



Noise-gating to Clean Astrophysical Image Data

C. E. DeForest

Southwest Research Institute, 1050 Walnut Street, Boulder, CO, USA

Received 2017 February 17; revised 2017 March 17; accepted 2017 March 17; published 2017 April 5

Abstract

I present a family of algorithms to reduce noise in astrophysical images and image sequences, preserving more information from the original data than is retained by conventional techniques. The family uses locally adaptive filters (“noise gates”) in the Fourier domain to separate coherent image structure from background noise based on the statistics of local neighborhoods in the image. Processing of solar data limited by simple shot noise or by additive noise reveals image structure not easily visible in the originals, preserves photometry of observable features, and reduces shot noise by a factor of 10 or more with little to no apparent loss of resolution. This reveals faint features that were either not directly discernible or not sufficiently strongly detected for quantitative analysis. The method works best on image sequences containing related subjects, for example movies of solar evolution, but is also applicable to single images provided that there are enough pixels. The adaptive filter uses the statistical properties of noise and of local neighborhoods in the data to discriminate between coherent features and incoherent noise without reference to the specific shape or evolution of those features. The technique can potentially be modified in a straightforward way to exploit additional *a priori* knowledge about the functional form of the noise.

Key words: methods: data analysis – methods: observational – techniques: photometric

Supporting material: animations, tar.gz file

1. Introduction

Images and image sequences are central to astrophysics and especially to the subfield of solar physics. Most such images are affected by shot noise or other additive noise, which limits sensitivity and indirectly spatial and temporal resolution of the data: small or short-duration image features may not rise above the noise floor even though they are, at least in principle, resolved by the instrument that collected the image.

The most commonly used denoising methods involve smoothing the data by direct convolution with a rectangular or smooth kernel, by median filtering, or by fitting a semi-empirical model to relevant image features. All of these methods remove more information from the data than is strictly necessary. For example, smoothing averages across pixels to beat down the noise, but it also discards all of the high-resolution (small-feature) aspects of the data. Semi-empirical fitting extracts a few parameters from a larger data set, but at the cost of ignoring all aspects of the data that do not match the structure of the fitted model.

In general, noise reduction and detection schemes attempt to isolate aspects of the data that contain mainly noise from those aspects of the data that contain signal. Convolutional smoothing is an example: it works because the Fourier transform concentrates image structure near the origin, while keeping additive noise spread across the entire Fourier domain. Convolution with a smoothing kernel attenuates the high frequencies, reducing the overall Fourier energy of the noise spectrum—at the cost of discarding any information contained in those high frequencies (e.g., Bracewell 2000).

But noise, despite being a random variable, has clearly defined statistical characteristics that can be used to separate it from signal on the basis of coherence within a small neighborhood of pixels (e.g., Lee 1980). Adaptive filters that take advantage of these statistics by changing parametrically with local image characteristics are well studied in the image processing literature (two good reviews are Shynk 1992 and

Buades et al. 2005). Adaptive filters have long been used to take advantage of known functional dependence between image values and noise level (Kuan et al. 1985); such filters remain a topic of active research and refinement (Shaick et al. 2000; Huang 2015, p. 1657; Lebrun et al. 2015). A particularly useful class of adaptive image filter is a 2D noise gate: a filter that identifies which Fourier components rise above a modeled noise spectrum, and attenuates or discards those that are weak while retaining those that are strong.

Noise-gating in the Fourier domain is a technique that takes advantage of the properties of the Fourier transform. In particular, image patterns that might be recognized as “features” typically have a concentrated Fourier spectrum. Hence, in each small region of a scientific image containing noisy but discernible structure, a few Fourier components typically rise above the noise floor. This concentration can be used to retain the “features” while attenuating or rejecting components containing more noise. This process is more selective than direct convolution by a single smoothing kernel, because the filter is locally adapted to preserve identified signal in each neighborhood of the image.

Noise-gating is commonly used in 1D to “clean” audio signals (Davis 1989). A time series of pressure levels (audio) is broken into segments, each of which is Fourier-transformed. Fourier components whose amplitudes rise above a predetermined noise floor spectrum are retained, and all others are either zeroed or attenuated. Then the segments are recombined to create the cleaned audio signal. The technique works to remove hum, buzz, crowd noise, tape hiss, and other additive noise from high-quality recordings.

Generalizing noise-gating to 2D (images) or 3D (image sequences) allows far better discrimination between information and noise than simple smoothing, simply by taking advantage of the redundancy inherent in visually distinguishable image structure: as in single images, coherent “features” tend to produce localized spatiotemporal spectra. The practical

result is that statistically significant features in solar, astrophysical, or other image sequences can be retained even at the full instrumental resolution, while the background noise is strongly attenuated by a zero-amplitude gate, a simple coefficient, or a component-wise adjustable Wiener filter.

Noise-gating in 2D has been used since the mid 1990s to process still images with additive noise (e.g., Yaroslavsky 1996), but 3D applications have largely relied on separate spatial and temporal filtering, using local motion tracking in video to allow co-adding with minimal motion blur. These deblurring techniques do not work well for images (such as solar magnetograms) that do not contain conventional objects as would a cinematographic scene, or where the direction of optical flow is ambiguous (Guo et al. 2007; Jovanov et al. 2009; Liu & Freeman 2010). Generalizing the Fourier noise gate to 3D affords the same basic benefit as the transition to 2D noise-gating: it helps to discriminate coherent structure in both space and time, without regard for the particular shape or evolution of the coherent pattern itself, thereby preventing the need to track motion and allowing denoising of astrophysical imagery.

Finally, local-neighborhood noise-gating permits the noise model to be adjusted on the basis of local image characteristics. This permits the gating of noise spectra that vary with location or with image characteristics, for example shot noise from a telescope with a nonflat vignetting function. Noise-gating of images and image sequences is sufficiently promising and spectacular, compared to smoothing, that it should be a standard part of the data reduction “toolbox” used for heliophysical and astrophysical image data. In this article, I present a method for noise-gating images and image sequences with variable (shot) or constant (fixed additive) noise spectra, and demonstrate its performance on commonly used image sources in solar physics.

In Section 2 I describe the technique of noise-gating in detail. In Section 3 I present results from four data sets: shot-noise-limited image sequences from *SDO/AIA*; additive-noise-limited sequences from *SDO/HMI*; shot- and film-grain-limited still images from the Normal Incidence X-ray Telescope (NIXT) sounding rocket; and a standard template image used by the image processing community. In Section 4 I discuss the results and their relevance to other data sets of astrophysical interest. Finally, in Section 5 I summarize some advantages and limitations of noise-gating and similar methods, for improving the use of astrophysical and solar remote-sensing data.

