

Removing noise from astronomical images using a pixel-specific noise model

Harold Christopher Burger, Bernhard Schölkopf, and Stefan Harmeling
Max Planck Institute for Biological Cybernetics, Tübingen, Germany

firstname.lastname@tuebingen.mpg.de

Abstract

For digital photographs of astronomical objects, where exposure times are usually long and ISO settings high, the so-called dark-current is a significant source of noise. Dark-current refers to thermally generated electrons and is therefore present even in the absence of light. This paper presents a novel approach for denoising astronomical images that have been corrupted by dark-current noise. Our method relies on a probabilistic description of the dark-current of each pixel of a given camera. The noise model is then combined with an image prior which is adapted to astronomical images. In a laboratory environment, we use a black and white CCD camera containing a cooling unit and show that our method is superior to existing methods in terms of root mean squared error. Furthermore, we show that our method is practically relevant by providing visually more appealing results on astronomical photographs taken with a single lens reflex CMOS camera.

1. Introduction

The problem of removing noise from images has been extensively studied. Methods to denoise images are numerous and diverse. However, most work focuses on the denoising of natural images or images of everyday scenes. Denoising algorithms are usually evaluated on their ability to remove additive white Gaussian noise with zero mean and uniform variance across the image. Often it is assumed that the variance of the noise is known.

The state-of-the-art image denoising methods try to leverage properties that are inherent to natural images. *E.g.* denoising with Fields of Experts [9] exploits the statistics of natural images: the algorithm relies on a set of small filters that have been trained on a dataset of natural images. Another successful denoising method, Bayesian least squares - Gaussian scale mixtures (BLS-GSM) [8] applies a wavelet transform on a noisy image and then exploits correlations between neighboring coefficients that are observed on natu-

ral images. A further recent method, called BM3D [2] uses the fact that in natural images, different patches are often similar in appearance.

The problem of attenuating noise in astronomical images is quite different, however. Astronomical images have statistics that are completely different from natural images. Often, such images contain little structure, such as a few stars against a black background. In addition, the characteristics on the dark-current noise that corrupts these images differs from the uniform additive white Gaussian noise most methods have been tuned to remove. We will show that pixels of a sensor behave differently. Therefore, we emphasize understanding and exploiting the statistics of the noise, rather than the image.

Assumption: The statistics of sensor noise is not adequately described by uniform additive white Gaussian noise (AWGN), see *e.g.* [6].

Hypothesis: Exploiting the noise statistics of each individual pixel of a camera's sensor leads to better denoising results.

Contribution: We present a novel method to remove noise from astronomical images that have been corrupted by dark-current noise. Our method relies on the combination of a statistical description of the noise of each pixel of a camera's sensor and an image prior. We show that the results achieved in that way are better than those obtained with state-of-the-art denoising methods.

Related work: Besides the state-of-the-art works on image denoising mentioned above, various authors have tried to create an artificial dark frame to decrease dark current noise. Such dark-frames are created in such a way as to optimize some image quality measure once the dark-frame has been subtracted from the noisy image. *E.g.*, Goesele *et al.* [4] have proposed to create an artificial dark-frame by scaling a given dark-frame in such a way as to minimize the entropy of the image. The method assumes that dark-current increases with increasing temperature, but do not take into account the random fluctuations for the dark-current. Gomez-Rodriguez *et al.* [5] have proposed to create

a convex combination of previously recorded dark-frames in such a way as to minimize the discrete gradient of the image at certain locations.

2. Dark-current noise

There exist many sources of noise in the imaging process. The different units of a CCD chip produce different voltages for the same amount of input light, a phenomenon sometimes called “fixed pattern noise” (FPN). This effect is due to the fact that the wells in a CCD chip slightly vary in size. An imaging sensor also produces “dark current”, which is due to thermal energy [12] and therefore also present in the absence of light. The analog to digital conversion process also introduces noise.

In the following we are not considering all possible sources of noise, but focus our attention on the image sensor’s dark current. The thermal energy of the imaging sensor frees electrons, which then accumulate in the chip’s wells. When an image is read out, the thermal electrons are indistinguishable from photoelectrons. We therefore assume the dark current to be additive and independent of the image signal. This assumption has been successfully exploited by others to denoise astronomical images [5]. Random samples of dark-current can be recorded with a so-called “dark-frame”, which is a photograph taken with closed lens and non-zero exposure time.

