

Image Inpainting for High-Resolution Textures using CNN Texture Synthesis

P. Laube*, M. Grunwald*, M. O. Franz* and G. Umlauf*

* Institute for Optical Systems, University of Applied Sciences Konstanz, Germany

Abstract

Deep neural networks have been successfully applied to problems such as image segmentation, image super-resolution, coloration and image inpainting. In this work we propose the use of convolutional neural networks (CNN) for image inpainting of large regions in high-resolution textures. Due to limited computational resources processing high-resolution images with neural networks is still an open problem. Existing methods separate inpainting of global structure and the transfer of details, which leads to blurry results and loss of global coherence in the detail transfer step. Based on advances in texture synthesis using CNNs we propose patch-based image inpainting by a CNN that is able to optimize for global as well as detail texture statistics. Our method is capable of filling large inpainting regions, oftentimes exceeding the quality of comparable methods for high-resolution images. For reference patch look-up we propose to use the same summary statistics that are used in the inpainting process.

CCS Concepts

•Computing methodologies → Neural networks; Image processing;

1. Introduction

Image inpainting is the process of filling missing or corrupted regions in images based on surrounding image information so that the result looks visually plausible. Most image inpainting approaches are based on sampling existing information surrounding the inpainting region, which is called exemplar-based [CPT04, WSI07, KEBK05, EL99, BSFG09] inpainting. Recently machine learning techniques have been applied successfully to the problem of texture synthesis and inpainting [LW16, GEB16, JAFF16, DB16]. First introduced by Gatys et al. in [GEB15] texture synthesis CNNs have been shown to surpass well-known methods like the one by Portilla et al. [PS00] for many textures. Wallis et al. [WFE*17] recently showed that artificial images produced from a parametric texture model closely match texture appearance for humans. Especially, the CNN texture model of [GEB15] and the extension by Liu [LGX16] are able to capture important aspects of material perception in humans. For many textures the synthesis results are indistinguishable under foveal inspection. Other methods like the ones by Phatak et al. [PKD*16] and Yang et al. [YLL*16] train auto-encoder-like networks, called context-encoders, for inpainting. Inpainting methods using neural networks still suffer two main drawbacks: Due to limited computational resources they are restricted to small inpainting regions and results often lack details and are blurry. For high-resolution textures the inpainting result not only needs to reproduce texture details but also global structure. Applying details after a first coarse inpainting step distorts global statistics. Fig. 1

shows some examples where well-known inpainting methods fail to reproduce global and local structure. To resolve the outlined is-




Figure 1: Inpainting results for some of the example textures from Fig. 5 using the methods of (a) Photoshop CS7 which is a combination of methods [BSFG09] and [WSI07], (b) the method by Yang et al. [YLL*16], (c) and by Criminisi et al. [CPT04].

sues we propose an inpainting approach that produces results that reproduce global statistics and contain blur-free details. We fill the inpainting region by synthesis of new texture patch by patch, which enables us to process high-resolution textures. Our inpainting approach creates a smooth transition between the sampling and the inpainting region as well as between patches. Our setup is able to shift focus from optimizing detail to global statistics on different levels of resolutions.

Sections of this paper are arranged as follows. The process of tex-

ture synthesis by CNNs is then explained in Sec. 2. In Sec. 3 we present our inpainting approach, followed by an experimental evaluation in Sec. 4. We conclude in Sec. 5.

2. Texture synthesis

First introduced by Gatys et al. [GEB15] CNN texture synthesis uses summary statistics derived from filter responses of convolutional layers, the feature maps, to synthesize new texture. In a first step some vectorized texture \mathbf{x} of size P is presented to the analysis CNN. Based on the resulting feature maps one can compute the Gramians which are spatial summary statistics. The Gramian of some network layer l is defined as

$$\mathbf{G}_{ij}^l = \sum_k \mathbf{F}_{ik}^l \mathbf{F}_{jk}^l, \quad (1)$$

where \mathbf{F}_{ik}^l is feature map i and \mathbf{F}_{jk}^l feature map j at location k given input \mathbf{x} . These inner products of filter activations of different layers are then used to define a synthesis loss

