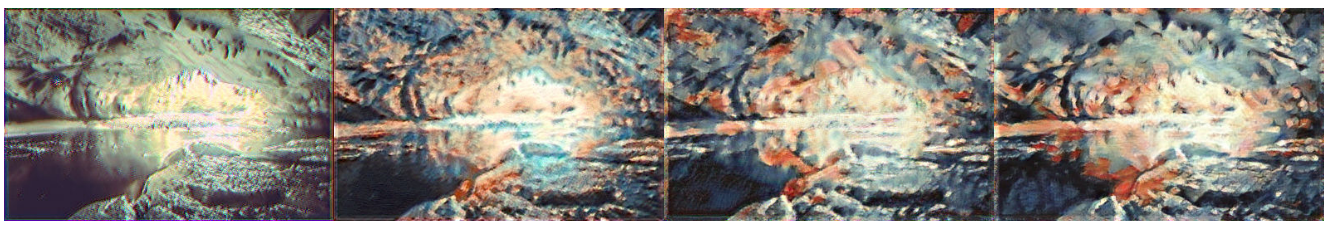

(i) NPND, 4 filters | (j) NPND, 16 filters | (k) NPND, 71 filters | (l) NPND, 512 filters

Figure 54: Removal of structure

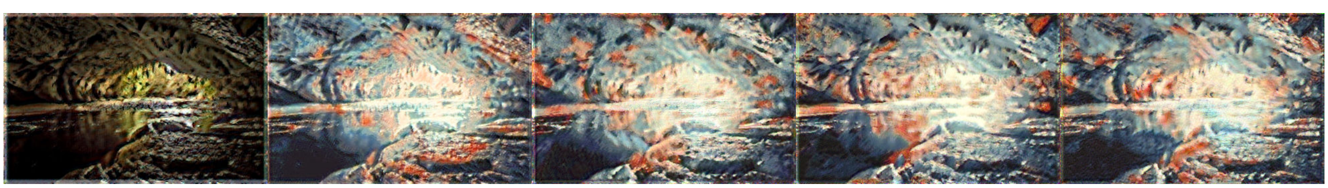

(a) Basic set 1, strategy 4 (32-channel) | (b) Basic set 2, strategy 1 (18-channel) | (c) Basic set 3, strategy 3 (64-channel) | (d) 2×2 kernels, random, 64-channel | (e) 3×3 kernels, random, 64-channel

Figure 55: Alternative first-layer convolution kernels

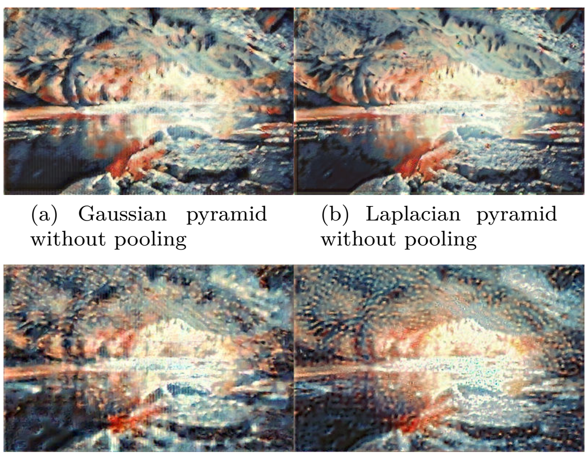

(a) Gaussian pyramid without pooling

(b) Laplacian pyramid without pooling

(c) Gaussian pyramid with pooling

(d) Laplacian pyramid with pooling

Figure 56: Using image pyramids instead of multi-layer nonlinearities

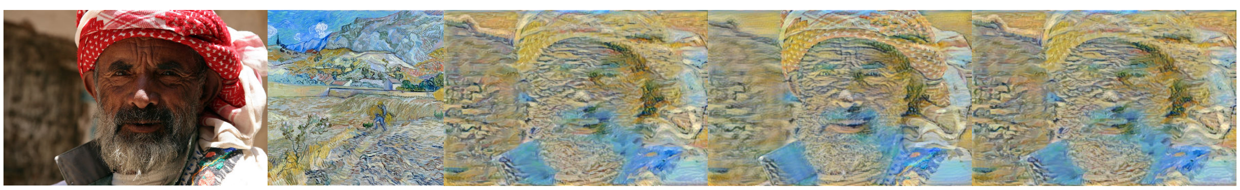

(a) Content Image (b) Style Image (c) Baseline (d) $\beta/\alpha = 10^2$ (e) Removing the content loss