Since dark-current noise is due to thermal electrons, it is possible to reduce the amount of dark-current by cooling the camera’s sensor. For cameras on which cooling the image sensor is not possible, a simple approach for denoising photographs is to subtract a dark-frame with matching camera settings from the image. An improvement over this method would be to subtract the average of many dark-frames [6], thereby reducing the random components of the dark current.

Dark-current depends on the exposure time of the image, as well as the ISO-setting of the camera and the temperature of the image sensor [14]. More generally speaking, the problem of removing dark-current noise can be seen as a decomposition problem: Given a noisy observation y , what is the most likely true image x and the most likely noise sample d such that $y = x + d$, see Fig. 1. The problem is solvable if we exploit knowledge about the statistics of the image and about the distribution of the noise.

To study the statistics of dark current, we use a camera with a CCD chip (pco.2000, image sensor KAI-4021, cooled) that has a fixed conversion factor of 2.2 electrons per pixel count. Each dark-frame contains 2048×2048 pixels. Each pixel value is encoded as a 16 bit value. The camera has a built-in cooling unit, allowing us to study the property of the noise at different temperature settings. We record 200 dark-frames with 10 second exposure time at chip temperatures 10°C , 5°C , 0°C , and -25°C .




Figure 1. The noisy image is modelled as a sum of the true image and the dark current.

We randomly select 1000 pixels which we analyse in the following. For each temperature setting we calculate the mean and variances of these pixels. The top left panel of Fig. 2 shows the sorted mean values of those 1000 pixels over the 200 frames recorded at 20°C . Using this sorting of the pixels we show in the top row panel the corresponding mean values for the other temperatures and in the bottom panel (again with the same ordering as the top left panel) the variances. First of all we observe that the variances increase with the mean pixel values. Also we see that for different temperatures the same pixels exhibit a large mean and variance. Furthermore, we observe that for lower temperatures the means and variances decrease.

3. Theory

For notational simplicity, the recorded image y , the dark frame d , and the true image x are column vectors of the same lengths. The entries are denoted by subscripts, *i.e.* y_i , d_i , and x_i .

Maximum likelihood estimator. Theoretically, the pixel values of a dark frame d should be modelled with a Poisson distribution, *e.g.* [10]. This would imply that for the correct (ISO-dependent) scaling of the dark frames their pixel-wise mean and variance should coincide. However, such a scaling factor did not exist for our library of dark frames. For that reason we model the pixel values of the dark frames with Gaussian distributions for which we can adjust pixel-wise the mean and variance independently. For simplicity we consider in this part only dark frames of a fixed temperature. For each pixel d_i in a dark frame d , we can estimate its mean vector μ and variance vector σ^2 from some large set of dark frames that have been recorded with a fixed temperature. Modelling the recorded noisy image as $y = x + d$, the negative log likelihood of observing some image y is:

$$-\log p(y|x) = \frac{1}{2} \|y - x - \mu\|_D^2 + c \quad (1)$$

$$= \frac{1}{2} (y - x - \mu)^T D^{-1} (y - x - \mu) + c \quad (2)$$

$$= \sum_i \frac{(y_i - x_i - \mu_i)^2}{2\sigma_i^2} + c \quad (3)$$




Figure 2. Mean (top) and variance (bottom) of 1000 randomly chosen pixels for decreasing temperatures (left to right).

where D is the diagonal matrix with the variances σ^2 along its diagonal. The constant c depends only on D but not on y or x . Note that this likelihood models the amount of noise individually for each pixel location. This is different from most other methods, *e.g.* [3], who assume a global value for the variance. Methods that allow different amounts of noise in different locations of the image, usually estimate that amount from the noisy image [7]. Instead, we determine the noise distribution from a library of dark frames.

The maximum likelihood estimate of x is equal to the difference of y and μ , because setting $x = y - \mu$ minimizes $-\log p(y|x)$. We will denote this approach by DF-ML.

