

A Tour of Modern Image Filtering

PHOTO BY DAVID WAGNER

[Peyman Milanfar]

[New insights and methods,
both practical and theoretical]

Recent developments in computational imaging and restoration have heralded the arrival and convergence of several powerful methods for adaptive processing of multidimensional data. Examples include moving least square (from graphics), the bilateral filter (BF) and anisotropic diffusion (from computer vision), boosting, kernel, and spectral methods (from machine learning), nonlocal means (NLM) and its variants (from signal processing), Bregman iterations (from applied math), kernel regression, and iterative scaling (from statistics). While these approaches found their inspirations in diverse fields of nascence, they are deeply connected.

Digital Object Identifier 10.1109/MSP.2011.2179329
Date of publication: 5 December 2012

In this article, I present a practical and accessible framework to understand some of the basic underpinnings of these methods, with the intention of leading the reader to a broad understanding of how they interrelate. I also illustrate connections between these techniques and more classical (empirical) Bayesian approaches.

The proposed framework is used to arrive at new insights and methods, both practical and theoretical. In particular, several novel optimality properties of algorithms in wide use such as block-matching and three-dimensional (3-D) filtering (BM3D), and methods for their iterative improvement (or non-existence thereof) are discussed.

A general approach is laid out to enable the performance analysis and subsequent improvement of many existing filtering

algorithms. While much of the material discussed is applicable to the wider class of linear degradation models beyond noise (e.g., blur,) to keep matters focused, we consider the problem of denoising here.

INTRODUCTION

Multidimensional filtering is the most fundamental operation in image and video processing, and low-level computer vision. In particular, the most widely used canonical filtering operation is one that removes or attenuates the effect of noise. As such, the basic design and analysis of image filtering operations form a very large part of the image processing literature; the resulting techniques often quickly spreading to the wider range of restoration and reconstruction problems in imaging. Over the years, many approaches have been tried, but only recently in the last decade or so, a great leap forward in performance has been realized. While largely unacknowledged in our community,

this phenomenal progress has been mostly thanks to the adoption and development of nonparametric point estimation procedures adapted to the local structure of the given multidimensional data. Viewed through the lens of the denoising application, here we develop a general framework for understanding the basic science and engineering behind these techniques and their generalizations. Surely this is not the first article to attempt such an ambitious overview, and it will likely not be the last; but the aim here is to provide a self-contained presentation that distills, generalizes, and puts into proper context many other excellent earlier works such as [1]–[5], and, more recently, [6]. It is fair to say that this article is, by necessity, not completely tutorial. Indeed it does contain several novel results; yet these are largely novel interpretations, formalizations, or generalizations of ideas already known or empirically familiar to the community. Hence, I hope that the enterprising reader will find this article not only a good overview, but as should be the case with any useful presentation, a source of new insights and food for thought.

So to begin, let us consider the measurement model for the denoising problem

$$y_i = z_i + e_i, \quad \text{for } i = 1, \dots, n, \quad (1)$$

where $z_i = z(x_i)$ is the underlying latent signal of interest at a position $x_i = [x_{i,1}, x_{i,2}]^T$, y_i is the noisy measured signal (pixel value), and e_i is zero-mean, white noise with variance σ^2 . We make no other distributional assumptions for the noise. The problem of interest then is to recover the complete set of samples of $z(x)$, which we denote vectorially as $\mathbf{z} = [z(x_1), z(x_2), \dots, z(x_n)]^T$ from the corresponding data set \mathbf{y} . To restate the problem more concisely, the complete measurement model in vector notation is given by (surely a

similar analysis to what follows can and should be carried out for more general inverse problems such as deconvolution, interpolation, etc.

$$\mathbf{y} = \mathbf{z} + \mathbf{e}. \quad (2)$$

It has been realized for some time now that effective restoration of signals will require methods which either model the signal a priori (i.e., are Bayesian) or learn the underlying characteristics of the signal from the given data (i.e., learning, nonparametric, or empirical Bayes' methods.) Most recently, the latter category of approaches has become exceedingly popular.

Perhaps the most striking recent example is the popularity of patch-based methods [7]–[10]. This new generation of algorithms exploit both local and nonlocal redundancies or “self-similarities” in the signals being treated. Earlier on, the BF [8] was developed with very much the same idea in mind, as were its

spiritually close predecessors: the Susan filter [11], normalized convolution [12], and the filters of Yaroslavsky [13]. The common philosophy among these and related techniques is the notion of measuring and making use of affinities between a given data point (or more generally patch or region) of interest, and others in the given measured signal \mathbf{y} . These similarities are then used in a filtering context to give higher weights to contributions from more similar data values, and to properly discount data points that are less similar. The pattern recognition literature has also been a source of parallel ideas. In particular, the celebrated mean-shift algorithm [14], [15] is in principle an iterated version of point-wise regression as also described in [1] and [2]. In the machine learning community, the general regression problem has been carefully studied, and deep connections between regularization, least-squares regression, and the support vector formalism have also been established [16]–[19].

Despite the voluminous recent literature on techniques based on these ideas, simply put, the key differences between the resulting practical filtering methods have been relatively minor, but rather poorly understood. In particular, the underlying framework for each of these methods is distinct only to the extent that the weights assigned to different data points is decided upon differently. To be more concrete and mathematically precise, let us consider the denoising problem (2) again. The estimate of the signal $z(x)$ at the position x is found using a (nonparametric) point estimation framework; specifically, the weighted least squares problem

$$\hat{z}(x_j) = \arg \min_{z(x_j)} \sum_{i=1}^n [y_i - z(x_j)]^2 K(x_i, x_j, y_i, y_j), \quad (3)$$

where the weight (or kernel) function $K(\cdot)$ is a symmetric function with respect to the indices i and j . $K(\cdot)$ is also a

**RECENT DEVELOPMENTS IN
COMPUTATIONAL IMAGING AND
RESTORATION HAVE HERALDED
THE ARRIVAL AND CONVERGENCE
OF SEVERAL POWERFUL METHODS
FOR ADAPTIVE PROCESSING OF
MULTIDIMENSIONAL DATA.**

positive valued and unimodal function that measures the “similarity” between the samples y_i and y_j , at respective positions x_i and x_j . If the kernel function is restricted to be only a function of the spatial locations x_i and x_j , then the resulting formulation is what is known as (classical, or not data-adaptive) kernel regression in the nonparametric statistics literature [20], [21]. Perhaps more importantly, the key difference between local and nonlocal patch-based methods lies essentially in the definition of the range of the sum in (3). Specifically, indices covering a small spatial region around a pixel of interest define local methods, and vice versa.

Interestingly, in the early 1980s, the essentially identical concept of moving least-squares emerged independently [22], [23] in the graphics community. This idea has since been widely adopted in computer graphics [24] as a very effective tool for smoothing and interpolation of data in three dimensions. Surprisingly, despite the obvious connections between moving least-squares and the adaptive filters based on similarity, their kinship has remained largely hidden so far.

EXISTING ALGORITHMS

Over the years, the measure of similarity $K(x_i, x_j, y_i, y_j)$ has been defined in a number of different ways, leading to a cacophony of filters, including some of the most well-known recent approaches to image denoising [7]–[9]. Figure 1 gives a graphical illustration of how different choices of similarity kernels lead to different classes of filters, some of which we discuss next.

CLASSICAL REGRESSION FILTERS

Naturally, the most naive way to measure the “distance” between two pixels is to simply consider their spatial Euclidean distance; specifically, using a Gaussian kernel,

$$K(x_i, x_j, y_i, y_j) = \exp\left(\frac{-\|x_i - x_j\|^2}{h_x^2}\right).$$

Such filters essentially lead to (possibly space-varying) Gaussian filters which are quite familiar from traditional image processing [13], [20], [21], [25]. It is possible to adapt the variance (or bandwidth) parameter h_x to the local image statistics, and obtain a relatively modest improvement in performance. But the lack of stronger adaptivity to the underlying structure of the signal of interest is a major drawback of these classical approaches.

THE HUMBLE WEIGHTED LEAST-SQUARES, WITH LOCALLY DATA-DEPENDENT WEIGHTS BASED ON SIMILARITY KERNELS, EXPLAINS AN ENORMOUS CLASS OF RECENTLY POPULARIZED AND VERY EFFECTIVE FILTERS.

THE BILATERAL FILTER


Another manifestation of the formulation in (3) is the BF [8], [13] where the spatial and photometric distances between two pixels are taken into account in separable fashion as follows:

$$\begin{aligned} K(x_i, x_j, y_i, y_j) &= \exp\left(\frac{-\|x_i - x_j\|^2}{h_x^2}\right) \exp\left(\frac{-(y_i - y_j)^2}{h_y^2}\right) \\ &= \exp\left\{\frac{-\|x_i - x_j\|^2}{h_x^2} + \frac{-(y_i - y_j)^2}{h_y^2}\right\}. \end{aligned} \quad (4)$$

As can be observed in the exponent on the right-hand side, and in Figure 1, the similarity metric here is a weighted Euclidean distance between the vectors (x_i, y_i) and (x_j, y_j) . This approach has several advantages. Specifically, while the kernel is easy to construct and computationally simple to calculate, it yields useful local adaptivity to the given data. In addition, it has only two control parameters (h_x, h_y), which make it very convenient to use. However, as is well known, this filter does not provide effective performance in low signal-to-noise scenarios [3].

NONLOCAL MEANS

The NLM algorithm [7], [26], [27], originally proposed in [28] and [29], has stirred a great deal of interest in the community in recent years. At its core, however, it is a relatively simple generalization of the BF; specifically, the photometric term in the bilateral similarity kernel, which is measured point-wise, is simply replaced with one that is patch-wise. A second difference is that (at least in theory) the geometric distance



[FIG1] Similarity metrics and the resulting filters.

between the patches (corresponding to the first term in the bilateral similarity kernel), is essentially ignored, leading to strong contribution from patches that may not be physically near the pixel of interest (hence, the name nonlocal). To summarize, the NLM kernel is

$$K(x_i, x_j, y_i, y_j) = \exp\left(-\frac{\|x_i - x_j\|^2}{h_x^2}\right) \exp\left(-\frac{\|y_i - y_j\|^2}{h_y^2}\right) \quad (5)$$

with $h_x \rightarrow \infty$, where y_i and y_j refer now to patches of pixels centered at pixels y_i and y_j , respectively. In practice, two implementation details should be observed. First, the patch-wise photometric distance $\|y_i - y_j\|^2$ in the above is in fact measured as $(y_i - y_j)^T \mathbf{G} (y_i - y_j)$, where \mathbf{G} is a fixed diagonal matrix containing Gaussian weights, which give higher importance to the center of the respective patches. Second, it is rather computationally impractical to compare all the patches y_i to y_j , so although the NLM approach in Buades et al. [28] theoretically forces h_x to be infinite, in practice typically the search is limited to a reasonable spatial neighborhood of y_j . Consequently, in effect, the NLM filter too is more or less local; or said another way, h_x is never infinite in practice. The method in Awate et al. [29], on the other hand, proposes a Gaussian-distributed sample that comes closer to the exponential weighting on Euclidean distances in (5).

Despite its popularity, the performance of the NLM filter leaves much to be desired. The true potential of this filtering scheme was demonstrated only later with the optimal spatial adaptation (OSA) approach of Boulanger and Kervrann [26]. In their approach, the photometric distance was refined to include estimates of the local noise variances within each patch. Specifically, they computed a local diagonal covariance matrix, and defined the locally adaptive photometric distance as $(y_i - y_j)^T \mathbf{V}_j^{-1} (y_i - y_j)$ in such a way as to minimize an estimate of the local mean squared error (MSE). Furthermore, they considered iterative application of the filter as discussed in the section “Improving the Estimate by Iteration.”

LOCALLY ADAPTIVE REGRESSION (STEERING) KERNELS

The key idea behind this measure of similarity, originally proposed in [9], is to robustly obtain the local structure of images by analyzing the photometric (pixel value) differences based on estimated gradients, and to use this structure information to adapt the shape and size of a canonical kernel. The locally adaptive regression kernel (LARK) is defined as follows:

$$K(x_i, x_j, y_i, y_j) = \exp\{- (x_i - x_j)^T \mathbf{C}_i (x_i - x_j)\}, \quad (6)$$

where the matrix $\mathbf{C}_i = \mathbf{C}(y_i, y_j)$ is estimated from the given data as

$$\mathbf{C}_i = \sum_j \begin{bmatrix} z_{x_{i,1}}^2(x_j) & z_{x_{i,1}}(x_j) z_{x_{i,2}}(x_j) \\ z_{x_{i,1}}(x_j) z_{x_{i,2}}(x_j) & z_{x_{i,2}}^2(x_j) \end{bmatrix}.$$

Specifically, $z_{x_{i,\alpha}}(x_j)$ are the estimated gradients of the underlying signal at point x_i , computed from the given measurements y_j in a patch around the point of interest. In particular, the gradients used in the above expression can be estimated from the given noisy image by applying classical (i.e., nonadaptive) locally linear kernel regression. Details of this estimation procedure are given in [30]. The reader may recognize the above matrix as the well-studied “structure tensor” [31]. The advantage of the LARK descriptor is that it is exceedingly robust to noise and perturbations of the data. The formulation is also theoretically well motivated since the quadratic exponent in (6) essentially encodes the local geodesic distance between the points (x_i, y_i) and (x_j, y_j) on the graph of the function $z(x, y)$ thought of as a two-dimensional (2-D) surface (a manifold) embedded in three dimensions. The geodesic distance was also used in the context of the Beltrami-flow kernel in [32] and [33] in an analogous fashion.

GENERALIZATIONS AND PROPERTIES

The above discussion can be naturally generalized by defining the augmented data variable $\mathbf{t}_i = [x_i^T, y_i^T]^T$ and a general Gaussian kernel as follows:

$$K(\mathbf{t}_i, \mathbf{t}_j) = \exp\{- (\mathbf{t}_i - \mathbf{t}_j)^T \mathbf{Q} (\mathbf{t}_i - \mathbf{t}_j)\}, \quad (7)$$

$$\mathbf{Q}_{i,j} = \begin{bmatrix} \mathbf{Q}_x & \mathbf{0} \\ \mathbf{0} & \mathbf{Q}_y \end{bmatrix} \quad (8)$$

where \mathbf{Q} is symmetric positive definite (SPD).

Setting $\mathbf{Q}_x = \frac{1}{h_x^2} \mathbf{I}$ and $\mathbf{Q}_y = \mathbf{0}$, we have classical kernel regression, whereas one obtains the BF framework when $\mathbf{Q}_x = \frac{1}{h_x^2} \mathbf{I}$ and $\mathbf{Q}_y = \frac{1}{h_y^2} \text{diag}[0, 0, \dots, 1, \dots, 0, 0]$. The latter diagonal matrix picks out the center pixel in the element-wise difference of patches $\mathbf{t}_i - \mathbf{t}_j$. When $\mathbf{Q}_x = \mathbf{0}$ and $\mathbf{Q}_y = \frac{1}{h_y^2} \mathbf{G}$, we have the NLM filter and its variants. Finally, the LARK kernel in (6) is obtained when $\mathbf{Q}_x = \mathbf{C}_i$ and $\mathbf{Q}_y = \mathbf{0}$. More generally, the matrix \mathbf{Q} can be selected so that it has nonzero off-diagonal blocks. However, no practical algorithms with this choice have been proposed so far. As detailed below, with an SPD \mathbf{Q} , this general approach results in valid SPD kernels, a property that is used throughout the rest of our discussion. The definition of \mathbf{t} given here is only one of many possible choices. Our treatment in this article is equally valid when, for instance, $\mathbf{t} = \mathbf{T}(x, y)$ is any feature derived from a convenient linear or nonlinear transformation of the original data.

