

Learning Nonlinear Spectral Filters for Color Image Reconstruction

Michael Moeller¹, Julia Diebold¹, Guy Gilboa² and Daniel Cremers¹

¹TU Munich, Germany*

²Technion - IIT, Israel

Abstract

This paper presents the idea of learning optimal filters for color image reconstruction based on a novel concept of nonlinear spectral image decompositions recently proposed by Guy Gilboa. We use a multiscale image decomposition approach based on total variation regularization and Bregman iterations to represent the input data as the sum of image layers containing features at different scales. Filtered images can be obtained by weighted linear combinations of the different frequency layers. We introduce the idea of learning optimal filters for the task of image denoising, and propose the idea of mixing high frequency components of different color channels. Our numerical experiments demonstrate that learning the optimal weights can significantly improve the results in comparison to the standard variational approach, and achieves state-of-the-art image denoising results.

1. Introduction

The great success of linear spectral decomposition methods such as the Fourier transform (FT) is based on their ability to represent the input data at different scales. The FT for instance represents a signal as the superposition of sine and cosine of different frequencies, such that one can enhance, damp, or eliminate certain frequencies differently by the design of Fourier filters. While this theory and its applications like high-pass, low-pass, band-pass, or band-stop filterings is well understood for linear transformations, recent works have extended such concepts to nonlinear variational techniques.

In [16, 17] Guy Gilboa proposed to use the total variation (TV) gradient flow to define a notion of nonlinear spectral representations of images. Burger *et al.* generalized this concept to arbitrary one-homogeneous regularizations in [4] and considered three different possible definitions of nonlinear spectral representations.

In all of the above works, the general idea of nonlinear spectral representations is to define a function $\psi(t)$ called

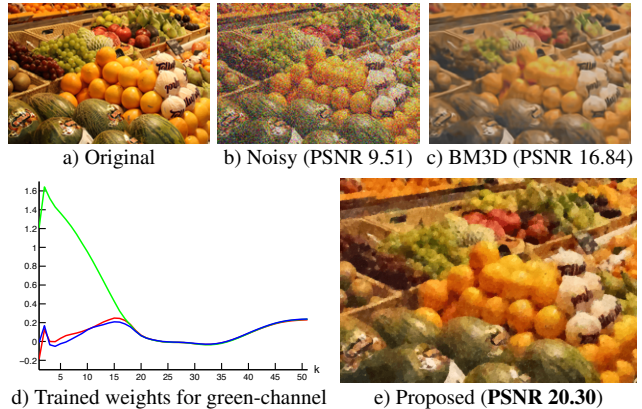


Figure 1. **Learned spectral filtering.** Enhancing low and filtering high frequencies of a spectral total variation image decomposition according to the learned filters shown in d), yields the denoised image in e), which compares favorably to the BM3D algorithm c) at high noise levels.

the frequency representation of the input data f , such that

$$f = \int_0^\infty \psi(t) dt.$$

Similar to the notion of the frequency in classical methods such as the FT, the size of the features in $\psi(t)$ decreases as t increases. The latter motivates the definition of filters $\omega(t)$ in the frequency domain to reconstruct a filtered version

$$u_\omega = \int_0^\infty \omega(t)\psi(t) dt \quad (1)$$

of the input data. The above approach has the flexibility to enhance ($\omega(t) > 1$), damp ($\omega(t) < 1$), or eliminate ($\omega(t) = 0$) different frequencies, where the meaning of the frequency depends on the particular type of decomposition.

In this paper we consider the application of nonlinear spectral image decompositions to *color image denoising*, *i.e.* the task of separating the input data f into the sum of a clean signal \hat{u} and undesirable noise n . In particular, we propose to learn optimal filters ω in (1) on training data sets. To account for inter-channel correlations, we learn the natural spectral relation between different color chan-

*This work was supported by the ERC Starting Grant ‘ConvexVision’.

nels by allowing the reconstruction of each channel to incorporate frequency information from other color channels. Figure 2 illustrates the proposed processing, and Figure 1 demonstrates the effectiveness of the proposed approach in comparison to the state-of-the-art technique of BM3D [12]. While the main focus of our experiments is image denoising, we demonstrate that the general idea can be extended to several image reconstruction problems including contrast enhancement, deblurring and compressed sensing.

In summary, we make the following contributions:

- We study the nonlinear spectral TV decomposition of color images.
- We learn spectral filters for color image denoising on a training data set.
- We propose to mix high frequency components of different color channels to account for inter-channel correlations.
- We demonstrate that the proposed formalism extends to reconstruction problems beyond image denoising.

2. Related work

2.1. Image Denoising

Due to the large amount of relevant work in the field of image denoising, we limit ourselves to a review of a small selection of denoising strategies.