$$\mathcal{L}_s(\mathbf{x}, \hat{\mathbf{x}}) = \sum_{l=0}^L \frac{1}{2N_l^2 M_l^2} \sum_{i,j} (\mathbf{G}_{ij}^l - \hat{\mathbf{G}}_{ij}^l)^2,$$

with N_l feature maps of size M_l at layer l . Here $\hat{\mathbf{G}}_{ij}^l$ are the Gramians of a synthesis CNN. Based on this loss some randomly initialized input vector $\hat{\mathbf{x}}$ of the synthesis CNN is optimized to satisfy statistics derived from the analysis CNN. Since Gramians average over feature map positions this leads to a loss of global texture coherence. Berger and Memisevic [BM16] introduce a second cross-correlation loss by computing Gramians between feature maps \mathbf{F}^l and a spatial translation T of the feature maps $T(\mathbf{F}^l)$. By discarding either rows or columns of feature maps one can now compute correlations of features at some location $k = (x, y)$ and a shifted location $T_{x,+\delta}(k) = (x + \delta, y)$ or $T_{y,+\delta}(k) = (x, y + \delta)$. The horizontally translated Gramian becomes

$$\mathbf{G}_{x,\delta,ij}^l = \sum_k T_{x,+\delta}(\mathbf{F}_{ik}^l) T_{x,-\delta}(\mathbf{F}_{jk}^l), \quad (2)$$

and $\mathbf{G}_{y,\delta,ij}^l$ analogous. The cross-correlation loss \mathcal{L}_{cc} for an arbitrary shift δ is defined as

$$\mathcal{L}_{cc}(\mathbf{x}, \hat{\mathbf{x}}) = \sum_{l,i,j} \frac{(\mathbf{G}_{x,\delta,ij}^l - \hat{\mathbf{G}}_{x,\delta,ij}^l)^2 + (\mathbf{G}_{y,\delta,ij}^l - \hat{\mathbf{G}}_{y,\delta,ij}^l)^2}{4N_l^2 M_l^2}.$$

The combined loss is then defined as

$$\mathcal{L}_{s,cc}(\mathbf{x}, \hat{\mathbf{x}}) = w_s \mathcal{L}_s + w_{cc} \mathcal{L}_{cc},$$

with weight factors w_s and w_{cc} . The interested reader is referred to [GEB15] and [BM16] for an in-depth explanation of texture synthesis.

3. Patch-based texture synthesis for image inpainting

3.1. Patch-based texture synthesis

Given some image with high-resolution, uncorrupted texture Φ we propose the application of the synthesis method introduced in Sec. 2 on different scales of resolution to fill the inpainting region Ω (Fig. 3a). A schematic overview of our setup is given in Fig. 2. We

propose to inpaint region Ω patch by patch with each patch satisfying global as well as detail statistics. For this purpose, we define a texture loss function that simultaneously evaluates the quality of the synthesized patch $\hat{\mathbf{x}}_d$ in native resolution as well as the quality of an embedding of $\hat{\mathbf{x}}_d$ into a pooled window of its surroundings $\hat{\mathbf{x}}_g$ capturing global information. $\hat{\mathbf{x}}_g$ is initialized with a Q -times average-pooled window of the image so that this window fully contains Ω and the boundary Ψ . Q average-pooling-layers are introduced in-between $\hat{\mathbf{x}}_d$ and $\hat{\mathbf{x}}_g$ so that $\hat{\mathbf{x}}_d$ can become a subtensor of $\hat{\mathbf{x}}_g$ at the correct (pooled) position. Depending on the size of Ω , Q needs to be adjusted as a parameter before inpainting. Before generating the next patch $\hat{\mathbf{x}}_d$ at a new location we update Ω with the synthesis result in $\hat{\mathbf{x}}_d$ and reinitialize $\hat{\mathbf{x}}_g$. Only $\hat{\mathbf{x}}_d$ is optimized in the synthesis process. For the synthesis as described in Sec. 2, suitable reference textures \mathbf{x}_d and \mathbf{x}_g are needed. We will describe the reference patch look-up in Sec. 3.2. While \mathbf{x}_g needs to be initialized only once at the beginning of the inpainting process, \mathbf{x}_d is reinitialized with a new reference for every new position of the inpainting patch $\hat{\mathbf{x}}_d$. We further define a boundary loss, that limits the optimization of region Ψ inside $\hat{\mathbf{x}}_d$ in the input domain. We define the boundary loss as

