

Denoising photographs using dark frames optimized by quadratic programming

Manuel Gomez-Rodriguez
Electrical Engineering Department
Stanford University
manuelgr@stanford.edu

Jens Kober
Max Planck Institute
for Biological Cybernetics
kober@tuebingen.mpg.de

Bernhard Schölkopf
Max Planck Institute
for Biological Cybernetics
bs@tuebingen.mpg.de

Abstract

Photographs taken with long exposure or high ISO setting may contain substantial amounts of noise, drastically reducing the Signal-To-Noise Ratio (SNR). This paper presents a novel optimization approach for denoising. It is based on a library of dark frames previously taken under varying conditions of temperature, ISO setting and exposure time, and a quality measure or prior for the class of images to denoise. The method automatically computes a synthetic dark frame that, when subtracted from an image, optimizes the quality measure. For specific choices of the quality measure, the denoising problem reduces to a quadratic programming (QP) problem that can be solved efficiently. We show experimentally that it is sufficient to consider a limited subsample of pixels when evaluating the quality measure in the optimization, in which case the complexity of the procedure does not depend on the size of the images but only on the number of dark frames.

We provide quantitative experimental results showing that our method automatically computes dark frames that are competitive with those taken under idealized conditions (controlled temperature, ISO setting, exposure time, and averaging of multiple exposures). We provide application examples in astronomical image denoising. The method is validated on two CMOS SLRs.

1. Introduction

Among the sources of noise in images taken using a CMOS digital camera we can distinguish photon shot noise, reset noise, dark current noise, MOS device noise (mainly thermal noise) and readout noise, and there are fixed pattern noise (FPN) and temporal components (random) [4].

A bias frame is a raw image taken with closed shutter and an exposure time of (almost) zero seconds. The bias value

is usually caused by the readout noise. A dark frame is a raw image taken with closed shutter and a nonzero exposure time. Dark frames are a method by which the thermal noise can be recorded. In a nutshell, each dark frame contains a bias frame plus a component that increases with exposure time, in a way that depends on several other factors, including temperature and ISO setting.

A standard method for denoising long exposure images, variants of which are implemented on many commercial cameras [2], records a dark frame of matching exposure time immediately after each long exposure. This dark frame is subtracted from the so-called light frame (the image to be denoised) to result in an image where a substantial part of the noise cancels out. This procedure doubles the amount of time that the imaging process takes, but it works reasonably well provided that the temperature of the sensor has not changed. In practice, however, the temperature tends to change during camera operation even if the outside temperature is constant. The method effectively uses a one-point sample from the joint distribution of the noise (jointly over all pixels). As it is sampled from the joint distribution, it does reflect pixel dependencies, but being a one point sample it can be a poor estimator if the intrinsic fluctuations of the noise are substantial.

In fields where exposure times are long and low noise is important (e.g., astrophotography), it is thus common practice to record a whole *set* of dark frames taken under conditions matching the ones of the light frame, and denoise the image using the mean of those.¹ This leads to better estimates of the expected noise and thus to better denoising results, particularly with professional cooled CCDs, which have precise temperature control.²

Alternatively, one can try to take into account statistical aspects of the denoising problem in rather different ways.

¹Instead of the mean, sometimes the median or a form of trimmed mean or sigma clipping is used. We have found that this leads to only slightly different results, without changing our conclusions in the experimental part of this paper.

²For temperature controlled CCDs, one can also get away with dark frames whose exposure does not match, by separately removing the bias, and scaling the de-biased dark frames.

Work done while Manuel Gomez-Rodriguez was with the Max Planck Institute for Biological Cybernetics.

Many denoising methods assume that the sum of all the sources of noise can be modeled as white Gaussian noise, and work with demosaicked images (after interpolating the Bayer array [5]). Wavelet methods [11], bilateral filtering [13], anisotropic diffusion [14] or NL-means [1] perform particularly well in practice, preserving image boundaries while smoothing flat regions. They are thus very well adapted to the statistics of natural images, however, they do not take into account the joint statistics of the sensor noise.

Moreover, after the demosaicking process, the noise is intensity dependent, as noted by [8]. The same reference proposes a model for noise estimation and denoising, taking into account the demosaicking process and the camera response function (CRF). However, the noise sources are considered zero mean Gaussians, an assumption that prohibits modeling dark current fixed pattern noise, one of the main sources of noise in long exposure images. In the field of astronomical images, [10] proposes a Bayesian framework that includes dark current fixed pattern noise.