The above concepts can be further extended using the powerful theory of reproducible kernels originated in functional analysis (and later successfully adopted in machine learning [34], [35]) to present a significantly more general framework for selecting the similarity functions. This will help us identify the wider class of admissible similarity kernels more formally, and to understand how to produce new kernels from ones already defined [35]. Formally, a scalar-valued function $K(\mathbf{t}, \mathbf{s})$ over a compact region of its domain \mathbb{R}^n is called an admissible kernel if

- K is symmetric: $K(t, s) = K(s, t)$
- K is positive definite. That is, for any collection of points t_i , $i = 1, \dots, n$, the Gram matrix with elements $K_{i,j} = K(t_i, t_j)$ is positive definite.

THE PERFORMANCE OF ANY KERNEL-BASED DENOISING METHOD CAN BE IMPROVED BY SOME TYPE OF ITERATION. THE KEY IS TO USE THE RIGHT ITERATION.

nonparametric approach in (3) can be further extended to include a more general expansion of the signal $z(x)$ in a desired basis. We briefly discuss this case in “Generalization to Arbitrary Bases,”

but leave its full treatment for future research.

Such kernels satisfy some useful properties such as positivity, $K(t, t) \geq 0$, and the Cauchy–Schwartz inequality $K(t, s) \leq \sqrt{K(t, t)K(s, s)}$.

With the above definitions in place, there are numerous ways to construct new valid kernels from existing ones. We list some of the most useful ones below, without proof [35]. Given two valid kernels $K_1(t, s)$ and $K_2(t, s)$, the following constructions yield admissible kernels:

- 1) $K(t, s) = \alpha K_1(t, s) + \beta K_2(t, s)$ for any pair $\alpha, \beta \geq 0$
- 2) $K(t, s) = K_1(t, s)K_2(t, s)$
- 3) $K(t, s) = k(t)k(s)$, where $k(\cdot)$ is a scalar-valued function
- 4) $K(t, s) = p(K_1(t, s))$, where $p(\cdot)$ is a polynomial with positive coefficients
- 5) $K(t, s) = \exp(K_1(t, s))$.

Regardless of the choice of the kernel function, the weighted least-square optimization problem (3) has a simple solution. In matrix notation we can write

$$\hat{z}(x_j) = \arg \min_{z(x_j)} [\mathbf{y} - z(x_j) \mathbf{1}_n]^T \mathbf{K}_j [\mathbf{y} - z(x_j) \mathbf{1}_n], \quad (9)$$

where $\mathbf{1}_n = [1, \dots, 1]^T$, and $\mathbf{K}_j = \text{diag}[K(x_1, x_j, y_1, y_j), K(x_2, x_j, y_2, y_j), \dots, K(x_n, x_j, y_n, y_j)]$. The closed-form solution to the above is

$$\begin{aligned} \hat{z}(x_j) &= (\mathbf{1}_n^T \mathbf{K}_j \mathbf{1}_n)^{-1} \mathbf{1}_n^T \mathbf{K}_j \mathbf{y} \\ &= \left(\sum_i K(x_i, x_j, y_i, y_j) \right)^{-1} \left(\sum_i K(x_i, x_j, y_i, y_j) y_i \right) \\ &= \sum_i \frac{K_{ij}}{\sum_i K_{ij}} y_i \\ &= \sum_i W_{ij} y_i \\ &= \mathbf{w}_j^T \mathbf{y}. \end{aligned} \quad (10)$$

$$(11)$$

So in general, the estimate $\hat{z}(x_j)$ of the signal at position x_j is given by a weighted sum of all the given data points $\mathbf{y}(x_i)$, each contributing a weight commensurate with its similarity as indicated by $K(\cdot)$, with the measurement $\mathbf{y}(x_j)$ at the position of interest. Furthermore, as should be apparent in (10), the weights sum to one. To control computational complexity, or to design local versus nonlocal filters, we may choose to set the weight for some “sufficiently far-away” pixels to zero or a small value, leading to a weighted sum involving a relatively small number of data points in a properly defined vicinity of the sample of interest. This is essentially the only distinction between locally adaptive processing methods (such as BL and LARK) and so-called nonlocal methods such as NLM. It is worth noting that in the formulation above, despite the simple form of (10), in general we have a nonlinear estimator since the weights $W_{ij} = W(x_i, x_j, y_i, y_j)$ depend on the noisy data. The

but leave its full treatment for future research.

To summarize the discussion so far, we have presented a general framework that absorbs many existing algorithms as special cases. This was done in several ways, including a general description of the set of admissible similarity kernels, which allows the construction of a wide variety of new kernels not considered before in the image processing literature. Next, we turn our attention to the matrix formulation of the nonparametric filtering approach. As we shall see, this provides a framework for more in-depth and intuitive understanding of the resulting filters, their subsequent improvement, and for their respective asymptotic and numerical properties.

Before we end this section, it is worth saying a few words about computational complexity. In general, patch-based methods are quite computationally intensive. Recently, many works have aimed at both efficiently searching for similar patches, and more cleverly computing the resulting weights. Among these, notable recent work appears in [36] and [37].

THE MARTIX FORMULATION AND ITS PROPERTIES

In this section, we analyze the filtering problems posed earlier in the language of linear algebra and make several theoretical and practical observations. In particular, we are able to not only study the numerical/algebraic properties of the resulting filters, but also to analyze some of their fundamental statistical properties.

To begin, recall the convenient vector form of the filters:

$$\hat{z}(x_j) = \mathbf{w}_j^T \mathbf{y}, \quad (12)$$

where $\mathbf{w}_j^T = [W(x_1, x_j, y_1, y_j), W(x_2, x_j, y_2, y_j), \dots, W(x_n, x_j, y_n, y_j)]$ is a vector of weights for each j . Writing the above at once for all j we have

$$\hat{\mathbf{z}} = \begin{bmatrix} \mathbf{w}_1^T \\ \mathbf{w}_2^T \\ \vdots \\ \mathbf{w}_n^T \end{bmatrix} \mathbf{y} = \mathbf{W} \mathbf{y}. \quad (13)$$

As such, the filters defined by the above process can be analyzed as the product of a (square, $n \times n$) matrix of weights \mathbf{W} with the vector of the given data \mathbf{y} . First, a notational matter: \mathbf{W} is in general a function of the data, so strictly speaking, the notation $\mathbf{W}(\mathbf{y})$ would be more descriptive. But as we will describe later in more detail, the typical process for computing these weights in practice involves first computing a preliminary denoised “pilot,” or “prefiltered” version of the image, from which the weights are then calculated. This preprocessing, done only for the purposes of computing the parameters

GENERALIZATION TO ARBITRARY BASES

The nonparametric approach in (3) can be further extended to include a more general model of the signal $z(x)$ in some appropriate basis. Specifically, expanding the regression function $z(x)$ in a desired basis ϕ_i , we can formulate the following optimization problem:

$$\hat{z}(x_j) = \arg \min_{\beta(x_j)} \sum_{i=1}^n \left[y_i - \sum_{l=0}^N \beta_l(x_j) \phi_l(x_i, x_j) \right]^2 K(y_i, y, x_i, x_j), \quad (S1)$$

where N is the model (or regression) order. For instance, with the basis set $\phi_l(x_i, x_j) = (x_i - x_j)^l$, we have the Taylor series expansion, leading to the local polynomial approaches of classical kernel regression [20], [71], [9]. Alternatively, the basis vectors could be learned from the given image using a method such as K-SVD [72]. In matrix notation we have

$$\hat{\beta}(x_j) = \arg \min_{\beta(x_j)} [\mathbf{y} - \Phi_j \beta(x_j)]^T \mathbf{K}_j [\mathbf{y} - \Phi_j \beta(x_j)], \quad (S2)$$

where

$$\Phi_j = \begin{bmatrix} \phi_0(x_1, x_j) & \phi_1(x_1, x_j) & \dots & \phi_N(x_1, x_j) \\ \phi_0(x_2, x_j) & \phi_1(x_2, x_j) & \dots & \phi_N(x_2, x_j) \\ \vdots & \vdots & \ddots & \vdots \\ \phi_0(x_n, x_j) & \phi_1(x_n, x_j) & \dots & \phi_N(x_n, x_j) \end{bmatrix} \quad (S3)$$

is a matrix containing the basis vectors in its columns; and where $\beta(x_j) = [\beta_0, \beta_1, \dots, \beta_N]^T(x_j)$ is the coefficient vector of the local signal representation in this basis. Meanwhile, \mathbf{K}_j is the diagonal matrix of weights as defined in (9). To maintain the ability to represent the “DC” pixel values, we can insist that the matrix Φ_j contain the vector $\mathbf{1}_n = [1, 1, \dots, 1]^T$ as one of its columns. Without loss of generality, we assume this to be the first column so that by definition $\phi_0(x_i, x_j) = 1$ for all i and j .

of the filter, reduces the effective variance of the input noise and stabilizes the calculation of the weights so that they are not a strong function of the input noise level. As such, it is practically more useful to think of the weights as depending more directly on the underlying (unknown) image $\mathbf{W}(z)$. As such, we shall simply denote the weight matrix going forward as \mathbf{W} . A more rigorous justification for this is provided later in the next section, and supported in “Stability of Filter Coefficients to Perturbations Due to Noise.” Further justification is provided in [75].

Next, we highlight some important properties of the matrix \mathbf{W} , which lend much insight to our analysis. Referring to (10), \mathbf{W} can be written as the product

$$\mathbf{W} = \mathbf{D}^{-1} \mathbf{K}, \quad (14)$$

where $[\mathbf{K}]_{ij} = K(x_i, x_j, y_i, y_j) = K_{ij}$ and \mathbf{D} is a positive definite diagonal matrix with the normalizing factors $\mathbf{D}_{jj} = \text{diag}\{\sum_i K_{ij}\}$ along its diagonals. We can write

$$\mathbf{W} = \mathbf{D}^{-1} \mathbf{K} = \mathbf{D}^{-1/2} \underbrace{\mathbf{D}^{-1/2} \mathbf{K} \mathbf{D}^{-1/2}}_{\mathbf{L}} \mathbf{D}^{1/2}. \quad (15)$$

Denoting the j th row of Φ_j as $\phi_j^T = [\phi_0(x_j, x_j), \phi_1(x_j, x_j), \dots, \phi_N(x_j, x_j)]$ we have the closed-form solution

$$\hat{z}(x_j) = \phi_j^T \hat{\beta}(x_j) \quad (S4)$$

$$= \underbrace{\phi_j^T (\Phi_j^T \mathbf{K}_j \Phi_j)^{-1} \Phi_j^T \mathbf{K}_j \mathbf{y}}_{w_j^T} = \mathbf{w}_j^T \mathbf{y}. \quad (S5)$$

This again is a weighted combination of the pixels, though a rather more complicated one than the earlier formulation. Interestingly, the filter vector \mathbf{w}_j still has elements that sum to one. This can be seen as follows. We have

$$\mathbf{w}_j^T \Phi_j = \phi_j^T (\Phi_j^T \mathbf{K}_j \Phi_j)^{-1} \Phi_j^T \mathbf{K}_j \Phi_j = \phi_j^T.$$

Writing this explicitly, we observe

$$\mathbf{w}_j^T \begin{bmatrix} 1 & \phi_1(x_1, x_j) & \dots & \phi_N(x_1, x_j) \\ 1 & \phi_1(x_2, x_j) & \dots & \phi_N(x_2, x_j) \\ \vdots & \vdots & \ddots & \vdots \\ 1 & \phi_1(x_n, x_j) & \dots & \phi_N(x_n, x_j) \end{bmatrix} = [1, \phi_1(x_j, x_j), \dots, \phi_N(x_j, x_j)].$$

Considering the inner product of \mathbf{w}_j^T with the first column of Φ_j , we have $\mathbf{w}_j^T \mathbf{1}_n = 1$ as claimed.

The use of a general basis as described above carries one disadvantage, since the coefficient vectors \mathbf{w}_j are influenced now not only by the choice of the kernel, but also by the choice of the basis, the elements of \mathbf{w}_j can no longer be guaranteed to be positive. While this can be useful in the context of filtering as it affords additional degrees of freedom, it does significantly complicate the analysis of the resulting algorithms. Most notably, the resulting weight matrix \mathbf{W} will no longer have nonnegative elements, and hence the powerful machinery afforded by the Perron–Frobenius theory is no longer directly applicable. As such, we must resort to a broader definition of a generalized stochastic matrix [73].

It is evident that since \mathbf{K} and \mathbf{D} are SPD, then so is \mathbf{L} . Meanwhile, \mathbf{W} and \mathbf{L} are related through a diagonal similarity transformation, and therefore have the same eigenvalues. Hence, interestingly, \mathbf{W} has real, nonnegative eigenvalues (even though it is not symmetric). An alternative, though not necessarily more intuitive, description of the action of \mathbf{W} is possible in terms of graphical models. Specifically, if we consider the data (x_i, y_i) to be nodes on a (finite, undirected, weighted) graph with weights between nodes i and j given by the kernel values $K(x_i, x_j, y_i, y_j)$, the matrix $\mathcal{L} = \mathbf{L} - \mathbf{I}$ is precisely the “graph Laplacian” [38], [5], [39]. The name is not surprising. In the canonical case, \mathbf{W} is a low-pass filter (though possibly nonlinear and space-varying.) Since \mathbf{W} and \mathbf{L} share the same spectrum, \mathbf{L} is also a low-pass filter. Therefore, $\mathcal{L} = \mathbf{L} - \mathbf{I}$ is generically a high-pass filter, justifying the Laplacian label from the point of view of filtering. This is illustrated in Figure 2.

Next, we note that \mathbf{W} is a positive row-stochastic matrix—specifically, by definition, the sum of each of its rows is equal to one, as should be clear from (10). Furthermore, the structure of the pixel data as samples of a function on a finite, fixed, lattice means that \mathbf{W} can also be interpreted [5], [39] as the transition matrix of an aperiodic and irreducible Markov

STABILITY OF FILTER COEFFICIENTS TO PERTURBATIONS DUE TO NOISE

The statistical analysis of the nonparametric filters \mathbf{W} is quite complicated when \mathbf{W} is considered random. Fortunately, however, the stability results in [47] allow us to consider the resulting filters as approximately deterministic. More specifically, we assume that the noise \mathbf{e} corrupting the data has i.i.d. samples from a symmetric (but not necessarily Gaussian) distribution, with zero-mean and finite variance σ^2 . The results in [47] imply that if the noise variance σ^2 is small relative to the clean data \mathbf{z} (i.e., when signal-to-noise ratio is high), the weight matrix $\overline{\mathbf{W}} = \mathbf{W}(\mathbf{y})$ computed from the noisy data $\mathbf{y} = \mathbf{z} + \mathbf{e}$ is near the latent weight matrix $\mathbf{W}(\mathbf{z})$. That is, with the number of samples n sufficiently large [47],

$$\|\overline{\mathbf{W}} - \mathbf{W}\|_F^2 \leq_p c_1 \sigma_e^{(2)} + c_2 \sigma_e^{(4)}, \quad (56)$$

for some constants c_1 and c_2 , where $\sigma_e^{(2)}$ and $\sigma_e^{(4)}$ denote the second and fourth order moments of $\|\mathbf{e}\|$, respectively, and " \leq_p " indicates that the inequality holds in probability. When

the noise variance is sufficiently small, the fourth-order moment is even smaller; hence, the change in the resulting coefficient matrix is bounded by a constant multiple of the small noise variance.

