

# Fast High-Dimensional Bilateral and Nonlocal Means Filtering

Pravin Nair<sup>ID</sup>, *Student Member, IEEE*, and Kunal N. Chaudhury<sup>ID</sup>, *Senior Member, IEEE*

**Abstract**—Existing fast algorithms for bilateral and nonlocal means filtering mostly work with grayscale images. They cannot easily be extended to high-dimensional data such as color and hyperspectral images, patch-based data, and flow-fields. In this paper, we propose a fast algorithm for high-dimensional bilateral and nonlocal means filtering. Unlike existing approaches, where the focus is on approximating the data (using quantization) or the filter kernel (via analytic expansions), we locally approximate the kernel using weighted and shifted copies of a Gaussian, where the weights and shifts are inferred from the data. The algorithm emerging from the proposed approximation essentially involves clustering and fast convolutions, and is easy to implement. Moreover, a variant of our algorithm comes with a guarantee (bound) on the approximation error, which is not enjoyed by existing algorithms. We present some results for high-dimensional bilateral and nonlocal means filtering to demonstrate the speed and accuracy of our proposal. Moreover, we also show that our algorithm can outperform the state-of-the-art fast approximations in terms of accuracy and timing.

**Index Terms**—High-dimensional filter, bilateral filter, nonlocal means, shiftability, kernel, approximation, fast algorithm.

## I. INTRODUCTION

**S**MOOTHING images while preserving structures (edges, corners, lines, etc.) is a fundamental task in image processing. A classic example in this regard is the diffusion framework of Perona and Malik [1]. In the last few decades, several filtering based approaches have been proposed for this task. Prominent examples include bilateral filtering [2]–[4], mean shift filtering [5], weighted least squares [6], domain transform [7], guided filtering [8], and nonlocal means [9]. The brute-force implementation of most of these filters is computationally prohibitive and cannot be used for real-time applications. To address this problem, researchers have come up with approximation algorithms that can significantly accelerate the filtering without compromising the quality. Unfortunately, for bilateral and nonlocal means filtering, most of these algorithms can be used only for grayscale images. It is difficult to use them even for color filtering, while preserving

their efficiency. The situation is more challenging for multispectral and hyperspectral images, where the range dimension is much larger. Of course, any algorithm for grayscale filtering can be used for channel-by-channel processing of high-dimensional images. However, by working in the combined intensity space, we can exploit the strong correlation between channels [4].

### A. High-Dimensional Filtering

The focus of the present work is on two popular smoothers, the bilateral and the nonlocal means filters, and their application to high-dimensional data. The former is used for edge-preserving smoothing in a variety of applications [10], while the latter is primarily used for denoising. Though nonlocal means has limited denoising capability compared to state-of-the-art denoisers [11], [12], it continues to be of interest due to its simplicity and the availability of low-complexity algorithms [13]–[15]. The connection between these filters is that they can be interpreted as a multidimensional Gaussian filter operating in the joint spatio-range space [16]. The term *high-dimensional filtering* is used when the dimension of the spatio-range space is large [14], [17], [18]. In this paper, we will use this term when the range dimension is greater than one.

An unified formulation of high-dimensional bilateral and nonlocal means filtering is as follows. Suppose that the input data is  $f : \Omega \rightarrow [0, R]^n$ , where  $\Omega \subset \mathbb{Z}^d$  is the spatial domain,  $[0, R]^n$  is the range space, and  $d$  and  $n$  are dimensions of the domain and range of  $f$ . Let  $p : \Omega \rightarrow \mathbb{R}^p$  be the *guide* image, which is used to control the filtering [10]. The output  $g : \Omega \rightarrow [0, R]^n$  is given by

$$g(i) = \frac{1}{\eta(i)} \sum_{j \in W} \omega(j) \varphi(p(i-j) - p(i)) f(i-j), \quad (1)$$

where

$$\eta(i) = \sum_{j \in W} \omega(j) \varphi(p(i-j) - p(i)). \quad (2)$$

The aggregations in (1) and (2) are performed over a window around the pixel of interest, i.e.,  $W = [-S, S]^d$ , where  $S$  is the window length. We call  $\omega : \mathbb{Z}^d \rightarrow \mathbb{R}$  the *spatial* kernel and  $\varphi : \mathbb{R}^p \rightarrow \mathbb{R}$  the *range* kernel [4], where we denote the dimension of the range space of  $p$  as  $p$ . The spatial kernel is used to measure proximity in the spatial domain  $\Omega$  (as in classical linear filtering). On the other hand, the range kernel is used to measure proximity in the range space of  $p$ .

Manuscript received March 5, 2018; revised August 24, 2018 and October 3, 2018; accepted October 27, 2018. Date of publication October 31, 2018; date of current version November 21, 2018. The work of K. N. Chaudhury was supported by a Startup Grant from the Indian Institute of Science and EMR under Grant SERB/F/6047/2016-2017 from the Department of Science and Technology, Government of India. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Xudong Jiang. (*Corresponding author: Kunal N. Chaudhury.*)

The authors are with the Department of Electrical Engineering, Indian Institute of Science, Bengaluru 560012, India (e-mail: kunal@iisc.ac.in).

Digital Object Identifier 10.1109/TIP.2018.2878955

1057-7149 © 2018 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.

See [http://www.ieee.org/publications\\_standards/publications/rights/index.html](http://www.ieee.org/publications_standards/publications/rights/index.html) for more information.

For (joint) bilateral filtering,  $\omega$  and  $\varphi$  are generally Gaussian. While  $\mathbf{f}$  and  $\mathbf{p}$  are identical for the bilateral filter ( $\rho = n$ ), they are different for the joint bilateral filter [10]. In the original proposal for nonlocal means [9],  $\mathbf{p}(\mathbf{i})$  are the spatial neighbors of pixel  $\mathbf{i}$  (extracted from a patch around  $\mathbf{i}$ ) and hence  $\rho$  is  $n$  times the size of image patch,  $\varphi$  is a multidimensional Gaussian, and  $\omega$  is a box filter. Later, it was shown in [21] that by reducing the dimension of each patch (e.g., by applying PCA to the collection of patches), we can improve the speed and denoising performance. Similar to [14], [17], and [18], we have also considered PCA-based nonlocal means in this work. Needless to say, the proposed algorithm can also work with full patches.

The brute-force computation of (1) and (2) clearly requires  $\mathcal{O}(S^d(n+\rho))$  operations per pixel. In particular, the complexity scales exponentially with the window length  $S$ , which can be large for some smoothing applications. In the last decade, several fast algorithms have been proposed for bilateral filtering of grayscale images [20], [22]–[26]. The complexity can be cut down from  $\mathcal{O}(S^d)$  to  $\mathcal{O}(1)$  using these algorithms. Similarly, fast algorithms have been proposed for nonlocal means [13] and [15] that can reduce the complexity from  $\mathcal{O}(S^d\rho)$  to  $\mathcal{O}(S^d)$ .