Figure 57: Removing the content loss

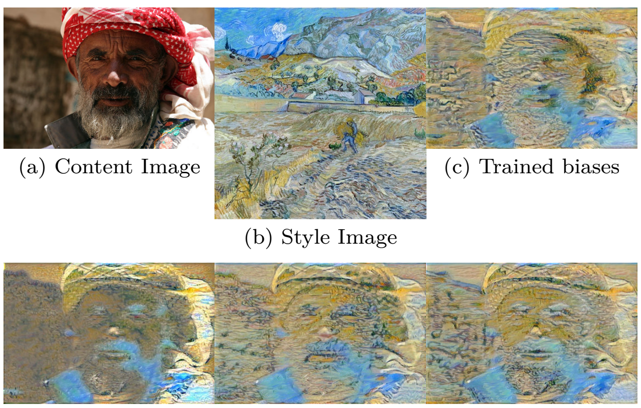

(a) Content Image (b) Style Image (c) Trained biases

(d) All zero biases (e) Fixed biases (0.5) (f) Reordered biases

Figure 58: Varing the biases while keeping trained weights

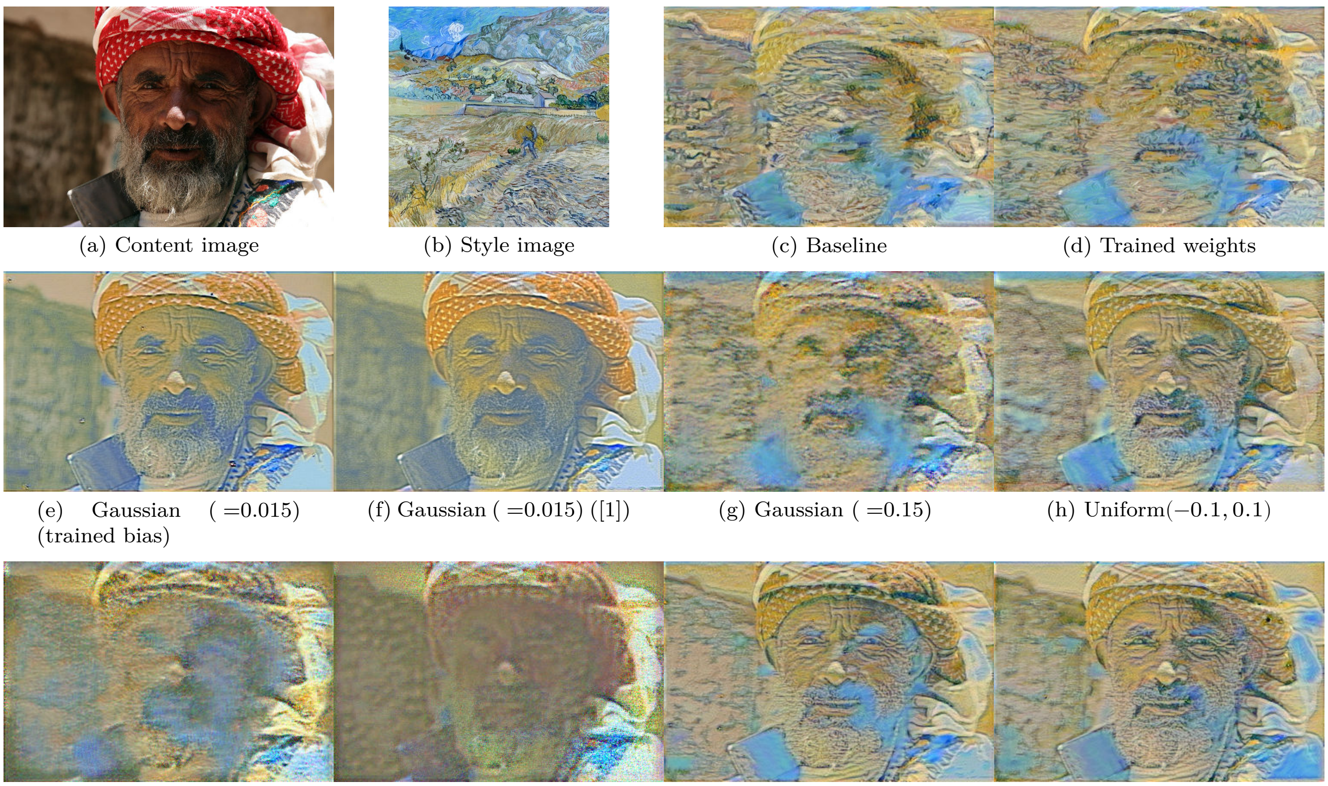

(a) Content image
(b) Style image
(c) Baseline
(d) Trained weights
(e) Gaussian (=0.015) (trained bias)
(f) Gaussian (=0.015) ([1])
(g) Gaussian (=0.15)
(h) Uniform($-0.1, 0.1$)
(i) Uniform($-0.5, 0.5$)
(j) Uniform($-1, 1$)
(k) Shuffled weights
(l) Shuffled weights (trained biases)

