Texture Networks: Supplementary material

Below, we provide several additional qualitative results demonstrating the performance of texture networks and comparing them to Gatys et al. approaches. The source code for the proposed method can be found at https://github.com/DmitryUlyanov/texture_nets.

1. A note on generator architecture

Since the generator is only restricted to produce good images in terms of texture loss, nothing stops it from generating samples with small variance between them. Therefore, if a model archives lower texture loss it does not implicate that this model is preferable. The generator should be powerful enough to combine texture elements, but not too complex to degrade to similar samples. The degrading effect is similar to overfitting but there is no obvious way to control it as there is no analogue of validation set available.

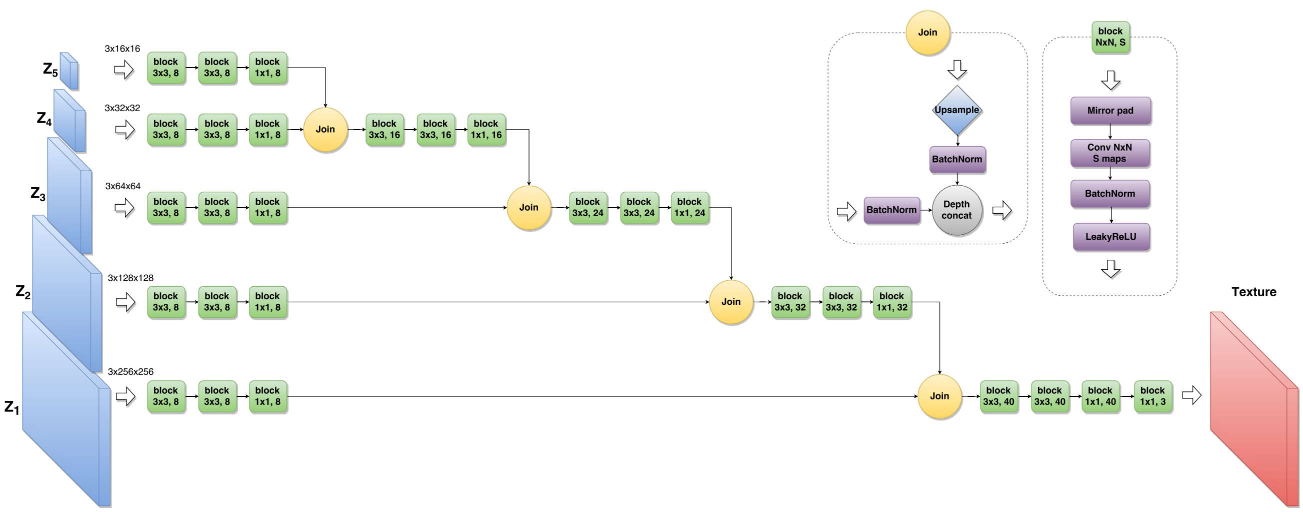

Figure 1. A more exact scheme of the architecture used for texture generation in our experiments.