$$\mathcal{L}_b(\mathbf{x}, \hat{\mathbf{x}}) = \frac{1}{p} \sum (\mathbf{m}\mathbf{x} - \mathbf{m}\hat{\mathbf{x}})^2, \quad (3)$$

where the binary mask \mathbf{m} equals 0, if $\mathbf{m}_i \in \Omega$, and 1 otherwise. The combined loss over both branches together with boundary loss becomes

$$\mathcal{L} = w_d \mathcal{L}_{s,cc}(\mathbf{x}_d, \hat{\mathbf{x}}_d) + w_g \mathcal{L}_{s,cc}(\mathbf{x}_g, \hat{\mathbf{x}}_g) + w_b \mathcal{L}_b(\mathbf{x}_b, \hat{\mathbf{x}}_d),$$

where w_d , w_g , and w_b are weight terms. \mathbf{x}_b is initialized with $\hat{\mathbf{x}}_d$ before optimization and does change for each new position of $\hat{\mathbf{x}}_d$.

3.2. Patch distance by Gramians

For the synthesis of patch $\hat{\mathbf{x}}_d$, suitable reference patches \mathbf{x}_d , and \mathbf{x}_g are needed. The initial $\hat{\mathbf{x}}_d$ is a window of the image containing parts of Ψ as well as parts of Ω while $\hat{\mathbf{x}}_g$ completely contains Ψ and Ω . One now has to find closest patches from Φ matching Ψ inside $\hat{\mathbf{x}}_d$, and $\hat{\mathbf{x}}_g$ as candidates for \mathbf{x}_d , and \mathbf{x}_g . Instead of the MSE, we propose to use the distance of texture Gramians as a similarity measure. Since values inside Ω are unknown we propose masking Ω for each individual feature map to remove Ω -related correlations from any of the resulting Gramians. Because the network input up to some layer l has passed through both pooling and convolutional layers we need to adapt the feature map masks to compensate for these operations. In a first step, the initial binary mask \mathbf{m} from Eq. (3) needs to be adapted in size to account for the pooling layers. This is done by applying each pooling step of the CNN that has been applied up to layer l to the mask \mathbf{m}^l which is responsible for masking feature maps \mathbf{F}^l . In a second step, one needs to account for the propagation of Ω into Ψ by convolutions. Masks \mathbf{m}^l also need to account for propagation of Ω into Ψ due to convolution. Simply discarding the affected values by setting them to zero in \mathbf{m}^l for each convolutional layer is too restrictive and would lead to masks with all values zero in later layers. We propose to expand Ω by a smaller individual number of pixels e^l for each convolutional layer (see Sec. 4). In our experiments this expansion has proven to be sufficient for compensation.