## B. Previous Work

1) *Quantization Methods*: Durand *et al.* [23] proposed a novel framework for approximating (1) and (2) using clustering and interpolation. In terms of our notations, their approximation of (1) and (2) can be expressed as

$$\frac{1}{\hat{\eta}(\mathbf{i})} \sum_{k=1}^K c_k(\mathbf{i}) \left( \sum_{j \in W} \omega(\mathbf{j}) \varphi(\mathbf{p}(\mathbf{i} - \mathbf{j}) - \boldsymbol{\mu}_k) \mathbf{f}(\mathbf{i} - \mathbf{j}) \right), \quad (3)$$

and

$$\hat{\eta}(\mathbf{i}) = \sum_{k=1}^K c_k(\mathbf{i}) \left( \sum_{j \in W} \omega(\mathbf{j}) \varphi(\mathbf{p}(\mathbf{i} - \mathbf{j}) - \boldsymbol{\mu}_k) \right). \quad (4)$$

The centers  $\boldsymbol{\mu}_k$  are obtained by uniformly sampling the range space of  $\mathbf{p}$ , whereas the interpolation coefficient  $c_k(\mathbf{i})$  is determined from the distance between  $\mathbf{p}(\mathbf{i})$  and  $\boldsymbol{\mu}_k$ . Notice that the inner summations in (3) and (4) can be expressed as convolutions, which are computed using FFT in [23]. Furthermore, the entire processing is performed on the subsampled image, following which the output image is upsampled. The approximation turns out to be effective for grayscale images. This is because the range space is one-dimensional in this case, and hence a good approximation can be achieved using a small the number of samples  $K$ . Moreover, it suffices to interpolate just the two neighboring pixels. Paris and Durand [20] showed that the accuracy of [23] can be improved by downsampling the intermediate images (instead of the input image) involved in the convolutions. Chen *et al.* [16] proposed to accelerate [20] by performing convolutions in the higher-dimensional spatio-range space. Yang *et al.* [26] observed that this framework can also be used with non-Gaussian range kernels, and that  $\mathcal{O}(1)$  convolutions can be used to improve the computational complexity. As the range dimension  $\rho$

increases, these methods however become computationally inefficient. In particular, an exponentially large  $K$  is required to achieve a decent approximation, and linear interpolation requires  $2^\rho$  neighboring samples. Yang *et al.* [27] extended this method to perform bilateral filtering of color images, and proposed a fast scheme for linear interpolation. However, [27] is based on uniform quantization, which is not optimal given the strong correlation between the color channels. In fact, a more efficient algorithm for color filtering was later proposed in [19], where clustering is used to find  $\boldsymbol{\mu}_k$ . Moreover, instead of interpolation, they used local statistics prior to improve the accuracy. However, the algorithm is used only for color bilateral filtering and its application to other forms of high-dimensional filtering is not reported in [19].

2) *Splat-Blur-Slice Method*: Another line of work is [14], [17], and [18], which is based on a slightly different form of approximation. More precisely, they are based on the “splat-blur-slice” framework, which involves data partitioning (clustering or tessellation), fast convolutions, and interpolation as the core operations. These are considered to be the state-of-the-art fast algorithms for high-dimensional bilateral and nonlocal means filtering. The idea, originally proposed in [16], is based on the observation that [20] corresponds to convolutions in the joint spatio-range space. The general idea is to sample the input pixels in a different space (splatting), perform Gaussian convolutions (blurring), and resample the result back to the original pixel space (slicing). Adams *et al.* tessellated the domain and performed blurring using the  $k$ -d tree in [17] and the permutohedral lattice in [18]. Gastal and Oliveira [14] divided and resampled the domain into non-linear manifolds, and performed blurring on them. This was shown to be faster than all other methods within the splat-blur-slice framework.

3) *Kernel Approximation*: In an altogether different direction, it was shown in [22], [24], [25], and [28] that fast algorithms for bilateral filtering can be obtained by approximating the range kernel  $\varphi$  using the so-called *shiftable* functions, which includes polynomials and trigonometric functions. The above mentioned algorithms provide excellent accelerations for grayscale images, and are generally superior to algorithms based on data approximation. Unfortunately, it is difficult to approximate a high-dimensional Gaussian using low-order shiftable functions. For example, the Fourier expansion in [25] is quite effective in one dimension (grayscale images), but its straightforward extension to even three dimensions (color images) results in exponentially many terms. This is referred to as the “curse of dimensionality.”

## C. Contributions

Existing algorithms have focussed on either approximating the data [14], [17]–[19], [26] or the range kernel [22], [24], [25], [29], [30]. In this paper, we combine the ideas of shiftable [22], [24], [28], [31] and data approximation [23], [26] within a single framework to approximate (1) and (2). In particular, we locally approximate the range kernel in (1) and (2) using weighted and shifted copies of a Gaussian, where the weights and shifts are determined from the data.




Fig. 1. Bilateral filtering of a color image [20], where the filter parameters are  $\sigma_s = 10$  and  $\sigma_r = 40$ . We have used 15 clusters for GCS and the proposed method. The number of manifolds was automatically set to 3 in AM. The run-time and PSNR (see definition in (21)) with 15 manifolds are 1.48 sec and 43.41 dB. In the second row, we show the error between a particular method (first row) and the brute-force implementation for just the red channel. (a) *Lena* ( $3 \times 512 \times 512$ ). (b) **Proposed, 274 ms, 55.36 dB**. (c) GCS [19], 278 ms, 53.31 dB. (d) AM [14], 272 ms, 40.58 dB. (e) Brute-force, 19 sec. (f) Error: (b)–(e). (g) Error: (c)–(e). (h) Error: (d)–(e).

More specifically, we use clustering to determine the shifts, whereby the correlation between data channels is taken into account. Once the shifts have been fixed, we determine the weights (coefficients) using least-squares fitting, where the data is again taken into account. An important technical point is that we are required to solve a least-squares problem at each pixel, which can be expensive. We show how this can be done using just one matrix inversion. In summary, the key distinctions of our method in relation to previous approaches are as follows:

- We use data-dependent methods to calculate both  $\mu_k$  and  $c_k$  in (3) and (4). The latter is done in a heuristic fashion in [14], [17], and [18]. We note that our approximation also reduces to the form in (3) and (4). However, the notion of shiftable approximation (in a sense which will be made precise in Section II) plays an important role. Namely, it allows us to interpret the coefficients in (3) and (4) from an approximation-theoretic point of view. This in turn forms the basis of the proposed optimization used to compute the coefficients.
- Unlike [19], [20], [23], [26], and [27], our approach is scalable in the dimensions of the input  $f$  and the guide  $p$ .
- An important difference with [22], [24], [25], and [28] is that a different shiftable approximation is used at each pixel in our proposal, while the approximation is global in [22], [24], [25], and [28]. This will be made precise in Section II.
- Similar to [14], [17], and [18], our method also involves splatting (clustering), blurring (convolutions) and slicing (weighted recombinations). However, the important

difference is that we perform slicing in an optimal manner, which is done in a heuristic fashion in these methods.

We explain the proposed shiftable approximation in detail in Section II. The algorithm emerging from the approximation is conceptually simple and easy to code. The flowchart of the processing pipeline is shown in Figure 2. The computation-intensive components are “Cluster Range Space” and “Fast Convolutions”, for which highly optimized routines are available. Excluding clustering, the complexity is  $\mathcal{O}(K(n + \rho))$ , where  $K$  is the number of clusters. Notice that there is no dependence on  $S$ . This is a big advantage compared to fast nonlocal means [13] and [15], where the scaling is  $\mathcal{O}(S^d)$ . The highlight of our proposal is that we are able to derive a bound on the approximation error for a particular variant of our algorithm [32]. In particular, the error is guaranteed to vanish if we use more clusters. This kind of guarantee is not offered by state-of-the-art algorithms for high-dimensional filtering [14], [17], [18], partly due to the complex nature of their formulation.