Figure 59: Continuously or densely distributed weights

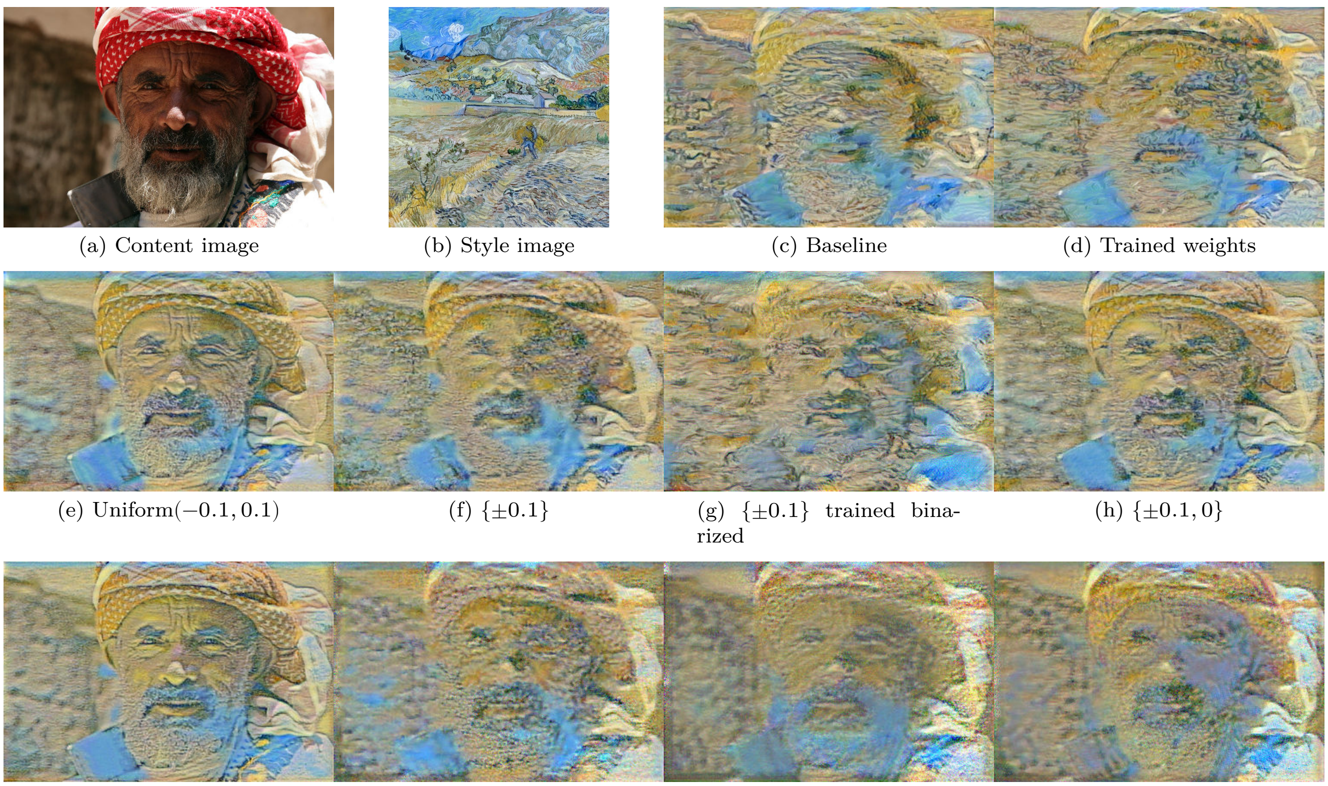

(a) Content image (b) Style image (c) Baseline (d) Trained weights

(e) Uniform$(-0.1, 0.1)$ (f) $\{\pm 0.1\}$ (g) $\{\pm 0.1\}$ trained binarized (h) $\{\pm 0.1, 0\}$

(i) $\{\pm 0.1, 0, 0\}$ (j) $\{\pm 0.2\}$ (k) $\{\pm 0.4\}$ (l) $\{\pm 0.4, \pm 0.2, \pm 0.1, 0\}$

Figure 60: Symmetric discrete weights with fixed biases (0.5)