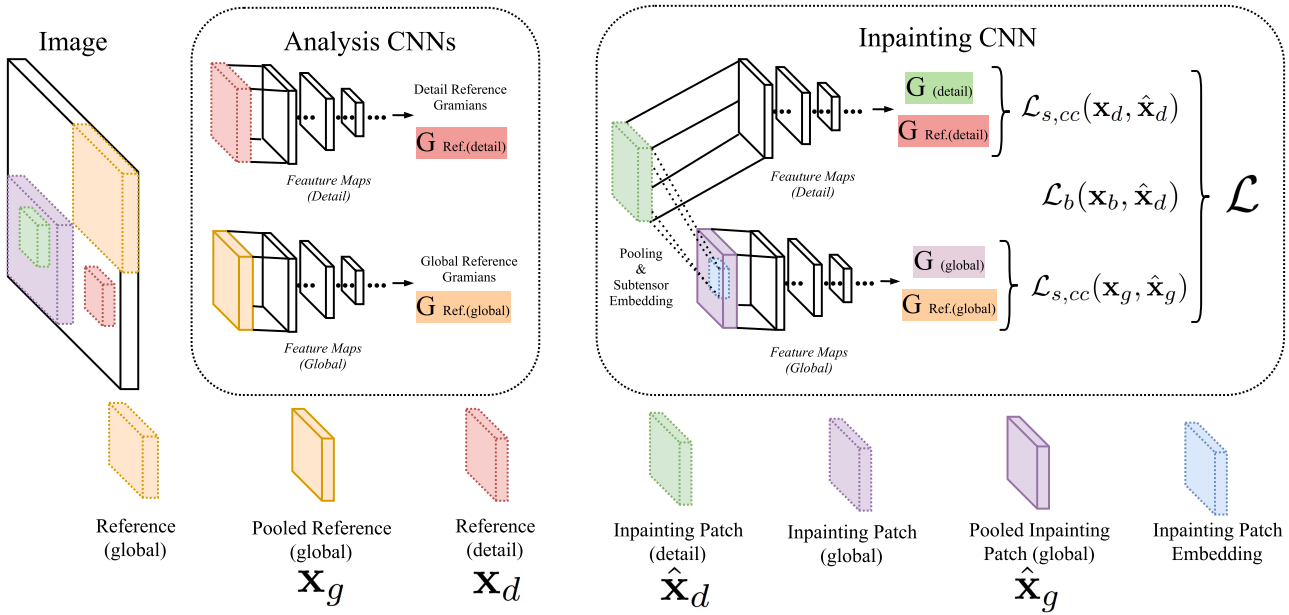


Figure 2: Scheme of our proposed inpainting setup. On the top left the inpainting image together with important image regions is shown. Under "Analysis CNN" the generation of detail as well as reference Gramians is shown. On the top right our "Inpainting CNN" together with the resulting loss terms is shown. The inpainting patch $\hat{\mathbf{x}}_d$ is input to the detail branch (top) as well as, after embedding, the global branch (bottom). A legend of the involved image elements is given at the bottom.

Taking these considerations into account we define our patch distance as

$$\Delta_{\mathbf{G}}(\mathbf{x}, \hat{\mathbf{x}}) = \sum_{l,i,j} \left(\sum_k \mathbf{m}^l \mathbf{F}_{ik}^l \mathbf{F}_{jk}^l - \sum_k \mathbf{m}^l \hat{\mathbf{F}}_{ik}^l \hat{\mathbf{F}}_{jk}^l \right)^2.$$

3.3. Inpainting

For inpainting we propose a coarse to fine inpainting process with two steps. At each stage $\hat{\mathbf{x}}_d$ is optimized by applying L-BFGS-B [ZBLN97]. We initialize each color channel in region Ω with the corresponding color channel mean from Φ . In the coarse inpainting step we focus on optimizing global statistics by setting $w_d = 0$, $w_g = 1$. This leads to $\hat{\mathbf{x}}_d$ satisfying global statistics but at low resolution. Pooling larger input regions introduces color artifacts since loss is shared among pooled pixels as can be seen in Fig. 3b. We eliminate these color artifacts by converting Ω to greyscale (see Fig. 3c) with RGB weights $r = 0.212$, $g = 0.7154$, $b = 0.0721$. Only this structure is used for initialization of the second stage. In the fine inpainting step we set $w_d = 1$ and w_g to a value in the range of $[0.01, 0.1]$. This ensures focus on the optimization for detail statistics through the *detail branch* while constraining the optimization to also maintain global texture statistics.