While the first denoising methods applied linear filters, nonlinear variational techniques computing

$$u(t) = \arg \min_u \frac{1}{2} \|u - f\|^2 + tJ(u) \quad (2)$$

for a suitable regularization functional J such as the total variation [25] have revolutionized the field. Today, many patch-based methods such as nonlocal means [3], nonlocal TV [18], dictionary learning (e.g. [13, 21]), the Zoran-Weiss EPLL model [33], patch-based Wiener filtering [9], or the BM3D algorithm [12] yield state-of-the-art denoising results, particularly if the image to be denoised is self-similar. Competitive non-patch based methods are for instance based on learning analysis operators [10], or learning iterative denoising schemes motivated from optimization methods [27].

2.2. Spectral Representations

Methods based on finding different representations which are better suited to separate certain desirable features from parts which ought to be suppressed are widely used in the literature. In addition to classical Fourier analysis, the design of more sophisticated orthogonal transformations based on wavelets has attracted a lot of attention (cf. [22]). While classical wavelets lead to linear transformations, nonlinear methods based on the variational formulation (2) are

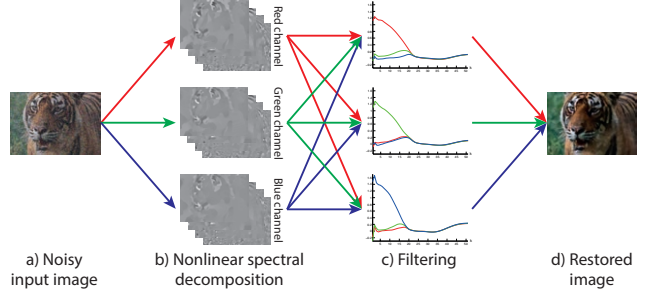


Figure 2. **Spectral filtering.** We determine the nonlinear spectral decomposition of the noisy input image a) as illustrated in b). The frequency layers are recombined with learned optimal filters c) that allow a mixing of frequency components of the different color channels and lead to the final reconstruction shown in d).

popular due to their versatility and ability to preserve sharp edges. Although multiscale decompositions based on (2) have been widely studied (e.g. [1, 30]), the works [4, 16, 17] were – to the best knowledge of the authors – the first to establish a clear analogy between linear and nonlinear spectral decompositions. In the next section we recall these works in more detail, as they are the foundation of the proposed filter learning.

3. Spectral Filtering

The work [4] proposed three different ways of defining nonlinear spectral decompositions, namely via (2), via a gradient flow (as originally studied in [16, 17]), or by considering the so-called *inverse scale space flow* (ISSF) [5, 6].

For the sake of brevity, we will limit our discussion to the ISSF formulation, since this is the method we used in the numerical implementation of our framework. Although the details of the relation between the three different approaches presented in [4] remains an open question, we expect them to yield similar results. For the special case of the data being a nonlinear eigenfunction, i.e. there exists a $\lambda \in \mathbb{R}$ such that $\lambda f \in \partial J(f)$, the exact equivalence of all three spectral decompositions was shown in [4].

Let us detail how to construct nonlinear spectral decompositions. As observed in [23] variational reconstructions via (2) contain a systematic error, bias, or loss of contrast, which can be avoided by considering the Bregman iteration

$$u^{k+1} = \arg \min_u \frac{1}{2} \|u - f\|^2 + \alpha(J(u) - \langle p^k, u \rangle), \quad (3)$$

with $p^k \in \partial J(u^k)$. For large α and $p^0 = 0$, Bregman iteration starts at an approximation u^1 of f such that $J(u^1)$ is small. For J being the TV this means a strong oversmoothing, possibly even a completely constant image. As the iteration proceeds, the iterates u^k converge to f by including finer and finer features. The latter makes Bregman iteration an ideal candidate for spectral decompositions.

The difference between two successive iterates corresponds to features of f that have a particular frequency, such that the k -th frequency component can be defined as

$$\psi^k = u^k - u^{k-1}.$$

As an example, Figure 3 shows three different image frequency components ψ^k . As we can see, the structures are rather large for small k , and rather fine for large k .




Figure 3. **Frequency representation.** Low, medium, and high frequencies of a spectral color TV image decomposition with amplified contrast.

By defining

$$u^0 = \arg \min_u \frac{1}{2} \|u - f\|^2 \text{ s.t. } J(u) = 0,$$

i.e. the orthogonal projection of f onto the kernel of J , and denoting $\psi^0 = u^0$, one obtains

$$f = \sum_{k=0}^{\infty} \psi^k.$$

The original continuous ISSF framework considered in [4] can be recovered from (3) in the limit of $\alpha \rightarrow \infty$.