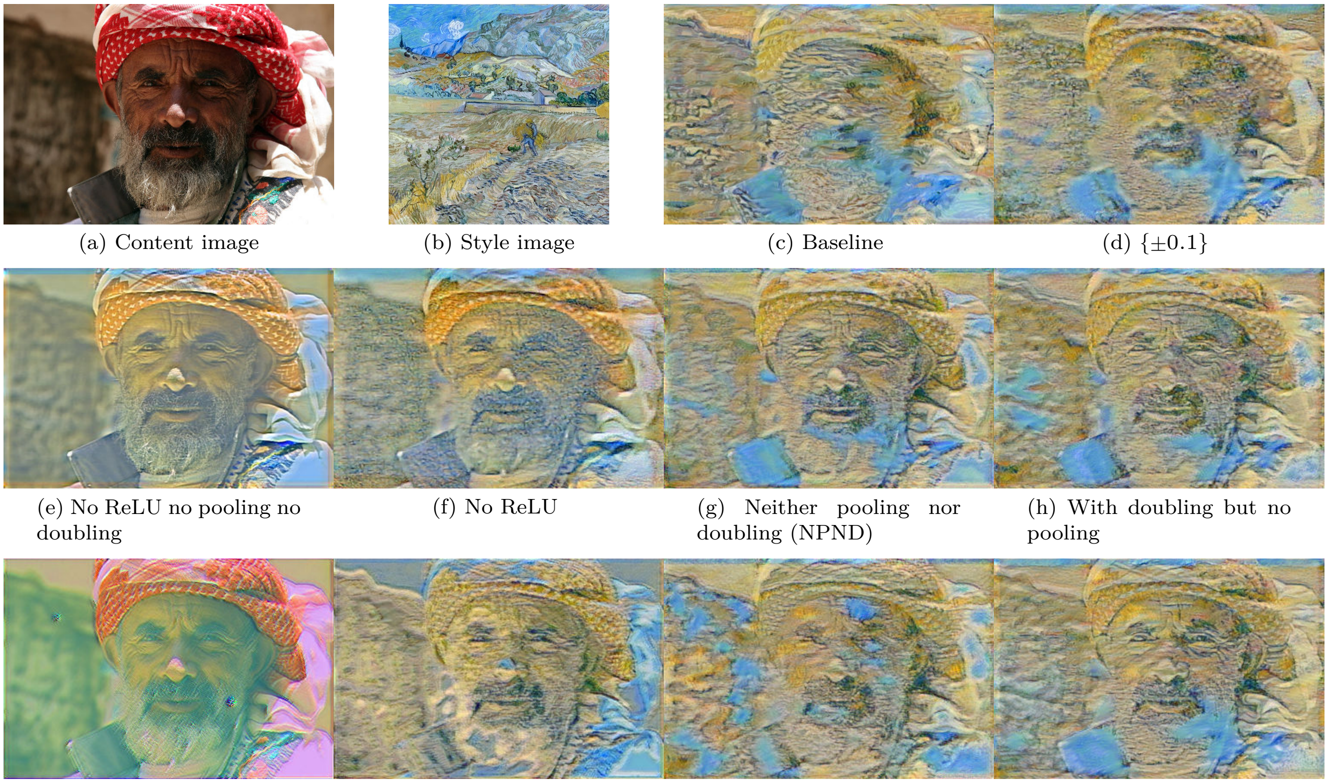

(a) Content image
(b) Style image
(c) Baseline
(d) $\{\pm 0.1\}$
(e) No ReLU no pooling no doubling
(f) No ReLU
(g) Neither pooling nor doubling (NPND)
(h) With doubling but no pooling
(i) NPND, 4 filters
(j) NPND, 16 filters
(k) NPND, 71 filters
(l) NPND, 512 filters

Figure 61: Removal of structure

(a) Basic set 1, strategy 4 (32-channel)
(b) Basic set 2, strategy 1 (18-channel)
(c) Basic set 3, strategy 3 (64-channel)
(d) 2×2 kernels, random, 64-channel