2. Noise-gating Technique

Monochrome image sequences are mappings $\text{Im}: \mathbb{Z}^3 \rightarrow \mathbb{R}$, where the domain runs over pixel coordinates and the range describes image brightness. In general, practical image data (after correction for fixed detector effects, i.e., “flat-fielding”) contain at least additive and shot noise:

$$\text{Im}(x, y, t) = \text{Im}_0(x, y, t) + N_a(x, y, t) + N_s(x, y, t) + N_{\text{other}}(x, y, t), \quad (1)$$

where Im is an image sequence data set, Im_0 is the “ideal” noise-free data set, N_a is an additive “background noise” term independent of Im_0 , N_s is an additive “shot noise” term that depends on Im_0 , and N_{other} is all other noise sources.

In high-fidelity audio processing and in some image applications, N_a is the dominant noise/background term. For example, in audio data (which can be described as a single-

pixel “image” sequence), N_a is commonly composed of white- or pink-noise “static,” highly spectrally peaked “hum,” crowd noises, or other uniform background noises that do not change character or amplitude across time. This noise is commonly removed from high-fidelity audio by “noise-gating.” The audio recording is broken into small segments that are individually Fourier-transformed, and the Fourier amplitude spectrum of N_a is determined from a segment that contains little or none of the actual signal to be reconstructed (“active silence”). The silence can be identified manually, or identified automatically by searching for the minimum spectrum from a long recording that contains at least one silent interval. Then all segments are processed by zeroing or greatly attenuating Fourier components that do not exceed the amplitude of the active silence by a predetermined factor.

Noise-gating to remove N_a can be adapted in a straightforward way to image sequences, by segmenting (and Fourier-transforming) in two (x, y) or three (x, y, t) dimensions rather than one. However, this adaptation is not readily applicable to many image applications, because uniform additive noise is often not the dominant contaminant of image data. In a large class of images, N_s , rather than N_a , dominates the noise field in the image: N_a and N_{other} can be neglected. Since N_s depends on the value of Im_0 , filtering based on a simple threshold of component amplitude is not sufficient as it is for N_a .

Estimating noise level. Shot noise arises from the Poisson-distribution statistics of counting discrete quanta—photons, photoelectrons, or other quantized elements that depend on the detection technology. N_s is then a random variable, sampled once per pixel, whose value depends also on the local value of Im_0 . When the number of quanta is high, the Poisson distribution is well approximated by a Gaussian (normal) distribution, and an estimate of the shot noise can be written directly:

$$N_s(x, y, t) \approx \alpha G(x, y, t) \sqrt{\text{Im}_0(x, y, t)}, \quad (2)$$

where α is an instrument-dependent constant, and $G(x, y, t)$ is a random variable with a fixed Gaussian distribution of mean 0 and variance 1. Equation (2) is particularly useful because it divides shot noise into three components, two of which can be characterized well. The α coefficient is a constant of the instrument (which could in principle vary with x and y , for example if the vignetting function of the instrument is nontrivial), and can be reconstructed from flat-field images or directly from image data. $G(x, y, t)$ is a standard tool of statistical analysis. By Fourier-transforming Equation (2) it is possible to estimate the spectral amplitude of the noise in the image:

$$N'_s(k_x, k_y, \omega) \approx \alpha G'(k_x, k_y, \omega) \otimes \mathbb{F}(\sqrt{\text{Im}_0(x, y, t)}), \quad (3)$$

where we have neglected any (x, y) dependence of α . But $G'(k_x, k_y, \omega)$ is constant across Fourier space, since $G(x, y, t)$ is a normal random variable; and for a very broad class of scenes the zero-frequency component of the Fourier transform dominates the spectrum, and the term $\mathbb{F}(\sqrt{\text{Im}_0(x, y, t)})$ can be treated as a delta function. Hence we arrive at the approximation:

$$|N'_s(k_x, k_y, \omega)| \approx \beta(k_x, k_y, \omega) \sum_{x, y, t} \sqrt{\text{Im}(x, y, t)}, \quad (4)$$

where β is a spectrum that is constant across images and is characteristic of the instrument that acquired the image, and we have used the fact that the sum of the noise term over many pixels is approximately zero to replace Im_0 with Im under the radical.

Equation (4) is useful because it estimates the noise amplitude in a given image or image subregion, provided that β is known. In principle, β can be determined a priori from the absolute sensitivity of the instrument, but it is also accessible via a posteriori analysis of the data themselves. This is accomplished by breaking a full data set $\text{Im}(x, y, t)$ into multiple small samples Im_i and searching for a minimum scaled spectrum for each one. Fourier-transforming Equation (1) (with N_a and N_{other} neglected) and substituting Equation (4) yields

$$\text{Im}'_i(k_x, k_y, \omega) \approx \text{Im}'_{0,i}(k_x, k_y, \omega) + \beta \sum_{x,y,t} \sqrt{\text{Im}_i(x, y, t)} \quad (5)$$

for each subimage index i . Solving for β gives

$$\beta_i(k_x, k_y, \omega) \approx \frac{|\text{Im}'_i(k_x, k_y, \omega) - \text{Im}'_{0,i}(k_x, k_y, \omega)|}{\overline{\text{Im}_i}}, \quad (6)$$

where the difference in the numerator is still just the (unknown) Fourier spectrum of the shot noise, and the bar over Im_i indicates summing the square root of each (positive-definite) pixel value. Across a large population of samples, the estimates of β_i at a given location in Fourier space will vary from a minimum where the sample of the noise spectrum at that particular point is near zero to a maximum where the local sampled value of the shot noise is much larger than the corresponding Fourier component Im'_0 . But structured images and image sequences containing coherent features are dominated by a few sparse Fourier components where $\text{Im}'_0(k_x, k_y, \omega) \approx \text{Im}'_i(k_x, k_y, \omega)$. Because the shot noise is a random variable, its Fourier amplitude is more nearly constant throughout the space, and most Fourier components are instead dominated by the noise: $|\text{Im}'_0| \ll |\text{Im}'_i(k_x, k_y, \omega)|$. Because the latter is the more common case, the median value of β_i across many image samples is a good estimator of the noise spectrum, and we can take

$$\beta_{\text{approx}}(k_x, k_y, \omega) = \text{median}_i \left(\frac{|\text{Im}'_i(k_x, k_y, \omega)|}{\overline{\text{Im}_i}} \right), \quad (7)$$

which depends only on the statistics across image subsamples of the Fourier spectrum in the original data set. The approximation in Equation (7) requires that a significant fraction of Fourier space be noise-dominated: at least half of all samples. This is typically the case in image sequences that have direct visual evidence of shot noise, but for images that are more highly structured the median could be replaced with a lower percentile value. The calculated value of β_{approx} allows estimation of the noise level across *all* regions of an image sequence dominated by conventional shot noise, per Equation (4).