The above representation of the input data as the sum over contributions of different frequencies motivates filtering approaches of the form

$$u_\omega = \sum_{k=0}^{\infty} \omega_k \psi^k, \quad (4)$$

for *weights* or *filter coefficients* $\omega_k \in \mathbb{R}$, which is a discretization of the time continuous filtering (1).

As an example, consider the ideal low pass filter

$$\omega_k = \begin{cases} 1 & k \leq K, \\ 0 & \text{else,} \end{cases}$$

which restores the K -th Bregman iterate as the filtered solution. In the special case of the input data being a generalized eigenfunction, one can show that the spectral representation just consists of a single peak. In this case even the solution of the variational regularization (2) can be restored by a particular choice of filter coefficients, namely those that decrease linearly to zero.

Considering the popularity of variational methods as well as of Bregman iterative methods, one might not only ask the question if a linearly decaying or a rectangular shaped spectral filter yield better reconstruction results, but also what the optimal shape of a filter is. In this manuscript we propose to learn such optimal filters for TV color image denoising based on a training set of natural images.

4. Learning Spectral Filters for RGB-Images

4.1. Color TV Regularization

We consider color image denoising by TV regularization as an effective and efficient regularization technique. More specifically, we use the color TV definition considered in [2] which originated from [26]. For this type of TV not only the derivatives but also the color channels are coupled in an ℓ^2 fashion, *i.e.*

$$TV(u) = \int_{\Omega} \sqrt{\sum_{i=1}^3 ((\partial_{x_1} u^i(x))^2 + (\partial_{x_2} u^i(x))^2)} dx \quad (5)$$

for $u^i : \Omega \rightarrow \mathbb{R}$ representing the different color channels, red, green, and blue.

Because TV based variational reconstruction methods often obtain improved results if the color channel correlation is avoided by considering transformations into luminance and chrominances (cf. [8, 11]), we additionally consider

$$Au(x) := \begin{pmatrix} 1 & 1 & 1 \\ -1 & 2 & -1 \\ -1 & 0 & 1 \end{pmatrix} \begin{pmatrix} u^1(x) \\ u^2(x) \\ u^3(x) \end{pmatrix},$$

normalize the columns of A and regularize the (uncoupled) total variation of the transformed image, where the weight of the luminance channel is reduced by a factor of 0.75. We refer to this approach as color transformed TV (CTTV).

4.2. Exploring Channel Correlations

The color channels of natural images are often highly correlated. While color differences occur at rather large scales, the high frequency features and textures are often shared by all three color channels. It can therefore be beneficial to consider a mixture of the high frequency components of all color channels for color image restoration. While the mixing surely reduces the standard deviation of the noise, the true texture does not change significantly due to the positive correlation. Naturally, the stronger the correlation between the color channels is, the stronger the mixing may be.

We propose to not only learn the optimal filter on each channel separately, but also learn the correlation between the channels at different frequencies by considering a reconstruction model of the form

$$u_\omega^c = \sum_{l \in \{\text{red, green, blue}\}} \sum_{k=0}^K \omega_k^l \psi_k^l, \quad (6)$$

for all colors $c \in \{\text{red, green, blue}\}$. This way the learning technique can automatically determine the correlation between different color channels at different frequencies k .




Figure 4. **Image set.** For our experiments we used a set of 24 natural images and split it by a ratio of 50/50 for training and testing. The first and second row show the training and test set respectively.

4.3. Learning Optimal Filters

We use a set of clean training images to learn optimal filters based on (6). Denoting the clean images by g_i , we generate noisy images $f_i = g_i + n_i^\sigma$ by adding noise of fixed standard deviation σ to each of the images and compute the spectral decomposition of each of the f_i according to (3). Our goal is to find weights ω_k^l such that the u_ω^c in (6) approximate the g_i^c as closely as possible.

If the total number of pixels in our entire training data set is N and we used $K - 1$ Bregman iterations, we write the weights ω into a $\mathbb{R}^{3K \times 3}$ matrix, arrange the clean images in a matrix $g \in \mathbb{R}^{N \times 3}$, and write the ψ_i^k into a single matrix $\psi \in \mathbb{R}^{N \times 3K}$. Now we can either consider the simple least squares problem $\min_\omega \|\psi\omega - g\|^2$, or – if we expect smooth filtering curves – rather consider a regularized least-squares problem of the form

$$\omega = \arg \min_\omega \|\psi\omega - g\|^2 + \gamma \|\nabla\omega\|^2. \quad (7)$$

Additionally, we considered a non-negativity constraint on the weights which, however, did not lead to improved results.

5. Implementation

We use the primal-dual hybrid gradient (PDHG) method [7, 14, 24, 32] with the adaptive time stepping scheme proposed in [19] and a fixed number of 500 iterations to solve the minimization problems in (3). Additionally, we found initializing the minimization algorithm with the previous u^k to improve the convergence.