Approximations to the moments of the perturbation $d\mathbf{W} = \overline{\mathbf{W}} - \mathbf{W}$ are also given in [47] and [78]

$$\begin{aligned} E[d\mathbf{W}] &\approx E[d\mathbf{D}]\mathbf{D}^{-2}\mathbf{K} - \mathbf{D}^{-1}E[d\mathbf{K}], \\ E[d\mathbf{W}^2] &\approx E[d\mathbf{D}^2]\mathbf{D}^{-4}\mathbf{K}^2 + \mathbf{D}^{-2}E[d\mathbf{K}^2] - E[d\mathbf{D}d\mathbf{K}]\mathbf{D}^{-3} \circ \mathbf{K}, \end{aligned}$$

where $d\mathbf{D} = \overline{\mathbf{D}} - \mathbf{D}$, and $d\mathbf{K} = \overline{\mathbf{K}} - \mathbf{K}$, and \circ denotes an element-wise product. A thorough description of the perturbation analysis for the class of Gaussian data-adaptive weights is given in [47] and [78]. To summarize, for a sufficiently small noise variance, the matrix of weights can be treated, with high confidence, as a deterministic quantity that depends only on the content of the underlying latent image \mathbf{z} .

chain defined on the entries of the data vector \mathbf{y} . Hence, the classical Perron–Frobenius theory is at our disposal [40], [41], which completely characterizes the spectra of such matrices. In particular [42], [43], we have that \mathbf{W} has spectral radius $0 \leq \lambda(\mathbf{W}) \leq 1$. Indeed, the largest eigenvalue of \mathbf{W} is simple (unrepeated) and is exactly $\lambda_1 = 1$, with corresponding eigenvector $\mathbf{v}_1 = (1/\sqrt{n})[1, 1, \dots, 1]^T = (1/\sqrt{n})\mathbf{1}_n$. Intuitively, this means that filtering by \mathbf{W} preserves the average grey level, and will leave a constant signal (i.e., a “flat” image) unchanged. In fact, with its spectrum inside the unit disk, \mathbf{W} is an ergodic matrix [40], [41], which is to say that its powers converge to a matrix of rank one, with identical rows. More explicitly, using the eigen-decomposition of \mathbf{W} we can write

$$\mathbf{W}^k = \mathbf{V}\mathbf{S}^k\mathbf{V}^{-1} = \mathbf{V}\mathbf{S}^k\mathbf{U}^T = \sum_{i=1}^n \lambda_i^k \mathbf{v}_i \mathbf{u}_i^T,$$

where we have defined $\mathbf{U}^T = \mathbf{V}^{-1}$, while denoting by \mathbf{u}_i^T the left eigenvectors of \mathbf{W} which are the columns of \mathbf{U} . Therefore,

$$\lim_{k \rightarrow \infty} \mathbf{W}^k = \mathbf{1}_n \mathbf{u}_1^T.$$

So \mathbf{u}_1 summarizes the asymptotic effect of applying the filter \mathbf{W} many times.


As we made clear earlier, $\mathbf{W} = \mathbf{D}^{-1}\mathbf{K}$ is generically not a symmetric matrix, though it always has real, positive eigenvalues. As such, it should not come as a surprise that \mathbf{W} is quite close to an SPD matrix. As we illustrate in “Symmetric Approximation of \mathbf{W} and its Properties,” this is indeed the case. As such, we shall find it useful to approximate \mathbf{W} with such an SPD matrix. To effect this approximation requires a bit of care, as the resulting matrix must still be row-stochastic. Fortunately, this is not difficult to do. Sinkhorn’s algorithm [44], [45] for matrix scaling provides a provably optimal [46] way to accomplish this task. We refer the reader to [75] and

“Symmetric Approximation of \mathbf{W} and its Properties” and citations therein, where we illustrate and justify both the approximation procedure and its fidelity in detail. The obtained matrix is a close approximation of a given \mathbf{W} which is both SPD, and now doubly (i.e., row- and column-) stochastic. Working with a symmetric (or rather “symmetrized”) \mathbf{W} makes possible much of the analysis that follows in the remainder of this article. Interestingly, there is an inherent and practical advantage in using symmetrization. As we will see in the experimental results, given any \mathbf{W} , employing its symmetrized version generally improves performance in the MSE sense. We do not present a proof of this observation here, but refer the interested reader for details in [75].

Therefore, from this point forward, we shall consider \mathbf{W} to be an SPD, and (doubly) stochastic matrix. The spectrum of \mathbf{W} determines the effect of the filter on the noisy signal \mathbf{y} . We write the eigen-decomposition

$$\mathbf{W} = \mathbf{V}\mathbf{S}\mathbf{V}^T, \quad (16)$$

where $\mathbf{S} = \text{diag}[\lambda_1, \dots, \lambda_n]$ contains the eigenvalues in decreasing order $0 \leq \lambda_n \leq \dots \leq \lambda_1 = 1$, and \mathbf{V} is an orthogonal matrix. As an example, we illustrate the spectrum of the LARK [9]



[FIG2] Construction of the graph Laplacian \mathcal{L} from the kernel weights \mathbf{K} .

SYMMETRIC APPROXIMATION OF \mathbf{W} AND ITS PROPERTIES

As elaborated in detail in [75], we approximate the matrix $\mathbf{W} = \mathbf{D}^{-1}\mathbf{K}$ with a doubly stochastic (symmetric) positive definite matrix. The algorithm we use to effect this approximation is due to Sinkhorn [44], [45]. He proved that given an irreducible matrix with nonnegative elements (such as \mathbf{W}), there exist diagonal matrices $\mathbf{R} = \text{diag}(\mathbf{r})$ and $\mathbf{C} = \text{diag}(\mathbf{c})$ such that

$$\widehat{\mathbf{W}} = \mathbf{R}\mathbf{W}\mathbf{C}$$

is doubly stochastic. That is,

$$\widehat{\mathbf{W}}\mathbf{1}_n = \mathbf{1}_n \quad \text{and} \quad \mathbf{1}_n^T \widehat{\mathbf{W}} = \mathbf{1}_n^T. \quad (S7)$$

Furthermore, the vectors \mathbf{r} and \mathbf{c} are unique to within a scalar (i.e., $\alpha \mathbf{r}$, $\alpha \mathbf{c}$.) Sinkhorn's algorithm for obtaining \mathbf{r} and \mathbf{c} in effect involves repeated normalization of the rows and columns (see Algorithm 1 below for details) so that they sum to one, and is provably convergent and optimal in the cross-entropy sense [46]. To see that the resulting matrix $\widehat{\mathbf{W}}$ is SPD, we note a corollary of Sinkhorn's result: When a symmetric matrix is considered, the diagonal scalings are in fact identical [44], [45]. Applied to \mathbf{K} , we see that the symmetric diagonal scaling $\widehat{\mathbf{K}} = \mathbf{M}\mathbf{K}\mathbf{M}$ yields a symmetric doubly stochastic matrix. But we have $\mathbf{W} = \mathbf{D}^{-1}\mathbf{K}$, so

$$\widehat{\mathbf{W}} = \mathbf{R}\mathbf{D}^{-1}\mathbf{K}\mathbf{C}.$$

Since the $\widehat{\mathbf{W}}$ and the diagonal scalings are unique (to within a scalar), we must have $\mathbf{R}\mathbf{D}^{-1} = \mathbf{C} = \mathbf{M}$, and therefore $\widehat{\mathbf{W}} = \widehat{\mathbf{K}}$. That is to say, applying Sinkhorn's algorithm to either \mathbf{K} or its row-sum normalized version \mathbf{W} yields the very same SPD doubly stochastic result, $\widehat{\mathbf{W}}$.

Algorithm 1: Scaling a matrix \mathbf{A} to a nearby doubly stochastic matrix $\widehat{\mathbf{A}}$

Given a matrix \mathbf{A} , let $(n, n) = \text{size}(\mathbf{A})$ and initialize $\mathbf{r} = \text{ones}(n, 1)$; for $k = 1 : \text{iter}$;

$$\mathbf{c} = 1./(\mathbf{A}^T \mathbf{r});$$

$$\mathbf{r} = 1./(\mathbf{A} \mathbf{c});$$

end

$$\mathbf{C} = \text{diag}(\mathbf{c}); \mathbf{R} = \text{diag}(\mathbf{r});$$

$$\widehat{\mathbf{A}} = \mathbf{R} \mathbf{A} \mathbf{C}$$

The convergence of this iterative algorithm is known to be linear [45], with rate given by the subdominant eigenvalue λ_2 of \mathbf{W} . It is of interest to know how much the diagonal scalings will perturb the eigenvalues of \mathbf{W} . This will indicate how accurate our approximate analysis will be, which uses the eigenvalues of $\widehat{\mathbf{W}}$ instead of \mathbf{W} . Fortunately, Kahan [74], [43] provides a nice result: Suppose that \mathbf{A} is a non-Hermitian matrix with eigenvalues $|\lambda_1| \geq |\lambda_2| \geq \dots \geq |\lambda_n|$. Let $\widehat{\mathbf{A}}$ be a

perturbation of \mathbf{A} that is Hermitian with eigenvalues $\widehat{\lambda}_1 \geq \widehat{\lambda}_2 \geq \dots \geq \widehat{\lambda}_n$, then

$$\sum_{i=1}^n |\lambda_i - \widehat{\lambda}_i|^2 \leq 2 \|\mathbf{A} - \widehat{\mathbf{A}}\|_F^2.$$

Specializing the result for \mathbf{W} , defining $\Delta = \mathbf{W} - \widehat{\mathbf{W}}$ and normalizing by n , we have

$$\frac{1}{n} \sum_{i=1}^n |\lambda_i - \widehat{\lambda}_i|^2 \leq \frac{1}{n} \|\Delta\|_F^2.$$


The right-hand side can be further bounded as described in [75].

Hence, for patches of reasonable size (say $n \geq 11$), the bound on the right-hand side and the eigenvalues of the respective matrices are quite close, as demonstrated in Figure S1. Bounds may be established on the distance between the eigenvectors of \mathbf{W} and $\widehat{\mathbf{W}}$ as well [75]. Of course, since both matrices are row-stochastic, they share their first right eigenvector $\widehat{\mathbf{v}}_1 = \mathbf{v}_1$, corresponding to a unit eigenvalue. The second eigenvectors are arguably more important, as they tell us about the dominant structure of the region being filtered, and the corresponding effect of the filters. Specifically, as indicated in [47], applying Theorem 5.2.8 from [43] gives a bound on the perturbation of the second eigenvectors

$$\|\widehat{\mathbf{v}}_2 - \mathbf{v}_2\| \leq \frac{\|4\Delta\|_F}{\nu - \sqrt{2}\|\Delta\|_F},$$

where ν is the gap between the second and the third eigenvalues. It is worth noting that the bound essentially does not depend on the dimension n of the data y .


As a side note, we mention the interesting fact that, when applied to a symmetric squared distance matrix, Sinkhorn's algorithm has a rather nice geometric interpretation [76], [77]. Specifically, if P is the set of points that generated the symmetric distance matrix A , then elements of its Sinkhorn scaled version $B = \mathbf{M}\mathbf{A}\mathbf{M}$ correspond to the square distances between the corresponding points in a set Q , where the points in Q are obtained by a stereographic projection of the points in P . The points in Q are confined to a hypersphere of dimension d embedded in \mathbb{R}^{d+1} , where d is the dimension of the subspace spanned by the points in P .



[FIGS1] Eigenvalues of row-stochastic (asymmetric) filter matrices and the corresponding doubly stochastic (symmetric) filter matrices symmetrized using Sinkhorn's algorithm: (a) an NLM filter example and (b) a LARK filter example.

STATISTICAL ANALYSIS OF THE FILTERS' PERFORMANCE

We are now in a position to carry out an analysis of the performance of filters defined earlier. As we just explained, to carry out this analysis, we first note that W is, strictly speaking, a function of the noisy data y . Indeed, the dependence on the given data is what gives strength to these adaptive filters. On the other hand, one may legitimately worry that the effect of noise on the computation of these weights may be dramatic, resulting in too much sensitivity for the resulting filters to be effective. This issue has indeed been considered in the context of most patch-based adaptive filters. The consensus approach, in keeping with similar thinking in the nonparametric statistics literature, is to compute the weights from a “pilot” (or prefiltered) version \hat{z}_0 of the noisy data; that is, $W(\hat{z}_0)$. The preprocessing step yielding



[FIG3] Spectrum of the LARK filter on different types of patches.

kernel in Figure 3 for different types of patches. As can be observed, the spectrum of W decays rather quickly for “flat” patches, indicating that the filter is very aggressive in denoising these types of patches. As we consider increasingly more complex anisotropic patches containing edges, corners, and various textures, the spectrum of the filter is (automatically) adjusted to be less aggressive in particular directions. To take the specific case of a patch containing a simple edge, the filter is adjusted in unsupervised fashion to perform strong denoising along the edge, and to preserve and enhance the structure of the patch across the edge. This type of behavior is typical of the adaptive nonparametric filters we have discussed so far, but it is particularly striking and stable in the case of LARK kernels shown in Figure 3.

Of course, the matrix W is a function of the noisy data y and hence it is strictly speaking a stochastic variable. But another interesting property of W is that when the similarity kernel $K(\cdot)$ is of the general Gaussian type as in (7), as described in “Symmetric Approximation of W and its Properties” the resulting filter coefficients W are quite stable to perturbations of the data by noise of modest size [47]. From a practical point of view, and insofar as the computation of the matrix W is concerned, it is always reasonable to assume that the noise variance is relatively small, because in practice we typically compute W on a “prefiltered” version of the noisy image y anyway. This small noise assumption is made only for the analysis of the filter coefficients, and is not invoked in rest of the article. Going forward, therefore, we consider W as essentially nonstochastic. Some experiments in the section “Improving the Estimate by Iteration” and in [75] confirm the validity of this assumption. We also elaborate a bit more on this point in the next section.