Of course, for image-independent additive noise, one can estimate the noise level with a simpler calculation of a constant level across image segments:

$$|N'_{i,a}(k_x, k_y, \omega)| = \text{median}_j(|\text{Im}'_j(k_x, k_y, \omega)|), \quad (8)$$

where (as in Equation (7)) one may replace the median with a lower percentile across image samples in the case that noise is low or the image is very highly structured.

Filtering. Having produced a noise model, one can generate an adaptive filter tuned to the estimated noise spectrum in each sample:

$$\tilde{\text{Im}}'_i(k_x, k_y, \omega) \equiv \text{Im}'_i(k_x, k_y, \omega) F'_i(k_x, k_y, \omega), \quad (9)$$

where F'_i is a filter function of some sort. For the simplest processing, F_i is the gating function:

$$F'_{i,\text{gate}}(k_x, k_y, \omega) \equiv \begin{cases} 0 & \text{if } \text{Im}'_i(k_x, k_y, \omega) < T_i(k_x, k_y, \omega) \\ 1 & \text{otherwise} \end{cases}, \quad (10)$$

where T_i is a threshold function based on the noise levels from Equation (4) or (8). The Wiener filter

$$F'_{i,\text{wiener}}(k_x, k_y, \omega) \equiv \frac{\text{Im}'_i(k_x, k_y, \omega) / T_i(k_x, k_y, \omega)}{1 + \text{Im}'_i(k_x, k_y, \omega) / T_i(k_x, k_y, \omega)} \quad (11)$$

rolls off filter response more gradually but may admit more noise at a given threshold level.

In practice, the threshold can be defined using an ad hoc factor γ to bias the filtering between the preference to preserve the most signal possible and the preference to reject the most noise possible:

$$T_i(k_x, k_y, \omega) \equiv \gamma N'_i(k_x, k_y, \omega) \quad (12)$$

for whichever noise model is most appropriate. For the solar applications in Section 3, $\gamma = 3$ provides a good balance between recognizable features and noise reduction.

Apodization and reconstruction. As with all Fourier methods, the noise-gating described here requires careful apodization of the Im_i image segments. Apodization brings the edges of each Im_i smoothly to zero in a way that minimizes edge effects on the Fourier spectrum. It is accomplished by multiplying each image sample by a windowing function $w(x, y, z)$:

$$\text{Im}_i(x, y, z) = w(x, y, z) \text{Im}_{i,\text{raw}}(x, y, z) \quad (13)$$

where $\text{Im}_{i,\text{raw}}$ represents a small sample of data in an image sequence and $\text{Im}_i(x, y, z)$ is the image function treated above. Multiplying by a windowing function convolves the Fourier spectrum:

$$\text{Im}'_i(k_x, k_y, \omega) = \text{Im}'_{i,\text{raw}}(k_x, k_y, \omega) \otimes w'(k_x, k_y, \omega). \quad (14)$$

The optimal windowing function is the Hanning window:

$$w_h(x, y, t) = \sin^2 \left((x + 0.5) \frac{\pi}{n_x} \right) \sin^2 \left((y + 0.5) \frac{\pi}{n_y} \right) \times \sin^2 \left((t + 0.5) \frac{\pi}{n_t} \right), \quad (15)$$

where n_x , n_y , and n_t are the sizes (in pixels) of each dimension of Im_i . The Hanning window has Fourier power only in the zero and first nonzero frequencies in each axis, and therefore minimally spreads the spectrum of the apodized image.

Hanning windows are also ideal for smooth image reconstruction, because $\sin^2(\alpha) + \cos^2(\alpha) = 1$: by overlapping adjacent image samples by 50%, one can simply sum

adjacent apodized samples to reconstruct the full original (pre-apodization) image sequence. This approach is, in fact, used in one dimension in popular Fourier audio compression schemes such as MP3. However, in this application the summation properties of the apodization function are lost because of the noise-gating itself. The Fourier-space product between the filter and the transformed, apodized image block is equivalent to a convolution in real space, which damages the profile of the window function.

There are various rigorous techniques for building windowing functions that avoid this problem. The simplest is to apply a double-Hanning window: set

$$\tilde{\text{Im}}_{i,\text{final}}(x, y, t) \equiv w_h \tilde{\text{Im}}_i = w_h G(w_h \text{Im}_{i,\text{raw}}), \quad (16)$$

where the G operator represents the entire noise-gating sequence described above. If the gating step in Fourier space were the no-op (no operation), this would result in a simple windowing by \sin^4 . The original image can be reconstructed by oversampling the width of Im_i by a factor of 4 in each axis and summing all the resulting image segments. This is because

$$\sum_{j=0}^3 \sin^4\left(x + \frac{j\pi}{4}\right) = 1.5, \quad (17)$$

for all x . The oversampling and additional windowing attenuate any edge artifacts caused by the filtered gating function.

Implementation. I implemented this calculation in the Perl/PDL language and executed it on several scientific data sets, notably EUV images of the Sun. Data sets were subsampled to $12 \times 12 \times 12$ pixel sequences staggered every three pixels in space and time, with the double-Hanning apodization described above. The code is available as part of the “solarpdl-tools” distribution at GitHub.com.¹ Several small adaptations have been made in application. For example, rather than the full complex discrete Fourier transform the code uses the real discrete Fourier transform or equivalently the discrete cosine transform. Several “wrapper” routines implement multitasking and/or stream processing of a potentially semi-infinite set of images.

I implemented the code in Perl/PDL because of that environment’s superior handling of high-dimensional objects, ready library of scientific application modules, simple process control for multitasking, and ease of adaptation to compiled code. The existing code makes heavy use of the PDL vectorization engine. This is convenient for prototyping, but (as with all vectorized languages, including Numeric Python and IDL) it is pessimal for cache maintenance. Hence, the code almost certainly runs an order of magnitude more slowly than an optimized application written in a fully compiled language such as C or FORTRAN.