Since many more changes of the time continuous flow discretized by (3) happen at small times, we use an adaptive time resolution. Because the reciprocal of the regularization parameter α in (3) acts like a time step we start with a large value of $\alpha = 20$ and decrease the value of α by a factor of 0.92 in each iteration. We compute a total number of 50 iterations according to (3) and set $u^{51} = f$.

Learning the optimal filters via (7) leads to a simple and small linear equation. In our experiments we used $\gamma = 1000$ as a regularization parameter for the 12 images g_i being on a scale from 0 to 1.

Interestingly, the decoupled CTTV regularization of luminance and chrominances did not improve the results of the learned optimal filters, such that we focused on (5).

The complexity of the spectral decomposition itself amounts to 50 TV minimization problems. While our Matlab implementation needs about 85 seconds per TV minimization on a 640×640 image, recent GPU implementations have demonstrated real time capabilities on similar problems (e.g. [28]), which means the full decomposition could be computed in a couple of seconds. Note that the filter learning as well as the application of a filter are extremely cheap and run in real-time. Thus, once the spectral decomposition is computed, even adapting the denoising strength by changing (or interpolating between) learned filters runs in real-time on a CPU.

6. Experimental Results

Our numerical experiments are conducted on a data set of 24 natural color images, which we divided into equally sized training and test sets. We add zero mean Gaussian noise of different standard deviations σ to the images and compute the spectral decomposition as well as the optimal weights as described in Sections 4 and 5.

For a qualitative evaluation, we compare our method to four different techniques. Firstly, we use the method and code from [31] (abbreviated by *DCT*) as a recent, fast, and effective denoising strategy. Secondly, we compare our approach to the results obtained by TV denoising because this method is most closely related to our approach. Since the CTTV yielded in better denoising results, we limit the presentation of the results to this definition. Finally, we include the block matching 3D (*BM3D*) algorithm [12] with code from [20] as a state-of-the-art technique into our comparison.

For each method, each image, and each noise level, we compute the peak signal to noise ratio (*PSNR*) as well as the structural similarity index (*SSIM*) [29] which better reflects the visual quality of the images.

Let us first consider the optimal filters found by the learning procedure (7) shown in Figure 5 for two different noise levels. The first column shows the learned filter coefficients used for the reconstruction of the red channel. The red

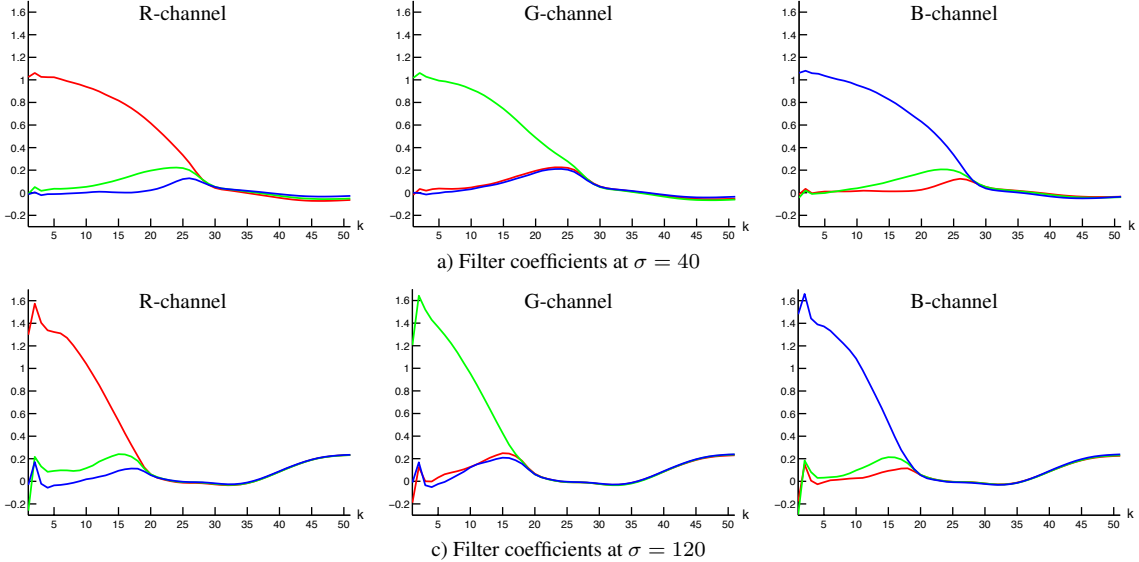


Figure 5. **Learned optimal filters.** Filter coefficients $w_k^l, l \in \{\text{red, green, blue}\}$ (cf. Equation (6)) used for the reconstruction of the R, G and B-channel at different standard deviations σ .

curve within the first column corresponds to the filter coefficients of the red frequencies used for red channel reconstruction. Respectively, the green and blue curves illustrate the weights with which the green and blue frequencies contribute to the red channel reconstruction. The second and third columns show the filter coefficients for the green and blue channel in a similar fashion. Note that the x -axis simply corresponds to the Bregman iterates of our algorithm. In a time continuous representation, the k -th iterate corresponds to the time $t^k = \sum_{i=1}^k \frac{1}{20 \cdot 0.92^i}$.