We use the proposed fast algorithm for various applications such as smoothing [4], denoising [9], low-light denoising [33], hyperspectral filtering [34], and flow-field denoising [35]. In particular, we demonstrate that our algorithm is competitive with existing fast algorithms for color bilateral and nonlocal means filtering (cf. Figures 1 and 11). We note that although our original target was high-dimensional filtering, our algorithm outperforms [26], which is considered as the state-of-the-art for bilateral filtering of grayscale images (cf. Figure 5). Finally, we note that we have not used the particular form of



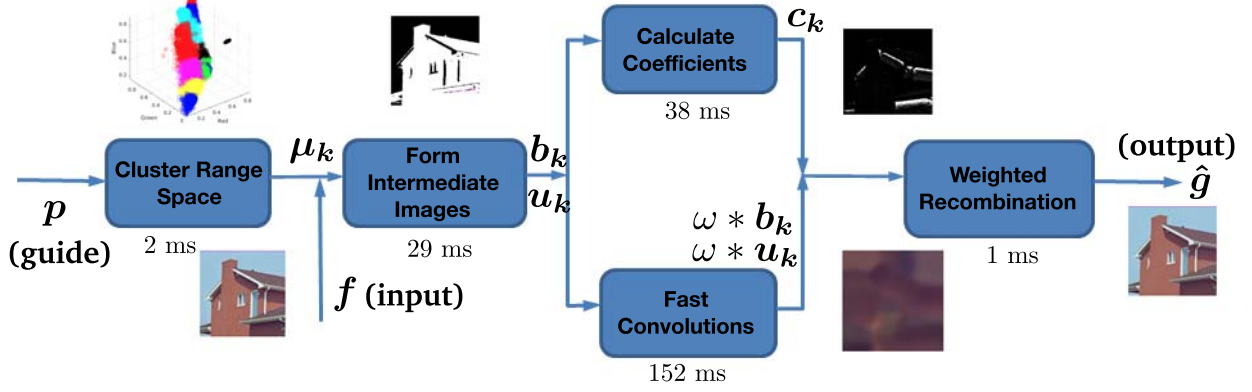


Fig. 2. Flowchart of our algorithm with the main modules. Also shown are the timings for the bilateral filtering of a 1MP color image. The filter parameters are  $\sigma_s = 10$  and  $\sigma_r = 75$ . When  $K = 4$  (number of clusters), total time required is 253 ms and the PSNR is 48.30 dB. The symbols used in the flowchart are defined in Section II (also see Algorithm 1).

the Gaussian kernel in the derivation of the fast algorithm. Therefore, it can also be used for non-Gaussian kernels, such as the exponential and the Lorentz kernel [23], [36].

#### D. Organization

This rest of the paper is organized as follows. In Section II, we describe the proposed approximation and the resulting fast algorithm. We also derive error bounds for a particular variant of the approximation. In Section III, we apply the fast algorithm for high-dimensional bilateral and nonlocal means filtering, and compare its performance (timing and accuracy) with state-of-the-art algorithms. We conclude the paper with a summary of the results in Section IV.

## II. PROPOSED METHOD

Notice that, if the guide  $p$  has constant intensity value at each pixel, then (1) and (2) reduce to linear convolutions. This observation is essentially at the core of the fast algorithms in [20], [23], and [26]. On the other hand, the fast algorithms in [22], [25], and [28] are derived using a completely different idea, where the univariate Gaussian kernel  $\varphi$  is approximated using a polynomial or a trigonometric function  $\psi$ . These functions are *shiftable* in the following sense: There exists basis functions  $\psi_1, \dots, \psi_K$  such that, for any  $\tau \in \mathbb{R}$ , we can write

$$\psi(x - \tau) = \sum_{k=1}^K c_k(\tau) \psi_k(x), \quad (5)$$

where the coefficients  $c_1, \dots, c_K$  depend only on  $\tau$ . One can readily see that polynomials and trigonometric functions,

$$\psi(x) = \sum_{n=0}^N p_n x^n \quad \text{and} \quad \psi(x) = \sum_{n=0}^N q_n \cos(n\omega_0 x), \quad (6)$$

are shiftable. The utility of expansion (5) is that it allows us to factor out  $p(i)$  from the kernel. In particular, by replacing  $\varphi$  by  $\psi$ , it was shown in [22] that we can compute the bilateral filter using fast convolutions. The Taylor polynomial of  $\varphi$  was used for  $\psi$  in [24], and the Fourier approximation of  $\varphi$  was used in [22], [25], and [28].

#### A. Shiftable Approximation

As mentioned earlier, it is rather difficult to obtain shiftable approximations for high-dimensional Gaussians that are practically useful. We can use separable extensions of (6) to generate such approximations. However, the difficulty is that the number of terms grows as  $N^D$  in this case, where  $N$  is the order in (6). We address this fundamental problem using a data-centric approach. First, we do not use a shiftable (trigonometric or polynomial) approximation of  $\varphi$  [22], [25]. Instead, for some fixed pixel  $i \in \Omega$ , we consider the multidimensional Gaussian  $\varphi(x - p(i))$  centered at  $p(i)$  appearing in (1) and (2). Similar to (5), we wish to find  $\psi_1, \dots, \psi_K$  such that

$$\varphi(x - p(i)) \approx \sum_{k=1}^K c_k(i) \psi_k(x), \quad (7)$$

where  $c_1, \dots, c_K$  depend only on  $p(i)$ . The key distinction with [22], [24], [25], and [28] is that the functional approximation is defined locally in (7), namely, the approximation is different at each pixel. Moreover, (7) is neither a polynomial nor a trigonometric function. In spite of this, we continue to use the term “shiftable” to highlight the fact that the approximation is based on “shifts” of the basis functions  $\psi_1, \dots, \psi_K$ . In this work, we propose to use the translates of the original kernel as the basis functions, i.e., we set  $\psi_k(x) = \varphi(x - \mu_k)$ . One can in principle use different basis functions, but we will not investigate this possibility in this paper.

The important consideration is that the shifts  $\{\mu_k\}$  are set globally. Since the range kernel is defined via the guide  $p$ , we propose to cluster the range space of  $p$  and use the cluster centroids for  $\{\mu_k\}$ . That is we partition the range space

$$\Theta = \{p(i) : i \in \Omega\}, \quad (8)$$

into clusters  $\mathcal{C}_1, \dots, \mathcal{C}_K$ . We set  $\mu_k$  to be the centroid of  $\mathcal{C}_k$ . In summary, for each  $i$ , we require that

$$\varphi(x - p(i)) \approx \sum_{k=1}^K c_k(i) \varphi(x - \mu_k). \quad (9)$$




Fig. 3. Proposed approximation of the Gaussian range kernel  $\varphi(x - p(i))$  in (9) for  $p(i) = 25$  (left) and  $p(i) = 100$  (right). The shifts  $\{\mu_k\}$  are obtained by clustering a grayscale image whose histogram is shown above. As expected, the approximation is better for order  $K = 8$  compared to  $K = 4$ . In particular, the approximation is quite accurate over the entire dynamic range  $[0, 255]$  when  $K = 8$ . However, notice that for  $K = 4$ , the approximation is better in the interval  $[100, 200]$ , where the density is higher. As the approximation is data-driven, the error gets distributed in tune with the underlying histogram.

At this point, notice the formal resemblance between (5) and (9). Hence, we will refer to (9) as a shiftable approximation in the rest of the discussion, even though it is not shiftable in the sense of (5).

Note that we need a good approximation in (9) only for  $x \in \Theta_i = \{p(i - j) : j \in W\}$ . This is because the samples appearing in (1) and (2) are  $\varphi(x - p(i))$ , where  $x$  takes values in  $\Theta_i$ . To determine the coefficients, it is therefore natural to consider the following problem:

$$\min_{c(i) \in \mathbb{R}^K} \sum_{x \in \Theta_i} \left( \varphi(x - p(i)) - \sum_{k=1}^K c_k(i) \varphi(x - \mu_k) \right)^2. \quad (10)$$

The difficulty is that this requires us to solve an expensive least-squares problem (the matrix to be inverted is large) for each  $i$ . More importantly, we have a different inversion, at each pixel. This is time consuming and impractical. On the other hand, notice that  $\Theta_i$  is included in  $\Theta$ . Hence, we could perform the fitting over  $\Theta$ . Instead, we set  $\Lambda = \{\mu_1, \dots, \mu_K\}$  and choose to fit over  $\Lambda$ , which are quantized representatives of the points in  $\Theta$ . In short, while the cluster centers  $\Lambda$  most likely do not belong to  $\Theta_i$ , they are representative of the larger set  $\Theta$ . Therefore, instead of (10), we consider the surrogate problem:

$$\min_{c(i) \in \mathbb{R}^K} \sum_{x \in \Lambda} \left( \varphi(x - p(i)) - \sum_{k=1}^K c_k(i) \varphi(x - \mu_k) \right)^2.$$

In matrix notation, we can compactly write this as

$$\min_{c(i) \in \mathbb{R}^K} \|\mathbf{A}c(i) - \mathbf{b}(i)\|^2, \quad (11)$$

where  $\mathbf{A} \in \mathbb{R}^{K \times K}$  and  $\mathbf{b}(i) \in \mathbb{R}^K$  are given by

$$A_{kl} = \varphi(\mu_k - \mu_l) \text{ and } b_k(i) = \varphi(\mu_k - p(i)). \quad (12)$$

The solution of (11) is  $c(i) = \mathbf{A}^\dagger \mathbf{b}(i)$ , where  $\mathbf{A}^\dagger$  is the pseudo-inverse of  $\mathbf{A}$ . In particular, we need to compute the pseudo-inverse just once; the coefficients at each pixel are obtained using a matrix-vector multiplication. In Figure 6, we show that solving (11) is much faster than solving (10), and that the coefficients from (11) are close to those from (10).