(e) 3×3 kernels, random, 64-channel

Figure 62: Alternative first-layer convolution kernels

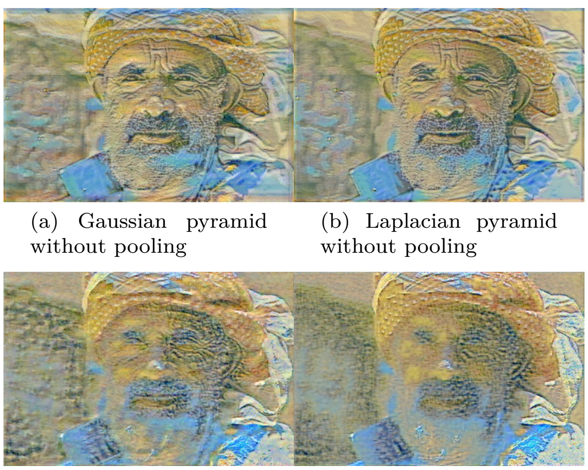

(a) Gaussian pyramid without pooling

(b) Laplacian pyramid without pooling

(c) Gaussian pyramid with pooling

(d) Laplacian pyramid with pooling

Figure 63: Using image pyramids instead of multi-layer nonlinearities

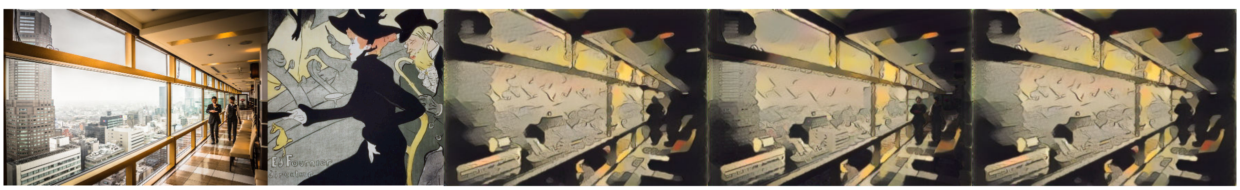

(a) Content Image (b) Style Image (c) Baseline (d) $\beta/\alpha = 10^2$ (e) Removing the content loss