Gating with the existing code is feasible even for large image sequences (such as full-frame *SDO/AIA* images) on current hardware. Stream processing *AIA* images with block size $12 \times 12 \times 12$ and $4\times$ oversampling requires 30–90 s per frame on typical recent (2016) hardware, with the longer time representing a laptop computer and the shorter a rackmount “CPU server” with high-speed internal bus. A cache-optimized implementation might be expected to run $10\times$ faster, and a

GPU-optimized or similar highly parallel implementation might run $10\times$ faster still.

3. Results

Here I present results of applying the noise-gating algorithm described in Section 2 to images and image sequences. Noise-gating works best on image sequences because of the additional feature coherence afforded by the time dimension, and examples in Sections 3.1 and 3.2 show noise reduction in different types of solar image sequences. The technique is also applicable to individual images, and Sections 3.3 and 3.4 demonstrate performance in those domains.

3.1. Shot Noise: Coronal Image Sequences from *SDO/AIA*

Extreme-ultraviolet images of the solar corona have revolutionized solar physics since the 1990s. These images are collected in short spectral lines emitted by very hot, multiply-ionized trace elements (such as iron). The Atmospheric Imaging Assembly on board the *Solar Dynamics Observatory* (*SDO/AIA*; Lemen et al. 2011) produces 4096×4096 pixel EUV images of the Sun in several spectral lines, once every 12 s. This high cadence is intended to capture the evolution and dynamics of both large- and small-scale features. At the smallest scales visible in *AIA* images (up to a few seconds of arc), most features are dominated by shot noise in the original images. This makes them susceptible to improvement by scaled noise-gating as described in Section 2.

Figure 1 shows the limitations imposed by shot noise on the *AIA* images. It is available as a still frame and also (in the digital version of this article) as a movie. The differences are apparent in the still figure, but are yet more visually striking in the movie. The upper left panel is an image of a mid-latitude coronal plume, acquired on 2016 July 4 in that instrument’s 17.1 nm passband. It is a subfield of the *SDO/AIA* “Level 1” data set that is available via the virtual solar observatory (<http://vso.stanford.edu>) or the mission’s primary data distribution center (<http://jsoc.stanford.edu>). The image has been “unsharp masked”: a smooth background made with the *minsmooth* operator (DeForest et al. 2016) has been subtracted from it. The smooth background is made from a local-minimum estimator with an aperture diameter of 30 arcsec, so that the final image is positive-definite but contains only features with a spatial scale (in at least one direction) less than 30 arcsec. The image is easily seen to be affected by shot noise from the quantization of the EUV flux coming from the feature.

The lower left panel of Figure 1 has been smoothed with a 5 pixel full-width spherical Gaussian kernel in both space and time. The shot noise has been reduced by a factor of just over 10, at the cost of a $5\times$ reduction in both spatial and temporal resolution (i.e., a reduction of the data set to under 1% of its original information content).

The upper right panel of Figure 1 shows the effect of noise-gating as described in Section 2, using $12 \times 12 \times 12$ pixel subregions with double-Hanning windowing, direct threshold gating of Fourier components and a γ factor of 3. The noise is reduced by approximately the same factor as the smoothed image, but high spatial and temporal resolution are preserved. In particular, small bright points and ejecta, many of which are not directly noticeable in the original data, are preserved in the noise-gated sequence even though they are lost in the original and/or in the smoothed sequence.

¹ An archival copy of the code may be found attached to the tar.gz file accompanying this article; as of 2017, the complete current code and its development history may be downloaded from <https://github.com/drzowie/solarpdl-tools>, in the “image” subdirectory of that repository.

3-D Noise Gating Improves AIA Images




Figure 1. Noise-gating significantly improves *SDO*/AIA EUV images, which are limited by conventional shot noise, as seen in this unsharp-masked closeup of a low-latitude coronal plume. Clockwise from top left: the original L1 image is strongly noise-limited; the noise-gated image preserves resolution while deeply suppressing the shot noise by more than $10\times$; residual differences between the original and gated images show no overall structure; a Gaussian-smoothed image achieves comparable suppression, but with $5\times$ degradation in both spatial and temporal resolution. See also the movie attached to the digital version of this article, which shows the improvement still more dramatically.

(An animation of this figure is available.)

Moreover, myriad small bright ejecta in the plume itself are visible in the gated sequence, traveling from the core near $(-680, 270)$ leftward along the bright fibrils of the plume structure. At least some of these ejecta, once spotted, can be identified in the original data, but they are lost in the smoothed data. Others can be spotted with frame-to-frame stepping of the gated sequence but are too far below the noise floor of the original data to spot by eye.

The lower right panel shows the residual (difference) between the original image and the noise-gated image in the top two panels. While the noise level varies as expected, being stronger in the brighter portions of the image, it contains no visually identifiable image structure. The normalized correlation coefficient between the difference image and the

noise-gated image fluctuates from frame to frame, is 0 on average, and has a value well under 1% for individual frames.

It is worth noting that the residual (difference) images do *not* show signs of the faint ejecta along the plume, indicating that they are present in the original data although they are difficult to discern against the noise floor.

3.2. Additive Noise: Solar Magnetogram Sequences from *SDO*/HMI

Solar magnetograms are image data products that represent the amount of magnetic flux penetrating the surface of the Sun. The Helioseismic and Magnetic Imager on board the *Solar Dynamics Observatory* (*SDO*/HMI; Scherrer et al. 2012)