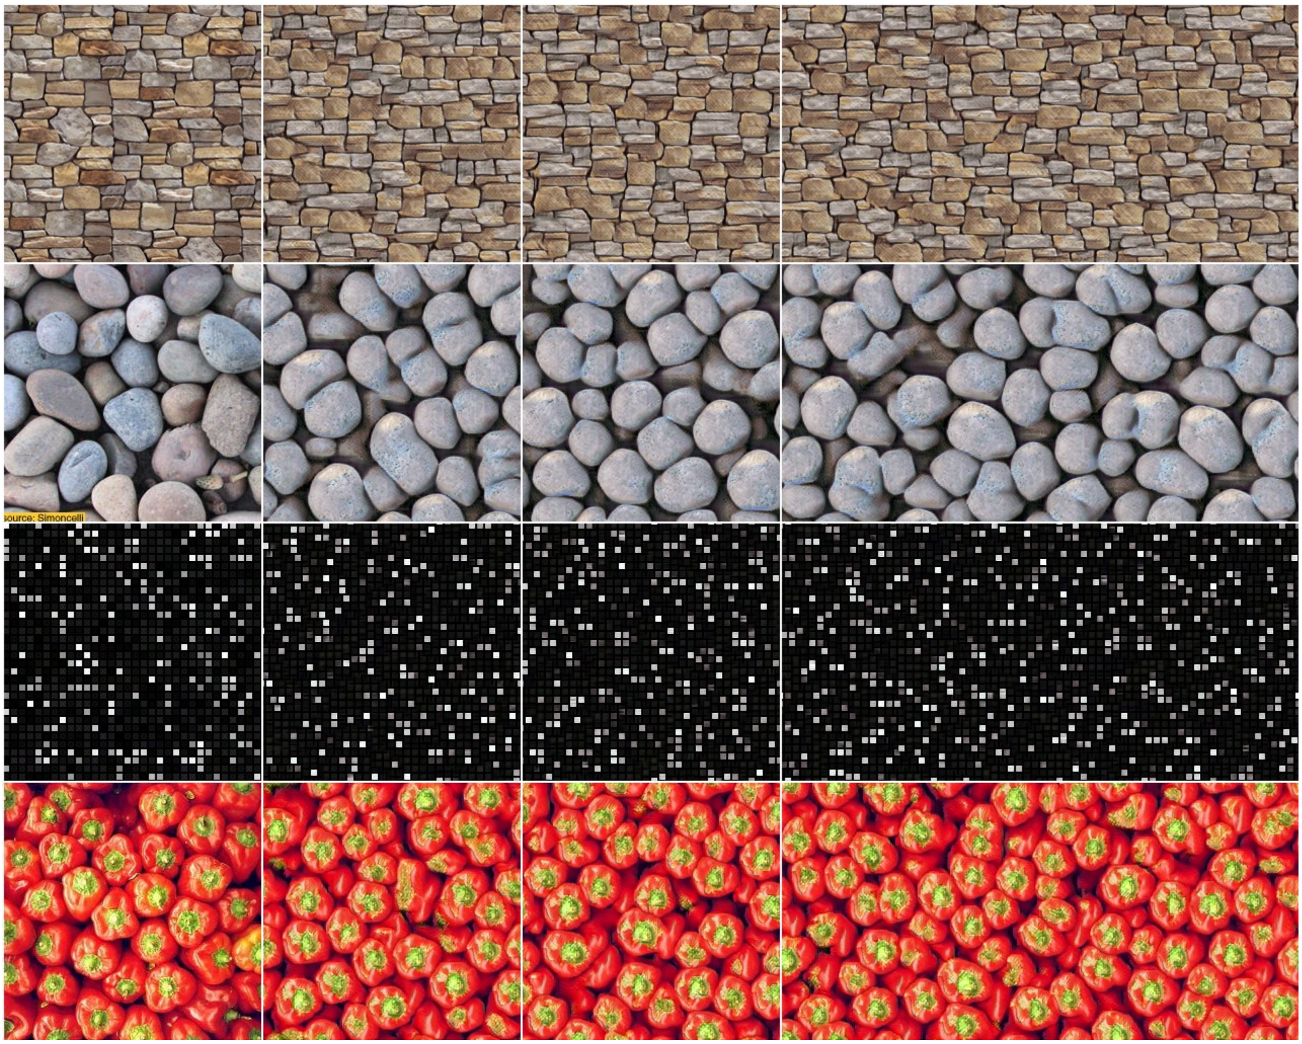

Input Sample 1 Sample 2 Sample 512x256

Figure 2. Various samples drawn from three texture networks trained for the samples shown on the left. While training was done for `256x256` samples, in the right column we show the generated textures for a different resolution.