Figure 64: Removing the content loss

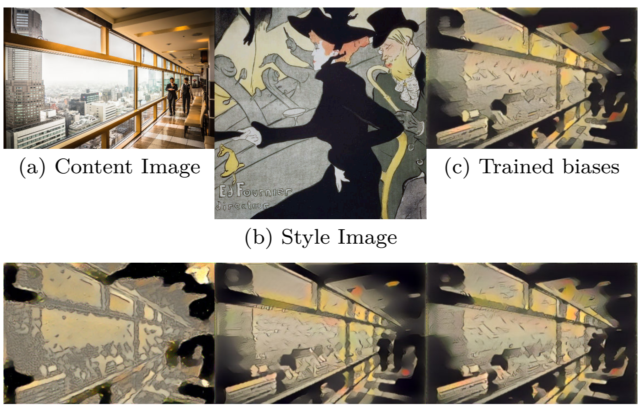

(a) Content Image (b) Style Image (c) Trained biases

(d) All zero biases (e) Fixed biases (0.5) (f) Reordered biases

Figure 65: Varing the biases while keeping trained weights

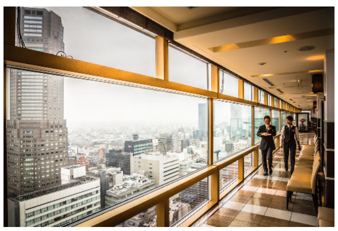

(a) Content image

(b) Style image

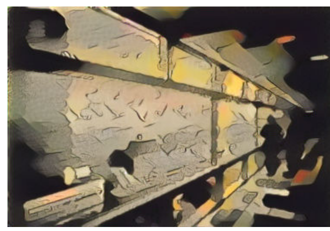

(c) Baseline

(d) Trained weights

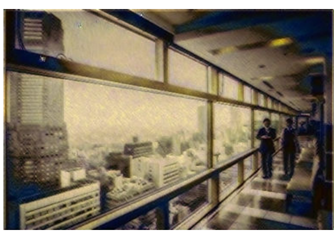

(e) Gaussian (=0.015) (trained bias)

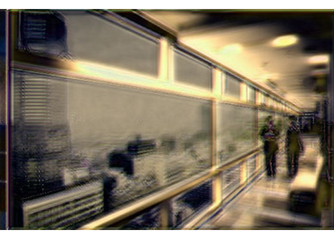

(f) Gaussian (=0.015) ([1])

(g) Gaussian (=0.15)

(h) Uniform$(-0.1, 0.1)$

(i) Uniform$(-0.5, 0.5)$

(j) Uniform$(-1, 1)$

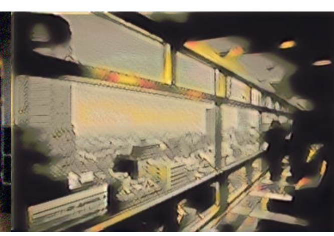

(k) Shuffled weights

(l) Shuffled weights (trained biases)

Figure 66: Continuously or densely distributed weights

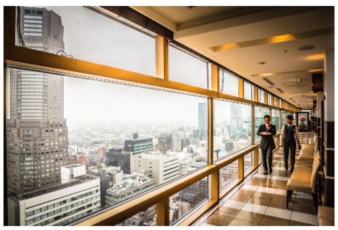

(a) Content image

(b) Style image

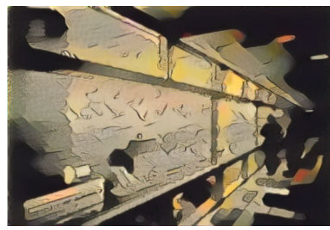

(c) Baseline

(d) Trained weights

(e) Uniform$(-0.1, 0.1)$

(f) $\{\pm 0.1\}$

(g) $\{\pm 0.1\}$ trained binarized

(h) $\{\pm 0.1, 0\}$

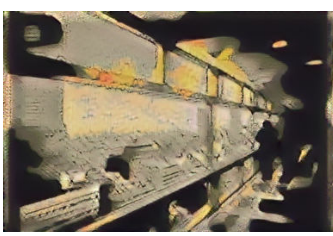

(i) $\{\pm 0.1, 0, 0\}$

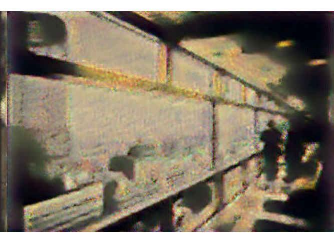

(j) $\{\pm 0.2\}$

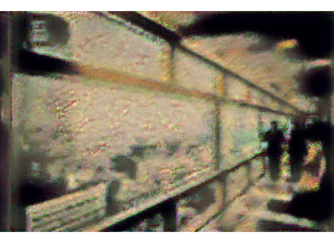

(k) $\{\pm 0.4\}$

(l) $\{\pm 0.4, \pm 0.2, \pm 0.1, 0\}$

Figure 67: Symmetric discrete weights with fixed biases (0.5)