We can see that, for a fixed noise level, the main filter curves, *i.e.* the red filter coefficients used for the reconstruction of the red channel, the green filter coefficients used for the reconstruction of the green channel, and the blue filter coefficients used for the reconstruction of the blue channel, look very similar. Interestingly, the optimal filter coefficients underline our assumption that at medium to high frequencies it makes sense to mix the different color channels. Note that the correlation between red and green as well as the correlation between blue and green is higher than the red-blue correlation, which is to be expected considering their respective distances in the electromagnetic spectrum.

As the noise level increases, the iterate at which the filters reach a value of zero moves from about 30 at $\sigma = 40$ to 20 at $\sigma = 120$ because a stronger filtering is required. It is interesting to see that some of the very high frequency components get reincorporated for higher noise levels.

A third interesting observation is the fact that the low frequencies are boosted with filter coefficients larger than one as the noise level increases. Since the noisy images used for our experiments were saved in the usual 8-bit format, values below 0 or above 255 are clipped. The latter leads

to the noise in a saved image not exactly following a zero-mean Gaussian distribution anymore. Particularly, the mean value of the noise is negative in bright image areas and positive in dark image areas. Since many denoising methods are (locally) mean value preserving, the denoised image will be too dark in bright areas and too bright in dark areas, hence leading to a reduced contrast. This effect can clearly be seen in the result of the BM3D algorithm in Figure 1 c). By boosting low and medium frequencies our method is able to restore the loss of contrast caused by noise clipping. Thus, spectral filtering possibly offers an alternative to incorporating additional transforms into denoising strategies to correct for the aforementioned bias as investigated in [15].

Let us now look at the actual evaluation of the denoising algorithms. The average PSNR and SSIM values over the 12 test images that all methods achieved for different standard deviations σ of the noise are shown in Figure 6 a) and b) respectively. We can see that while the BM3D algorithm yields the best results for low noise levels such as $\sigma = 40$, its performance drops as the noise level increases. TV denoising on the other hand starts with rather low PSNR and SSIM values but does not show an equally fast decay of the quality metric values, such that it yields better results for noise levels above $\sigma = 100$. The DCT denoising pays for its efficiency by showing the weakest denoising performance.

While the proposed approach yields slightly worse results than the BM3D algorithm at $\sigma = 40$, it can handle high noise levels very well, leading to PSNR values about 3dB higher than BM3D at the highest noise level. Moreover, the SSIM metric indicates a significantly higher visual quality of our approach as the noise level increases. The qualitative results shown in Figure 8 underline this indication. The


σ	Image name	Noisy input	DCT [31]	BM3D [12, 20]	CTTV [8, 11]	Proposed	Ideal filters
40	playground	16.40	25.63	26.70	25.66	26.00	26.10
		16.40	26.14	27.60	26.38	26.60	26.80
60	flowers	13.90	21.26	22.60	22.08	23.20	23.70
		13.60	20.48	21.70	21.13	21.50	21.80
80	fruits	12.00	18.80	19.90	19.72	21.90	22.20
		12.20	20.88	21.80	21.45	24.20	25.10
100	facade	10.20	20.54	21.60	21.47	22.90	23.30
		10.40	18.53	19.80	19.61	21.40	21.80
120	zebra	9.30	20.08	20.70	20.68	23.30	24.10
		9.30	20.10	21.10	20.87	24.00	24.40

a) PSNR values

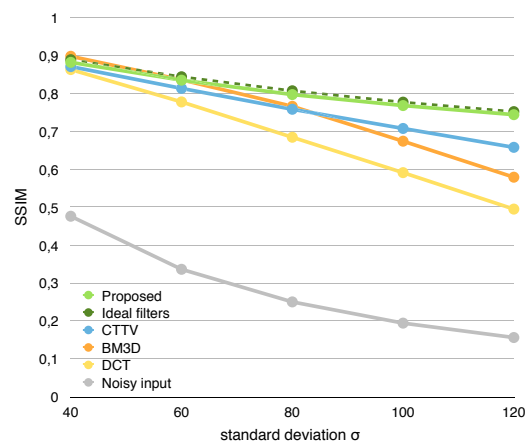
σ	Image name	Noisy input	DCT [31]	BM3D [12, 20]	CTTV [8, 11]	Proposed	Ideal filters
40	playground	0.47	0.85	0.88	0.85	0.87	0.87
		0.35	0.84	0.89	0.85	0.86	0.87
60	flowers	0.43	0.78	0.83	0.82	0.86	0.88
		0.36	0.64	0.73	0.72	0.74	0.75
80	fruits	0.51	0.82	0.85	0.85	0.91	0.92
		0.18	0.65	0.71	0.72	0.78	0.77
100	facade	0.17	0.77	0.77	0.78	0.84	0.84
		0.18	0.62	0.62	0.64	0.71	0.71
120	zebra	0.21	0.68	0.73	0.77	0.85	0.86
		0.15	0.70	0.76	0.78	0.86	0.87