For our approach inpainting order is not important as long as the first patch overlaps with Ψ and consecutive patches overlap. Overlapping a patch by $\frac{1}{4}$ of its own size with surrounding texture has proven to be sufficient for smooth boundary transition. We chose to fill Ω in a top to bottom, left to right fashion. To ensure a smooth transition in-between patches we apply image quilting [EF01] on overlaps. As a result of our experiments we set $w_s = 1e6$ and $w_{cc} = 1e7$ for inpainting of 8-bit color images. Choosing w_b in

the range $[5, 25]$ has shown to be sufficient. The large difference between Gramian-based loss weights and weights related to loss in pixel space results from different value ranges.




Figure 3: (a) Example image (2048x2048px) with inpainting region Ω , boundary Ψ and texture Φ . (b) First patch of the coarse inpainting step. (c) Fine inpainting of texture detail after coarse inpainting.

4. Experimental Evaluation

We present inpainting results of exemplar high resolution textures. All textures have a resolution of 2048x2048px while the inpainting region Ω is of size 512x512px. We use ImageNet pre-trained VGG-19 CNNs for analysis as well as synthesis with input size 256x256px. We use layers *conv1_1*, *pool1*, *pool2*, *pool3* and *pool4* for computing global as well as detail statistics. For very stochastic textures we propose to use *pool3*, *pool4* and *pool5* to compute global statistics since this leads to improved texture scale in

the coarse inpainting step. For patch distance computation we define pixel expansions $\mathbf{e} = (1, 1, 2, 3, 2)$, and for shift δ of translated Gramians $\mathbf{G}_{x,\delta}^l$ and $\mathbf{G}_{y,\delta}^l$ we define $\delta = (6, 6, 5, 4, 3)$. We use $Q = 2$ pooling layers. To find suitable reference patches \mathbf{x}_d and \mathbf{x}_g region Φ is searched at a step size of 64px. Inpainting of the exemplar textures was done using a Nvidia GeForce 1080 Ti and took roughly 8 min strongly depending on the number of iterations of the L-BFGS-B optimization. In Fig. 6 we present results of our inpainting approach for inpainting Ω of the example textures in Fig. 5. While many methods have difficulties maintaining global as well as local texture characteristics our results look reasonable on both scales. Using the difference of masked Gramians as a metric for patch distance has major benefits for our inpainting approach over using simple MSE. Since we are not dependent on reference textures \mathbf{x}_d or \mathbf{x}_g exactly matching Ψ inside $\hat{\mathbf{x}}_d$ or $\hat{\mathbf{x}}_g$ in terms of MSE, we can reduce the number of samples taken from Φ in reference patch look-up. Due to the averaging of feature information inside Gramians, global spacial information is lost. This enables the Gramian to represent texture invariant to rotation and translation to some degree (see Fig. 4). Because our loss term \mathcal{L} is based on the difference of Gramians this further ensures that Ψ inside $\hat{\mathbf{x}}$ already satisfies target statistics to some extent. When choosing w_d and w_g




Figure 4: (a) Inpainting patch $\hat{\mathbf{x}}_d$. (b) Closest reference patch from Φ . (c) Inpainting result.



Figure 5: Examples for evaluation with inpainting region Ω .

one needs to be aware of the trade-off introduced. While higher w_g ensures persistence of global statistics it also introduces artifacts as a result of pooling $\hat{\mathbf{x}}_d$ before subtensor embedding and vice versa. Higher w_d lays larger emphasis on details while possibly violating



Figure 6: Closeup on results using our method to inpaint region Ω from Fig. 5.

global structure. This trade-off is further influenced by the number of poolings Q .

5. Conclusion

In this work, we presented a new CNN-based method for inpainting that can be applied to large-scale, high-resolution textures. Texture analysis and inpainting are done on two scales, one for global structure and one for details. This avoids the problems of blurry or missing details from which previous CNN approaches suffered while plausibly continuing global image structure. In principle, our network architecture can be extended to include a hierarchy of more than two interacting scales. The design of such a multi-resolution architecture could be an interesting line of research that we plan to pursue in the future. We would also like to apply this concept to the inpainting of height-maps for highly detailed 3D surfaces.