Maximum a posteriori estimator. The state of the art denoising methods have successfully introduced priors for natural images, *e.g.* [13]. Unfortunately, these priors do not apply to astronomical images as we will see in the experimental section. Instead we employ a generalization of a simple image prior used by Gomez-Rodriguez *et al.* [5] which is based on the idea that an astronomical image should be smooth and not grainy, *i.e.* the prior assumes that neighboring pixels should have similar values,

$$-\log p(x) = \lambda \frac{1}{|N_i|} \sum_{j \in N_i} |x_i - x_j|^p + c, \quad (4)$$

where c is a constant independent of x , and N_i are the indices of the eight neighbors of pixel i . In case the camera has a color filter array (CFA), the set of neighbors refers to the closest pixels on the same color channel. Gomez-Rodriguez *et al.* [5] applied this prior for $p = 2$. However, in our experiments it turned out that we get the best results for $p = 1.4$. The factor λ depends on the inverse variance of the image prior, *i.e.* it controls to which degree a pixel can be different from its neighbors. For the maximum a posteriori estimator, λ is a hyper parameter which

controls the trade-off between the image prior and the likelihood. Throughout all reported experiments it was fixed to $\lambda = 100$. This value was determined on artificial training images.

Together with the likelihood we can write down the posterior distribution for x ,

$$-\log p(x|y) = -\log p(y|x) - \log p(x) + c, \quad (5)$$

where c is again a constant independent of x and y . To determine the maximum a posteriori estimator, we have to minimize the negative log-likelihood $-\log p(x|y)$ in x . For this, we initialize x with its maximum likelihood estimate and proceed with gradient descent steps minimizing the negative log posterior $-\log p(x|y)$ in x . We employ an early stopping strategy to avoid over-smoothing. In the following, we denote this method by DF-MAP_p.

4. Experiments

4.1. Artificial stars with ground truth

To evaluate our method, we create artificial stars in a dark sky by employing a black surface containing small holes through which dim light shines. We obtain a low-noise ground truth image by the following procedure: we take 200 photos with a camera (pco.2000, image sensor KAI-4021, cooled down to -25°C) and average the resulting images to reduce the noise. From the resulting image, we subtract the average of 200 dark-frames recorded with the same chip temperature and with the same exposure times. The exposure time of all images is 10 seconds. Besides the ground truth image, we take also 100 noisy test images at 10°C chip temperature. The images are shot in very low light conditions, resulting in noticeable noise, see Fig. 3. Our goal is to recover the clean image as well as possible, given a single noisy image.




Figure 3. The results obtained by denoising using various methods. BM3D, which is particularly well suited for natural images, is not suitable for denoising astronomical images and results in a black image.

To compare a reconstructed image with the ground truth image, we calculate the root mean squared error (RMSE) between the ground true image x^* and the reconstructed image x , defined as $\text{RMSE}(x^*, x) = \sqrt{\frac{1}{n} \sum_{i=1}^n (x_i^* - x_i)^2}$, where n is the number of pixels in the image.

As competitors to our method we apply state-of-the-art methods for image denoising: BM3D [2], BLS-GSM [8], Bilateral filtering [11], Fields of Experts (FoE) [9] and dictionary-based denoising with an overcomplete DCT dictionary (DCT) [3]. All those methods are initialized with the maximum likelihood solution, *i.e.* with the difference between the noisy frame and the mean dark-frame. We then apply the various denoising algorithms. Usually, these denoising algorithms require a parameter describing the standard deviation of the noise present in the noisy image and it is assumed that this standard deviation is the same for each pixel in the image. We set this parameter value to be the average of all σ_i s.

A method that is tailored for astronomical images is the method proposed by Gomez-Rodriguez *et al.* [5], which creates an artificial dark-frame as a convex combination of available dark-frames. The method—called QP in the following—attempts to minimize the same image penalty function as we do (with $p = 2$) on a small selection of pixels (so-called “evaluation points”), which leads to a tractable quadratic programming problem. Following [5], we apply their method using 1000 evaluation points. We use two different settings regarding the library of dark-frames: once we use 160 dark-frames, all recorded with 10 second ex-

posure time but with temperatures ranging from -25°C to 10°C (20 frames per temperature, using 5°C increments). In a second setting, we use 200 dark-frames recorded with 10 second exposure time and the temperature matching the temperature with which the image was taken.