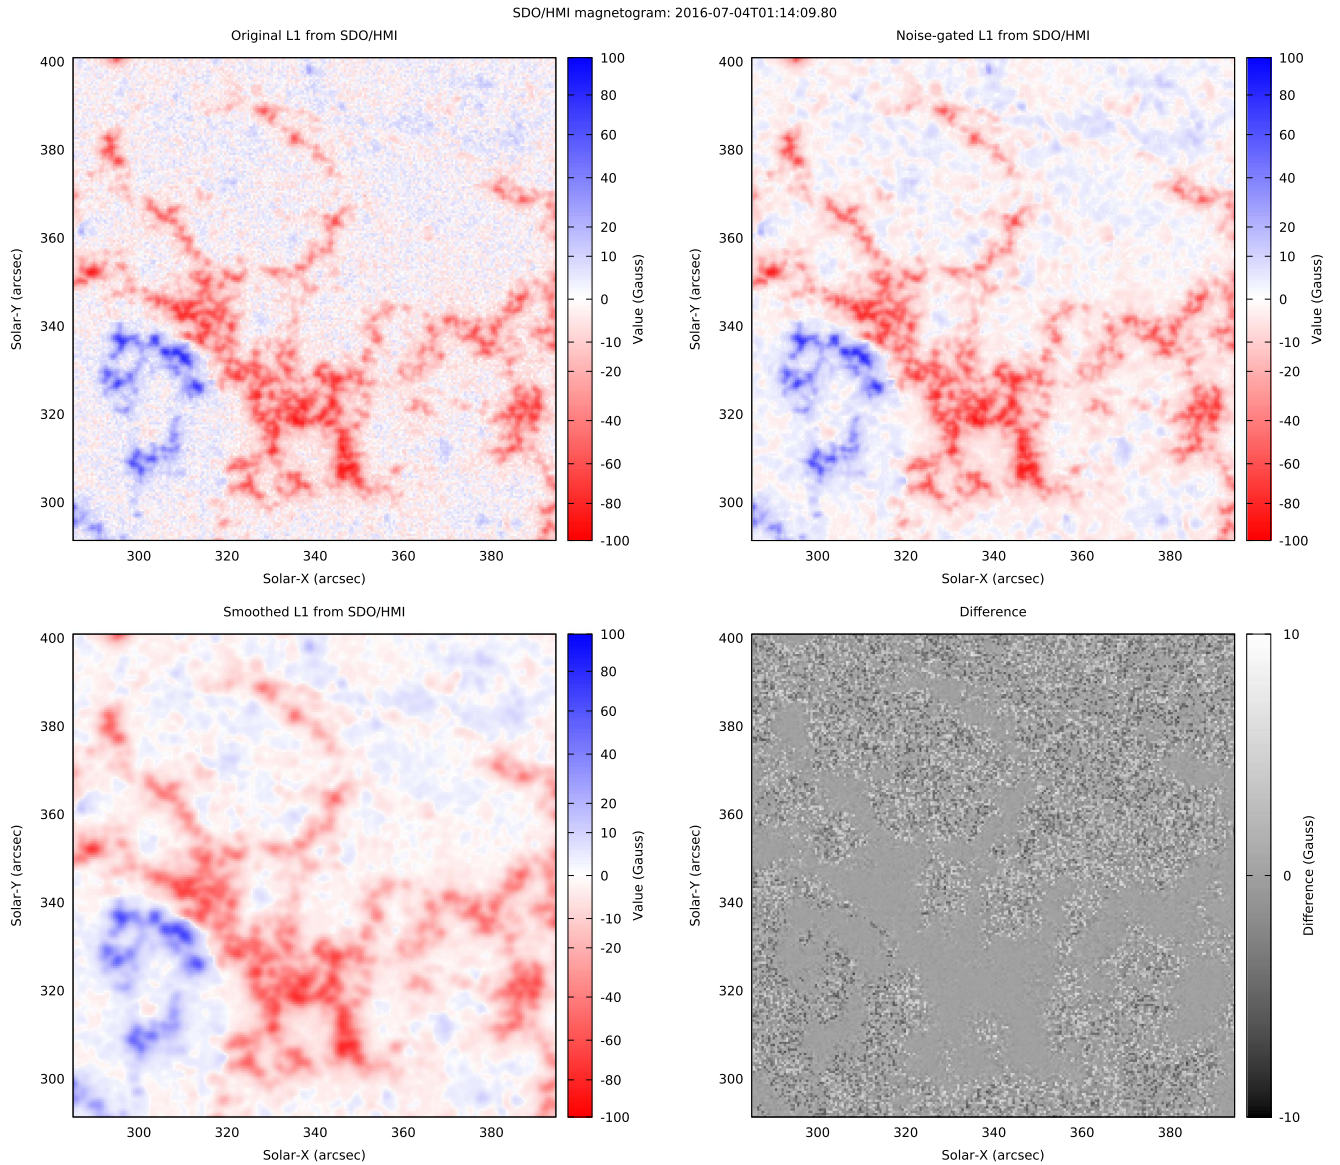


Figure 2. Noise-gating improves *SDO*/HMI magnetograms, which are limited by image-independent additive noise. Clockwise from top left: closeup of an original HMI 45 s magnetogram; the same magnetogram, gated as described in the text; difference image between the gated and original magnetograms; and a smoothed magnetogram with similar noise spectrum to the gated image. Gating reduces noise in the quiet regions of the image, without incurring the spatial or temporal resolution “hit” of smoothing where the signal is strong.

(An animation of this figure is available.)

downlinks polarized images in multiple narrow spectral bands (“filtergrams”), spanning a single spectral absorption line. These images are assembled post facto into a data product that represents the line-of-sight magnetic field on the surface of the Sun, exploiting the Zeeman effect. The noise spectrum of resulting “magnetogram” images is dominated by shot noise in the individual filtergrams but is not dependent on the local value of the magnetogram data product. Hence it can be modeled as additive noise on the data product itself. The method outlined in Section 2 is therefore applicable, using Equation (8) to define a threshold spatial spectrum.

Figure 2 shows processing of an example image sequence in extreme closeup (240×240 pixels compared to the native 4096×4096) of a sequence of 128 magnetograms from *SDO*/HMI. These data are available from the same locations as *SDO*/AIA data (Section 3.1, above). The fixed spectrum

threshold was set as $4 \times$ the 25th percentile value of each spectral component across all $8 \times 8 \times 8$ data samples in the magnetogram sequence (5×10^5 samples in all).

Because magnetograms are subject to additive noise and have most values near zero, the noise level can be seen through direct visual inspection of the probability distribution function (e.g., Schrijver et al. 1997) as shown in Figure 3. The additive noise in the magnetograms is well approximated as a normal (Gaussian) additive source with $\sigma = 14.5$ G, which appears on the semilog plot as an inverted parabola, centered on zero, as the core of the magnetogram’s probability distribution function. The width of the parabola represents the noise level. Smoothing the original with a spherical Gaussian kernel with full width of three pixels reduces σ of this distribution core to 5.1 G, at the cost of reducing the effective number of data points by a factor of roughly 10. Noise-gating removes noise from regions where




Figure 3. The probability distribution function (PDF) of the image value in Figure 2 shows the effects of noise and its reduction. Strong magnetic features form broad tails in the PDF. Additive shot noise forms a Gaussian core (inverted parabola on this semilog plot) with $\sigma = 14.5$ G. Both noise-gating and smoothing reduce the Gaussian core to $\sigma = 6.0$ G.

it would dominate the magnetic signal, both revealing weak magnetic features and preserving the full instrumental resolution in strong features.

The difference image at lower right in Figure 2 shows the effectiveness of the technique at preserving strong magnetic features: the portions of the image that are dominated by strong magnetic signal are preserved exactly as in the original data. Weak-field regions have high spatial frequencies removed, preserving only those elements of image structure that rise above the noise “floor.” While the photospheric magnetic field evolves relatively slowly compared to the overlying corona, at the very smallest scales flux emergence and/or cancellation is visible on timescales of 3–5 frames. As with Figure 1, the attached digital movie is yet more visually striking than the still frame.