In the present paper, we develop a new method to denoise long exposure raw photographs. It uses a library of dark frames to obtain information about the joint distribution of the noise for a given camera. This distribution depends on various conditions, including temperature, ISO settings and exposure time. If we knew the conditions for the image to be denoised, we should ideally use a library that matches the conditions of the image as closely as possible. However, it will turn out that our method can generalize across varying conditions, e.g., by denoising images whose exposure time or temperature does not match the one in the dark frames. This is convenient not only if the exact temperature is unknown, but also if our space to store dark frames is limited.

Given a set of dark frames taken under different conditions of temperature, ISO settings and exposure time, and a quality measure or prior for the class of images to denoise, our method computes a convex combination of dark frames such that subtracting it from the given image optimizes the quality measure. Empirically, we find that the quality measure need only be evaluated on a relatively small set of pixels (whose size is set independently of the image size), which are chosen a priori to correspond to locations where the noise has large variance (i.e., where the different dark frames in the library will perform differently well).

The paper is organized as follows. Section 2 is devoted to the formulation of the convex optimization problem. In section 3, a criterion to select the pixels used in our optimization problem is explained. Results and performance analysis, with an application to photography, are shown in section 4, and the paper finishes with some conclusions in section 5.

2. Problem formulation

Removing the noise from a long exposure image smoothes the original noisy (grainy) image. We set our quality measure for our class of images to denoise to be the smoothness of the image. A good measure of smoothness is the discrete derivative; smooth regions have small derivatives. Our aim is to find an elements of the convex hull of the dark frames $D^{(1)} \dots D^{(N)}$,

$$C = \{\hat{D} \in \mathbb{R}^{m \times n} : \hat{D} = \sum_{i=1}^N \alpha_i D^{(i)}, \alpha \succeq 0, \mathbf{1}^T \alpha = 1\}, \quad (1)$$

that — when subtracted from the given image — minimizes a convex cost function over our measure of smoothness, where we use \succeq between vectors to mean componentwise inequality. Note that the constraint $\mathbf{1}^T \alpha = 1$ ensures that the bias noise (which is present in all dark frames) is removed.

Then, a mathematical formulation of our problem is to

$$\begin{aligned} & \text{minimize} && \sum_{x_i \in S} \sum_{x_j \in L_{x_i}} \phi(\hat{I}_{x_i} - \hat{I}_{x_j}) \\ & \text{subject to} && \alpha \succeq 0 \\ & && \mathbf{1}^T \alpha = 1, \end{aligned} \quad (2)$$

where $I_z \in \mathbb{R}$ is the intensity level of the raw image at location z , $D_z^{(k)} \in \mathbb{R}$ is the intensity level of the raw dark frame k at location z ,

$$\hat{I}_z = I_z - \sum_k \alpha_k D_z^{(k)} \quad (3)$$

is the denoised intensity, $\alpha \in \mathbb{R}^N$ is the variable, ϕ is a real convex cost function, $S \subset \mathbb{N}$ is a set of *evaluation* points of I and $L_z \subset \mathbb{N}$ is the 8-neighbor set of the location z in the raw image. For a raw image recorded with a Bayer array, 8-neighbor set means the 8 closest points with the same type of filter (R, G or B).

If a quadratic penalty function, $\phi(x) = x^2$, is chosen, (2) is equivalent to the following quadratic programming (QP) problem, reminiscent of a support vector machine [12]:³

$$\begin{aligned} & \text{minimize} && \frac{1}{2} \alpha^T H \alpha + \alpha^T f \\ & \text{subject to} && \alpha \succeq 0 \\ & && \mathbf{1}^T \alpha = 1, \end{aligned} \quad (4)$$

³Note that the size of the QP is determined by the number of dark frames. Using a few hundred or thousand dark frames would constitute a relatively small QP that can be solved fast [9]; in machine learning, similar problems whose size is of the order of millions are solved routinely using methods exploiting the structure (e.g., sparsity) of the QP.

where

$$\begin{aligned} H &= \sum_{x_i \in S} \sum_{x_j \in L_{x_i}} (D_{x_i} - D_{x_j}) (D_{x_i} - D_{x_j})^T \\ f &= \sum_{x_i \in S} \sum_{x_j \in L_{x_i}} (I_{x_i} - I_{x_j}) (D_{x_i} - D_{x_j}) \\ D_z &= [D_z^{(1)} D_z^{(2)} \dots D_z^{(N)}]^T. \end{aligned}$$

H is positive semidefinite and thus (4) is bounded from below.