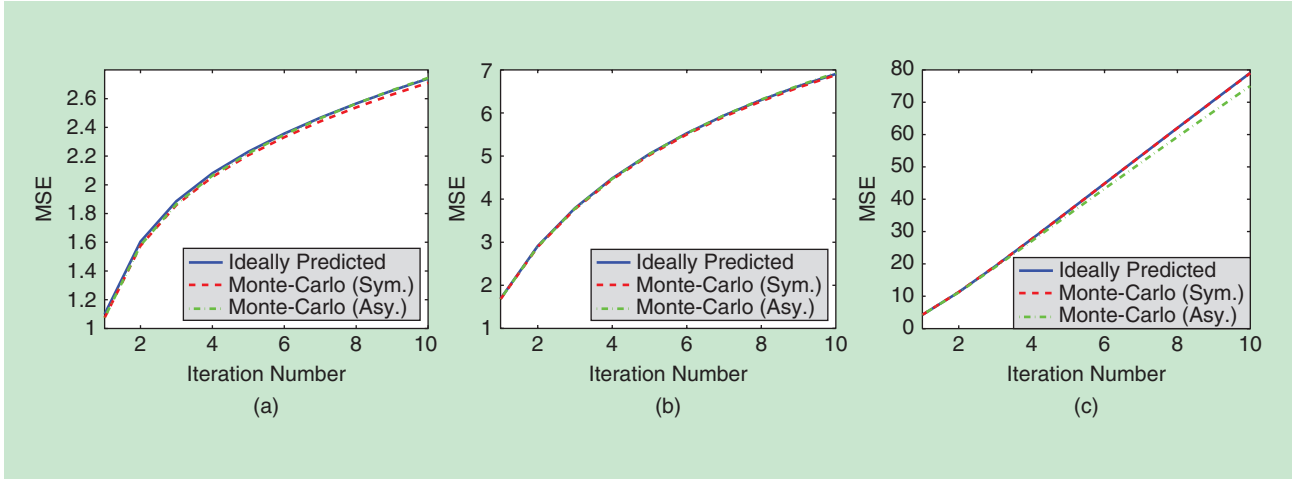
the pilot estimate \hat{z}_0 can be carried out with any one of many simpler (nonadaptive) filters, and is done only for the purpose of stably computing the parameters of the filter. The net effect of this step, however, is significant: it reduces the effective variance of the input noise and stabilizes the calculation of the weights so that they are not a strong function of the input noise level any longer. In this setting, the noise affecting the calculation of the weights can, for all practical intents, be assumed small. As such, the stability results in [47] (summarized in “Stability of Filter Coefficients to Perturbations Due to Noise”) give strong guarantees that the weights are consequently disturbed relatively little. That is, $W(\hat{z}_0) \approx W(z)$. For the purposes of performance analysis below, it is therefore reasonable (and practically verified in Figures 4–7) to consider the weights as depending more directly on the underlying (unknown) image $W(z)$, rather than $W(y)$. As such, we also drop the dependence on its argument for the sake of simplifying the notation. But the reader is advised to always remember that the weights are adapted to the (estimated) structure of the underlying image.

Now, let us compute an approximation to the bias and variance of the estimator $\hat{z} = Wy$. The bias in the estimate is

$$\begin{aligned} \text{bias} &= E(\hat{z}) - z = E(Wy) - z \\ &\approx Wz - z = (W - I)z. \end{aligned}$$

Recalling that the matrix W has a unit eigenvalue corresponding to a (constant) vector, we note that \hat{z} is an unbiased estimate if z is a constant image, but is biased for all other underlying latent images. The squared magnitude of the bias is given by

$$\|\text{bias}\|^2 = \|(W - I)z\|^2. \quad (17)$$



[FIG4] Parts (a)–(c) show plots of predicted MSE in diffusion based on ideally estimated symmetric LARK filters, and MSE of noisy symmetric/asymmetric LARK filters through Monte-Carlo simulations. The input noise variance $\sigma^2 = 0.5$, and for each simulation 100 independent noise realizations are implemented.

Writing the image \mathbf{z} in the column space of the orthogonal matrix \mathbf{V} (which contains the eigenvectors as its columns) as $\mathbf{z} = \mathbf{V}\mathbf{b}$, we can rewrite the squared bias magnitude as

$$\begin{aligned} \|\text{bias}\|^2 &= \|(\mathbf{W} - \mathbf{I})\mathbf{z}\|^2 = \|\mathbf{V}(\mathbf{S} - \mathbf{I})\mathbf{b}\|^2 \\ &= \|(\mathbf{S} - \mathbf{I})\mathbf{b}\|^2 = \sum_{i=1}^n (\lambda_i - 1)^2 b_i^2. \end{aligned} \quad (18)$$

We note that the last sum can be expressed over the indices $i = 2, \dots, n$, because the first term in fact vanishes since $\lambda_1 = 1$. Next, we consider the variance of the estimate. We have

$$\begin{aligned} \text{cov}(\hat{\mathbf{z}}) &= \text{cov}(\mathbf{W}\mathbf{y}) \approx \text{cov}(\mathbf{W}\mathbf{e}) = \sigma^2 \mathbf{W}\mathbf{W}^T \implies \\ \text{var}(\hat{\mathbf{z}}) &= \text{tr}(\text{cov}(\hat{\mathbf{z}})) = \sigma^2 \sum_{i=1}^n \lambda_i^2. \end{aligned}$$

Overall, the MSE is given by


$$\text{MSE} = \|\text{bias}\|^2 + \text{var}(\hat{\mathbf{z}}) = \sum_{i=1}^n (\lambda_i - 1)^2 b_i^2 + \sigma^2 \lambda_i^2. \quad (19)$$

THE IDEAL SPECTRUM: BM3D


Let us pause for a moment and imagine that we can optimally design the matrix \mathbf{W} from scratch in such a way as to minimize the MSE expression above. That is, consider the set of eigenvalues for which (19) is minimized. Differentiating the expression for MSE with respect to λ_i and setting equal to zero leads to a familiar expression

$$\lambda_i^* = \frac{b_i^2}{b_i^2 + \sigma^2} = \frac{1}{1 + \text{snr}_i^{-1}}. \quad (20)$$

This is, of course, nothing but a manifestation of the Wiener filter with $\text{snr}_i = \frac{b_i^2}{\sigma^2}$ denoting the signal-to-noise ratio at each component i of the signal. This observation naturally raises the question of whether any existing filters in fact get close to this optimal performance [49]. Two filters that are designed to approximately achieve this goal are BM3D [48] and patchwise locally optimal Wiener (PLOW) [50], which are currently considered to be the state of the art in denoising. The BM3D



[FIG5] Parts (a)–(c) show plots of predicted MSE in diffusion based on ideally estimated symmetric NLM filters, and MSE of noisy symmetric/asymmetric NLM filters through Monte-Carlo simulations. The input noise variance $\sigma^2 = 0.5$ and for each simulation 100 independent noise realizations are implemented.



[FIG6] Parts (a)–(c) show plots of predicted MSE in twicing based on ideally estimated symmetric LARK filters, and MSE of noisy symmetric/asymmetric LARK filters through Monte-Carlo simulations. The input noise variance $\sigma^2 = 0.5$ and for each simulation 100 independent noise realizations are implemented.

algorithm can be briefly summarized by the following three steps:

- 1) Patches from an input image are classified according to their similarity, and are grouped into 3-D clusters accordingly.
- 2) A so-called “3-D collaborative Wiener filtering” is implemented to process each cluster.
- 3) Filtered patches inside all the clusters are aggregated to form an output.

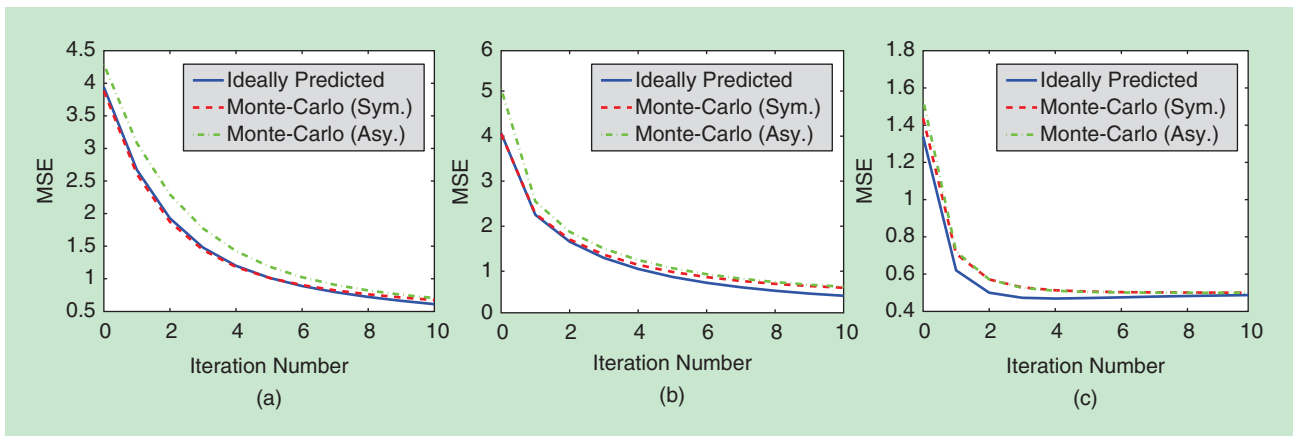
The core process of the BM3D algorithm is the collaborative Wiener filtering that transforms a patch cluster through a fixed orthonormal basis V (e.g., DCT) [48], where the coefficients of each component are then scaled by some “shrinkage” factor $\{\lambda_i\}$, which are precisely chosen according to (20). Following this shrinkage step, the data are transformed back to the spatial domain to generate denoised

patches for the aggregation process. In practice, of course, one does not have access to $\{b_i^2\}$, so they are estimated from the given noisy image y , as is done also in [48]. The BM3D is not based on the design of a specific kernel $K(\cdot)$, or even on a regression framework, as are other patch-based methods (such as NLN and LARK). In fact, there need not exist any closed-form kernel that gives rise to the BM3D filter. Still, with the present framework, the diagonal Wiener shrinkage matrix

$S = \text{diag}[\lambda_1, \dots, \lambda_n]$ and the overall 3-D collaborative Wiener filtering process can now be understood using a symmetric, positive-definite filter matrix W with orthonormal

ONE WAY TO REMEDY THE LACK OF OPTIMALITY OF THE CHOICE OF KERNEL IS TO APPLY THE RESULTING FILTERS ITERATIVELY.

eigen-decomposition VSV^T . If the orthonormal basis V contains a constant vector $v_i = \frac{1}{\sqrt{n}}\mathbf{1}_n$, one can easily make W doubly stochastic by setting its corresponding shrinkage factor $\lambda_i = 1$.



[FIG7] Parts (a)–(c) show plots of predicted MSE in diffusion based on ideally estimated symmetric NLM filters, and MSE of noisy symmetric/asymmetric NLM filters through Monte-Carlo simulations. The input noise variance $\sigma^2 = 0.5$ and for each simulation 100 independent noise realizations are implemented.

To summarize this section, we have presented a widely applicable matrix formulation $\hat{z} = \mathbf{W}y$ for the denoising problem, where \mathbf{W} is henceforth considered symmetric, positive-definite, and doubly stochastic. This formulation has thus far allowed us to characterize the performance of the resulting filters in terms of their spectra and to obtain insights as to the relative statistical efficiency of existing filters such as BM3D. One final observation worth making here is that the Birkhoff-von Neumann theorem [41] notes that \mathbf{W} is doubly stochastic if and only if it is a convex combination of permutation matrices. Specifically, $\mathbf{W} = \sum_{l=1}^{\bar{n}} \alpha_l \mathbf{P}_l$, where \mathbf{P}_l are permutation matrices; α_l are nonnegative scalars with $\sum \alpha_l = 1$, and \bar{n} is at most $(n-1)^2 + 1$. This means that regardless of the choice of (an admissible) kernel, each pixel of the estimated image

$$\hat{z} = \mathbf{W}y = \sum_{l=1}^{\bar{n}} \alpha_l (\mathbf{P}_l y)$$

is always a convex combination of at most $\bar{n} = (n-1)^2 + 1$ (not necessarily nearby) pixels of the noisy input.

Going forward, we first use the insights gained thus far to study further improvements of the nonparametric approach. Then, we describe the relationship between the nonparametric approaches and more familiar parametric (empirical) Bayesian methods.

IMPROVING THE ESTIMATE BY ITERATION

As we described above, the optimal spectrum for the filter matrix \mathbf{W} is dictated by the Wiener condition (20), which requires clairvoyant knowledge of the signal or at least careful estimation of the SNR at each spatial frequency. In practice, however, nearly all similarity-based methods we described so far (with the notable exception of the BM3D) design a matrix \mathbf{W} based on a choice of the kernel function, and without regard to the ultimate consequences of this choice on the spectrum of the resulting filter. Hence, at least in the MSE sense, such choices of \mathbf{W} are always deficient. As former U.S. Secretary of Defense Donald Rumsfeld might have said, we filter with the spectrum that we have, not the spectrum that we would like to have! One way to remedy the lack of optimality of the choice of kernel is to apply the resulting filters iteratively, and that is the subject of this section.

To put the problem in more practical terms, the estimate derived from the application of a nonparametric denoising filter \mathbf{W} will not only remove some noise, but invariably some of the underlying signal as well. A number of different approaches have been proposed to reintroduce this “lost” component back to the estimate. Interestingly, as we describe later in the section “Relationship to (Empirical) Bayes’ and Regularization Approaches,” similar to the steepest descent (SD) methods employed in the solution of more classical Bayesian optimization approaches, the iterative method employed here involve multiple applications of a filter (or sequence of filters) on the data, and somehow pooling the results. Next we discuss the two most well-known and frequently used approaches. The first is anisotropic diffusion

[51], whereas the second class includes the recently popularized Bregman iterations [52], which is closely related to ℓ_2 -boosting [53]. The latter is a generalization of twicing introduced by Tukey [54] more than 30 years ago. In fact, as we will illustrate later, this latter class of iterations is also related to biased anisotropic diffusion [55], otherwise known as anisotropic reaction-diffusion.

DIFFUSION

So far, we have described a general formulation for denoising as a spatially adaptive, data-dependent filtering procedure (3) amounting to

$$\hat{z} = \mathbf{W}y.$$

Now consider applying the filter multiple times. That is, define $\hat{z}_0 = y$, and the iteration

$$\hat{z}_k = \mathbf{W}\hat{z}_{k-1} = \mathbf{W}^k y. \quad (21)$$

The net effect of each application of \mathbf{W} is essentially a step of anisotropic diffusion [51], [5]. Rewriting the iteration in (21), we have

$$\hat{z}_k = \mathbf{W}\hat{z}_{k-1}, \quad (22)$$

$$= \hat{z}_{k-1} - \hat{z}_{k-1} + \mathbf{W}\hat{z}_{k-1}, \quad (23)$$

$$= \hat{z}_{k-1} + (\mathbf{W} - \mathbf{I})\hat{z}_{k-1}, \quad (24)$$

which gives

$$\hat{z}_k - \hat{z}_{k-1} = (\mathbf{W} - \mathbf{I})\hat{z}_{k-1}. \quad (25)$$

Recall that $\mathbf{W} = \mathbf{D}^{-1/2} \mathbf{L} \mathbf{D}^{1/2}$. Hence, we have

$$\mathbf{W} - \mathbf{I} = \mathbf{D}^{-1/2} (\mathbf{L} - \mathbf{I}) \mathbf{D}^{1/2} = \mathbf{D}^{-1/2} \mathcal{L} \mathbf{D}^{1/2}, \quad (26)$$

where \mathcal{L} is the graph Laplacian operator mentioned earlier [39], [5]. Defining the normalized variable $\bar{z}_k = \mathbf{D}^{1/2} \hat{z}_k$, (25) embodies a discrete version of anisotropic diffusion


$$\bar{z}_k - \bar{z}_{k-1} = \mathcal{L} \bar{z}_{k-1}, \leftrightarrow \frac{\partial \bar{z}(t)}{\partial t} = \nabla^2 \bar{z}(t) \quad (\text{Diffusion Eqn}). \quad (27)$$

where the left-hand side of the above is a discretization of the derivative operator $\frac{\partial z(t)}{\partial t}$.