b) SSIM values

Table 1. We achieve competitive PSNR and SSIM values for all standard deviations σ with state-of-the-art approaches. The best results are given in bold. The respective qualitative results are given in Figure 8.



a) Average PSNR values



b) Average SSIM values

Figure 6. For increasing σ the proposed approach outperforms the comparative methods. Comparison of the average PSNR and SSIM values over the 12 test images for different standard deviations σ of the noise. Qualitative results are shown in Figure 8.

decreasing performance of BM3D at increasing noise levels is mainly due to two factors. Firstly, the BM3D method expects the noise to have zero mean. The aforementioned noise-clipping therefore leads to a loss of contrast. Secondly, the reliable identification of similar patches becomes significantly more difficult as the noise level increases. For illustration purposes, Figure 7 shows the similarity of a particular patch (highlighted in a) to all other patches in the image at different noise levels. As we can see, the accuracy with which relevant patches are identified decreases drastically as the noise level increases. Although the similarity measures of BM3D are improved by prior denoising and thresholding steps (cf. [20]), the general difficulty of finding matching patches in noisy images is unavoidable.

To be able to compare to what extent the results of the image quality metrics PSNR and SSIM coincide with the perceived visual quality, Table 1 a,b) shows the PSNR and SSIM values corresponding to the images in Figure 8.

In both, Table 1 and Figure 6 we included the results with *ideal filters*, i.e. when the spectral filters are learned on the

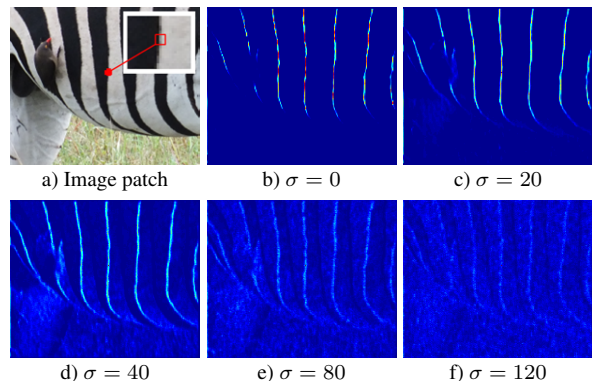


Figure 7. Noise harms reliable patch matching. Illustrated are the similarities (with red indicating a high similarity) of the patch highlighted in a) to all other patches in the image for different noise levels σ .

corresponding (unknown) noise-free image. While these images are of course impossible to compute in a realistic scenario the PSNR and SSIM values indicate that the proposed learned filters match the ideal filters closely.




Figure 8. **Qualitative results** for different standard deviations σ . The respective PSNR and SSIM values are given in Table 1.

7. Extensions

Finally, we would like to point out that the proposed framework of learning optimal filters for image reconstruction tasks based on nonlinear spectral image decompositions is not limited to image denoising.

7.1. Contrast Enhancement

Consider the problem of image sharpening or contrast enhancement. Since nonlinear spectral TV decompositions separate edges and features at different scales, our framework allows to learn filters to boost certain frequencies for visual quality enhancement as shown in Figure 9. We generate a low contrast test image by applying a bicubic downscaling followed by a bicubic upscaling of an image by a factor of four and mixing the original and blurry image to equal parts. The resulting image has the same resolution but a reduced contrast of small features.

We use (7) to compute the optimal smooth filters to restore the original image from its blurry decomposition. Figure 9 shows a) the original image, b) the image with reduced contrast, and c) the restoration using spectral filtering. Note that some of the lost contrast is restored, leading to a gain of 3.39 in the PSNR. Figure 9 d) shows the learned sharpening filters resulting from optimizing (7), where we omitted a color channel coupling for the sake of easier illustration.

7.2. Image Recovery

Although the theory of spectral decompositions developed in [4, 17] does not include additional linear operators, *i.e.* the reconstruction of u from $f = Au + n$, the numerical methods including the proposed filter learning are straight forward to apply to this case, too. Despite the missing spectral interpretation, we apply Bregman iteration to generate iterates u^k that approximate the TV minimizing solution to $Au = f$ at different scales and learn optimal filters for obtaining a good representation of the true underlying u that was used to generate the data.