Among all the available dark frames, only the ones that were taken under similar conditions as the noisy image should be used for denoising. We thus expect that a solution that generalizes well to the full image is sparse. Adding a L1-regularization to the objective function to enforce sparsity would not change the value of the optimal α because $\|\alpha\|_1$ is constant due to the constraints $\alpha \succeq 0$ and $\mathbf{1}^T \alpha = 1$.

As it turns out, our method also allows to estimate in an indirect way the exposure time, temperature and ISO of a photograph based on the selected dark frames, as the coefficients α_i of those darks frames that have matching parameters tend to be nonzero in the experiments.

3. Evaluation points

As evaluation points, we use points that have high variance between dark images (high level noise). Although points that have noise with low variance between dark frames (*e.g.* points of high readout noise) are thus not used in the optimization problem, they are still denoised: although the quality function is only evaluated at a small set of points, our effective dark frame $\sum_i \alpha_i D^{(i)}$ is subtracted from the whole image. We are denoising using a convex combination of samples from the full joint distribution of the noise, with expansion coefficients chosen based on a small set of pixels.

The selection of evaluation points has to be done only once for a specific camera (in principle, it could also depend on exposure time and ISO setting, although we have not found a strong dependence). Furthermore, H in (4) has to be computed only once.

In case of having a large amount of high resolution dark frames and noisy images, memory requirements may limit the scalability of our scheme if we load all the noisy images and dark frames in memory. For our cameras, 200 high resolution dark frames in raw format require ~ 4 Gb of memory. It is thus reasonable to load only the evaluation points and their neighbors because they are the only values of the dark images used in the optimization problem. In this way, for 200 dark frames and 1000 evaluation points, ~ 400 Kb of main memory are required. After the optimization problem is solved, we only need to load a few dark frames to denoise the original image because the solution is usually




Figure 1: 100 points with high variance between dark frames were selected; each curve shows the intensity value of one point for 25 subsequent exposures (ISO 1000 and 16 seconds exposure time), with breaks of ten seconds in between exposures. One can see that noise increases with the index of the exposure, due to an increase in the sensor temperature.

sparse. Moreover, subtraction of the individual dark frames occurring in the expansion of our synthetic dark frame can be done sequentially if memory is scarce.

4. Results and evaluation

An ISO 12233 test chart [7] is used for a quantitative analysis of the performance. Our method is compared to single dark frame denoising as well as to denoising using an average of dark frames (with the same exposure time and ISO setting as the noisy image). All photos for the quantitative analysis were taken with a Canon EOS 1Ds.

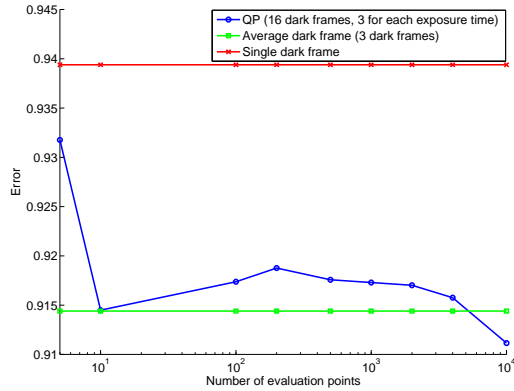
As common in machine learning, we use the same evaluation metric in the training set S (*i.e.* the optimization problem) and the test set T to numerically evaluate the performance of our method,

$$E = \frac{1}{8|T|} \sum_{x_i \in T} \sum_{x_j \in L_{x_i}} (\hat{I}_{x_i} - \hat{I}_{x_j})^2, \quad (5)$$


where T is a random set of size $|T| = 10^6$, disjoint from the set of evaluation points S . We normalize the value of E for the denoised image by the value of E for the noisy image.

A set of dark frames with an ISO of 800, 1000 and 1250 and exposure times from 1 to 128 seconds in powers of 2 and 21 seconds for different temperature conditions are used for the analysis. The dark frames were taken in two ways:

- *Constant temperature:* Given an exposure time, a set of 3 dark frames was taken, letting the camera cool down for 10 minutes between dark frames, ensuring



(a) Error (denoised image gradients relative to noisy image gradients) with respect to $|S|$. Our algorithm outperforms the single matching dark frame solution; for larger numbers of evaluation points, we perform as well as the mean of matching dark frames.