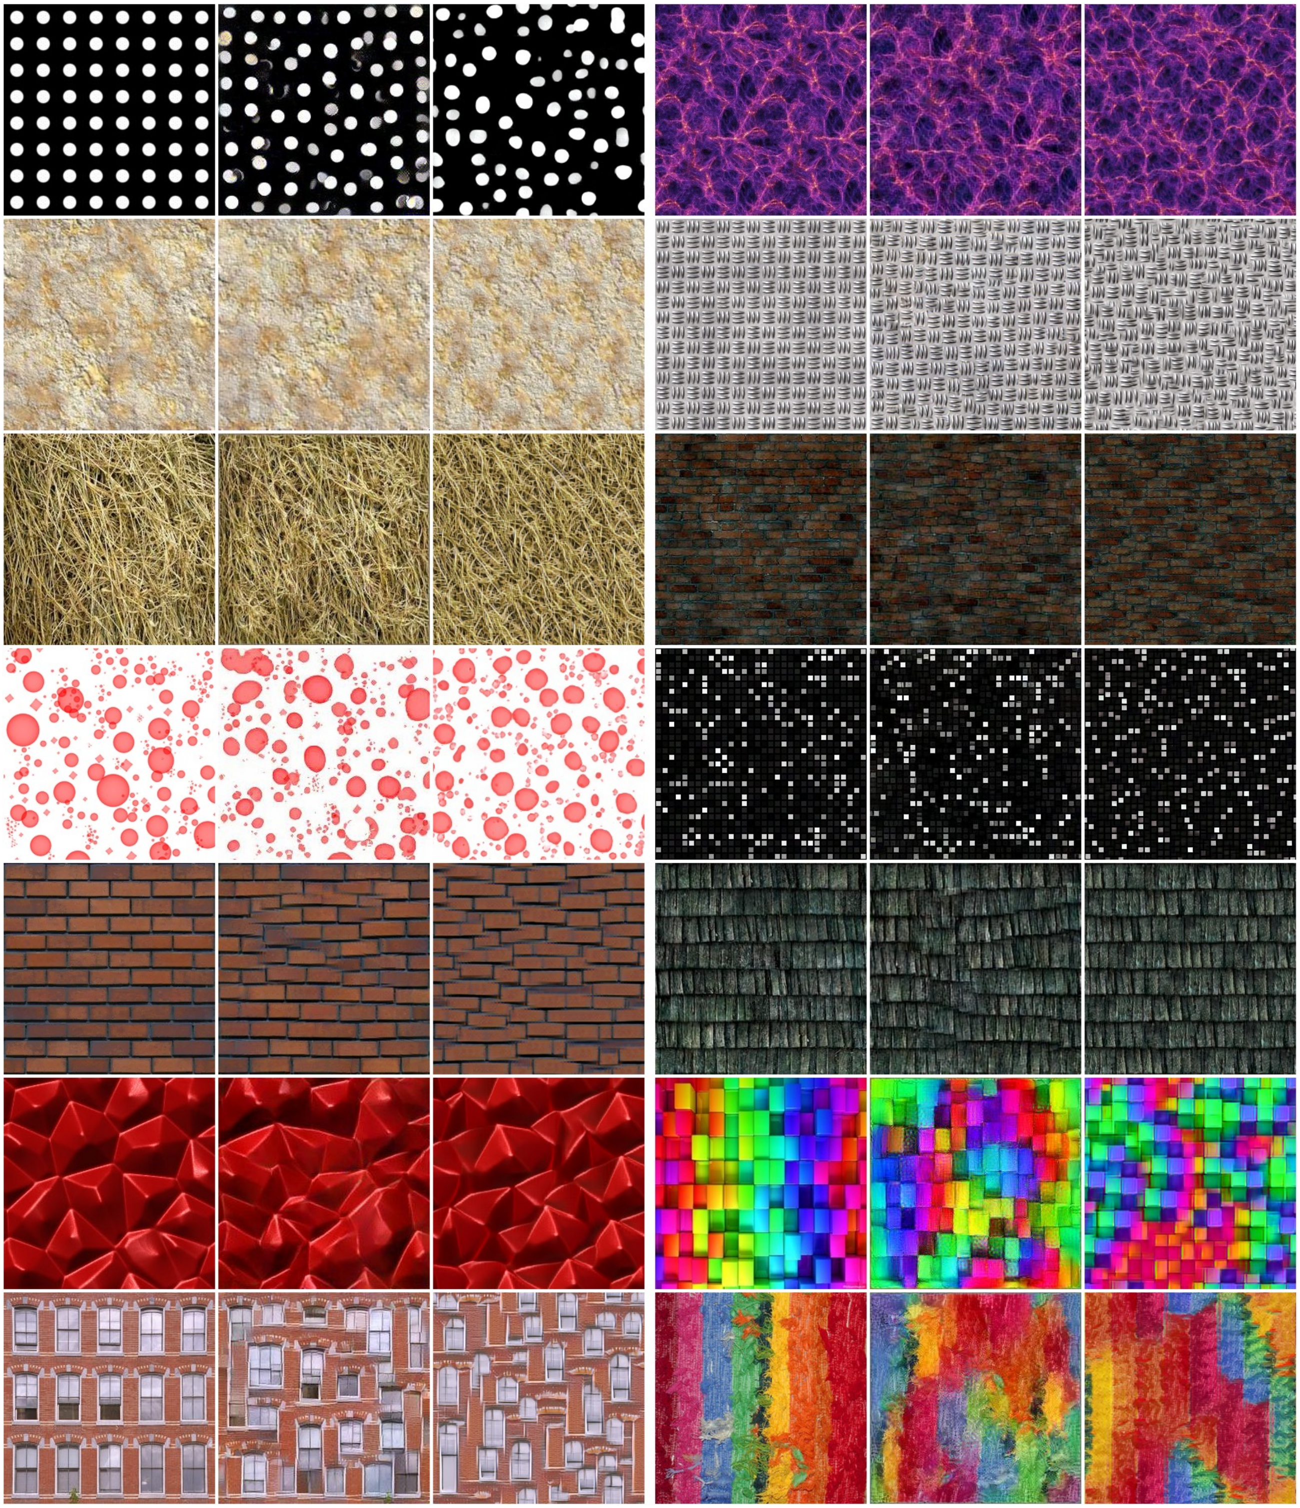

Input | Gatys et al. | Texture nets | Input | Gatys et al. | Texture nets

Figure 3. More comparison with Gatys et al. for texture synthesis.

Figure 4. Style transfer results for more styles. As before, each row represents a network trained with particular style with the original style image on the left, and columns correspond to different content images, with the original image on the top.

Figure 5. Being fully convolutional, our architecture is not bound to image resolution it was trained with. Above, a network trained to stylize 256×256 images was applied to 1024×768 reproduction of Pieter Bruegel's "The Tower of Babel".