Figure 10 shows an exemplary result for the reconstruction of an image which has been corrupted with a Gaus-




Figure 9. **Learning sharpening filters.** a) Original sharp image, b) image with reduced contrast, c) image with restored contrast based on learned ideal filters d).




Figure 10. **Details are recovered** when reconstructing blurred images. The lower row shows a zoom of the upper right image parts.

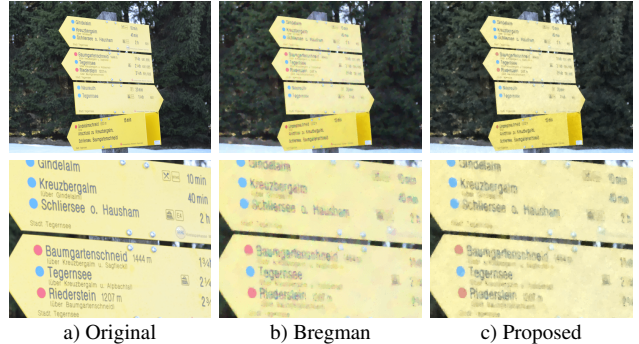


Figure 11. **Improved reconstruction quality** compared to the PSNR-optimal Bregman iteration for compressed image recovery.

sian blur of size 9×9 with standard deviation 2 and additive white Gaussian noise with standard deviation 25.5. For comparison purposes the PSNR-optimal Bregman iteration is shown as well. As we can see, the proposed scheme is able to recover finer details.

As a second image reconstruction example, Figure 11 shows the results we obtained on a compressed sensing problem. We generate a sparse matrix A that compresses the clean image to 10% of its original size by taking linear combinations of 10 random elements with random coefficients. To simulate the data, we additionally add white Gaussian noise with standard deviation 25.5. Note that the resulting data is not an image and thus cannot be visualized in a nice way. As we can see in Figure 11 c) the proposed learned spectral filters again yield an improved reconstruction quality and particularly suppresses color artifacts.

8. Conclusion

In this paper we have studied the nonlinear spectral TV decomposition of color images. We proposed to learn noise level specific filters that explore the natural inter-channel correlation of color images. Numerical results on image denoising show that learning filters for non-linear image decomposition yields state-of-the-art results at high noise levels. Additionally, the proposed framework demonstrates a great flexibility in adapting to additional tasks like image enhancement or reconstruction.