Returning to the diffusion estimate (21), the bias in the estimate after k iterations is

$$\text{bias}_k = \mathbf{E}(\hat{z}_k) - z = \mathbf{E}(\mathbf{W}^k y) - z = \mathbf{W}^k z - z = (\mathbf{W}^k - \mathbf{I})z.$$

Recalling that the matrix \mathbf{W} has a unit eigenvalue corresponding to a (constant) vector, we note that since $\mathbf{W}^k - \mathbf{I}$ has the same eigenvectors as \mathbf{W} , the above sequence of estimators produce unbiased estimates of constant images.



[FIG8] Example patches: (a) flat, (b) edge, and (c) texture.

The squared magnitude of the bias is given by

$$\|\text{bias}_k\|^2 = \|(W^k - I)z\|^2. \quad (28)$$

As before, writing the image z in the column space of V as $z = Vb$, we have

$$\|\text{bias}_k\|^2 = \|(W^k - I)z\|^2 = \|V(S^k - I)b\|^2 = \sum_{i=1}^n (\lambda_i^k - 1)^2 b_i^2. \quad (29)$$

As k grows, so does the magnitude of the bias, ultimately converging to $\|b\|^2 - b_1^2$. This behavior is consistent with what is observed in practice; specifically, increasing iterations of diffusion produce more and more blurry (biased) results. Next, we derive the variance of the diffusion estimator

$$\text{cov}(\hat{z}_k) = \text{cov}(W^k y) = \text{cov}(W^k e) = \sigma^2 W^k (W^k)^T,$$


which gives

$$\text{var}(\hat{z}_k) = \text{tr}(\text{cov}(\hat{z}_k)) = \sigma^2 \sum_{i=1}^n \lambda_i^{2k}.$$

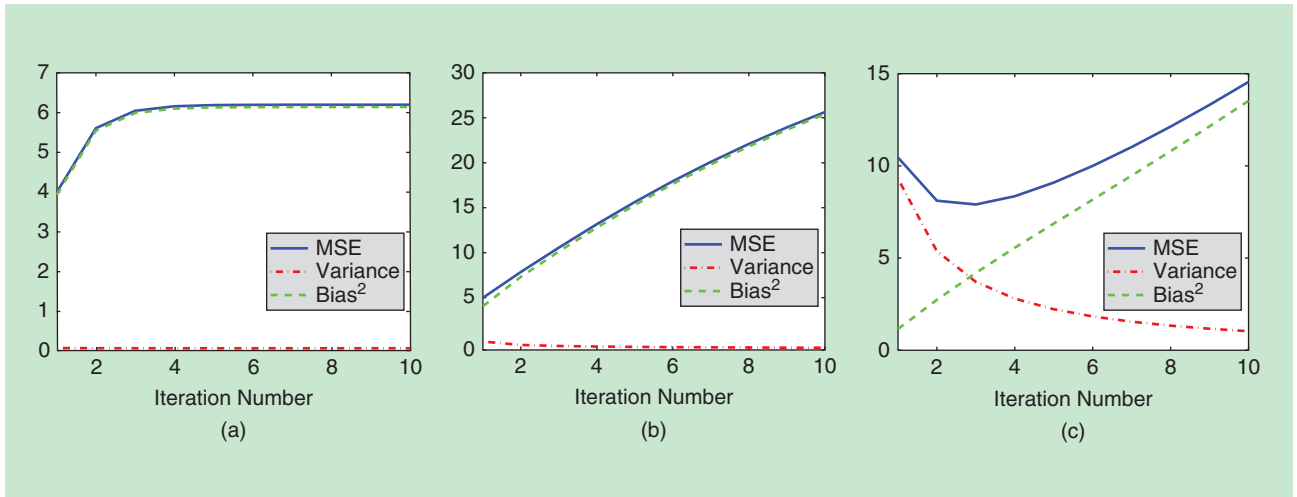
As k grows, the variance tends to a constant value of σ^2 . Overall, the MSE is given by

$$\text{MSE}_k = \|\text{bias}_k\|^2 + \text{var}(\hat{z}_k) = \sum_{i=1}^n (\lambda_i^k - 1)^2 b_i^2 + \sigma^2 \lambda_i^{2k}. \quad (30)$$

Diffusion experiments for denoising three types of patches (shown in Figure 8) are carried out to illustrate the evolution of MSE. The variance of input (Gaussian) noise is set to $\sigma^2 = 25$. Filters based on both LARK [9] and NLM [7] estimated from the latent noise-free patches are tested, and the Sinkhorn algorithm (see “Symmetric Approximation of W and its Properties”) is implemented to make the filter matrices doubly stochastic. Predicted MSE_k , $\text{var}(\hat{z}_k)$, and $\|\text{bias}_k\|^2$ according to (30) through the diffusion iteration are shown in Figure 9 (LARK) and Figure 10 (NLM), where we can see that diffusion monotonically reduces the estimation variance and increases the bias. In some cases [such as Figure 9(a) and (b)] diffusion further suppresses MSE and improves the estimation performance. True MSEs for the standard (asymmetric) filters



[FIG9] Parts (a)–(c) show plots of predicted MSE_k , $\text{var}(\hat{z}_k)$ and $\|\text{bias}_k\|^2$ using (30) for LARK [9] filters in diffusion process.



[FIG10] Parts (a)–(c) show plots of predicted MSE_k , $\text{var}(\hat{z}_k)$ and $\|\text{bias}_k\|^2$ using (30) for NLM [7] filters in diffusion process.

and their symmetrized versions computed by Monte-Carlo simulations are also shown, where in each simulation 100 independent noise realizations are averaged. We see that the estimated MSEs of the symmetrized filters match quite well with the predicted ones (see Figures 11 and 12). The asymmetric filters, meanwhile generate slightly higher MSE, but behave closely to the symmetrized ones especially in regions around the optimal MSEs.

In the next set of Monte-Carlo simulations, we illustrate the effect of \mathbf{W} being computed from clean versus noisy images (see Figures 4 and 5). Noise variance $\sigma^2 = 0.5$, which is relatively small, simulating the situation where “pre-processing” has been applied to suppress estimation variance. It can be observed that the MSEs estimated from Monte-Carlo simulations are quite close to the ones predicted from the ideal filters, which confirms the assumption in the section “The Martix Formulation and its Properties” that the filter matrix \mathbf{W} can be treated as deterministic under most circumstances.

To further analyze the change of MSE in the diffusion process, let us consider the contribution of MSE in each component (or mode) separately

$$\text{MSE}_k = \sum_{i=1}^n \text{MSE}_k^{(i)}, \quad (31)$$

where $\text{MSE}_k^{(i)}$ denotes the MSE of the i th mode in the k th diffusion iteration, which is given by


$$\text{MSE}_k^{(i)} = (\lambda_i^k - 1)^2 b_i^2 + \sigma^2 \lambda_i^{2k}. \quad (32)$$

Equivalently, we can write


$$\overline{\text{MSE}}_k^{(i)} = \frac{\text{MSE}_k^{(i)}}{\sigma^2} = \text{snr}_i (\lambda_i^k - 1)^2 + \lambda_i^{2k}, \quad (33)$$

where $\text{snr}_i = \frac{b_i^2}{\sigma^2}$ is defined as the signal to noise ratio in the i th mode.

One may be interested to know at which iteration $\text{MSE}_k^{(i)}$ is minimized. As we have seen experimentally, with an arbitrary



[FIG11] Parts (a)–(c) show plots of predicted MSE and Monte-Carlo estimated MSE in diffusion process using LARK [9] filters. In the Monte-Carlo simulations, results for row-stochastic (asymmetric) filters and their symmetrized versions are both shown.



[FIG12] Parts (a)–(c) show plots of predicted MSE and Monte-Carlo estimated MSE in diffusion process using NLM [7] filters. In the Monte-Carlo simulations, result for row-stochastic (asymmetric) filters and their symmetrized versions are both shown.

filter W , there is no guarantee that even the very first iteration of diffusion will improve the estimate in the MSE sense. The derivative of (33) with respect to k is

$$\frac{\partial \overline{\text{MSE}}_k^{(i)}}{\partial k} = 2\lambda_i^k \log \lambda_i [(\text{snr}_i + 1)\lambda_i^k - \text{snr}_i]. \quad (34)$$

For $i \geq 2$, since $\lambda_i < 1$, we note that the sign of the derivative depends on the term inside the brackets, specifically, $(\text{snr}_i + 1)\lambda_i^k - \text{snr}_i$. To guarantee that in the k th iteration the MSE of the i th mode decreases (i.e., negative derivative), we must have

$$(\text{snr}_i + 1)\lambda_i^q - \text{snr}_i > 0 \text{ for any } 0 \leq q < k. \quad (35)$$

In fact, the condition

$$(\text{snr}_i + 1)\lambda_i^k - \text{snr}_i > 0 \quad (36)$$

guarantees that the derivative is always negative for $q \in (0, k)$ because given any scalar $t = \frac{q}{k} \in (0, 1)$

$$\begin{aligned} (\text{snr}_i + 1)\lambda_i^k - \text{snr}_i > 0 &\Rightarrow \lambda_i^{tk} > \left(\frac{\text{snr}_i}{\text{snr}_i + 1}\right)^t \Rightarrow \\ (\text{snr}_i + 1)\lambda_i^{tk} > \text{snr}_i \left(\frac{\text{snr}_i}{\text{snr}_i + 1}\right)^{t-1} > \text{snr}_i \end{aligned} \quad (37)$$

The condition (36) has an interesting interpretation. Rewriting (36), we have

$$\log(1 + \text{snr}_i) < \underbrace{\log\left(\frac{1}{1 - \lambda_i^k}\right)}_{\epsilon_i}, \quad (38)$$

where $\lambda_i^k = \lambda_i^k$ denotes the i th eigenvalue of W^k . The left-hand side of the inequality is Shannon's channel capacity [56] of the i th mode for the problem $y = z + e$. We name the right-hand side expression ϵ_i the entropy of the i th mode of the filter W^k . The larger ϵ_i , the more variability can be expected in the output

image produced by the i th mode of the denoising filter. The above condition implies that if the filter entropy exceeds or equals the channel capacity of the i th mode, then the k th iteration of diffusion will in fact produce an improvement in the corresponding mode. One may also be interested in identifying an approximate value of k for which the overall MSE_k is minimized. This depends on the signal energy distribution ($\{b_i^2\}$) over all the modes, which is generally not known, but for which we may have some prior in mind. But for a given mode i , the minimum $\text{MSE}^{(i)}$ is achieved in the k^* th iteration, where

$$k^* = \log\left(\frac{\text{snr}_i}{\text{snr}_i + 1}\right) / \log \lambda_i. \quad (39)$$

Of course, in practice we can seldom find a single iteration number k^* minimizing $\text{MSE}^{(i)}$ for all the modes simultaneously; but in general, the take-away lesson is that optimizing a given filter through diffusion can [under proper conditions (38)] make its corresponding eigenvalues closer to those of an ideal Wiener filter that minimizes MSE.

Another interesting question is this: given a particular W , how much further MSE reduction can be achieved by implementing diffusion? For example, can we use diffusion to improve the 3-D collaborative Wiener filter in BM3D [48]? Let us assume that this filter has already approximately achieved the ideal Wiener filter condition in each mode, specifically

$$\lambda_i^* \approx \frac{\text{snr}_i}{\text{snr}_i + 1}. \quad (40)$$

Then we can see that

$$(\text{snr}_i + 1)(\lambda_i^*)^k - \text{snr}_i < 0 \text{ for any } k > 1, \quad (41)$$

which means for all the modes, diffusion will definitely worsen (increase) the MSE. In other words, multi-iteration diffusion is not useful for filters where the Wiener condition has been (approximately) achieved.

EXTENSIONS TO MORE GENERAL RESTORATION PROBLEMS

Many of the ideas discussed in the context of denoising here can be extended to more general image restoration problems. As a simple example, we briefly note the case where the measurement model includes a linear distortion such as blur, expressed as follows:

$$\mathbf{y} = \mathbf{A}\mathbf{z} + \mathbf{e}. \quad (S8)$$

One approach would be to employ the framework of the residual iterations defined earlier. More specifically, define

$$\hat{\mathbf{z}}_{k+1} = \hat{\mathbf{z}}_k + \mathbf{A}^T \mathbf{W}(\mathbf{y} - \mathbf{A}\hat{\mathbf{z}}_k), \quad (S9)$$

where we observe that the term $\mathbf{W}(\mathbf{y} - \mathbf{A}\hat{\mathbf{z}}_k)$ is filtering (i.e., denoising) the residuals from the measurement model in a spatially adaptive, nonparametric manner using \mathbf{W} . The operator \mathbf{A}^T is the adjoint to the blur, and applied to these filtered residuals,

acts as a “backprojection” operation. The asymptotic behavior of this iterative process is interesting. Specifically, as $k \rightarrow \infty$,

$$\mathbf{A}^T \mathbf{W} \mathbf{y} = \mathbf{A}^T \mathbf{W} \mathbf{A} \hat{\mathbf{z}}_\infty \quad (S10)$$

or equivalently

$$\hat{\mathbf{z}}_\infty = (\mathbf{A}^T \mathbf{W} \mathbf{A})^{-1} \mathbf{A}^T \mathbf{W} \mathbf{y}. \quad (S11)$$

Without the adaptive weights in \mathbf{W} , this asymptotic solution is the familiar (*global*) least squares solution, which is unstable due to the direct inverse. The inclusion of the kernel weights, and stopping the process after a finite number of iterations can result in a stabilized solution with smaller MSE. The analysis of the stability and statistical characteristics of such estimates, along with the choice of the optimal stopping point will require in depth understanding of the spectrum of the operator $\mathbf{A}^T \mathbf{W}$, much in the same way as we did with the spectrum of \mathbf{W} in the denoising framework.