(b) Contribution of dark frames depending on the number of evaluation points. Not that for a sufficiently large set of evaluation points, the correct exposure time is identified automatically in the sense that the majority of dark frames that are used correspond to the correct exposure time.

Figure 2: Error performance and relative contribution of selected dark frames (*i.e.*, $\alpha_i \neq 0$) for a photograph of the test chart (ISO 800 and 16 seconds), and different sizes of the set of evaluation points S with dark frames with the same ISO settings (ISO 800) and temperature conditions but different exposure times.

that all photos were taken under the same temperature conditions in the CMOS sensor.

- *Variable temperature:* Given an exposure time, a set of 25 dark frames was taken, with an interval of ten seconds in between. In ten seconds, the CMOS sensor does not cool down, and the temperature in the sensor thus ends up increasing with the chronological order. Figure 1 shows the increase of the intensity values of the 100 points with the highest variance between these dark frames.

All the images were taken in a room whose temperature was approximately constant.

Three problem instances are proposed to validate our method. In the first and easiest one, we use dark frames taken under the same temperature conditions and ISO setting as the noisy image but different exposure times, including dark frames with the same exposure time as the noisy image. In the second instance, we make things a little harder by having variable temperature conditions. For the last case, we reuse the set of dark frames of the second problem instance but we denoise an image with an exposure time that does not match any of the exposure times of the dark frames.⁴ All photos were taken with a Canon EOS 1Ds using a full-frame CMOS sensor [6], and processed with dcrw [3] and MATLAB (for our dark frame

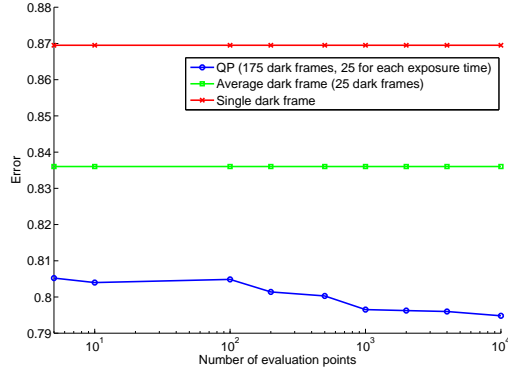
⁴We also performed experiments where dark frame libraries included photos taken at varying ISO settings, with qualitatively similar results to what is shown below.

libraries, the optimization problems can be solved using the standard Matlab optimizer in a fraction of a second).


For the first instance, figure 2a shows the performance in terms of normalized E and figure 2b shows the distribution of dark frames used by our method with respect to the number of points in S . A photograph of the test chart with ISO 800 and 16 seconds of exposure time was denoised using 3 dark frames of each exposure time with ISO 800 and the same temperature as the noisy photograph. Our method, as expected, outperforms the single dark frame solution and performs as well as the average of the dark matching temperature, for sufficiently large S , without giving information about the exposure time. The optimal α is sparse, and the majority of the matching dark frames are the ones corresponding to 16 seconds of exposure time.

From these results, it seems that our method works well but it does not outperform an average of dark frames with the same exposure time, ISO setting and temperature as the noisy image. However, requiring a set of dark frames taken under the same temperature conditions (in the CMOS sensor) and exposure time as the noisy image is inconvenient because it would necessitate a database of dark frames for every possible temperature and exposure time. Furthermore, controlling or measuring the temperature in the CMOS sensor is not an easy task without additional hardware.

For our second problem instance, where the temperature is no longer constant, figure 3a shows the error performance and figure 3b shows the weights α_i for $|S| = 10^4$ for our scheme. In this case, a photograph of the test chart with



(a) Error (denoised image gradients relative to noisy image gradients) as a function of the size of the evaluation set S . Our method outperforms the single matching dark frame and the mean of matching dark frames.



(b) Coefficient vector α for $|S| = 10^4$. The weights are maximal for the correct exposure time.

Figure 3: Error and weight distribution, for a photograph of the test chart (ISO 1000 and 16 seconds), and different sizes of the set of evaluation points S . Dark frames with the same ISO settings (ISO 1000) but different exposure times and temperature conditions were used.

ISO 1000 and 16 seconds of exposure time was denoised using 25 dark frames of each exposure time with ISO 1000 and variable temperature. Our method outperforms the single dark frame solution and the average of the dark frames with 16 seconds of exposure time and ISO 1000 but variable temperatures. As in the previous case, the optimal α is sparse, and the majority of matching dark frames are the ones with 32 and 16 second of exposure time.