Overall “photometry” and feature preservation are further demonstrated in Figure 4, which is a plot of the results of noise-gating and Gaussian smoothing on a single horizontal slice of the image in Figure 2. The gated data and smoothed data have roughly equal attenuations of the high-spatial-frequency “noise”—but the noise-gated trace preserves strong mixed-polarity flux and magnetic features at the instrumental resolution, while the smoothed line does not.

3.3. Single-image Shot Noise-gating: EUV Images from NIXT

Although the noise-gating works best on image sequences, it is also suitable for enhancement of still-frame images. The NIXT was a sounding rocket payload flown in the late 1980s through the mid 1990s that, together with its sister rocket the MSSTA, prototyped all modern EUV solar telescopes (e.g., Walker et al. 1988; Golub et al. 1990). These early sounding rockets used photographic film and were strongly limited by both film grain and photon shot noise throughout most of the field of view, though NIXT did achieve full subarcsecond resolution in bright features. Applying the noise-gating to NIXT images reveals more fine coronal structure than is apparent in the images themselves. Figure 5 shows the enhancement. An active region in the northeast (upper left) quadrant of the Sun is dominated by film grain (upper right panel). Noise-gating reduces the effect of the film grain (center left), revealing myriad faint fine-scale loops (center right; several are visible near $X \sim -700$, $Y \sim 100$). The features could in principle be artifacts, but they are visible in an




Figure 4. A single slice through the HMI image shown in Figure 2 reveals that individual features’ “photometry” and detail are preserved despite rejection of high-spatial-frequency noise. Traces are offset vertically by 50 G for clarity. For comparison to Figure 2, HMI pixels subtend 0.5 arcsec.

unprocessed longer exposure from the same rocket flight (lower right panel), verifying that they are real solar features being exposed by noise reduction.

The processed NIXT image has a different “texture” than many modern digital EUV images acquired in this spectral line, but that is attributable to remnant film grain and to the highly nonlinear response of photographic film to solar fluence. Images from the *SDO/AIA* 193 Å channel reveal similar fine-scale structure in the corona (lower left panel) when the contrast is adjusted appropriately—in this case by taking the 12th root of the reported image values (raising them to the power 0.083).

The NIXT images are a challenging target for this algorithm: they are individual stills with different exposure characteristics, limiting analysis to 2D rather than 3D; the detector response is highly nonlinear (the best-match AIA contrast profile scales as the 12th root of the instrument-reported radiance); and the images are affected by a mixture of shot noise and film grain. Nevertheless, these historical images are noticeably improved and reveal features in the corona that were present, but invisible, before processing.

3.4. A Conventional Photograph: the Standard Lena Test Image

I applied noise-gating not only to astrophysical images but also, for reference, to “Lena,” an industry-standard image adopted by the image processing community for over four decades. Figure 6 shows the effect of noise-gating on a noise-degraded version of the “Lena” image.

To highlight the difference between 2D and 3D processing, I simulated a slowly varying scene using the Lena image. I duplicated the image 16 times with a slow parametric shift of 0.5° rotation and 0.5 pixel displacement per frame, with the middle frame (no. 8) in its original position. I then added a different instance of the shot noise model to each copy. Noise-gating in 2D improved the signal-to-noise ratio (S/N) by a factor of 6, and in 3D by a factor of 14. Both applications recovered subtle and small image features that are apparently lost in the low-S/N degraded copy.

The additional noise reduction for 3D processing arises because the Fourier method exploits feature coherence across both time and space. It requires only that the moving/evolving object have a discrete Fourier spectrum that is distinguishable in amplitude from the background noise field, without regard