The approximation result for univariate Gaussians are shown in Figure 3. Notice that our approximation depends on the distribution of  $p$  (because of the clustering), and the coefficients are obtained by solving a least square problem at each pixel  $i$ . Hence, in Figure 3, a particular distribution is selected and




Fig. 4. Error rates for the proposed approximation (see text for the definition of  $\mathcal{E}_K$ ), where  $\rho$  is the dimension of the range space of the guide image  $p$ . We notice that the approximation error falls off with the increase in order  $K$ . The error is averaged over the color images ( $\rho = 3$ ) from the Kodak dataset and hyperspectral images ( $\rho = 33$ ) from [37].

the approximations for  $p(i) = 25$  and  $p(i) = 100$  are shown. The error rates (error vs.  $K$ ) for multidimensional Gaussians (covariance  $\sigma_r^2 \mathbf{I}$ ) corresponding to dimensions  $\rho = 3$  and 33 are shown in Figure 4. The error  $\mathcal{E}_K$  for a fixed  $K$  is simply (10) averaged over all the pixels.

To summarize, the steps in the proposed shiftable approximation are as follows:

- Cluster the range space  $\Theta$  and use the centers for  $\mu_k$ .
- Set the basis functions as  $\psi_k(x) = \varphi(x - \mu_k)$ .
- Set up  $\mathbf{A}$  using (12) and compute its pseudo-inverse  $\mathbf{A}^\dagger$ .
- At each pixel  $i$ , set  $\mathbf{b}(i)$  using (12) and  $\mathbf{c}(i) = \mathbf{A}^\dagger \mathbf{b}(i)$ .

We note that one can freely choose different shifts (and basis functions) in (9) for different applications. That is, (9) offers a broad approximation framework, where one can consider other rules for fixing the parameters (shifts and coefficients).

### B. Fast Algorithm

We now develop a fast algorithm by replacing the kernel in (1) and (2) with the approximation in (9). For  $1 \leq k \leq K$ , define  $\mathbf{v}_k : \Omega \rightarrow \mathbb{R}^n$  and  $r_k : \Omega \rightarrow \mathbb{R}$  to be

$$\mathbf{v}_k(i) = \sum_{j \in W} \omega(j) \varphi(p(i - j) - \mu_k) \mathbf{f}(i - j), \quad (13)$$

and

$$r_k(i) = \sum_{j \in W} \omega(j) \varphi(p(i - j) - \mu_k). \quad (14)$$

For  $\mathbf{x} = p(i - j)$ ,  $j \in W$ , we replace  $\varphi(\mathbf{x} - p(i))$  in (1) and (2) with the approximation in (9). After some manipulations (see Appendix), (1) and (2) are approximated as,

$$\hat{\mathbf{g}}(i) = \frac{1}{\hat{\eta}(i)} \sum_{k=1}^K c_k(i) \mathbf{v}_k(i), \quad (15)$$

**Algorithm 1** Fast High-Dimensional Filtering

---

```

1 Input:  $f : \Omega \rightarrow \mathbb{R}^n$  and  $p : \Omega \rightarrow \mathbb{R}^\rho$ , kernels  $\omega$  and  $\varphi$ ;
2 Parameter: Number of clusters  $K$ ;
3 Output: Approximation in (15);
  // CLUSTER RANGE SPACE
4  $\{\mu_k\} \leftarrow$  cluster  $\Theta$  in (8) using bisecting  $K$ -means;
  // FORM INTERMEDIATE IMAGES
5 Set up  $\mathbf{A}$  in (12) using  $\{\mu_k\}$  and  $\varphi$ , and compute  $\mathbf{A}^\dagger$ ;
6 for  $i \in \Omega$  do
7   for  $k = 1, \dots, K$  do
8      $b_k(i) \leftarrow \varphi(\mu_k - p(i));$ 
9      $u_k(i) \leftarrow b_k(i)f(i);$ 
10  end
11 end
  // CALCULATE COEFFICIENTS
12 for  $i \in \Omega$  do
13    $c(i) \leftarrow \mathbf{A}^\dagger b(i);$ 
14 end
  // FAST CONVOLUTIONS
15 Initialize  $v : \Omega \rightarrow \mathbb{R}^n$  and  $r : \Omega \rightarrow \mathbb{R}$  with zeros;
16 for  $k = 1, \dots, K$  do
17    $v \leftarrow v \oplus [c_k \otimes (\omega * u_k)];$ 
18    $r \leftarrow r \oplus [c_k \otimes (\omega * b_k)];$ 
19 end
20  $\hat{g} \leftarrow v \oslash r.$ 

```

---

and

$$\hat{\eta}(i) = \sum_{k=1}^K c_k(i) r_k(i). \quad (16)$$

The advantage with (15) and (16) is clear once we notice that (13) and (14) can be expressed as convolutions. Indeed, by defining  $u_k : \Omega \rightarrow \mathbb{R}^n$  to be  $u_k(i) = b_k(i)f(i)$ , where  $b_k$  is given by (12), we can write

$$v_k(i) = (\omega * u_k)(i) = \sum_{j \in W} \omega(j) u_k(i - j),$$

and

$$r_k(i) = (\omega * b_k)(i) = \sum_{j \in W} \omega(j) b_k(i - j),$$

where  $*$  denotes linear convolution. In summary, using (9), we have been able to approximate the high-dimensional filtering using weighted combinations of fast convolutions. The overall process is summarized in Algorithm 1. Note that we have used  $\oplus$ ,  $\otimes$  and  $\oslash$  to represent pointwise addition, multiplication, and division. In steps 17 and 18,  $c_k$  denotes the mapping  $c_k : \Omega \rightarrow \mathbb{R}$ . The dominant computation in Algorithm 1 are the  $(n+1)K$  convolutions in steps 17 and 18. Clustering and inversion of  $\mathbf{A}$  is performed just once. Notice that  $b_k$ , which is used for computing the coefficients in step 13, is anyways required to form the intermediate images in step 9. The flowchart of our algorithm, along with typical timings for a megapixel image, is shown in Figure 2. Note that the steps in Figure 2 (resp. loops in Algorithm 1) can

be parallelized: clustering the range space (loop 6), finding coefficients (loop 12), and performing convolutions (loop 16).

### C. Implementation

We have used bisecting  $K$ -means [38] for iteratively clustering the range space in step 4. In particular, the cluster with largest variance is divided into two parts (using 2-means clustering) at every iteration. This ensures that we do not end up with few large clusters at the end. We initialize the 2-means clustering using a pair of points with the largest separation. We picked bisecting  $K$ -means over other clustering algorithms as its cost is linear in  $K$ , and is generally faster than other clustering algorithms [38]. Importantly, its computational overhead is negligible compared to the time required for the convolutions. Needless to say, we can use any efficient clustering method and the fast algorithm would still go through.