The method that only subtracts the mean dark-frame, corresponding to the maximum likelihood estimate, is denoted by DF-ML. The results of our proposed maximum a posteriori methods are denoted by DF-MAP $_p$.

denoising method	mean RMSE
no denoising	416.3
BM3D [2]	103.7
BLS-GSM [8]	36.7
QP (all temperatures) [5]	35.2
DF-ML (mean dark frame)	32.1
QP (only one temperature) [5]	29.6
DF-MAP₂ (our method)	27.8
Matlab’s wiener2	27.5
Bilateral Filtering [11]	27.4
Fields of Experts [9]	27.2
DCT [3]	26.9
DF-MAP_{1.4} (our method)	20.0

Table 1. Average results obtained by denoising 100 noisy test images with different methods

The results in Tab. 1 show the state-of-the-art denoising methods are not suitable for astronomical images. *E.g.*,

BM3D and BLS-GSM yield results that are worse than those obtained by removing the mean dark-frame from the noisy image. Since the mean dark-frame subtracted image is provided as input, one can say that these methods deteriorate the results. Fig. 3 compares the results obtained by denoising strategies we have tried. The result of BLS-GSM contains artifacts that resemble the “ringing” phenomenon. We presume that both BM3D and BLS-GSM make assumptions about images which are violated by the images we are trying to denoise. Further evidence is that simpler denoising methods, such as Matlab’s `wiener2` function and bilateral filtering [11] work relatively well: these methods make fewer assumptions about the statistics of images. Another method that works well is the dictionary-based denoising method with an overcomplete DCT Dictionary [3]. The method assumes that patches in an image can be well represented using a sparse combination of predefined patches, which appears to work well for astronomical images.

If we set the parameter p of our image prior to 2, we use the same image penalty function as the method proposed by Gomez-Rodriguez *et al.* [5]. Yet, our results are better. This is presumably due to the fact that our model is more flexible to repair individual pixels, whereas Gomez-Rodriguez *et al.* are forced to subtract a convex combination of dark frames.

Modifying the image prior by setting $p = 1.4$, the results obtained with our method are even better (last line in Tab. 1). It is not clear if the method proposed by Gomez-Rodriguez *et al.* could be modified to use a different image prior.

4.2. Real astronomical images

Orion constellation. To test our approach under real-world conditions we applied it to a noisy image of part of the constellation Orion. The image was taken by a Canon EOS 5D with 60 seconds exposure time and ISO 1600 and Canon lens 35/1.4 at aperture $f/2.8$. To minimize motion blur due to celestial movements a tracking mount was used. Demosaicing and gamma correction was performed using `dcrw` [1].

In addition to the noisy astronomical image, we had a library of dark-frames recorded with the same camera. We had no control over the temperature of the image sensor. The library of dark-frames was composed of 16 dark-frames at exposure time 10 seconds, 32 at exposure time 60 seconds, 32 at exposure time 120 seconds as well as 16 bias-frames (dark-frames with the shortest possible exposure time).

Fig. 4 shows the results of different denoising approaches for an enlarged cropped version (800×1000 pixels). The presented images were reconstructed by first denoising the raw images, then demosaicing with `dcrw` [1] and finally gamma correction. Because BLS-GSM is not meant to be applied to raw images, we apply it separately to

the four color channels. The approach proposed by Gomez-Rodriguez *et al.* [5] and our method are also able to treat raw images: the image prior is not calculated by considering a pixel’s immediate neighbors, but rather the closest neighbors on the same color channel.

Subtracting the mean dark-frame, *i.e.* the maximum likelihood estimate, does not significantly improve the visual quality of the image. In fact, artifacts are introduced: some pixels seem to be too dark. A possible explanation for this phenomenon would be if the camera’s sensor was warmer at the time the image was recorded than at the time the dark-frames were recorded. Dark-current increases with increasing temperature, so the dark-current in the image would be weaker than in the dark-frames.

Applying BLS-GSM to the dark-frame subtracted image provides little improvement over the dark-frame subtracted image. We found the results obtained by BLS-GSM to be representative of results obtained with the other state-of-the-art image denoising methods.

Using our method with $p = 2$ also did not create satisfactory results. Visually, it is not clear whether using our method with $p = 2$ provides better results than the mean dark-frame subtracted result or using BLS-GSM.

We apply the approach proposed by Gomez-Rodriguez *et al.* [5] using all dark-frames contained in our library. The result obtained in this way is visually better than both the original noisy image and the mean dark-frame subtracted image. The dark pixels that are present in the mean dark-frame subtracted image do not exist in the image denoised by the method proposed by Gomez-Rodriguez *et al.* It is possible that the method was able to select dark-frames that were recorded at a matching sensor temperature. However, noise is still present in the image. The background looks grainy.