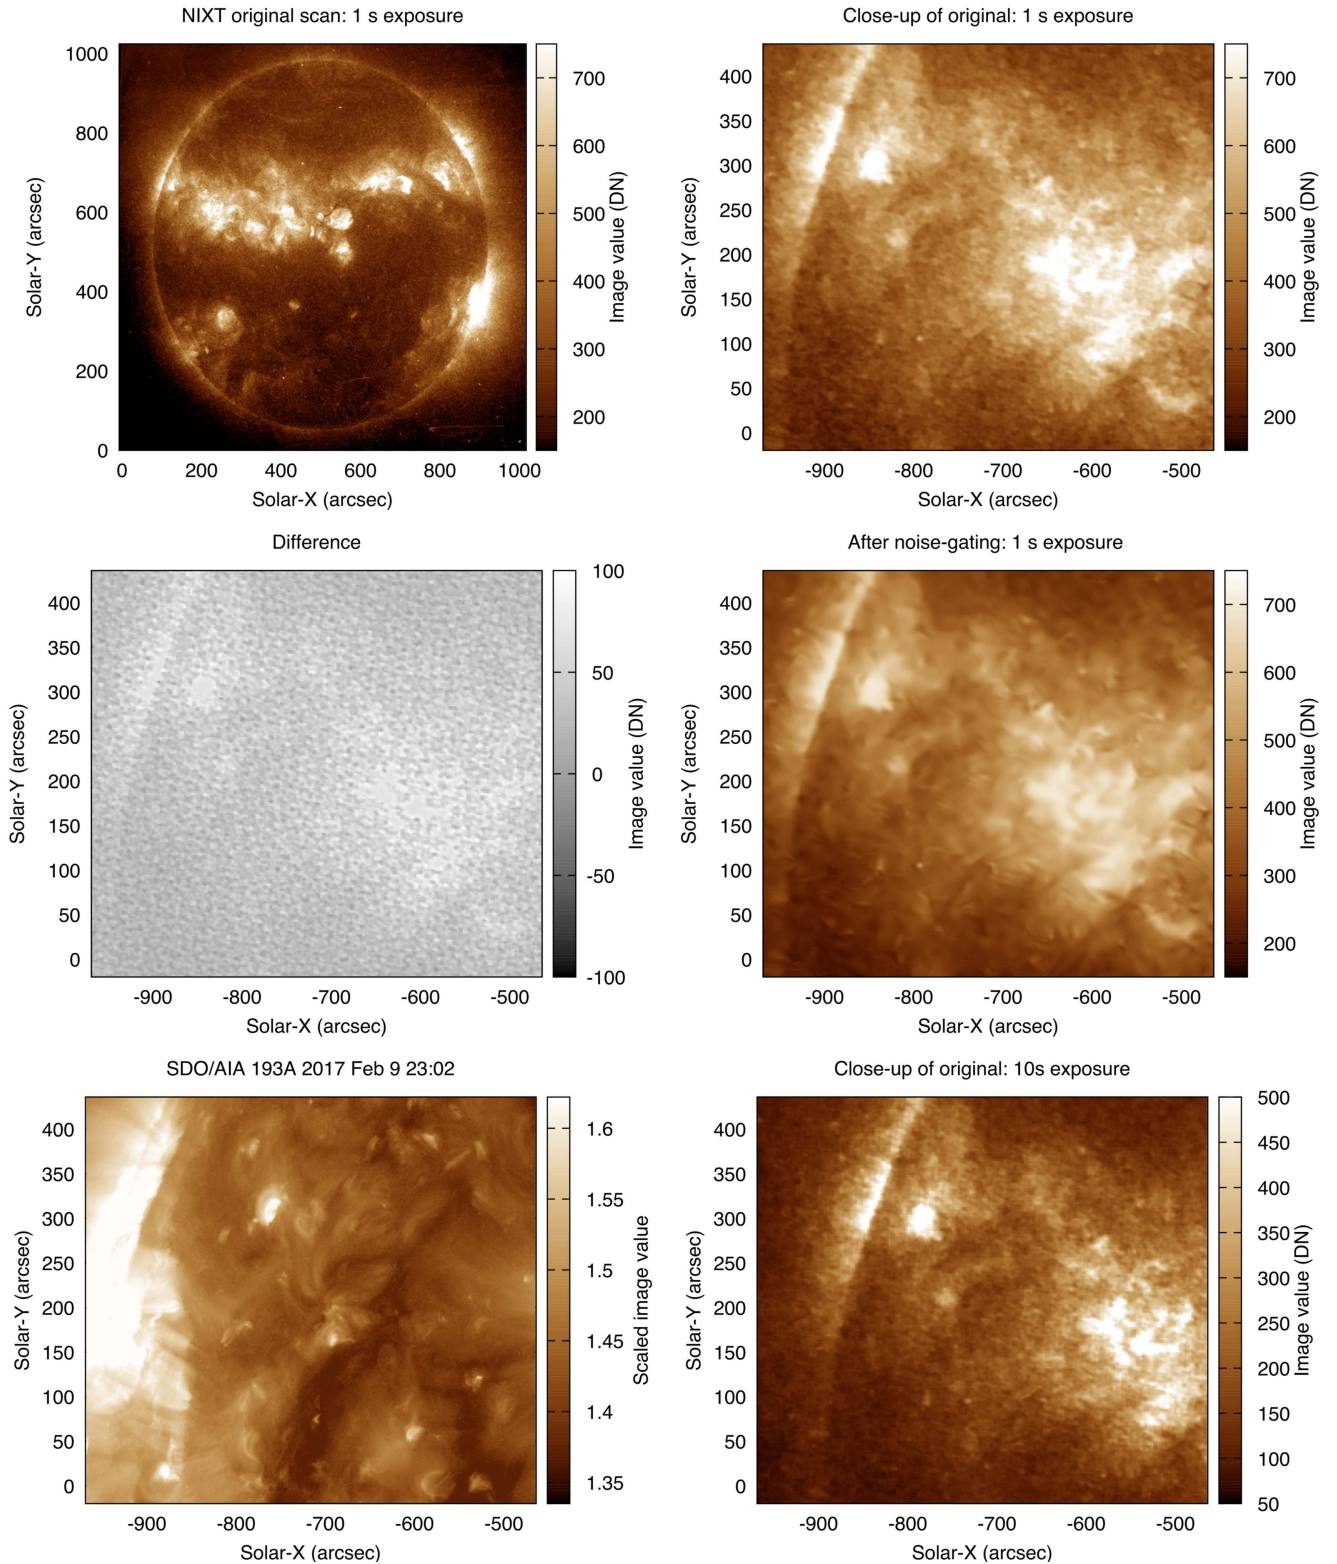


Figure 5. Enhancement of a 1 s exposure NIXT image of the Sun in the 193 \AA Fe emission line, from a sounding rocket flight on 1989 September 11, reveals fine-scale coronal loops not visible in the original image. Clockwise from top left: full-Sun image; closeup of an active region at upper left; the noise-gated image of the active region reveals small loop structures in the nearby corona; a longer exposure (unenhanced) from NIXT reveals the same loops; a modern 193 \AA image from *SDO/AIA*, scaled to match the NIXT film response, reveals similar small, faint structures in the corona; the difference image between the processed and unprocessed NIXT images reveals high-spatial-frequency components removed from the image.

for the shape of the feature or the particulars of its evolution across image frames.

While the output of noise-gating does not have the visual cleanliness of many of the more cosmetic denoisers (for a good

overview, see Buades et al. 2005), it does have direct roots in photometric noise reduction, and (as demonstrated in Section 3.2) preserves photometry of small features. The remaining shot noise is commensurate with the actual




Figure 6. Shot-noise degradation and restoration of the standard “Lena” image shows the advantage of 3D *sequence* processing for additional coherence. Clockwise from top left: original “Lena” image; degraded with shot noise: $\langle S/N \rangle = 1$; restored with 2D noise-gating: $\langle S/N \rangle = 6$; restored with 3D noise-gating: $\langle S/N \rangle = 14$.

sensitivity obtained by conventional averaging over the same size as the neighborhoods used for local Fourier transformation. In particular, the shot noise in the lower left panel of Figure 6 is reduced by $14\times$ by the 3D noise-gating, i.e., within a factor of two of the ideal photometric improvement obtained simply by averaging the noisy value across each $8 \times 8 \times 8$ pixel region of the synthetic image sequence. The improvement over smoothing arises from preservation of fine detail in places where the image contrast is sufficiently high. In particular the sharp edges of the hat’s brim and feather boa ribbon are preserved at full spatial resolution despite the low S/N in the degraded source image.

4. Discussion

Normal-distribution random additive noise with fixed or variable characteristics (generally photon shot noise, though there are other sources as well) affects nearly all scientific images. The most commonly used approaches to overcome it are smoothing of various sorts, and better approaches are available even without specific reference to the content of the data. 3D noise-gating can greatly improve data affected by this type of noise, by discriminating more precisely between likely noise and likely signal than can conventional content-blind methods.