We have used the Matlab routine *pinv* to compute  $\mathbf{A}^\dagger$  in step 5. As is well-known, when  $\omega$  is a box or Gaussian, we can convolve using  $\mathcal{O}(1)$  operations, i.e., the cost does not depend on the size of spatial filter [39], [40]. The Gaussian convolutions in steps 17 and 18 are implemented using Young's  $\mathcal{O}(1)$  algorithm [40]. For box convolutions, we used a standard recursive procedure which require just four operations per pixel for any box size. Since the leading computations are the convolutions, and there are  $\mathcal{O}(Kn)$  convolutions, the overall complexity is  $\mathcal{O}(K(n + \rho))$ ; the additional  $\mathcal{O}(K\rho)$  term is for computing (12). Importantly, we are able to obtain a speedup of  $S^d/K$  over the brute-force computations of (1) and (2). The storage complexity of Algorithm 1 is clearly linear in  $n$ ,  $K$ , as well as the image size. The MATLAB implementation of our algorithm is available here: <https://github.com/pravin1390/FastHDFilter>.

### D. Approximation Accuracy

We note that the algorithm in [32] is a variant of our method. Indeed, for  $1 \leq k \leq K$ , if we set the coefficients in (9) to be

$$c_k(i) = \begin{cases} 1 & \text{if } p(i) \in \mathcal{C}_k, \\ 0 & \text{else,} \end{cases} \quad (17)$$

then we obtain the fast algorithm in [32]. In other words, the filtering is performed on a cluster-by-cluster basis in this case. Correspondingly, (15) reduces to

$$\hat{g}(i) = \frac{v_s(i)}{r_s(i)}, \quad (18)$$

assuming that  $p(i)$  is in cluster  $\mathcal{C}_s$ .

In (15) and (16), we weight the results from the  $K$  clusters, where the weights are obtained using the optimization in (11). We will demonstrate in Section III that the weighting helps in reducing the approximation error. Unfortunately, it also makes the analysis of the algorithm difficult. Intuitively though, one would expect the error from (15) to be less than that from (18). For the latter approximation, we have the following result [32].

**Theorem 1 (Error Bound):** Suppose that  $\omega$  and  $\varphi$  are non-negative, and that the latter is Lipschitz continuous, i.e., for

some  $L > 0$ ,  $|\varphi(\mathbf{x}) - \varphi(\mathbf{y})| \leq L\|\mathbf{x} - \mathbf{y}\|$ , for any  $\mathbf{x}$  and  $\mathbf{y}$  in  $\mathbb{R}^p$ . Then, for some constant  $C > 0$ ,

$$\sum_{i \in \Omega} \|\hat{\mathbf{g}}(i) - \mathbf{g}(i)\|^2 \leq CL^2 n |W|^2 E_K, \quad (19)$$

where  $\hat{\mathbf{g}}$  is given by (18), and  $E_K$  is the clustering error:

$$E_K = \sum_{k=1}^K \sum_{p(i) \in C_k} \|\mathbf{p}(i) - \boldsymbol{\mu}_k\|^2. \quad (20)$$

We note that the assumptions in Theorem 1 are valid when  $\omega$  is a box or Gaussian, and  $\varphi$  is a multidimensional Gaussian. For completeness, we provide the derivation of Theorem 1 in Appendix. For several clustering methods, the clustering error  $E_K$  vanishes as  $K \rightarrow \infty$  [38]. For any of these methods, we can approximate (1) with arbitrary accuracy provided  $K$  is sufficiently large. We note that such a guarantee is not available for existing fast approximation algorithms for high-dimensional filtering [14], [17], [18].

### III. NUMERICAL RESULTS

We now report some numerical results for high-dimensional filtering. We have used the Matlab implementation of Algorithm 1. For fair timing comparisons with Adaptive Manifolds (AM) [14] and Global Color Sparseness (GCS) [19], we have used the Matlab code provided by the authors. The timing analysis was performed on an Intel 4-core 3.4 GHz machine with 32 GB memory. The grayscale and color images used for our experiments were obtained from standard databases.<sup>1</sup> The infrared and hyperspectral images are the ones used in [14] and [37], [41]. To compare the filtering accuracy with existing methods, we have fixed the timings by adjusting the respective approximation order. The objective of this section is to demonstrate that, in spite of its simple formulation, our algorithm is competitive with existing fast approximations of bilateral and nonlocal means filtering.

We have used an isotropic Gaussian range kernel for bilateral and nonlocal means filtering :

$$\varphi(\mathbf{x}) = \exp\left(-\|\mathbf{x}\|^2 / 2\sigma_r^2\right),$$

where the standard deviation  $\sigma_r$  is used to control the effective width (influence) of the kernel. We recall from (1) and (2) that the input to  $\varphi$  is the difference  $\mathbf{x} = \mathbf{p}(i) - \mathbf{p}(i - \mathbf{j})$ . In one of the experiments,  $\varphi(\mathbf{x})$  is an anisotropic Gaussian; we have explicitly mentioned this at the right place. The spatial kernel  $\omega$  for bilateral filtering is also Gaussian:

$$\omega(i) = \exp\left(-\|i\|^2 / 2\sigma_s^2\right),$$

where  $\sigma_s$  is again the standard deviation. For nonlocal means (NLM),  $\omega$  is a box filter, namely, no spatial weighting is used. The filter width for bilateral filtering is  $S = 3\sigma_s$ , while that for NLM is explicitly mentioned (typically,  $S = 10$ ).

Following [24] and [26], the error between the outputs of the brute-force implementation and the fast algorithm is measured

<sup>1</sup><https://goo.gl/821N2G>, <https://goo.gl/2fcNmu>, <https://goo.gl/MvxCMX>.




Fig. 5. Visual comparison of bilateral filtering for the *House* image. Filter parameters:  $\sigma_s = 10$  and  $\sigma_r = 30$ . For a fair comparison, we have used four clusters for the proposed method and four bins for Yang's method [26]. Notice that our result is better than Yang's, both visually (compare boxed areas) and in terms of PSNR. This is because we use non-uniform quantization (clustering) and data-driven approximation (see Figure 3). (a) Input ( $256 \times 256$ ). (b) Brute-force filtering. (c) **Proposed (43.5 dB)**. (d) Yang (34 dB).

TABLE I

COMPARISON OF APPROXIMATION ACCURACY (PSNR) FOR BILATERAL FILTERING OF *Barbara* ( $512 \times 512$ ). AVERAGE TIMING IS 530 MS FOR YANG AND 550 MS FOR PROPOSED METHOD

$\sigma_s \backslash \sigma_r$		10	20	30	40	50	60	70	80	90	100
<b>Yang [26]</b>											
10	<b>PSNR</b>	51.48	53.01	53.48	53.65	53.48	53.17	53.01	52.71	52.28	52.01
30	<b>PSNR</b>	51.23	53.17	53.81	53.98	53.98	53.81	53.81	53.98	53.98	53.98
30	<b>PSNR</b>	50.86	53.32	53.81	53.98	53.82	53.65	53.65	53.82	53.98	54.33
<b>Proposed</b>											
10	<b>PSNR</b>	56.08	64.61	61.69	59.50	57.50	56.09	54.88	53.81	53.32	52.71
30	<b>PSNR</b>	56.53	68.13	64.61	62.11	60.53	59.19	58.03	57.25	56.77	56.09
30	<b>PSNR</b>	57.01	70.07	67.30	64.61	62.56	60.89	60.17	59.19	58.88	58.59

using the peak signal-to-noise ratio (PSNR). This is given by  $\text{PSNR} = -10 \log_{10}(\text{MSE})$ , where

$$\text{MSE} = |\Omega|^{-1} \sum_{i \in \Omega} \|\hat{\mathbf{g}}(i) - \mathbf{g}(i)\|^2, \quad (21)$$

and  $|\Omega|$  is the number of pixels. Note that PSNR is also used to measure denoising performance; in this case,  $\hat{\mathbf{g}}(i)$  and  $\mathbf{g}(i)$  are the denoised and clean images in (21). It should be evident from the context as to what is being measured using PSNR. We also use SSIM [42] for measuring visual quality.