For our method, we estimate the pixel means μ and variances σ^2 using the dark-frames that have been recorded with an exposure time of 60 seconds. Ideally, we would have estimated μ and σ^2 on dark-frames whose recording temperature matches that of the noisy image. Nonetheless, the result obtained by our method is smoother than any of the previously applied methods. Our method was able to strongly reduce the graininess that was visible in all previous results. Our method provides a very smooth background, yet does not cause even faint stars to disappear. Also, the visual quality of the nebula is not deteriorated.

Milky way. Finally, we present results obtained on an image of the Milky Way, recorded at an ISO setting of 3200, with a Canon EOS 5D, see top panel of Fig. 5. For this image we had only six dark-frames of matching settings available to us. We use these six dark-frames for the method proposed by Gomez-Rodriguez *et al.* as well as for our method. On the left and right of the images we provide more detailed

views of parts of the image. The red inset contains four hot pixels in the noisy image (the bright green pixels), which were successfully removed by both the method proposed by Gomez-Rodriguez *et al.* and ours. However, the result obtained by our method appears much less grainy, which makes individual stars more discernible.

5. Conclusions

We presented a new method for denoising astronomical images containing dark current noise. The method relies on a probabilistic description of a given camera's dark-current, as well as an image prior appropriate for astronomical images. Our method treats every pixel of a camera's sensor individually. Our image prior is similar to the one used by Gomez-Rodriguez *et al.* [5] and attempts to capture the roughness or graininess of an image. However, different from [5] we are not limited to quadratic functions, which allows us to use an image prior that is better suited for astronomical images, and moreover, we are not restricted to subtracting a convex combination of dark frames.

In laboratory conditions, we have shown that our method provides better results than state-of-the-art denoising methods that are intended for use on natural images. Our method also outperformed the recent method by Gomez-Rodriguez *et al.*, which is designed to denoise astronomical images.

On real astronomical images, we have shown that our method provides visually more appealing results than other methods. Images appear much less grainy after applying our method than when applying other methods. Fine image structure such as faint stars and nebula are preserved. It should be added that our evaluation was on single images, which is the hardest case in the sense that their noise is higher than for averages over several images as often used in astrophotography. Moreover, some of the graininess that we remove can also be removed by using a more sophisticated image acquisition pipeline including dithering (combining multiple exposures offset with respect to each other). We would expect that this would further improve our results, but make the difference to the other methods smaller.

Our method is limited in that we assume that appropriate dark-frames are provided with the image to be denoised. We assume the exposure time, ISO setting and temperature of the camera's sensor to approximately match the conditions at which the noisy image was recorded. The method proposed by Gomez-Rodriguez *et al.* overcomes this difficulty: given a library of dark-frames recorded under varying conditions, the optimization problem selects dark-frames that were recorded under the same conditions as the image. It is therefore conceivable to combine the two methods: the quadratic optimization problem described by Gomez-Rodriguez *et al.* could be used to select a set of dark-frames from a larger library. Our method would then use the selected dark-frames to infer an appropriate probabilis-

tic dark-current model for the denoising process.