References

- [1] J.-F. Aujol, G. Gilboa, T. Chan, and S. Osher. Structure-Texture Image Decomposition—Modeling, Algorithms, and Parameter Selection. *Int. Journal of Computer Vision*, 67(1):111–136, 2006. [2](#)
- [2] X. Bresson and T. Chan. Fast Dual Minimization of the Vectorial Total Variation Norm and Applications to Color Image Processing. *Inverse Problems and Imaging*, 2(4):255–284, 2008. [3](#)
- [3] A. Buades, B. Coll, and J.-M. Morel. A Review of Image Denoising Algorithms, with a New One. *Multiscale Modeling & Simulation*, 4(2):490–530, 2005. [2](#)
- [4] M. Burger, L. Eckart, G. Gilboa, and M. Moeller. Spectral Representation of 1-Homogeneous Functionals. To appear at SSVN 2015. Preprint at <http://arxiv.org/abs/1503.05293>. [1](#), [2](#), [3](#), [8](#)
- [5] M. Burger, G. Gilboa, S. Osher, and J. Xu. Nonlinear inverse scale space methods. *Communications in Mathematical Sciences*, 4(1):179–212, 2006. [2](#)
- [6] M. Burger, S. Osher, J. Xu, and G. Gilboa. Nonlinear inverse scale space methods for image restoration. In *Variational, Geometric, and Level Set Methods in Computer Vision*, pages 25–36. Springer, 2005. [2](#)
- [7] A. Chambolle and T. Pock. A First-Order Primal-Dual Algorithm for Convex Problems with Applications to Imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145, 2011. [4](#)
- [8] T. Chan, S. Kang, and J. Shen. Total Variation Denoising and Enhancement of Color Images Based on the CB and HSV Color Models. *Journal of Visual Communication and Image Representation*, 12(4):422–435, 2001. [3](#), [6](#), [7](#)
- [9] P. Chatterjee and P. Milanfar. Patch-Based Near-Optimal Image Denoising. *IEEE Trans. on Image Processing*, 21(4):1635–1649, 2012. [2](#)
- [10] Y. Chen, R. Ranftl, and T. Pock. Insights into analysis operator learning: From patch-based sparse models to higher order MRFs. *IEEE Trans. on Image Processing*, 23(3):1060–1072, 2014. [2](#)
- [11] C. Condat and S. Mosaddegh. Joint Demosaicking and Denoising by Total Variation Minimization. In *IEEE Int. Conf. on Image Processing*, pages 2781–2784, 2012. [3](#), [6](#), [7](#)
- [12] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian. Image denoising by sparse 3-D transform-domain collaborative filtering. *IEEE Trans. on Image Processing*, 16(8):2080–2095, 2007. [2](#), [4](#), [6](#), [7](#)
- [13] M. Elad and M. Aharon. Image Denoising Via Sparse and Redundant Representations Over Learned Dictionaries. *IEEE Trans. on Image Processing*, 15(12):3736–3745, 2006. [2](#)
- [14] E. Esser, X. Zhang, and T. Chan. A General Framework for a Class of First Order Primal-Dual Algorithms for Convex Optimization in Imaging Science. *SIAM Journal on Imaging Sciences*, 3(4):1015–1046, 2010. [4](#)
- [15] A. Foi. Clipped noisy images: Heteroskedastic modeling and practical denoising. *Signal Processing*, 89(12):2609–2629, 2009. [5](#)
- [16] G. Gilboa. A Spectral Approach to Total Variation. In *Scale Space and Variational Methods in Computer Vision*, pages 36–47. Springer, 2013. [1](#), [2](#)
- [17] G. Gilboa. A total variation spectral framework for scale and texture analysis. *SIAM Journal on Imaging Sciences*, 7(4):1937–1961, 2014. [1](#), [2](#), [8](#)
- [18] G. Gilboa and S. Osher. Nonlocal Operators with Applications to Image Processing. *Multiscale Modeling & Simulation*, 7(3):1005–1028, 2008. [2](#)
- [19] T. Goldstein, E. Esser, and R. Baraniuk. Adaptive Primal-Dual Hybrid Gradient Methods for Saddle-Point Problems. *ArXiv preprint (arXiv:1305.0546)*, 2013. [4](#)
- [20] M. Lebrun. An Analysis and Implementation of the BM3D Image Denoising Method. *Image Processing On Line*, 2:175–213, 2012. <http://dx.doi.org/10.5201/ipol.2012.1-bm3d>. [4](#), [6](#), [7](#)
- [21] J. Mairal, M. Elad, and G. Sapiro. Sparse Representation for Color Image Restoration. *IEEE Trans. on Image Processing*, 17(1):53–69, 2008. [2](#)
- [22] S. Mallat. *A wavelet tour of signal processing*. Academic press, 1999. [2](#)
- [23] S. Osher, M. Burger, D. Goldfarb, J. Xu, and W. Yin. An Iterative Regularization Method for Total Variation Based Image Restoration. *SIAM Journal on Multiscale Modeling and Simulation*, 4:460–489, 2005. [2](#)
- [24] T. Pock, A. Chambolle, H. Bischof, and D. Cremers. A Convex Relaxation Approach for Computing Minimal Partitions. In *IEEE Int. Conf. on Computer Vision and Pattern Recognition*, pages 810–817, 2009. [4](#)
- [25] L. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992. [2](#)
- [26] G. Sapiro and D. Ringach. Anisotropic Diffusion of Multi-valued Images with Applications to Color Filtering. *IEEE Trans. on Image Processing*, 5(11):1582–1586, 1996. [3](#)
- [27] U. Schmidt and S. Roth. Shrinkage fields for effective image restoration. In *IEEE Int. Conf. on Computer Vision and Pattern Recognition*, pages 2774–2781, 2014. [2](#)
- [28] E. Strelakovsky and D. Cremers. Real-Time Minimization of the Piecewise Smooth Mumford-Shah Functional. In *European Conf. on Computer Vision*, pages 127–141, 2014. [4](#)
- [29] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. on Image Processing*, 13(4):600–612, 2004. [4](#)
- [30] L. Xu, C. Lu, Y. Xu, and J. Jia. Image Smoothing via L0 Gradient Minimization. *ACM Transactions on Graphics (TOG)*, 30(6):174:1–174:12, 2011. [2](#)
- [31] G. Yu and G. Sapiro. DCT image denoising: a simple and effective image denoising algorithm. *Image Processing On Line*, 2011. <http://dx.doi.org/10.5201/ipol.2011.ys-dct>. [4](#), [6](#), [7](#)
- [32] M. Zhu and T. Chan. An Efficient Primal-Dual Hybrid Gradient Algorithm for Total Variation Image Restoration. Technical Report 08-34, 2008. [4](#)
- [33] D. Zoran and Y. Weiss. From learning models of natural image patches to whole image restoration. In *IEEE Int. Conf. on Computer Vision*, pages 479–486, 2011. [2](#)