#### A. Grayscale Bilateral Filtering

As mentioned earlier, although the proposed method is targeted at high-dimensional filtering, we can outperform [26], which is considered the state-of-the-art for bilateral filtering of grayscale images. We demonstrate this in Table I and Figure 5 using a couple of examples. In particular, notice in Figure 5 that artifacts are visible in Yang's result (when  $K$  is small.). Our PSNR is almost 10 dB higher, and our output does not show any artifacts. A detailed PSNR comparison is provided in Table I. We notice that, for the same timing, our PSNR is consistently better than [26]. Note that we have set  $K = 4$  as



TABLE II  
TIMINGS FOR BILATERAL FILTERING OF *Barbara* FOR  
VARYING  $\sigma_s$  ( $\sigma_r = 100$ ,  $K = 16$ )

$\sigma_s$	10	20	30	40	50	60	70	80
Time (ms)	496	498	475	486	500	489	485	481

TABLE III  
PSNR FOR COLOR IMAGES ( $\sigma_s = 10$ ,  $\sigma_r = 50$ )

size \ $K$	4	8	16
1080P	33.54	48.68	53.86
4K	31.29	47.42	53.16
8K	34.25	46.90	54.03




Fig. 6. Bilateral filtering of a grayscale image. Parameters:  $\sigma_s = 2$ ,  $\sigma_r = 20$ , and  $K = 4$ . The coefficients are obtained from (10) and (11) in (c) and (d). (a) Input ( $256 \times 256$ ). (b) Brute-force (495 ms). (c) Using (10) (5 sec, 35.3 dB). (d) Using (11) (27 ms, 32.6 dB).

the number of clusters (resp. quantization bins) for our method (resp. [26]).

For completeness, we have reported the timings for different  $\sigma_s$  in Table II, where we have used Young's fast algorithm for performing the Gaussian convolutions [40]. As expected, the timing essentially does not change with  $\sigma_s$ .

In Figure 6, we have compared the bilateral filtering results obtained using (10) and (11). Since (11) is an approximation of (10), we see a drop in PSNR using (10), but importantly there is a significant speed up using (11). In fact, the time required using (10) is more than that of the brute-force implementation.

### B. Color Bilateral Filtering

We next perform fast bilateral filtering of color images, for which the state-of-the-art methods are AM [14] and GCS [19]. We have also compared with our previous work [32], which we simply refer to as "Clustering." In Figure 7, we have studied the effect of the number of clusters on the filtering




Fig. 7. Scaling of filtering accuracy (average PSNR) with number of clusters (manifolds for AM), for bilateral filtering of color images from the Kodak dataset. The parameters are  $\sigma_s = 20$  and  $\sigma_r = 40$ .




Fig. 8. Timing (seconds) vs. resolution (MP) for color filtering.




Fig. 9. Approximation accuracy (PSNR) for bilateral filtering of color *Barbara* ( $3 \times 512 \times 512$ ), at different  $\sigma_r$  values averaged over  $\sigma_s = 10, 20, 30$ . The average timings are: Proposed: 475 ms, Global Color Sparseness: 490 ms, Adaptive Manifolds: 370 ms and Clustering [32]: 370 ms.

accuracy, for our method, GCS, AM and [32]. While GCS performs better for small  $K$  (coarse approximation), our method outperforms GCS when  $K > 15$ . Note that the number of manifolds can only take dyadic (discrete) values [14].

In Table III, we report the filtering accuracy for full HD, 4K ultra HD, and 8K ultra HD by varying the number of clusters  $K$ . For a given resolution and  $K$ , we have averaged the PSNR over six high resolution images<sup>2</sup> with that resolution. Notice that, irrespective of the image resolution, the PSNR values are above 40 dB with just 8 clusters. In Figure 8, the timings for filtering color images are plotted for varying image sizes (1.5 MP to 32 MP). This plot verifies the linear dependence of the timing on the image size, irrespective of the number of clusters used.

In Figure 9, we have compared the PSNR and timings for bilateral filtering a color image, at different  $\sigma_s$  and  $\sigma_r$  values. We have used  $K = 16$  clusters for all three methods: proposed, GCS, and [32]. We use the default settings for the number of manifolds in AM. Notice that we are better by 5-10 dB, while the timings are roughly the same. In particular, notice the PSNR improvement that we obtain over [32] by optimizing the coefficients, i.e., by using (15) in place of (18). Interestingly, [32] can outperform GCS and AM at large  $\sigma_r$  values.

A visual comparison for color bilateral filtering was already shown in Figure 1, along with the error images, timings, and

<sup>2</sup><https://hdqwalls.com>





Fig. 10. Bilateral filtering of *Peppers* for different values of  $K$  ( $\sigma_s = 10$  and  $\sigma_r = 40$ ). Notice how the accuracy improves with  $K$ . (a)  $K = 2$  (15 ms, 22 dB). (b)  $K = 4$  (28 ms, 29 dB). (c)  $K = 8$  (52 ms, 37 dB). (d)  $K = 16$  (100 ms, 44 dB). (e) Brute-force.

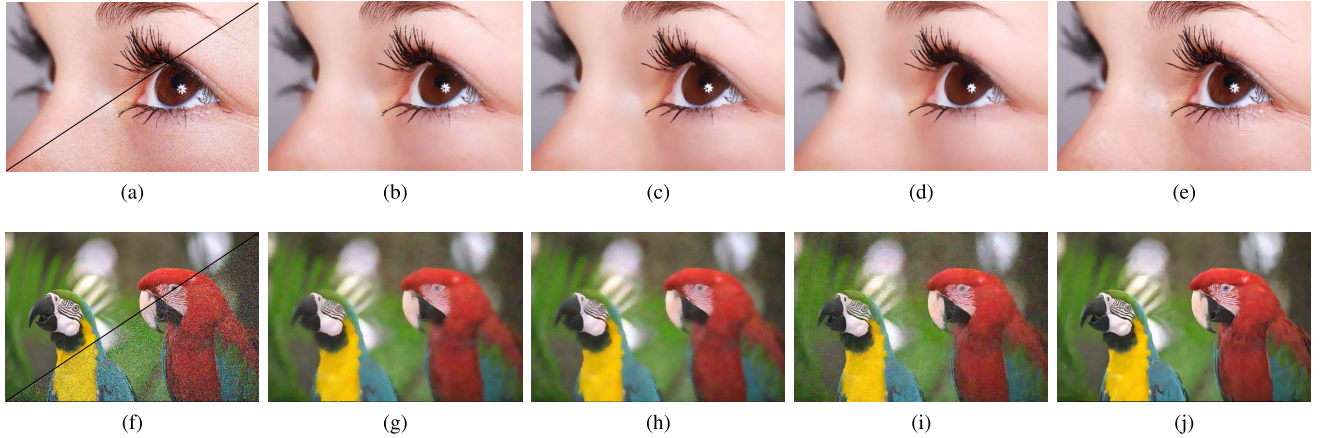


Fig. 11. Color image denoising (Gaussian noise) using PCA-NLM [21], where the patch and search sizes are  $7 \times 7$  and  $21 \times 21$ . For PCA-NLM, AM and our method, PCA was used to reduce the range dimension from  $3 \times 7 \times 7$  to 25. We used 31 clusters for our method. The number of manifolds was automatically set to 15 in the AM code [14]. The run-time, PSNR and SSIM for AM with 31 manifolds are 6 sec, 31.1 dB and 0.90 (top image) and 8.3 sec, 23.7 dB and 0.69 (bottom image). For comparison, we have also shown the result from a state-of-the-art Gaussian denoiser [12]. (a) Clean/Noisy (22 dB,  $\sigma = 20$ ). (b) [21] (8 min, 33.3 dB, 0.92). (c) **Ours** (1.7 sec, 31.2 dB, 0.92). (d) AM (3 sec, 30.7 dB, 0.90). (e) [12] (3.5sec, 36.2 dB, 0.93). (f) Clean/Noisy (12 dB,  $\sigma = 63$ ). (g) [21] (9 min, 26.7 dB, 0.77). (h) **Ours** (2.7 sec, 26.8 dB, 0.82). (i) AM (4.2 sec, 25 dB, 0.70). (j) [12] (5 sec, 31 dB, 0.86).