TWICING, ℓ_2 -BOOSTING, REACTION-DIFFUSION, AND BREGMAN ITERATIONS

An alternative to repeated applications of the filter \mathbf{W} is to consider the residual signals, defined as the difference between the estimated signal and the measured signal. The use of the residuals in improving estimates has a rather long history, dating at least back to the work of John Tukey, who, in his landmark book [54], termed the idea *twicing*. More recently, the idea has been proposed in the applied mathematics community for application to image restoration problems under the rubric of Bregman iterations [52]. And in the machine learning and statistics literature, [53] recently proposed the idea of ℓ_2 -boosting, which may be viewed as an (ℓ_2) regression counterpart of the more familiar notion of boosting in classification, originally proposed by Freund and Schapire [57]. All of the aforementioned ways of looking at the use of residuals are independently quite rich and worth exploring on their own. Indeed, for a wide-angle view of the boosting approach to regression and regularization, and its connections to functional gradient descent interpretations, we can refer the interested reader to the nice treatment in [58]. The general idea of Bregman divergence, and its use in solving large-scale optimization problems is also a very deep area of work, for which we can refer the interested reader to the book [59] in which both the foundations and some applications to signal processing are discussed. Also, in [60], the ideas of Bregman divergence and boosting are connected. More general inverse problems such as deconvolution can also be addressed in this framework. We give a very brief hint of this direction in “Extensions to More General Restoration Problems.”

In whatever guise, the basic notion—to paraphrase Tukey—is to use the filtered residuals to enhance the estimate by adding some “roughness” to it. Put another way, if the residuals contain some of the underlying signal, filtering them should recover this leftover signal at least in part.

Formally, the residuals are defined as the difference between the estimated signal and the measured signal: $\mathbf{r}_k = \mathbf{y} - \hat{\mathbf{z}}_{k-1}$, where here we define the initialization $\hat{\mathbf{z}}_0 = \mathbf{W}\mathbf{y}$. Note that this

initialization is a once-filtered version of the noisy data using some initial filter \mathbf{W} , and different from ($\mathbf{z}_0 = \mathbf{y}$) used in the diffusion iterations. With this definition, we write the iterated estimates as

$$\hat{\mathbf{z}}_k = \hat{\mathbf{z}}_{k-1} + \mathbf{W} \mathbf{r}_k = \hat{\mathbf{z}}_{k-1} + \mathbf{W}(\mathbf{y} - \hat{\mathbf{z}}_{k-1}). \quad (42)$$

These estimates also trade off bias against variance with increasing iterations, though in a fundamentally different way than the diffusion estimator we discussed earlier. Indeed, we immediately notice that this iteration has a very different asymptotic behavior than diffusion. Specifically, as $k \rightarrow \infty$, the estimate here tends back toward the original data \mathbf{y} , whereas the diffusion tends to a constant!

For an intuitive interpretation, note that for $k = 1$, the (residual) iteration above gives

$$\hat{\mathbf{z}}_1 = \hat{\mathbf{z}}_0 + \mathbf{W}(\mathbf{y} - \hat{\mathbf{z}}_0) = \mathbf{W}\mathbf{y} + \mathbf{W}(\mathbf{y} - \mathbf{W}\mathbf{y}) = (2\mathbf{W} - \mathbf{W}^2)\mathbf{y}.$$

This first iterate $\hat{\mathbf{z}}_1$ is in fact precisely the original “twicing” estimate of Tukey [54]. More recently, this particular filter has been suggested (i.e., rediscovered) in other contexts. In [39], for instance, Coifman et al. suggested it as an ad hoc way to enhance the effectiveness of diffusion. The intuitive justification given [5] was that $(2\mathbf{W} - \mathbf{W}^2) = (2\mathbf{I} - \mathbf{W})\mathbf{W}$ can be considered as a two-step process; specifically blurring, or diffusion (as embodied by \mathbf{W}), followed by an additional step of inverse diffusion or sharpening (as embodied by $2\mathbf{I} - \mathbf{W}$). As can be seen, the ad hoc suggestion has a clear interpretation here. Furthermore, we see that the estimate $\hat{\mathbf{z}}_1$ can also be thought of as a kind of nonlinear, adaptive unsharp masking process, further clarifying its effect. In [61], this procedure and its relationship to the Bregman iterations were studied in detail. For the sake of completeness, we also mention in passing that in the nonstochastic setting and where \mathbf{W} is rank one, the residual-based iterations (42) are known as matching pursuit [62]. This approach has found many applications and extensions for fitting overcomplete dictionaries.

To further understand the iteration (42) in comparison to the earlier diffusion framework, consider again a very similar algebraic manipulation of (42)

$$\hat{\mathbf{z}}_k = \hat{\mathbf{z}}_{k-1} + \mathbf{W}(\mathbf{y} - \hat{\mathbf{z}}_{k-1}), \quad (43)$$

$$= (\mathbf{I} - \mathbf{W})\hat{\mathbf{z}}_{k-1} + \mathbf{W}\mathbf{y}, \quad (44)$$

$$\hat{\mathbf{z}}_k - \hat{\mathbf{z}}_{k-1} = (\mathbf{I} - \mathbf{W})\hat{\mathbf{z}}_{k-1} + (\hat{\mathbf{z}}_0 - \hat{\mathbf{z}}_{k-1}). \quad (45)$$

Redefining the graph Laplacian [note that this definition of the graph Laplacian differs by a negative sign from the earlier one in (26)]

$$\mathbf{I} - \mathbf{W} = \mathbf{D}^{-1/2}(\mathbf{I} - \mathbf{L})\mathbf{D}^{1/2} = \mathbf{D}^{-1/2} \mathcal{L}\mathbf{D}^{1/2},$$

with a change of variable $\bar{\mathbf{z}}_k = \mathbf{D}^{1/2}\hat{\mathbf{z}}_k$, we can write

$$\underbrace{\bar{\mathbf{z}}_k - \bar{\mathbf{z}}_{k-1}}_{\text{diffusion}} = \underbrace{\mathcal{L}\bar{\mathbf{z}}_{k-1}}_{\text{diffusion}} + \underbrace{(\bar{\mathbf{z}}_0 - \bar{\mathbf{z}}_{k-1})}_{\text{reaction term}},$$

$$\downarrow$$

$$\underbrace{\frac{\partial \bar{\mathbf{z}}(t)}{\partial t}}_{\text{diffusion}} = \nabla^2 \bar{\mathbf{z}}(t) + \underbrace{(\bar{\mathbf{z}}(0) - \bar{\mathbf{z}}(t))}_{\text{reaction term}}. \quad (46)$$

The above is a modified version of the diffusion equation studied earlier, where the modification takes the form of a “forcing” function, which is in fact the residual between the estimate at time t and a prefiltered version of the original data. This additional term reacts against the strict smoothing effect of the standard diffusion, and results in a nontrivial steady-state. Not coincidentally, the notion of a “biased” anisotropic diffusion [55] is essentially identical to the above formulation. More recently, Farbman et al. [63] proposed a variational formulation which also leads precisely to the present setup. This will become more clear as we discuss the connection to (empirical) Bayesian approaches in the section “Relationship to (Empirical) Bayes’ and Regularization Approaches.”

The residual-based iterations also have fundamentally different statistical properties than diffusion. To study the statistical

behavior of the estimates, we rewrite the iterative process in (42) more explicitly in terms of the data \mathbf{y} . We have [40] and [53]

$$\hat{\mathbf{z}}_k = \sum_{j=0}^k \mathbf{W}(\mathbf{I} - \mathbf{W})^j \mathbf{y} = (\mathbf{I} - (\mathbf{I} - \mathbf{W})^{k+1})\mathbf{y}. \quad (47)$$

Clearly the above iteration does not monotonically blur the data; but a rather more interesting behavior for the bias-variance tradeoff emerges. Analogous to our earlier derivation, the bias in the estimate after k iterations is

$$\text{bias}_k = \mathbf{E}(\hat{\mathbf{z}}_k) - \mathbf{z} = (\mathbf{I} - (\mathbf{I} - \mathbf{W})^{k+1})\mathbf{E}(\mathbf{y}) - \mathbf{z} = -(\mathbf{I} - \mathbf{W})^{k+1}\mathbf{z}.$$

The squared magnitude of the bias is given by

$$\|\text{bias}_k\|^2 = \|(\mathbf{I} - \mathbf{W})^{k+1}\mathbf{z}\|^2 = \sum_{i=1}^n (1 - \lambda_i)^{2k+2} b_i^2. \quad (48)$$

The behavior of the (decaying) bias in this setting is in stark contrast to the (increasing) bias of the diffusion process.

The variance of the estimator and the corresponding MSE are

$$\begin{aligned} \text{cov}(\hat{\mathbf{z}}_k) &= \text{cov}[(\mathbf{I} - (\mathbf{I} - \mathbf{W})^{k+1})\mathbf{y}] \\ &= \sigma^2 (\mathbf{I} - (\mathbf{I} - \mathbf{W})^{k+1})(\mathbf{I} - (\mathbf{I} - \mathbf{W})^{k+1})^T, \end{aligned}$$

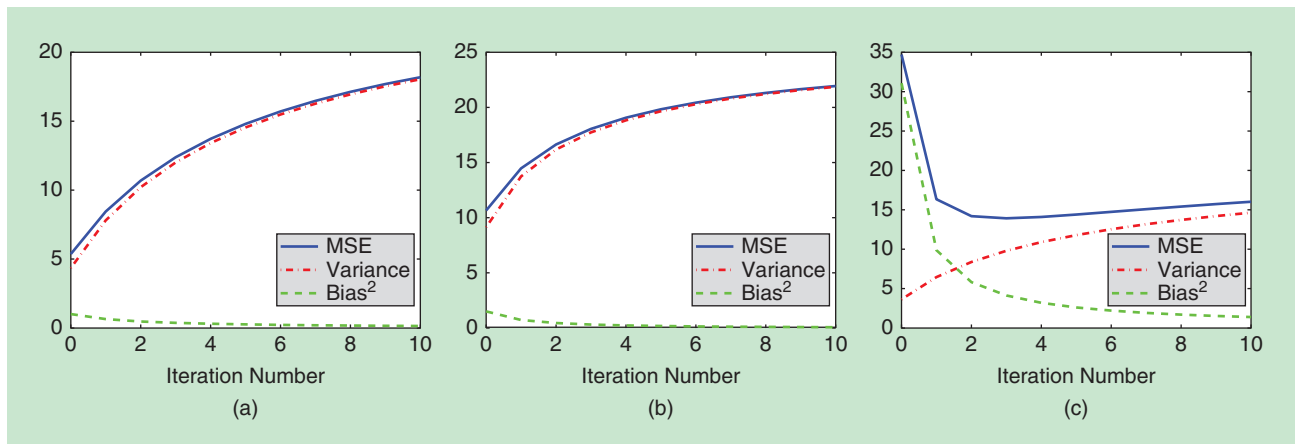
which gives

$$\text{var}(\hat{\mathbf{z}}_k) = \text{tr}(\text{cov}(\hat{\mathbf{z}}_k)) = \sigma^2 \sum_{i=1}^n (1 - (1 - \lambda_i)^{k+1})^2.$$


As k grows, the variance tends to a constant value of σ^2 . Overall, the MSE is

$$\begin{aligned} \text{MSE}_k &= \|\text{bias}_k\|^2 + \text{var}(\hat{\mathbf{z}}_k) \\ &= \sum_{i=1}^n (1 - \lambda_i)^{2k+2} b_i^2 + \sigma^2 (1 - (1 - \lambda_i)^{k+1})^2. \end{aligned} \quad (49)$$

Experiments for denoising the patches given in Figure 8 using the “twicing” approach are carried out, where the same doubly stochastic LARK and NLM filters used for the diffusion tests in Figures 9 and 10 are employed. Plots of predicted MSE_k , $\text{var}(\hat{\mathbf{z}}_k)$ and $\|\text{bias}_k\|^2$ with respect to iteration are illustrated in Figure 13 (LARK) and Figure 14 (NLM). It can be observed that



[FIG13] Parts (a)–(c) show plots of predicted MSE_k , $\text{var}(\hat{\mathbf{z}}_k)$ and $\|\text{bias}_k\|^2$ using (49) for LARK [9] filters in the twicing process.




[FIG14] Parts (a)–(c) show plots of predicted MSE_k , $var(\hat{z}_k)$ and $\|bias_k\|^2$ using (49) for NLM [7] filters in twicing process.

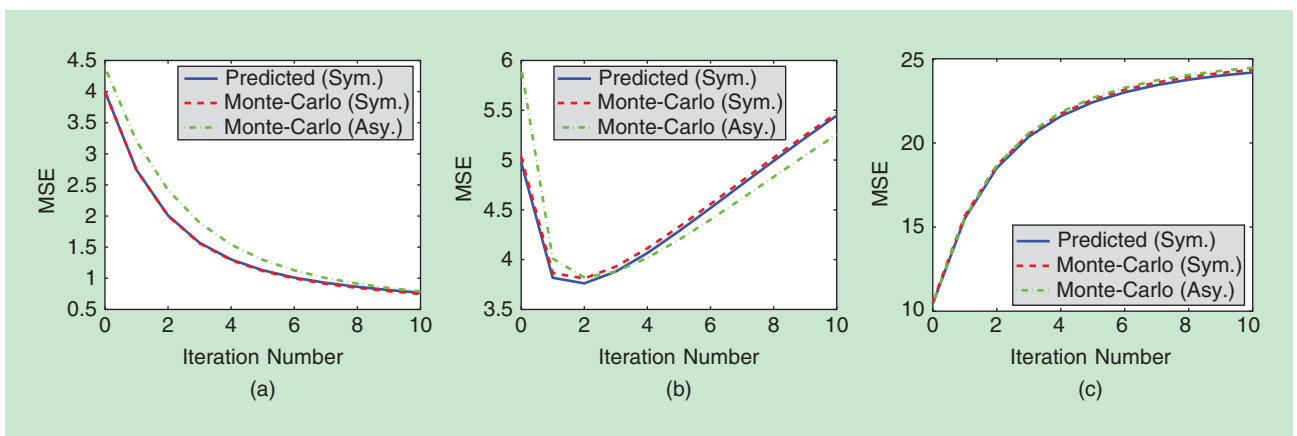
in contrast to diffusion, twicing monotonically reduces the estimation bias and increases the variance. So twicing may in fact reduce the MSE in cases where diffusion fails to do so [such as Figure 9(c)]. True MSEs, for the asymmetric filters and their symmetrized versions, are estimated through Monte-Carlo

simulations, and they all match very well with the predicted ones (see Figures 15 and 16).

Similar to the diffusion analysis before, we plot MSE resulting from the use of W filters directly estimated from noisy patches through Monte-Carlo simulations with noise variance



[FIG15] Parts (a)–(c) show plots of predicted MSE and Monte-Carlo estimated MSE in twicing process using LARK [9] filters. In the Monte-Carlo simulations, row-stochastic (asymmetric) filters and their symmetrized versions are all tested.



[FIG16] Parts (a)–(c) show plots of predicted MSE and Monte-Carlo estimated MSE in twicing process using NLM [7] filters. In the Monte-Carlo simulations, row-stochastic (asymmetric) filters and their symmetrized versions are all tested.

$\sigma^2 = 0.5$, and compare them with the predicted MSEs from the ideally estimated filters. Again, the MSEs estimated from Monte-Carlo simulations are quite close to the predicted data (see Figures 6 and 7).

