Image Restoration Using Deep Learning

Jonas De Vylder, Simon Donne, Hiep Luong, Wilfried Philips   JONAS.DEVYLDER@UGENT.BE

dept. of Telecommunications and Information Processing, iMinds, IPI, Ghent University, Belgium

Keywords: deep learning, sharpening, denoising

Abstract

We propose a new image restoration method that reduces noise and blur in degraded images. In contrast to many state of the art methods, our method does not rely on intensive iterative approaches, instead it uses a pre-trained convolutional neural network.

1. Image degradation and restoration

Due to physical and technical drawbacks of imaging devices the acquisition process introduces artifacts. Common examples are blur due to limited lens setups or noise when taking a picture in bad light conditions, e.g. at nightfall. To mitigate these restrictions image reconstruction algorithms aim at removing these artifacts. In this paper we will present a new method to overcome both blur and noise. To limit the scope of this work, we restrict our self to the following image degradation model

$$\mathbf{y} = H\mathbf{x} + \mathbf{n}, \tag{1}$$

where $\mathbf{y}$ corresponds to the observed image, $\mathbf{x}$ symbolizes the unknown non-degraded image H represents a blur kernel and $\mathbf{n}$ is additive white Gaussian noise. In (Roels et al., 2016) we discuss a more general degradation model, but for many applications the model in eq. (1) is adequate. Many state of the art methods estimate $\mathbf{x}$ using iterative methods, which typically result in computational intensive algorithms (Goldstein et al., 2010; Chambolle & Pock, 2011).

The increase in computational power and the existence of large annotated data-sets has led to remarkable results in the field of deep learning, especially with convolutional neural networks (CNNs). CNNs consist of

a set of layers, each typically combining a linear operation, i.e. a convolution, and a non-linear operation such as a rectified linear unit (ReLu) (Glorot & Bengio, 2011). In this work we investigate the use of such a CNN not on the typical classification task, but on image restoration, i.e. estimating the ideal image $\mathbf{x}$ in eq. (1).

In Fig. 1 we show the actual network: we start by expanding using 13 filters, after which we gradually converge to less dimensions. Note that the second layer has no spatial influence, and simply serves for a large degree of non-linearity using the ReLU activation functions. In contrast to methods such as deepdream or deepart (Gatys et al., 2015; Hayes, 2015), we aim at generating a realistic image as output from our network, in contrast to updating the input for our network, since this would result in another iterative computational demanding reconstruction algorithm.

2. Training the network

The network was trained using stochastic gradient descent (Bottou, 2010). In order not to optimize towards local minima, we added additional momentum terms and a shrinkage rate introduced on the filters (Ngiam et al., 2011; Sutskever et al., 2013). The used cost function is the Mean Squared Error (MSE) between the output image and the ground truth. The Berkeley training data-set (BSDS500) was used to train the network (Arbelaez et al., 2011). The training data-set of 200 images were first converted to gray-scale images, filtered with a Gaussian kernel ($\sigma = 2$, filter size = 7), and normalized such that the image intensities fall within $[0, 1]$ before being fetched to the network.

The network was trained using a deep learning framework implemented in the Quasar language for heterogeneous programming (Goossens et al., 2015). Compared to other machine learning frameworks such as caffe, Theano, Torch, etc. Quasar has the advantage that not just the training steps but all calculation in-
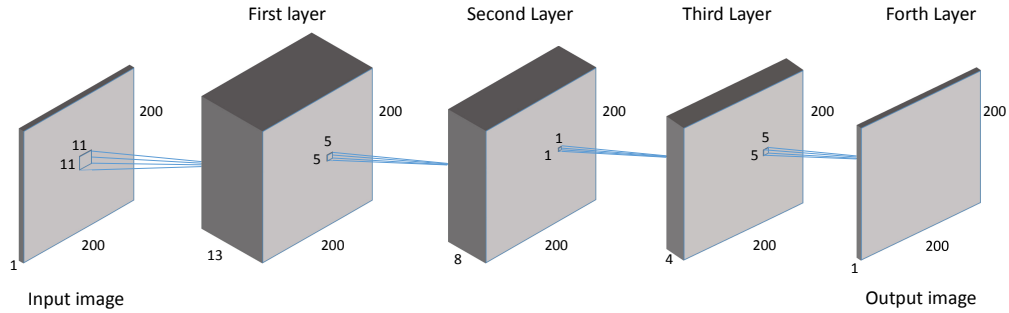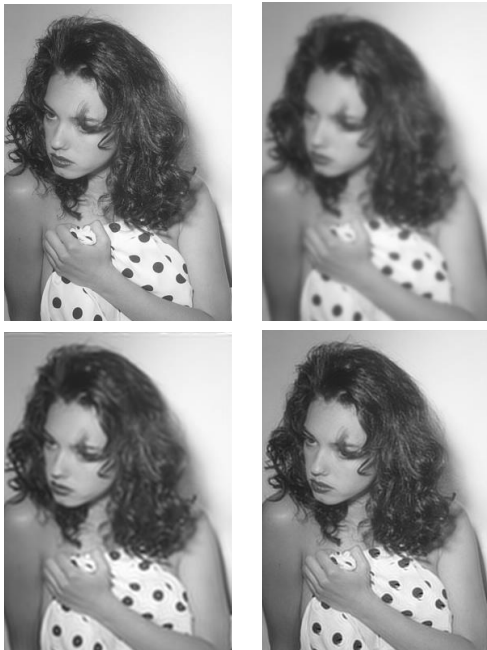
Figure 1. An overview of the proposed CNN.