Fig. 12. Comparison of the denoising performances of PCA-NLM [21], Adaptive Manifolds, and the proposed method. The noise level is denoted by  $\sigma (\times 255)$ . The PCA dimension was set to 6 for all methods. The PSNR and SSIM [42] values are averaged over the images from the Kodak dataset. The parameters used are  $K = 31$ ,  $S = 10$ , and  $m = 3$ .

PSNRs. In particular, notice that the error from our method is smaller than AM and GCS. This is interesting as our method is conceptually simpler than these algorithms — AM has a complex formulation and GCS is based on two-pass filtering.

In Figure 10, a visual result is provided to highlight how the approximation accuracy scales with  $K$ . Notice that the accuracy improves significantly as  $K$  increases from 2 to 16. In fact, the approximation already resembles the result of brute-force filtering when  $K$  is 8.

### C. Color Nonlocal Means

In NLM,  $f(i)$  is the noisy image (corrupted with additive Gaussian noise) and  $p(i)$  is a square patch of pixels around  $i$ .

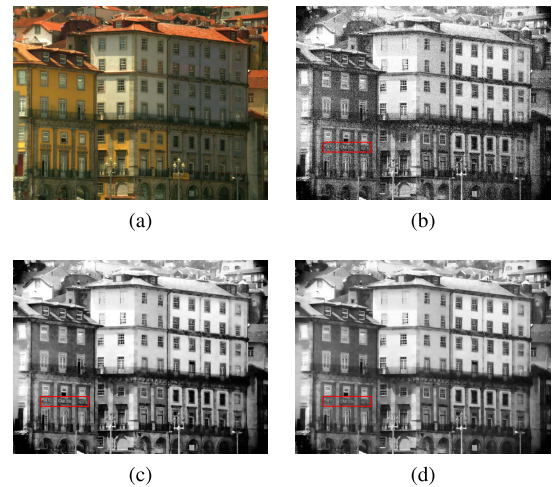


Fig. 13. Denoising of a hyperspectral image ( $1340 \times 1017 \times 33$ ) using bilateral filtering ( $\sigma_s = 5$  and  $\sigma_r = 100$ ). The image in (b) shows one of the noisy bands; the same band after filtering is shown in (c) and (d). We have shown one of the bands just for visualization; the filtering was performed on the entire hyperspectral image and not on a band-by-band basis. (a) Input [37]. (b) 440nm band. (c) Brute-force (27 min). (d) **Proposed** (15 sec).

In particular, if the patch is  $m \times m$ , then the dimension of  $p(i)$  is  $\rho = 3m^2$  for a color image. A popular trick to accelerate NLM (called PCA-NLM) is to project the patch onto a lower-dimensional space using PCA [21]. We note that




Fig. 14. Nonlocal means denoising of a low-light image using [43] and the proposed method, with and without infrared data [33]. Notice that we obtain better denoising by taking the infrared information into account (see zoomed sections). (a) Infrared input. (b) Noisy low-light input. (c) Without infrared. (d) **With infrared data.** (e) [43].

PCA was used for reducing the patch dimension (as explained earlier) for both AM and our algorithm.

A visual comparison is provided in Figure 11 for low and high noise, where the superior performance of our method over AM is evident. We have also compared with a state-of-the-art Gaussian denoiser (CPU implementation<sup>3</sup>) based on deep neural nets [12]. It is not surprising that the result from [12] is better both in terms of PSNR and visual quality. However, we are somewhat faster. Figure 11 also suggests that our method is visually closer to PCA-NLM at both low and high noise levels, albeit with a significant speedup (10 min to 2.7 sec).

In Figure 12, we have reported Gaussian denoising results on the Kodak dataset for different noise levels ( $\sigma$ ). While the proposed method and AM perform similarly for small  $\sigma$ , the former is more robust when  $\sigma$  is large. In fact, the denoising performance of AM degrades rapidly with the increase in  $\sigma$ . Another important point is that our result is close to PCA-NLM in terms of PSNR and SSIM for all noise levels.

#### D. Hyperspectral Denoising

We next perform denoising of hyperspectral images using the bilateral filter. For such images, the range dimension (i.e., the number of spectral bands) is high. In Figure 13, denoising results are exclusively shown to highlight the speedup of our approximation over brute-force implementation for a 33-band image. Notice the dramatic reduction in timing compared to the brute-force implementation. Moreover, notice that the denoising quality is in fact quite good for our method (compare the text in the boxed regions).

Moreover, we also compare with recent optimization-based denoising methods [44], [45], where parameters are tuned accordingly. Visual and quantitative comparisons are shown in Figure 15 (Pavia dataset). For quantitative comparisons, we have used MPSNR and MSSIM, which are simply the PSNR and SSIM values averaged over the spectral bands. We notice that the restoration obtained using our method is better than [44] and [45] (source code made public by authors), which is supported by the metrics shown in the figures. The same is visually evident from a comparison of the boxed regions in Figure 15. In particular, the color is not restored properly in [44], and grains can be seen in [45]. As expected,

<sup>3</sup><https://github.com/cszn/FFDNet>

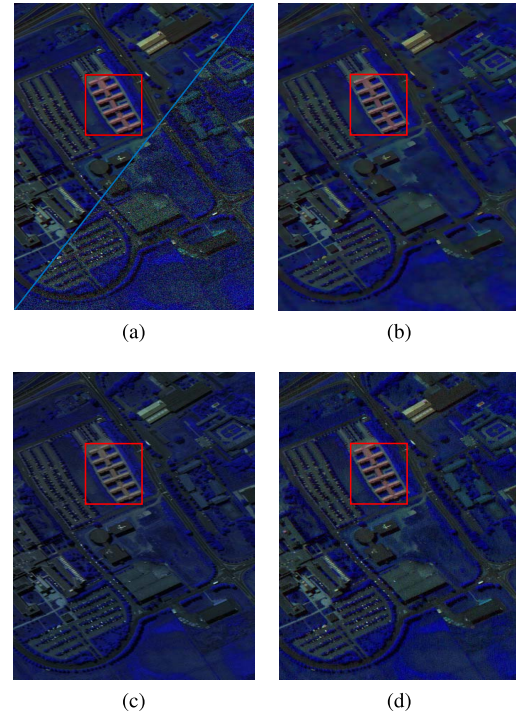


Fig. 15. Hyperspectral denoising of a natural image corrupted with Gaussian noise. Image size:  $(610 \times 340) \times 200$  bands. For our method,  $\sigma_s = 3$ ,  $\sigma_r = 100$ , and  $K = 32$ . We used 4 iterations for [44] and 1 iteration for [45]. The (timing, MPSNR, MSSIM) for (b), (c) and (d) are **(36 sec, 31.29 dB, 0.87)**, (192 sec, 30.54 dB, 0.89) and (1250 sec, 30.09 dB, 0.74). (a) Clean/Noisy ( $\sigma = 0.1$ ). (b) **Proposed.** (c) [44]. (d) [45].

we are much faster than these iterative methods, since we perform the filtering in one shot.