Convolutional smoothing—in space, time, or both—sacrifices resolution for noise mainly because the method is not signal-aware. Model fitting is generally extremely content-aware and therefore requires prior knowledge of the particular subject being recorded. Noise-gating is a “happy medium” in signal recovery: it is one of a large class of adaptive filters that respond to the characteristics of a data set to improve discrimination between signal and noise, without prior knowledge of the signal itself.

Noise-gating works because the prime discriminant between “signal” and “noise” is structural coherence, which yields peaks in the Fourier spectrum of each neighborhood in the data. Because discernible features have spatial and temporal extent, they are generally concentrated by the Fourier transform into a smaller region in Fourier space than in the direct space of independent variables, which affords easier discrimination between the feature and the background “noise.”

For a broad range of applications, noise-gating is therefore a large improvement over traditional noise reduction techniques. By discriminating noise from signal in the Fourier domain, 3D noise-gating greatly reduces the noise level without significant loss of spatial or temporal resolution.

There is further room for improvement. I have presented here a simple adaptive filter for rejecting variable or fixed additive noise, whose character depends on local signal brightness but

not on location in the data. Obvious extensions include filters that are aware of instrument focal variation or vignetting, of multiple (compound) noise sources such as shot noise mixed with non-negligible read noise, or of multiplicative or non-normally distributed noise sources. Although they are not considered here, there are also obvious generalizations to color or hyperspectral images, which might take advantage of correlation between different wavelength channels or filter images in indirect color spaces such as the hue–saturation–value system, rather than in direct wavelength–channel intensity spaces such as red–green–blue.

Noise-gating does have limitations. In particular, either form I have presented here (strict gating and Wiener filtering) either obliterates or strongly attenuates any weak signal that cannot be discriminated from noise using the neighborhood Fourier amplitudes. This highlights and reveals features and signals that may not have been apparent before processing, but prevents subsequent lock-in or related measurements that could “drill” well below the noise floor, acting on the original data. Therefore noise-gating should be applied as the *last* step in a multistep image reduction pipeline. In this respect, its limitations are similar to those of convolutional smoothing and related denoising techniques, but with the significant advantage of retaining full spatiotemporal resolution in the cleaned images.

I have demonstrated noise-gating primarily in solar applications, in part because data-rich image sequences are so readily available and so important to that subfield of astrophysics. But the technique is applicable to other astrophysical targets and techniques, including wide-field imaging of the solar wind itself, stellar-coronagraph imaging of exoplanets, studies of nebula evolution, time-lapse studies of supernova remnants such as Eta Carina, and time-lapse orbital analysis of globular clusters. The technique has broad applicability to subjects outside astronomy and astrophysics, though I have focused on astrophysical applications in this article.

5. Conclusions

I have introduced and demonstrated a novel denoising technique, 3D noise-gating, that is suitable for use in scientific-grade astrophysical images. The technique removes shot noise from image sequences while retaining photometry of coherent image features at the full resolution of the original data. Noise-gating can also be applied in 2D for still images, and has been known for some time in that context although it appears not to have been adopted by the astrophysical community. Noise-gating is useful in part because it relies primarily on feature coherence as a discriminant, without regard to any particular shape, motion profile, or evolutionary characteristic of the

desired signal. In that regard it is a powerful general-purpose technique for data improvement. In the applications I have demonstrated, it reduces shot noise by a factor of order 10, and is therefore equivalent to an increase in exposure time between $10\times$ and $100\times$ for revealing faint features in noise-limited astrophysical image sequences. The cost is that very faint features below the reduced noise floor, even those that could in principle be recovered from more targeted techniques, are obliterated. The algorithm has been implemented in Perl Data Language with C libraries, is easily encapsulated for general-purpose use, and is readily portable to other computing environments.

This work was funded by NASA grants NNX-16AF98G and NNX-16AG98G. Thanks are owed to several individuals for insightful and helpful discussion as the article was being prepared, in particular: T. Lauer, C. Kankelborg, K. Kobayashi, V. Martinez-Pillet, A. Parker, and S. Savage. The author is also grateful to L. Golub for kindly providing NIXT data, and to the *SDO* imaging teams for the use of AIA and HMI data. Image analysis and processing were performed using the freeware Perl Data Language (<http://pdl.perl.org>).

References

- Bracewell, R. N. 2000, *The Fourier Transform and its Applications* (Boston, MA: McGraw-Hill)
- Buades, A., Coll, B., & Morel, J. 2005, *Multiscale Model. Simul.*, 4, 490
- Davis, G. D. 1989, *The Sound Reinforcement Handbook* (Iwata: Yamaha)
- DeForest, C. E., Matthaeus, W. H., Viall, N. M., & Cranmer, S. R. 2016, *ApJ*, 828, 66
- Golub, L., Herant, M., Kalata, K., et al. 1990, *Natur*, 344, 842
- Guo, L., Au, O. C., Ma, M., & Liang, Z. 2007, *IEEE Trans. Circ. Syst. Video Tech.*, 17, 1423
- Huang, C.-T. 2015, *ITIP*, 24, 4299
- Jovanov, L., Pizurica, A., Schulte, S., et al. 2009, *IEEE Trans. Circ. Syst. Video Tech.*, 19, 417
- Kuan, D. T., Sawchuk, A. A., Strand, T. C., & Chavel, P. 1985, *ITPAM*, 7, 165
- Lebrun, M., Colom, M., & Morel, J.-M. 2015, *Image Processing On Line*, 5, 1
- Lee, J. S. 1980, *ITPAM*, 2, 165
- Lemen, J. R., Title, A. M., Akin, D. J., et al. 2011, in *The Solar Dynamics Observatory*, ed. P. Chamberlin, W. D. Pesnell, & B. Thompson (Berlin: Springer), 17
- Liu, C., & Freeman, W. T. 2010, in *Computer Vision—ECCV 2010*, ed. K. Daniilidis, P. Maragos, & N. Paragios (Berlin: Springer), 706
- Scherrer, P. H., Schou, J., Bush, R. I., et al. 2012, *SoPh*, 275, 207
- Schrijver, C. J., Title, A. M., van Ballegoijen, A. A., Hagenaar, H. J., & Shine, R. A. 1997, *ApJ*, 487, 424
- Shaick, B. Z., Ridel, L., & Yaroslavsky, L. 2000, in *2000 10th European Signal Processing Conf., A Hybrid Transform Method for Image Denoising*, 1
- Shynk, J. J. 1992, *ISPM*, 9, 14
- Walker, A. B. C., Jr., Lindblom, J. F., Barbee, T. W., Jr., & Hoover, R. B. 1988, *Sci*, 241, 1781
- Yaroslavsky, L. P. 1996, *Proc. SPIE*, 2825, 2