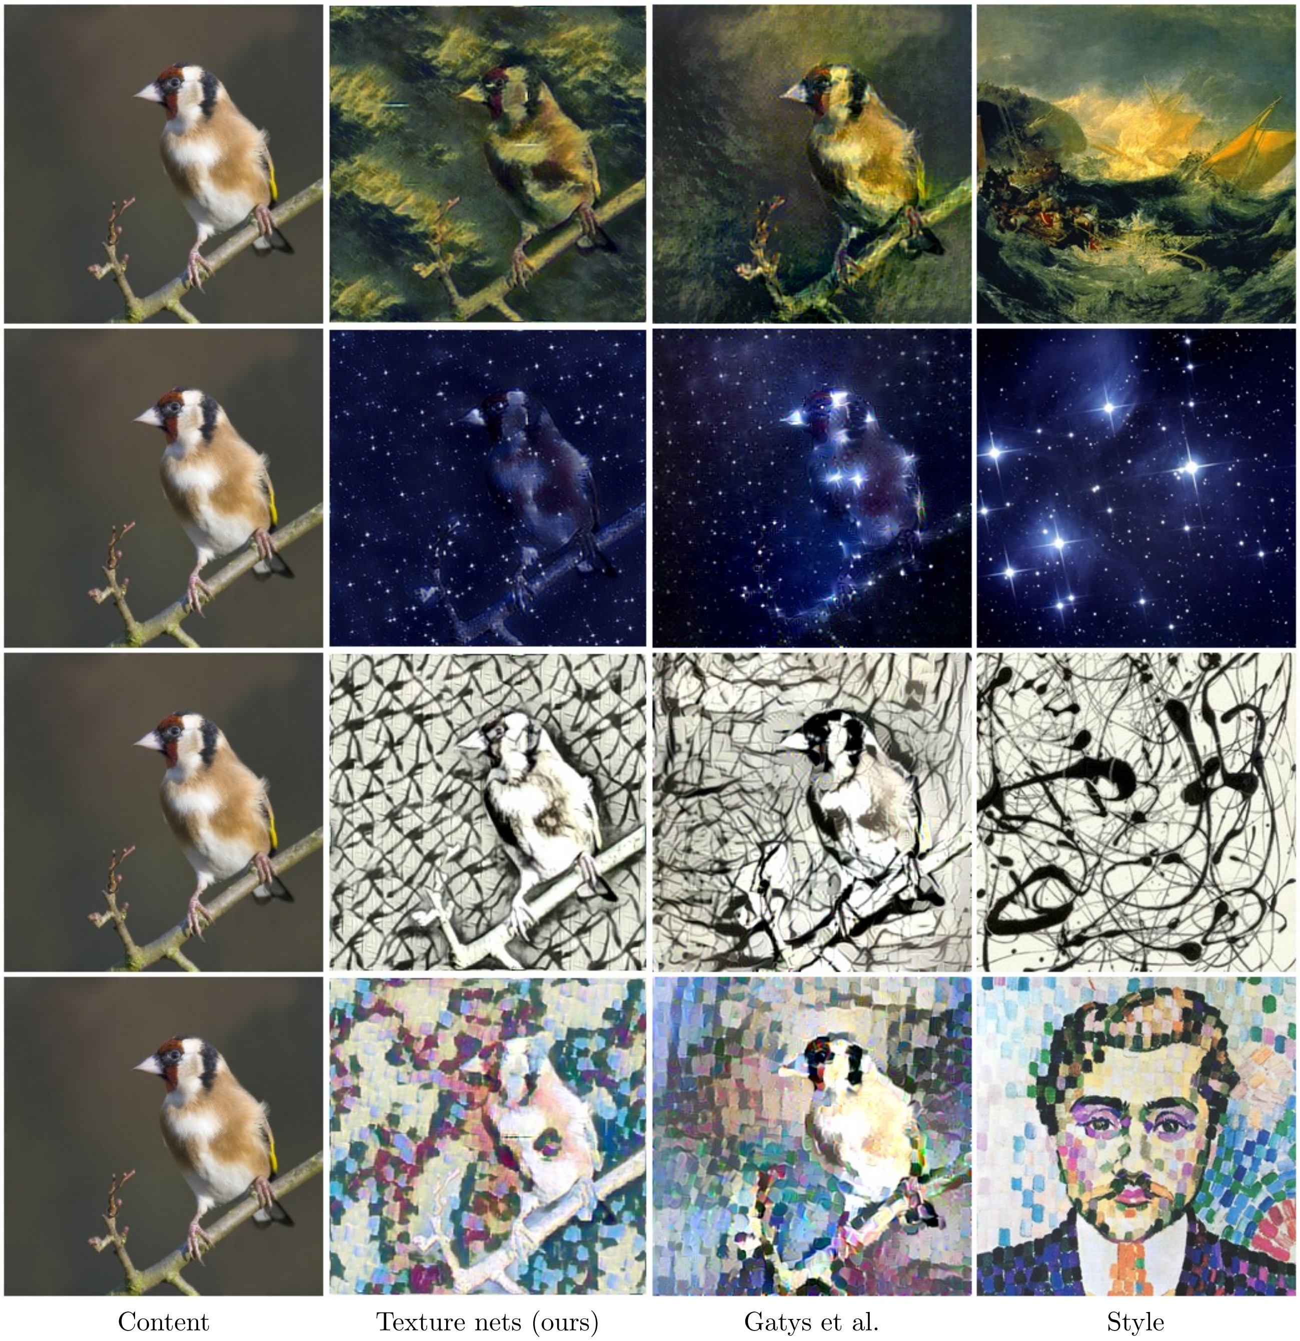

Figure 6. More style transfer comparisons with Gatys et al. For the styles above, the results of our approach are inferior to Gatys et al. It remains to be seen if more complex losses or deeper networks can narrow down or bridge this performance gap.