Here again, the contribution to MSE_k of the i th mode can be written as

$$\text{MSE}_k^{(i)} = (1 - \lambda_i)^{2k+2} \hat{\sigma}_i^2 + \sigma^2 (1 - (1 - \lambda_i)^{k+1})^2. \quad (50)$$

Proceeding in a fashion similar to the diffusion case, we can analyze the derivative of $\overline{\text{MSE}}_k^{(i)}$ with respect to k to see whether the iterations improve the estimate. We have

$$\begin{aligned} \frac{\partial \overline{\text{MSE}}_k^{(i)}}{\partial k} &= 2(1 - \lambda_i)^{k+1} \log(1 - \lambda_i) \\ &\times [(\text{snr}_i + 1)(1 - \lambda_i)^{k+1} - 1]. \end{aligned} \quad (51)$$

Reasoning along parallel lines to the earlier discussion, the following condition guarantees that the k th iteration improves the estimate:

$$(\text{snr}_i + 1)(1 - \lambda_i)^{k+1} > 1. \quad (52)$$

Rewriting the above, we have

$$\log(1 + \text{snr}_i) > \log\left(\frac{1}{1 - \lambda_i'}\right), \quad (53)$$

where now $\lambda_i' = 1 - (1 - \lambda_i)^{k+1}$ is the i th eigenvalue of $\mathbf{I} - (\mathbf{I} - \mathbf{W})^{k+1}$, which is the equivalent filter matrix for the k th twicing iteration. This is the reverse of the condition we derived earlier; specifically, here we observe that if the entropy does not exceed the channel capacity, then iterating with the residuals will indeed produce an improvement. This makes reasonable sense because if the filter removes too much detail from the denoised estimate, this “lost” detail is contained in the residuals, which the iterations will attempt to return to the estimate. Again, the minimum $\text{MSE}_k^{(i)}$ is achieved at the k^* th iteration, where

$$k^* = -\log(1 + \text{snr}_i) / \log(1 - \lambda_i') - 1, \quad (54)$$

which is when the i th eigenvalue of the matrix $\mathbf{I} - (\mathbf{I} - \mathbf{W})^{k^*+1}$ becomes $\text{snr}_i / (\text{snr}_i + 1)$, the Wiener condition. In most cases we cannot optimize all the modes with a uniform number of iterations, but under proper conditions (53) twicing can make the eigenvalues closer to those of the ideal Wiener filter.

From the above analysis, we can see that it is possible to improve the performance of many existing denoising algorithms in the MSE sense by implementing iterative filtering (see Figure 17). However, to choose either diffusion or twicing and to determine the optimal iteration number require prior

SPECIFYING THE KERNEL FUNCTION OR THE CORRESPONDING WEIGHTS IS ESSENTIALLY EQUIVALENT TO ESTIMATING A PARTICULAR TYPE OF EMPIRICAL PRIOR FROM THE GIVEN DATA.

knowledge (or estimation) of the latent signal energy $\{b_i^2\}$ (or the SNR,) and this is an ongoing challenge. Research on estimating the SNR of image or video data without a reference “ground truth” is of broad interest in the community as of late [64]. In particular,

this problem points to potentially important connections with ongoing work in no-reference image quality assessment [65] as well.


RELATIONSHIP TO (EMPIRICAL) BAYES' AND REGULARIZATION APPROACHES

It is interesting to contrast the adaptive nonparametric framework described so far with the more classical Bayesian estimation approaches. Of course, Bayesian thinking requires that a “prior” be specified independently of the measured data, which characterizes the aggregate behavior of the underlying latent image of interest. In the context we have described so far, no specific prior information was used. Instead, data-dependant weights were specified in a least-squares framework, yielding locally adaptive filters. On the surface then, it may seem that there is no direct relationship between the two approaches. However, as we will illustrate below, specifying the kernel function or the corresponding weights is essentially equivalent to estimating a particular type of “prior” from the given data. (One cringes to call this estimate a “prior” any longer, but we will do so for lack of a better word.) This empirical Bayesian approach has in fact a long and useful history [66] in the statistics literature. Furthermore, another class of modern patch-based filtering methods based on the empirical Bayesian framework were recently pioneered in [67] and extended in [68] and [69] based on the notion of “adaptive” (i.e., data-dependent) priors. In the same spirit, we illustrate below that an empirical version of the popular maximum a posteriori method has a data-dependent counterpart in the nonparametric techniques we described so far.

A canonical Bayesian approach is to specify a prior for the unknown \mathbf{z} , or to equivalently use a regularization functional. The so-called maximum a posteriori (MAP) estimate is then given by the solution of the following optimization problem:

$$\text{MAP: } \hat{\mathbf{z}} = \arg \min_{\mathbf{z}} \frac{1}{2} \|\mathbf{y} - \mathbf{z}\|^2 + \frac{\lambda}{2} \mathcal{R}(\mathbf{z}), \quad (55)$$

where the first term on the right-hand side is the data-fidelity (log-likelihood) term, and the second (regularization, or log-prior) term essentially enforces a soft constraint on the global structure and smoothness of the signal. In the MAP approach, the regularization term is a functional $\mathcal{R}(\mathbf{z})$, which is typically (but not always) convex, to yield a unique minimum for the overall cost. Particularly popular recent examples include $\mathcal{R}(\mathbf{z}) = \|\nabla \mathbf{z}\|$, and $\mathcal{R}(\mathbf{z}) = \|\mathbf{z}\|_1$. The above approach implicitly contains a global (Gaussian) model of the noise (captured by the



[FIG17] Denoising example using iterative LARK filtering: Parts (a)–(c) show input noisy patches; noise variance $\sigma^2 = 25$. (d)–(f) show output patches by filtering once. Parts (g)–(i) show MSE optimized outputs: part (g) is the sixth diffusion output; (h) is the fourth diffusion output; and (i) is the third twicing output.

quadratic data-fidelity term) and an explicit model of the signal (captured by the regularization term.)

Regardless of what choice we make for the regularization functional, this approach is global in the sense that the explicit assumptions about the noise and the signal constrain the degrees of freedom in the solution, hence limiting the global behavior of the estimate. This often results in well-behaved (but not always desirable) solutions. This is precisely the main

motivation for using (local/nonlocal) adaptive nonparametric techniques in place of global regularization methods. Indeed, though the effect of regularization is not explicitly present in the nonparametric framework, its work is implicitly done by the design of the kernel function $K(\cdot)$, which affords us local control, and therefore more freedom and often better adaptation to the given data, resulting in more powerful techniques with broader applicability.

Another important advantage of the local nonparametric methods is that the resulting algorithms inherently operate in relatively small windows, and can hence be efficiently implemented and often easily parallelized. With global parametric and variational methods, the solutions are often implemented using large-scale global optimization techniques, which can be hard (sometimes even impossible) to “kernelize,” for algorithmic efficiency [33].

Though we do not discuss them here, of course hybrid methods are also possible, where the first term in (55) is replaced by the corresponding weighted term (3), with regularization also applied explicitly [70]. As we shall see below, the iterative methods for improving the nonparametric estimates (such as diffusion, twicing, etc.) mirror the behavior of SD iterations in the empirical Bayesian case, but with a fundamental difference. Specifically, the corresponding implied regularization functionals deployed by the nonparametric methods are significantly different, adapting directly to the given data.

To connect MAP to the nonparametric setting, we proceed in similar lines as Elad [3] by considering the simplest iterative approach: the SD method. The SD iteration for MAP, with fixed step size μ is

$$\text{MAP: } \hat{z}_{k+1} = \hat{z}_k - \mu[(\hat{z}_k - y) + \lambda \nabla \mathcal{R}(z_k)]. \quad (56)$$

EMPIRICAL BAYES' INTERPRETATION OF DIFFUSION

Let us compare the diffusion iterations (25) to the SD MAP iterations above in (56). Specifically, we equate the right-hand sides of

$$\hat{z}_{k+1} = \hat{z}_k - \mu[(\hat{z}_k - y) + \lambda \nabla \mathcal{R}(\hat{z}_k)] \quad (57)$$

$$\hat{z}_{k+1} = \hat{z}_k + (W - I)\hat{z}_k, \quad (58)$$

which gives

$$-\mu[(\hat{z}_k - y) + \lambda \nabla \mathcal{R}(\hat{z}_k)] = (W - I)\hat{z}_k.$$

Solving for $\nabla \mathcal{R}(\hat{z}_k)$ we obtain

$$\nabla \mathcal{R}(\hat{z}_k) = \frac{1}{\mu\lambda}(W - (1 - \mu)I)(y - \hat{z}_k) - \frac{1}{\mu\lambda}(I - W)y. \quad (59)$$

When W is symmetric, integrating both sides we have (to within an additive constant)

$$\mathcal{R}_{\text{diff.}}(z) = \frac{1}{2\mu\lambda}(y - z)^T((1 - \mu)I - W)(y - z) + \frac{1}{\mu\lambda}y^T(I - W)z. \quad (60)$$

(Since the left-hand side of (59) is the gradient of a scalar-valued function, the right-hand side must be curl free. This is possible

if and only if W is symmetric.) We observe therefore that the equivalent regularization term is a quadratic function of the data and residuals. In particular, the function $\mathcal{R}_{\text{diffusion}}(\hat{z}_k)$ can be thought of as an estimated log-prior at each iteration as it depends on the data y , the last estimate \hat{z}_k , and weights W , which are computed from the data. More specifically, the implied empirical Bayes' “prior” would be

$$\hat{p}(z) = c \exp[-\mathcal{R}(z)],$$

where c is a normalization constant. The empirical (MAP) Bayesian interpretation of the residual-based method follows similarly, but is somewhat surprisingly more straightforward, as we illustrate below.

ANOTHER IMPORTANT ADVANTAGE OF THE LOCAL NONPARAMETRIC METHODS IS THAT THE RESULTING ALGORITHMS INHERENTLY OPERATE IN RELATIVELY SMALL WINDOWS.

EMPIRICAL BAYES' INTERPRETATION OF THE RESIDUAL ITERATIONS

We arrive at the empirical Bayes' interpretation of the residual iterations (42) in a similar manner. In particular, comparing (42) to the SD MAP iterations in (56) we have

$$\hat{z}_{k+1} = \hat{z}_k - \mu[(\hat{z}_k - y) + \lambda \nabla \mathcal{R}(\hat{z}_k)] \quad (61)$$

$$\hat{z}_{k+1} = \hat{z}_k + W(y - \hat{z}_k). \quad (62)$$

Again, equating the right-hand sides, we get

$$-\mu[(\hat{z}_k - y) + \lambda \nabla \mathcal{R}(z_k)] = W(y - \hat{z}_k).$$

Solving for $\nabla \mathcal{R}(\hat{z}_k)$ we obtain

$$\nabla \mathcal{R}(\hat{z}_k) = \frac{1}{\mu\lambda}(W - \mu I)(y - \hat{z}_k). \quad (63)$$

With W symmetric, we have (to within an additive constant)

$$\mathcal{R}_{\text{residual}}(z) = \frac{1}{2\mu\lambda}(y - z)^T(\mu I - W)(y - z). \quad (64)$$

Therefore, the (implicit) regularization term is an adaptive function of the data and residuals. Once again, the empirically estimated prior is of the same form $\hat{p}(z) = c \exp[-\mathcal{R}(z)]$, where $\mathcal{R}(z)$ is computed directly on the residuals $y - \hat{z}$ of the filtering process at every iteration. It is especially interesting to note that this empirical prior is a quadratic form in the residuals alone where, not coincidentally, the matrix $\mu I - W$ is closely related to the Laplacian operator defined earlier. This hints directly at the edge-preserving behavior of this type of iteration; a property not shared by the diffusion process, which monotonically blurs the data.

CONCLUSIONS

It has been said that in both literature and the arts, there are only seven basic plots (comedy, tragedy, etc.)—that all stories

are combinations or variations on these basic themes. Perhaps it is taking the analogy too far to say that an exact parallel exists in our field as well. But it is fair to argue that the basic tools of our trade in recent years have revolved around a small number of key concepts as well, which I tried to highlight in this article.

DEEP CONNECTIONS BETWEEN TECHNIQUES USED COMMONLY IN COMPUTER VISION, IMAGE PROCESSING, GRAPHICS, AND MACHINE LEARNING EXIST.

AUTHOR

Peyman Milanfar (milanfar@ucsc.edu) received the B.S. degree in electrical engineering and mathematics from the University of California, Berkeley, and the M.S. and Ph.D. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology, respectively. Until 1999, he was a senior research engineer at SRI International, Menlo Park, California. He is a professor of electrical engineering and was an associate dean for research at the Baskin School of Engineering at the University of California, Santa Cruz. He is currently a visiting scientist at Google[x]. His research interests include statistical signal, image and video processing, computational photography, and computer vision. He won the 2000 U.S. National Science Foundation CAREER Award and the 2010 IEEE Signal Processing Society (SPS) Best Paper Award. He is a member of the SPS's Awards Board as well as the Image, Video, and Multidimensional Signal Processing Technical Committee. He is an IEEE Fellow.

- The most successful modern approaches in image and video processing are nonparametric. We have drifted away from model-based methods, which have dominated signal processing for decades.

- In one-dimensional signal processing, there is a long history of design and analysis of adaptive filtering algorithms. A corresponding line of thought has only recently become the dominant paradigm in processing higher-dimensional data. Indeed, adaptivity to data is a central theme of all the algorithms and techniques discussed here, which represent a snapshot of the state of the art.

- The traditional approach that every graduate student of our field learns is to carefully design a filter, apply it to the given data, and call it a day. In contrast, many of the most successful recent approaches involve repeated applications of a filter or sequence of filters to data, and aggregating the results. Such ideas have been around for some time in statistics, machine learning, and elsewhere, but we have just begun to make careful use of sophisticated iteration and boosting mechanisms.

As I highlighted here, deep connections between techniques used commonly in computer vision, image processing, graphics, and machine learning exist. The practitioners of these fields have been using each other's techniques either implicitly or explicitly for a while. The pace of this convergence has quickened, and this is not a coincidence. It has come about through many years of scientific iteration in what my late friend and colleague Gene Golub called the "serendipity of science," an aptitude we have all developed for making desirable discoveries by happy accident.

ACKNOWLEDGMENTS

I thank the editorial board of *IEEE Signal Processing Magazine*, and Prof. Antonio Ortega in particular, for their support and valuable feedback. I am grateful to all my students in the Multidimensional Signal Processing research group at the University of California, Santa Cruz for their insights and comments on earlier versions of this manuscript. In particular, I acknowledge the diligent help of Xiang Zhu for many of the simulations contained here. I am also thankful for feedback and encouragement of my friends and colleagues Prof. Michael Elad of the Technion and Prof. Bernd Girod of Stanford. This work was supported in part by the U.S. Air Force grant FA9550-07-1-0365 and the National Science Foundation grant CCF-1016018.

neering and computer science from the Massachusetts Institute of Technology, respectively. Until 1999, he was a senior research engineer at SRI International, Menlo Park, California. He is a professor of electrical engineering and was an associate dean for research at the Baskin School of Engineering at the University of California, Santa Cruz. He is currently a visiting scientist at Google[x]. His research interests include statistical signal, image and video processing, computational photography, and computer vision. He won the 2000 U.S. National Science Foundation CAREER Award and the 2010 IEEE Signal Processing Society (SPS) Best Paper Award. He is a member of the SPS's Awards Board as well as the Image, Video, and Multidimensional Signal Processing Technical Committee. He is an IEEE Fellow.