Figure 4 shows the error performance for the last and hardest case. This time, the photograph of the test chart was taken with an exposure time of 21 seconds (ISO 1000), an exposure time which does not occur in our dark frame library. The photograph was denoised using 25 dark frames of each available exposure time, of variable temperature. Although no matching dark frames were present in the library, our method continues outperforming the single dark frame solution as well as the average of dark frames with 21 seconds exposure time and ISO 1000 but variable temperatures.

As an application of our method, a set of astronomical raw photographs were denoised. The photographs were recorded with ISO 1600, exposure times from 1 to 240 seconds and different unknown temperatures. All the astronomical photos were taken with a Canon EOS 5D containing a full-frame CMOS sensor [6].

A qualitative comparison with traditional methods (wavelets [11], bilateral filtering [13] or anisotropic diffusion [14]) is shown in figure 5. We found that none of the traditional methods removed the hot pixels (e.g., the red spot in the inset), and the denoised images are too blurred if we tune the parameters of the method in order to remove

significant noise. Note that knowledge of the dark frames, as utilized by our method, makes it trivial to remove hot pixels.

As the information used by our method is to some extent complementary to that used by the other methods, a combination is worthwhile. If we first apply first our method and afterwards a denoising based on wavelets with the parameters tuned less aggressively (*i.e.*, less blurring and less noise removal, because our method already removed a significant amount of noise), the result is satisfactory. Still better results could of course be obtained by additionally averaging over multiple light frames, as is common practice in astrophotography.

Figure 6 shows another noisy image and its denoised version using our algorithm and the wavelet denoising algorithm provided by dcrw.

5. Conclusions

In this paper, we have shown how a relatively simple method with low complexity can help denoise long exposure images in raw format in which dark current noise is the dominant source of noise.

Given a noisy raw image, we automatically compute a synthetic master dark frame from a library of dark frames recorded under different ISO settings, exposure times and temperature conditions, avoiding the necessity of acquiring dark frames every time a new photograph is taken. We have shown that to some extent, our approach interpolates across changes in exposure times and temperature — in other words, it is sufficient to have a sparse library where not all conceivable conditions are present.

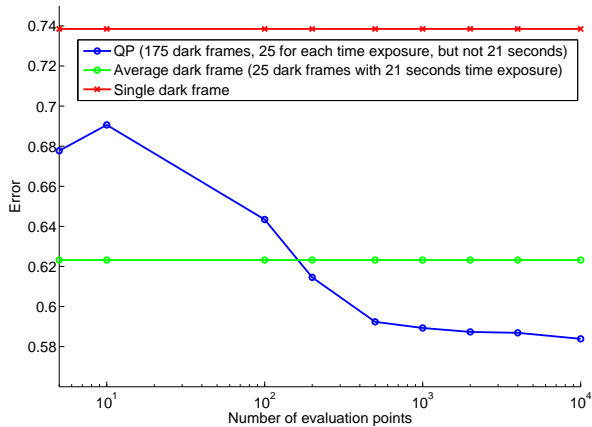


Figure 4: Error (denoised image gradients relative to noisy image gradients) as a function of $|S|$ for a photograph of the test chart (ISO 1000 and 21 seconds), using dark frames with the same ISO settings (ISO 1000) but different exposure times, not including 21 seconds, and temperature conditions. Our method outperforms the single matching (21 sec) dark frame, and it outperforms the mean of matching dark frames starting with 200 evaluation points.

Note that the software package Neat Image⁵ also uses different noise profiles depending on the camera model, ISO and exposure time (but not temperature conditions). However, like other noise reduction methods such as the ones built into Adobe Camera Raw or Digital Photo Professional, Neat Image is not camera specific in the sense that it does not take into account information about the sensor pattern noise (e.g., the location of hot pixels) in the specific exemplar of the camera used to capture the pictures. This is where our approach gets additional mileage, and it could beneficially be combined with other noise reduction methods.

Under our experimental conditions, we have seen the following:

- Looking at the solution of the quadratic programming problem that our method solves tells us which dark frames are used for denoising a given image. In the case of ISO setting (results not included) and exposure time, those settings turn out to resemble those of the photograph to be denoised.
- Provided the number of evaluation points is of the order of $10^3 - 10^4$, our approach does as well as the ideal scenario where a sizeable set of dark frames taken under exactly matching conditions are used. We have moreover found that our method always beats the one

of using a single dark frame, as implemented in commercial cameras.