References

- [1] D. J. Coffin. Decoding raw digital photos in Linux. <http://www.cybercom.net/~dcoffin/dcraw/>.
- [2] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian. Image denoising by sparse 3-D transform-domain collaborative filtering. *IEEE Transactions on Image Processing*, 16(8):2080–2095, 2007.
- [3] M. Elad and M. Aharon. Image denoising via sparse and redundant representations over learned dictionaries. *IEEE Transactions on Image Processing*, 15(12):3736–3745, 2006.
- [4] M. Goesele, W. Heidrich, and H. Peter Seidel. Entropy-based dark frame subtraction. In *Proceedings of PICS: Image Processing, Image Quality, Image Capture, Systems Conference*, 2001.
- [5] M. Gomez-Rodriguez, J. Kober, and B. Schölkopf. Denoising photographs using dark frames optimized by quadratic programming. In *IEEE International Conference on Computational Photography (ICCP)*, pages 1–9, 2009.
- [6] G. Healey and R. Kondepudy. Radiometric CCD camera calibration and noise estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 267–276, 1994.
- [7] C. Liu, W. Freeman, R. Szeliski, and S. Kang. Noise estimation from a single image. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 1, pages 901–908, 2006.
- [8] J. Portilla, V. Strela, M. Wainwright, and E. Simoncelli. Image denoising using scale mixtures of Gaussians in the wavelet domain. *IEEE Transactions on Image Processing*, 12(11):1338–1351, 2003.
- [9] S. Roth and M. Black. Fields of experts. *International Journal of Computer Vision*, 82(2):205–229, 2009.
- [10] H. Talbot, H. Phelippeau, M. Akil, and S. Bara. Efficient Poisson denoising for photography. In *Proceedings of the 16th International Conference on Image Processing (ICIP)*, pages 3881–3884, 2009.
- [11] C. Tomasi and R. Manduchi. Bilateral filtering for gray and color images. In *Proceedings of the Sixth International Conference on Computer Vision*, pages 839–846, 1998.
- [12] Y. Tsin, V. Ramesh, and T. Kanade. Statistical calibration of CCD imaging process. In *Proceedings of the Eighth International Conference on Computer Vision (ICCV)*, volume 1, pages 480–487, 2001.
- [13] Y. Weiss and W. Freeman. What makes a good model of natural images? In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2007.
- [14] R. Widenhorn, M. Blouke, A. Weber, A. Rest, and E. Bodegom. Temperature dependence of dark current in a CCD. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 4669, pages 193–201, 2002.