REFERENCES

- [1] D. Barash, "A fundamental relationship between bilateral filtering, adaptive smoothing and the nonlinear diffusion equation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, no. 6, pp. 844–847, June 2002.
- [2] D. Barash and D. Comaniciu, "A common framework for nonlinear diffusion, adaptive smoothing, bilateral filtering and mean shift," *Image Vis. Comput.*, vol. 22, no. 1, pp. 73–81, 2004.
- [3] M. Elad, "On the origin of the bilateral filter and ways to improve it," *IEEE Trans. Image Process.*, vol. 11, no. 10, pp. 1141–1150, Oct. 2002.
- [4] O. Scherzer and J. Weickert, "Relations between regularization and diffusion filtering," *J. Math. Imaging Vision*, vol. 12, no. 1, pp. 43–63, Feb. 2000.
- [5] A. Singer, Y. Shkolinsky, and B. Nadler, "Diffusion interpretation of nonlocal neighborhood filters for signal denoising," *SIAM J. Imaging Sci.*, vol. 2, no. 1, pp. 118–139, 2009.
- [6] L. Pizarro, P. Mrazek, S. Didas, S. Grewenig, and J. Weickert, "Generalised nonlocal image smoothing," *Int. J. Comput. Vision*, vol. 90, no. 1, pp. 62–87, 2010.
- [7] A. Buades, B. Coll, and J. M. Morel, "A review of image denoising algorithms, with a new one," *Multiscale Modeling Simulation (SIAM Interdisciplinary J.)*, vol. 4, no. 2, pp. 490–530, 2005.
- [8] C. Tomasi and R. Manduchi, "Bilateral filtering for gray and color images," in *Proc. IEEE Int. Conf. Comput. Vision*, Bombay, India, pp. 836–846, Jan. 1998.
- [9] H. Takeda, S. Farsiu, and P. Milanfar, "Kernel regression for image processing and reconstruction," *IEEE Trans. Image Process.*, vol. 16, no. 2, pp. 349–366, Feb. 2007.
- [10] K. Dabov, A. Foi, V. Katkovnic, and K. Egiazarian, "Image denoising by sparse 3D transform-domain collaborative filtering," *IEEE Trans. Image Process.*, vol. 16, no. 8, pp. 2080–2095, Aug. 2007.
- [11] S. M. Smith and J. M. Brady, "Susan: A new approach to low level image processing," *Int. J. Comput. Vision*, vol. 23, no. 1, pp. 45–78, 1997.
- [12] H. Knutsson and C. F. Westin, "Normalized and differential convolution: Methods for interpolation and filtering of incomplete and uncertain data," in *Proc. IEEE Comput. Society Conf. Comput. Vision Pattern Recognit. (CVPR)*, June 1993, pp. 515–523.
- [13] L. Yaroslavsky, *Digital Picture Processing*. Berlin, Germany: Springer-Verlag, 1987.
- [14] K. Fukunaga and L. D. Hostetler, "The estimation of the gradient of a density function, with applications in pattern recognition," *IEEE Trans. Inform. Theory*, vol. 21, no. 1, pp. 32–40, 1975.
- [15] D. Comaniciu and P. Meer, "Mean shift: A robust approach toward feature space analysis," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, no. 5, pp. 603–619, May 2002.
- [16] A. Smola and B. Scholkopf, "A tutorial on support vector regression," *Statistics Comput.*, vol. 14, no. 3, pp. 199–222, 2004.
- [17] A. Smola, B. Scholkopf, and K. Muller, "The connection between regularization operators and support vector kernels," *Neural Netw.*, vol. 11, no. 4, pp. 637–649, 1998.

- [18] R. Rosipal and L. J. Trejo, "Kernel partial least squares regression in reproducing kernel Hilbert space," *J. Mach. Learning Res.*, vol. 2, pp. 97–123, Dec. 2001.
- [19] F. Girosi, "An equivalence between sparse approximation and support vector machines," *Neural Comput.*, vol. 10, no. 6, pp. 1455–1480, 1998.
- [20] M. P. Wand and M. C. Jones, *Kernel Smoothing* (Monographs on Statistics and Applied Probability). London: Chapman & Hall, 1995.
- [21] W. Hardle, *Applied Nonparametric Regression*. Cambridge, U.K.: Cambridge Univ. Press, 1990.
- [22] P. Lancaster and K. Salkauskas, "Surfaces generated by moving least squares methods," *Math. Comput.*, vol. 87, no. 155, pp. 141–158, 1981.
- [23] D. Levin, "The approximation power of moving least-squares," *Math. Comput.*, vol. 67, no. 224, pp. 1517–1531, 1998.
- [24] M. Alexa, J. Behr, D. Cohen-or, S. Fleishman, and D. Levin, "Computing and rendering point set surfaces," *IEEE Trans. Visualization Comput. Graph.*, vol. 9, no. 1, pp. 3–15, 2003.
- [25] R. Gonzalez and J. Woods, *Digital Image Processing*, 3rd ed. Englewood Cliffs, NJ: Prentice-Hall, 2008.
- [26] C. Kervrann and J. Boulanger, "Optimal spatial adaptation for patch-based image denoising," *IEEE Trans. Image Process.*, vol. 15, no. 10, pp. 2866–2878, Oct. 2006.
- [27] S. P. Awate and R. T. Whitaker, "Unsupervised, information-theoretic, adaptive image filtering for image restoration," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 3, pp. 364–376, Mar. 2006.
- [28] A. Buades, B. Coll, and J. Morel, "A non-local algorithm for image denoising," in *Proc. CVPR*, June 2005, pp. 60–65.
- [29] S. Awate and R. Whitaker, "Higher-order image statistics for unsupervised, information-theoretic, adaptive, image filtering," in *Proc. CVPR*, June 2005, pp. 44–51.
- [30] H. Takeda, S. Farsiu, and P. Milanfar, "Kernel regression for image processing and reconstruction," *IEEE Trans. Image Process.*, vol. 16, no. 2, pp. 349–366, Feb. 2007.
- [31] H. Knutsson and C.-F. Westin, "Normalized and differential convolution: Methods for interpolation and filtering of incomplete and uncertain data," in *Proc. IEEE Comput. Soc. Conf. Comput. Vision Pattern Recognit.*, June 16–19 1993, pp. 515–523.
- [32] N. Sochen, R. Kimmel, and A. Bruckstein, "Diffusions and confusions in signal and image processing," *J. Math. Imaging Vision*, vol. 14, no. 3, pp. 195–209, 2001.
- [33] A. Spira, R. Kimmel, and N. Sochen, "A short time Beltrami kernel for smoothing images and manifolds," *IEEE Trans. Image Process.*, vol. 16, no. 6, pp. 1628–1636, 2007.
- [34] J. Ham, D. Lee, S. Mika, and B. Schölkopf, "A kernel view of the dimensionality reduction of manifolds," in *Proc. 21st Int. Conf. Mach. Learning*, 2004.
- [35] T. Hofmann, B. Schölkopf, and A. J. Smola, "Kernel methods in machine learning," *The Ann. Statistics*, vol. 36, no. 3, pp. 1171–1220, 2008.
- [36] N. Dowson and O. Salvado, "Hashed nonlocal means for rapid image filtering," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 33, no. 3, pp. 485–499, 2011.
- [37] J. Baek and D. Jacobs, "Accelerating spatially varying Gaussian filters," *ACM Trans. Graph.*, vol. 29, no. 6, Dec. 2010.
- [38] M. Belkin and P. Niyogi, "Towards a theoretical foundation for Laplacian-based manifold methods," in *Proc. 18th Conf. Learning Theory (COLT)*, 2005, pp. 486–500.
- [39] A. Szlám, M. Maggioni, and R. Coifman, "Regularization on graphs with function-adapted processes," *J. Mach. Learning Res.*, vol. 9, pp. 1711–1739, Aug. 2008.
- [40] E. Seneta, *Non-Negative Matrices and Markov Chains* (Springer Series in Statistics). Berlin, Germany: Springer, 1981.
- [41] R. A. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge, U.K.: Cambridge Univ. Press, 1991.
- [42] H. J. Landau and A. M. Odlyzko, "Bounds for eigenvalues of certain stochastic matrices," *Linear Algebra Its Applicat.*, vol. 38, pp. 5–15, June 1981.
- [43] W. J. Stewart, *Introduction to the Numerical Solution of Markov Chains*. Princeton, NJ: Princeton Univ. Press, 1994.
- [44] P. Knight, "The Sinkhorn-Knopp algorithm: Convergence and applications," *SIAM J. Matrix Anal. Appl.*, vol. 30, no. 1, pp. 261–275, 2008.
- [45] R. Sinkhorn, "A relationship between arbitrary positive matrices and doubly stochastic matrices," *Annals Math. Statist.*, vol. 35, no. 2, pp. 876–879, 1964.
- [46] J. Darroch and D. Ratcliff, "Generalized iterative scaling for log-linear models," *Ann. Math. Statist.*, vol. 43, no. 5, pp. 1470–1480, 1972.
- [47] L. Huang, D. Yan, M. I. Jordan, and N. Taft, "Spectral clustering with perturbed data," in *Proc. Neural Information Process. Syst. (NIPS)*, 2008.
- [48] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3-D transform-domain collaborative filtering," *IEEE Trans. Image Process.*, vol. 16, no. 8, pp. 2080–2095, Aug. 2007.
- [49] P. Chatterjee and P. Milanfar, "Is denoising dead?" *IEEE Trans. Image Process.*, vol. 19, no. 4, pp. 895–911, Apr. 2010.
- [50] P. Chatterjee and P. Milanfar, "Patch-based near-optimal denoising," *IEEE Trans. Image Process.*, vol. 21, no. 4, pp. 1635–1649, Apr. 2011.
- [51] P. Perona and J. Malik, "Scale-space and edge detection using anisotropic diffusion," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 12, no. 9, pp. 629–639, July 1990.
- [52] S. Osher, M. Burger, D. Goldfarb, J. Xu, and W. Yin, "An iterative regularization method for total variation-based image restoration," *Multiscale Modeling Simulation*, vol. 4, no. 2, pp. 460–489, 2005.
- [53] P. Buhlmann and B. Yu, "Boosting with the l_2 loss: Regression and classification," *J. Amer. Statist. Assoc.*, vol. 98, no. 462, pp. 324–339, 2003.
- [54] J. W. Tukey, *Exploratory Data Analysis*. Reading, MA: Addison Wesley, 1977.
- [55] N. Nordstrom, "Biased anisotropic diffusion: A unified regularization and diffusion approach to edge detection," *Image Vision Comput.*, vol. 8, no. 4, pp. 318–327, 1990.
- [56] D. P. Palomar and S. Verdu, "Gradient of mutual information in linear vector Gaussian channels," *IEEE Trans. Inform. Theory*, vol. 52, no. 1, pp. 141–154, 2006.
- [57] Y. Freund and R. E. Schapire, "A short introduction to boosting," *J. Japan. Soc. Artif. Intell.*, vol. 14, no. 5, pp. 771–780, Sept. 1999.
- [58] P. Buhlmann and T. Hothorn, "Boosting algorithms: regularization, prediction and model fitting," *Statist. Sci.*, vol. 22, no. 4, pp. 477–505, 2007.
- [59] Y. Censor and S. Zenios, *Parallel Optimization* (Numerical Mathematics and Scientific Computation). New York: Oxford Univ. Press, 1997.
- [60] M. Collins, R. Schapire, and Y. Singer, "Logistic regression, Adaboost and Bregman distances," *Mach. Learning*, vol. 48, no. 1–3, 2002.
- [61] M. Charest and P. Milanfar, "On iterative regularization and its application," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 18, no. 3, pp. 406–411, Mar. 2008.
- [62] S. G. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Trans. Signal Process.*, vol. 41, no. 12, pp. 3397–3415, Dec. 1993.
- [63] Z. Farbman, R. Fattal, D. Lischinski, and R. Szeliski, "Edge-preserving decompositions for multi-scale tone and detail manipulation," in *Proc. ACM SIGGRAPH ACM Trans. Graph.*, vol. 27, no. 3, Aug. 2008.
- [64] F. Luisier, T. Blu, and M. Unser, "A new SURE approach to image denoising: Interscale orthonormal wavelet thresholding," *IEEE Trans. Image Process.*, vol. 16, no. 3, pp. 593–606, Mar. 2007.
- [65] Z. Wang and A. Bovik, "Mean squared error: Love it or leave it? A new look at signal fidelity measures," *IEEE Signal Process. Mag.*, vol. 26, no. 1, pp. 98–117, Jan. 2009.
- [66] B. Carlin and T. Louis, *Bayes and Empirical Bayes Methods for Data Analysis* (Texts in Statistical Science). London: Chapman & Hall, 2000.
- [67] T. Weissman, E. Ordentlich, G. Seroussi, S. Verdu, and M. Weinberger, "Universal discrete denoising: Known channel," *IEEE Trans. Info. Theory*, vol. 51, no. 1, pp. 5–28, 2005.
- [68] S. Awate and R. Whitaker, "Feature-preserving MRI denoising: A nonparametric empirical Bayes approach," *IEEE Trans. Med. Imaging*, vol. 26, no. 9, pp. 1242–1255, 2007.
- [69] K. Sivaramakrishnan and T. Weissman, "A context quantization approach to universal denoising," *IEEE Trans. Signal Process.*, vol. 57, no. 6, pp. 2110–2129, 2009.
- [70] A. Bruckstein, "On globally optimal local modeling: From moving least square to overparametrization," in *Proc. SIAM Conf. Imaging Sci.*, Chicago, IL, 2000.
- [71] L. Wasserman, *All of Nonparametric Statistics*. Berlin, Germany: Springer, 2005.
- [72] M. Aharon, M. Elad, and A. Bruckstein, "The K-SVD: An algorithm for designing of overcomplete dictionaries for sparse representation," *IEEE Trans. Signal Process.*, vol. 54, no. 11, pp. 4311–4322, Nov. 2006.
- [73] R. Khoury, "Closest matrices in the space of generalized doubly stochastic matrices," *J. Math. Anal. Applicat.*, vol. 222, no. AY985970, pp. 562–568, 1998.
- [74] W. Kahan, "Spectra of nearly Hermitian matrices," in *Proc. Amer. Math. Soc.*, vol. 48, pp. 11–17, 1975.
- [75] P. Milanfar, "Symmetrizing smoothing filters," *SIAM J. Imaging Sci.*, to be published. DOI: 10.1137/120875843
- [76] H. Jegou, C. Schmid, H. Harzallah, and J. Verbeek, "Accurate image search using the contextual dissimilarity measure," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 32, no. 1, pp. 2–11, Jan. 2010.
- [77] C. Johnson, R. Masson, and M. Trosset, "On the diagonal scaling of Euclidean distance matrices to doubly stochastic matrices," *Linear Algebra Its Applicat.*, vol. 319, no. 1, pp. 253–264, 2005.
- [78] D. Yan, L. Huang, and M. I. Jordan, "Fast approximate spectral clustering," in *Proc. 15th ACM Conf. Knowledge Discovery Data Mining (SIGKDD)*, 2009, pp. 907–916.