Acknowledgments

This research is funded by the Federal Ministry of Education and Research (BMBF) of Germany (project number 02P14A035).

References

- [BM16] BERGER G., MEMISEVIC R.: Incorporating long-range consistency in cnn-based texture generation. *arXiv preprint arXiv:1606.01286* (2016). 2
- [BSFG09] BARNES C., SHECHTMAN E., FINKELSTEIN A., GOLDMAN D. B.: Patchmatch: A randomized correspondence algorithm for structural image editing. *ACM Trans. Graph.* 28, 3 (2009), 24–1. 1
- [CPT04] CRIMINISI A., PÉREZ P., TOYAMA K.: Region filling and object removal by exemplar-based image inpainting. *IEEE Transactions on image processing* 13, 9 (2004), 1200–1212. 1
- [DB16] DOSOVITSKIY A., BROX T.: Generating images with perceptual similarity metrics based on deep networks. In *Advances in Neural Information Processing Systems* (2016), pp. 658–666. 1
- [EF01] EFROS A. A., FREEMAN W. T.: Image quilting for texture synthesis and transfer. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques* (2001), ACM, pp. 341–346. 3
- [EL99] EFROS A. A., LEUNG T. K.: Texture synthesis by non-parametric sampling. In *Computer Vision, 1999. The Proceedings of the Seventh IEEE International Conference on* (1999), vol. 2, IEEE, pp. 1033–1038. 1
- [GEB15] GATYS L., ECKER A. S., BETHGE M.: Texture synthesis using convolutional neural networks. In *Advances in Neural Information Processing Systems* (2015), pp. 262–270. 1, 2
- [GEB16] GATYS L. A., ECKER A. S., BETHGE M.: Image style transfer using convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2016), pp. 2414–2423. 1
- [JAFF16] JOHNSON J., ALAHI A., FEI-FEI L.: Perceptual losses for real-time style transfer and super-resolution. In *European Conference on Computer Vision* (2016), Springer, pp. 694–711. 1
- [KEBK05] KWATRA V., ESSA I., BOBICK A., KWATRA N.: Texture optimization for example-based synthesis. *ACM Transactions on Graphics (ToG)* 24, 3 (2005), 795–802. 1
- [LGX16] LIU G., GOUSSEAU Y., XIA G.-S.: Texture synthesis through convolutional neural networks and spectrum constraints. In *Pattern Recognition (ICPR), 2016 23rd International Conference on* (2016), IEEE, pp. 3234–3239. 1
- [LW16] LI C., WAND M.: Combining markov random fields and convolutional neural networks for image synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2016), pp. 2479–2486. 1
- [PKD*16] PATHAK D., KRAHENBUHL P., DONAHUE J., DARRELL T., EFROS A. A.: Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2016), pp. 2536–2544. 1
- [PS00] PORTILLA J., SIMONCELLI E. P.: A parametric texture model based on joint statistics of complex wavelet coefficients. *International journal of computer vision* 40, 1 (2000), 49–70. 1
- [WFE*17] WALLIS T. S., FUNKE C. M., ECKER A. S., GATYS L. A., WICHMANN F. A., BETHGE M.: A parametric texture model based on deep convolutional features closely matches texture appearance for humans. *Journal of Vision* 17, 12 (2017), 5–5. 1
- [WSI07] WEXLER Y., SHECHTMAN E., IRANI M.: Space-time completion of video. *IEEE Transactions on pattern analysis and machine intelligence* 29, 3 (2007). 1
- [YLL*16] YANG C., LU X., LIN Z., SHECHTMAN E., WANG O., LI H.: High-resolution image inpainting using multi-scale neural patch synthesis. *arXiv preprint arXiv:1611.09969* (2016). 1
- [ZBLN97] ZHU C., BYRD R. H., LU P., NOCEDAL J.: Algorithm 778: L-bfgs-b: Fortran subroutines for large-scale bound-constrained optimization. *ACM Transactions on Mathematical Software (TOMS)* 23, 4 (1997), 550–560. 3