(a) Content image (b) Style image (c) Baseline (d) $\{\pm 0.1\}$

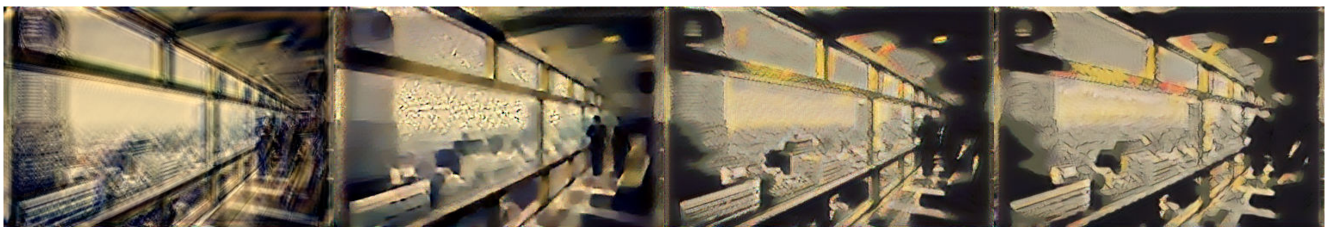

(e) No ReLU no pooling no doubling (f) No ReLU (g) Neither pooling nor doubling (NPND) (h) With doubling but no pooling

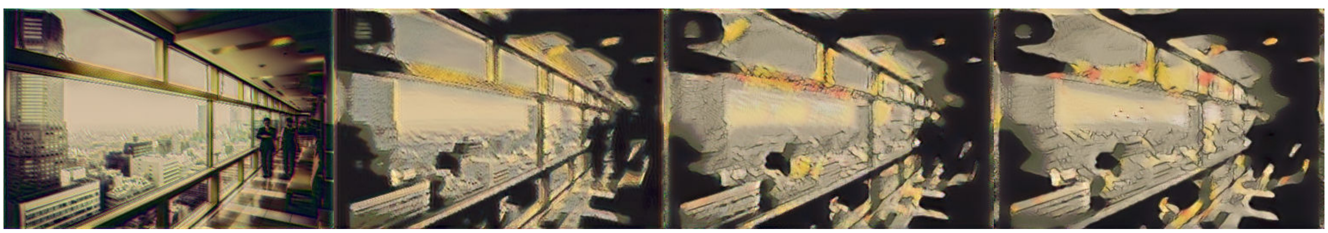

(i) NPND, 4 filters (j) NPND, 16 filters (k) NPND, 71 filters (l) NPND, 512 filters

Figure 68: Removal of structure

(a) Basic set 1, strategy 4 (32-channel) (b) Basic set 2, strategy 1 (18-channel) (c) Basic set 3, strategy 3 (64-channel) (d) 2×2 kernels, random, 64-channel (e) 3×3 kernels, random, 64-channel

Figure 69: Alternative first-layer convolution kernels

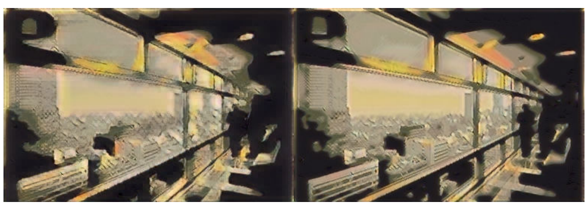

(a) Gaussian pyramid without pooling

(b) Laplacian pyramid without pooling

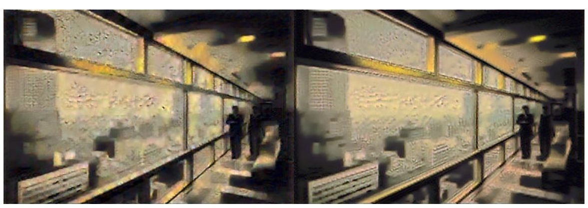

(c) Gaussian pyramid with pooling

(d) Laplacian pyramid with pooling

Figure 70: Using image pyramids instead of multi-layer nonlinearities