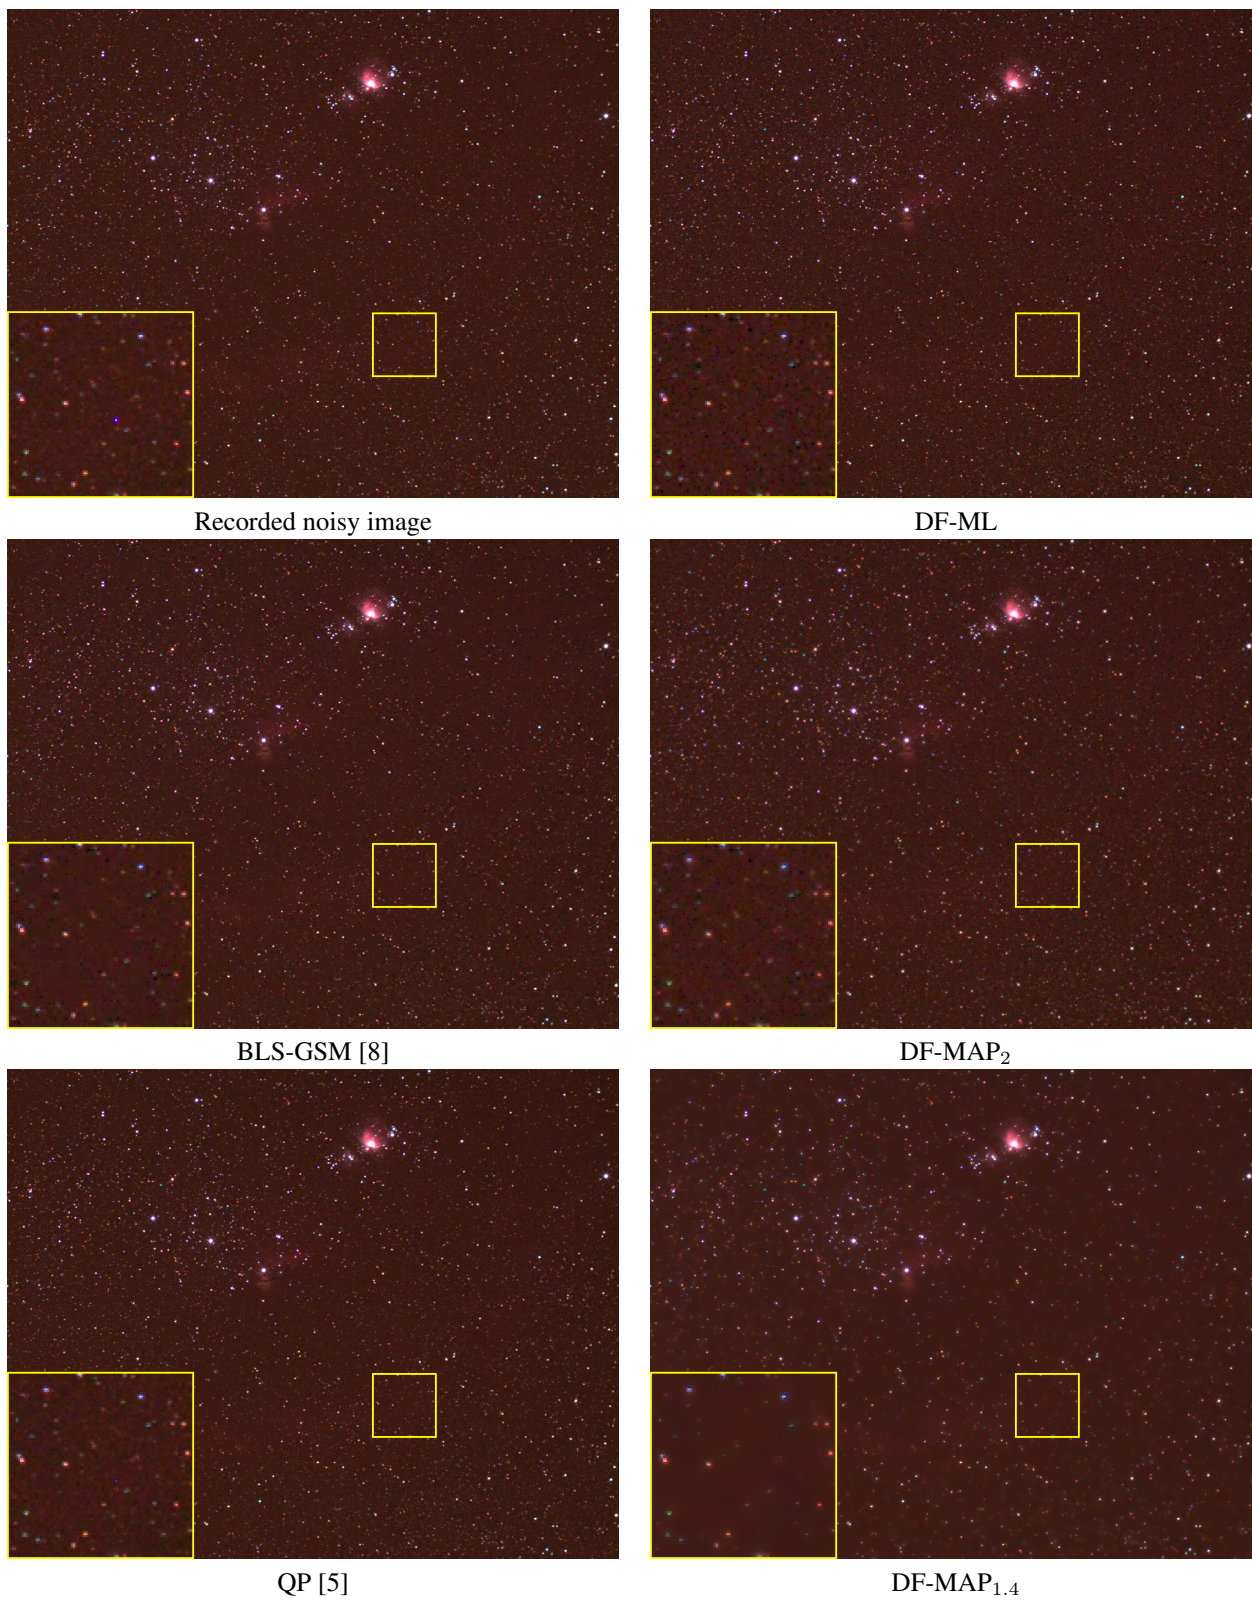


Figure 4. Comparison of various denoising techniques on a small section of the image of the constellation Orion.



Recorded noisy image



QP [5]



DF-MAP_{1.4} (our method)

Figure 5. Comparison of QP [5] and DF-MAP_{1.4} (our method) on a real astronomical image (taken with a Canon EOS 5D at ISO 3200 and 60 sec. exposure time with Zeiss Distagon 28mm lens at f/2.8, courtesy of Gomez-Rodriguez, Kober, and Schölkopf [5]).