Figure 2. The result of the propsed method: top left the original non-degraded image, top right shows the blurred input image, the bottom row shows the restored images. Left using iterative restoration from (Chambolle & Pock, 2011), right restoration using the proposed machine learning method.

	input	primal-dual	proposed
blur	24.8 dB	27.0 dB	30.2 dB
blur+noise	24.3 dB	26.2 dB	26.5 dB

Table 1. Quantitative evaluation of the proposed method.

We quantitatively tested two different degradations: blur and the combination of noise and blur. Blur was applied by convolving the input image with a Gaussian kernel ($\sigma = 2$, filter size = 7). For the noise, we added white Gaussian noise with $\sigma = 2$. We applied the same degradation parameters for training and validation. Table 3 summarizes the results. The case with blur clearly results in good restoration. The combination with blur and noise performs similar to state of the art methods (Chambolle & Pock, 2011), albeit less well than for restoration with only blur: while still improving the image with 2 dB, this result is significantly less than the case with only blur. This will require further research where both the impact of larger training data-sets will be tested and the added value of a deeper network will be investigated.

Another relevant aspect is the computational time. The proposed method only takes 493 ms to restore a single image of 512×512 pixels, compared to 5226 ms for a primal-dual based optimization method (Chambolle & Pock, 2011). Both methods are GPU accelerated using a Nvidia GeFORCE 770m.

tensive steps such as image blurring are optimized towards GPU acceleration.

3. Results

We tested the trained network using the Berkeley test data-set, i.e. 200 new images that where not used during training. Fig. 3 shows an example of a degraded input image and the restored result. Note that this input image was only degraded using blur and not by noise.

4. Conclusion

We proposed a new and fast image restoration method. The proposed approach makes use of a pre-trained CNN. Based on the results for deblurring we can state that this approach seems promising, but requires further research for other degradations or combination of image degradations.

References

Arbelaez, P., Maire, M., Fowlkes, C., & Malik, J. (2011). Contour detection and hierarchical image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, *33*, 898–916.

Bottou, L. (2010). Large-scale machine learning with stochastic gradient descent. *Proceedings of compstat 2010* (pp. 177–186). Springer.

Chambolle, A., & Pock, T. (2011). A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, *40*, 120–145.

Gatys, L. A., Ecker, A. S., & Bethge, M. (2015). A neural algorithm of artistic style. *CoRR, abs/1508.06576.*

Glorot, X.and Bordes, A., & Bengio, Y. (2011). Deep sparse rectifier neural networks. *International Conerence on Artificial Intelligence and Statistics* (pp. 315–323).

Goldstein, T., Bresson, X., & Osher, S. (2010). Geometric applications of the split bregman method: Segmentation and surface reconstruction. *Journal of Scientific Computing*, *45*, 272–293.

Goossens, B., Vylder, J. D., & Philips, W. (2015). Quasar : a new heterogeneous programming framework for image and video processing algorithms on cpu and gpu. *Proceedings of the International Conference on Image Processing.*

Hayes, B. (2015). Computer vision and computer hallucinations. *AMERICAN SCIENTIST*, *103*, 380–383.

Ngiam, J., Coates, A., Lahiri, A., Prochnow, B., Le, Q. V., & Ng, A. Y. (2011). On optimization methods for deep learning. *Proceedings of the 28th International Conference on Machine Learning (ICML-11)* (pp. 265–272).

Roels, J., Aelterman, J., De Vylder, J., Luong, Q., Lippens, S., Guerin, C., & Philips, W. (2016). Image degradation in microscopic images: avoidance, artifacts and solutions. *Focus on bio-image informatics*, *219.*

Sutskever, I., Martens, J., Dahl, G., & Hinton, G. (2013). On the importance of initialization and momentum in deep learning. *Proceedings of the 30th international conference on machine learning (ICML-13)* (pp. 1139–1147).