- Consumer cameras generally have no temperature control, in which case it is unrealistic to have dark frame libraries of matching temperature. In this case, our approach actually managed to beat the method of averaging dark frames taken with matching ISO and exposure (but varying temperature), even if our library did not contain dark frames of the correct exposure time.
- If we have a set of matching dark frames, and add some additional dark frames whose exposure or temperature do not match, our approach can still take advantage of them and improve upon the solution obtained with only the matching dark frames.

Implementation of the system on a given camera requires a certain amount of logistical overhead:

1. A library of dark frames needs to be recorded. This could be done either by the user, by the manufacturer, or automatically whenever the camera is turned on and not being used.
2. A small quadratic program needs to be solved for each image to be denoised. In our experiments, with off-the-shelf optimization code, this usually took a small fraction of a second.

We believe that if these issues are addressed, the proposed method can become a practical tool for digital photography, especially for sensors with small pixel sizes as used in an increasing number of today's digital cameras.

Acknowledgment

Manuel Gomez Rodriguez is supported by a Fundacion Caja Madrid Graduate Fellowship.

References

- [1] A. Buades, B. Coll, and J. Morel. A Review of Image Denoising Algorithms, with a New One. *Multiscale Modeling and Simulation*, 4(2):490, 2005. 2
- [2] Canon's full-frame CMOS sensor: the finest tools for digital photography. Full frame CMOS, Canon White paper. 1
- [3] D. Coffin. dcraw. <http://www.cybercom.net/~dcoffin/dcraw/>. 4, 7
- [4] H. P. Components Group. Noise Sources in CMOS Image Sensors. *Imaging Products Operations*, 1998. 1
- [5] E. Dubois. Frequency-domain methods for demosaicking of Bayer-sampled color images. *IEEE Signal Processing Letters*, 12(12):847–850, 2005. 2
- [6] E. Fossum. CMOS image sensors: electronic camera-on-a-chip. *IEEE Transactions on Electron Devices*, 44(10):1689–1698, 1997. 4, 5

⁵<http://www.neatimage.com/>



(a) Noisy image with magnified inset



(b) Denoised using the proposed QP method



(c) Denoised using the wavelet denoising scheme provided by dcrw with threshold = 1000 [3]




(d) Denoised using bilateral filter (standard deviations (4, 0.2) and half-window size = 5) [13], using matlab code of <http://mesh.brown.edu/dlanman/>




(e) Denoised using QP and dcrw wavelet denoising scheme with threshold = 600

Figure 5: Comparison between our denoising method and standard approaches for an image with ISO 1600 and exposure time 60 seconds (horsehead nebula Barnard 33 in nebula IC 434, flame nebula NGC 2024, Canon EOS 5D with 300mm f/2.8 lens).



(a) Noisy image



(b) Denoised using QP and dcrw wavelet denoising, threshold = 600. Hot pixels and band noise are significantly reduced.

Figure 6: Comparison between a noisy image and its denoised version with our algorithm and wavelet denoising. The image was taken with ISO 3200 and 60 seconds of exposure time (summer milky way, Canon EOS 5D with Zeiss Distagon 28mm lens at f/2.8).

[7] ISO 12233. <http://www.iso.org>. 3

[8] C. Liu, W. T. Freeman, R. Szeliski, and S. B. Kang. Noise Estimation from a Single Image. *IEEE Computer Society*

Conference on Computer Vision and Pattern Recognition, 1, 2006. 2

[9] H. Mittelmann. Benchmarks for Optimization Software.

<http://plato.asu.edu/bench.html>. 2

- [10] R. Molina, J. Nunez, F. J. Cortijo, and J. Mateos. Image restoration in astronomy: a Bayesian perspective. *IEEE Signal Processing Magazine*, 18(2):11–29, 2001. 2
- [11] J. Portilla, V. Strela, M. J. Wainwright, and E. P. Simoncelli. Image denoising using scale mixtures of Gaussians in the wavelet domain. *IEEE Transactions on Image Processing*, 12(11):1338–1351, 2003. 2, 5
- [12] B. Schölkopf and A. Smola. *Learning with kernels*. MIT Press, MA, 2002. 2
- [13] C. Tomasi and R. Manduchi. Bilateral filtering for gray and color images. *Proceedings of the Sixth International Conference on Computer Vision*, pages 839–846, 1998. 2, 5, 7
- [14] D. Tschumperlé. Fast Anisotropic Smoothing of Multi-Valued Images using Curvature-Preserving PDE's. *International Journal of Computer Vision*, 68(1):65–82, 2006. 2, 5