#### E. Low-Light Denoising

Finally, we use our fast algorithm for NLM denoising of low-light images using additional infrared data [33]. In Figure 14, we have shown a visual result which compares our method, both with and without infrared data, and [43] (an optimization method). We set the patch and search window sizes as  $7 \times 7$  and  $17 \times 17$ . We used an anisotropic Gaussian kernel, where  $\sigma_r = 50$  for the low-light data and  $\sigma_r = 10$  for the infrared data. The dimension was reduced from  $3 \times 7 \times 7$  to 6 using PCA, and the infrared data was added as the seventh dimension. We used 32 clusters for our method. In Figure 14, notice that some salient features are lost if we



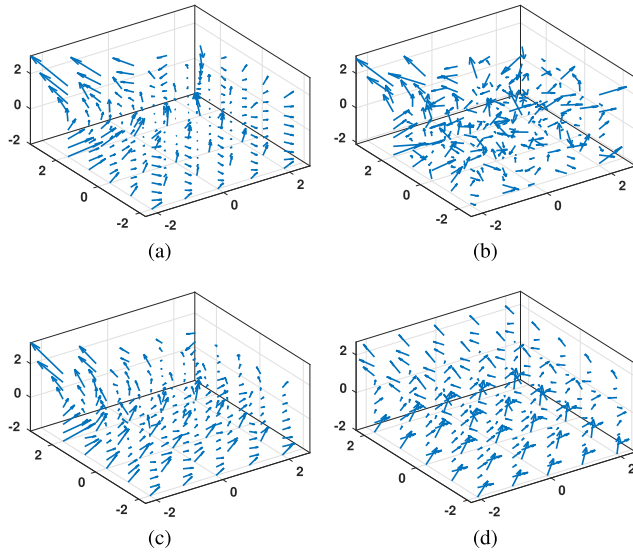


Fig. 16. Denoising of a synthetic flow-field using proposed NLM approximation and Gaussian smoothing. The synthetic flow-field used is  $u(x, y, z) = x^2 - ze^y$ ,  $v(x, y, z) = y^3 - xze^y$ , and  $w(x, y, z) = z^4 - xe^y$ . The patch and search window sizes for NLM are  $7 \times 7 \times 7$  and  $21 \times 21 \times 21$ , and  $\sigma_r = 7.5$ . We have used PCA to reduce the dimension of the patch from  $3 \times 7^3$  to 6. Notice that isotropic Gaussian smoothing fails to preserve sharp changes in the flow direction. (a) Clean flow-field. (b) Noisy flow-field ( $\sigma = 7.5$ ). (c) NLM output. (d) Gaussian smoothing.

solely use the low-light input as the guide. Instead, if we add the infrared data to the guide, then the restored image is sharper and the features are more apparent compared to [43] and NLM without infrared data.

#### F. Flow-Field Denoising

Finally, we apply the proposed method for flow-field denoising. In particular, sharp directional changes in the flow can be preserved much better using NLM, while simultaneously removing the noise (with Gaussian filtering as the baseline). An instance of flow-field denoising [35] is shown in Figure 16 and compared with Gaussian smoothing.

### IV. CONCLUSION

We proposed a framework for fast high-dimensional filtering by approximating both the data and the kernel. In particular, we derived an algorithm that fuses the scalability of the former with the approximation capability of the latter. At the core of our algorithm is the concept of shiftable approximation, which allows us to interpret the coefficients in the framework of Durand and Dorsey [23] from an approximation-theoretic point of view. We proposed an efficient method for determining the shifts (centers) and the coefficients using  $K$ -means clustering (inspired by [19]) and data-driven optimization. Though the proposed algorithm is conceptually simple and easy to implement (about 15 lines of code), it was shown to yield promising results for diverse applications. In particular, our algorithm was shown to be competitive with state-of-the-art methods for fast bilateral and nonlocal means filtering of color images.

### APPENDIX

#### A. Derivation of (15) and (16)

We recall that the basis of the approximation is the following: For  $\mathbf{x} = \mathbf{p}(\mathbf{i} - \mathbf{j})$ ,  $\mathbf{j} \in W$ , we replace  $\varphi(\mathbf{x} - \mathbf{p}(\mathbf{i}))$  in (1) and (2) with  $\sum_{k=1}^K c_k(\mathbf{i})\varphi(\mathbf{x} - \boldsymbol{\mu}_k)$ . That is, we approximate the numerator of (1) with

$$\sum_{\mathbf{j} \in W} \omega(\mathbf{j}) \left( \sum_{k=1}^K c_k(\mathbf{i}) \varphi(\mathbf{p}(\mathbf{i} - \mathbf{j}) - \boldsymbol{\mu}_k) \right) \mathbf{f}(\mathbf{i} - \mathbf{j}), \quad (22)$$

$$\text{and (2) with } \sum_{\mathbf{j} \in W} \omega(\mathbf{j}) \left( \sum_{k=1}^K c_k(\mathbf{i}) \varphi(\mathbf{p}(\mathbf{i} - \mathbf{j}) - \boldsymbol{\mu}_k) \right). \quad (23)$$

Exchanging the sums, we can write (22) as

$$\sum_{k=1}^K c_k(\mathbf{i}) \left( \sum_{\mathbf{j} \in W} \omega(\mathbf{j}) \varphi(\mathbf{p}(\mathbf{i} - \mathbf{j}) - \boldsymbol{\mu}_k) \mathbf{f}(\mathbf{i} - \mathbf{j}) \right) = \sum_{k=1}^K c_k(\mathbf{i}) \mathbf{v}_k(\mathbf{i}),$$

where  $\mathbf{v}_k$  is defined in (13). Similarly, we can write (23) as  $\sum_{k=1}^K c_k(\mathbf{i}) r_k(\mathbf{i})$ , where  $r_k$  is defined in (14). This completes the derivation of (15) and (16).

#### B. Proof of Theorem 1

For notational convenience, let  $\mathbf{g}(\mathbf{i}) = \boldsymbol{\xi}(\mathbf{i})/\eta(\mathbf{i})$ , where

$$\boldsymbol{\xi}(\mathbf{i}) = \sum_{\mathbf{j} \in W} \omega(\mathbf{j}) \varphi(\mathbf{p}(\mathbf{i} - \mathbf{j}) - \mathbf{p}(\mathbf{i})) \mathbf{f}(\mathbf{i} - \mathbf{j}),$$

and  $\eta$  is as defined in (2). For some fixed  $\mathbf{i} \in \Omega$ , assume that  $\mathbf{p}(\mathbf{i}) \in \mathcal{C}_s$  where  $1 \leq s \leq K$ . Then, following (18),

$$\hat{\mathbf{g}}(\mathbf{i}) = \frac{\mathbf{v}_s(\mathbf{i})}{r_s(\mathbf{i})}. \quad (24)$$

Note that we can write

$$\hat{\mathbf{g}}(\mathbf{i}) - \mathbf{g}(\mathbf{i}) = \frac{1}{\eta(\mathbf{i})} \left( \hat{\mathbf{g}}(\mathbf{i}) (\eta(\mathbf{i}) - r_s(\mathbf{i})) + (\mathbf{v}_s(\mathbf{i}) - \boldsymbol{\xi}(\mathbf{i})) \right).$$

By triangle inequality, we can bound  $\|\hat{\mathbf{g}}(\mathbf{i}) - \mathbf{g}(\mathbf{i})\|$  using

$$\frac{1}{|\eta(\mathbf{i})|} \left( \sqrt{nR} |\eta(\mathbf{i}) - r_s(\mathbf{i})| + \|\mathbf{v}_s(\mathbf{i}) - \boldsymbol{\xi}(\mathbf{i})\| \right), \quad (25)$$

where we have used the fact that  $\hat{\mathbf{g}}(\mathbf{i}) \in [0, R]^n$ . This follows from (13), (14), and (24). Now

$$\mathbf{v}_s(\mathbf{i}) - \boldsymbol{\xi}(\mathbf{i}) = \sum_{\mathbf{j} \in W} \omega(\mathbf{j}) \delta(\mathbf{i}, \mathbf{j}) \mathbf{f}(\mathbf{i} - \mathbf{j}),$$

where  $\delta(\mathbf{i}, \mathbf{j}) = \varphi(\mathbf{p}(\mathbf{i} - \mathbf{j}) - \mathbf{p}(\mathbf{i})) - \varphi(\mathbf{p}(\mathbf{i} - \mathbf{j}) - \boldsymbol{\mu}_s)$ .

By Lipschitz property,  $\|\delta(\mathbf{i}, \mathbf{j})\| \leq L \|\mathbf{p}(\mathbf{i}) - \boldsymbol{\mu}_s\|$ .

Note that we can assume that each  $\omega(\mathbf{j})$  is in  $[0, 1]$ . This is simply because the weights appear both in the numerator and denominator in (1) and (24). Moreover, since the range of  $\mathbf{f}$  is  $[0, R]^n$ , using triangle inequality again, we obtain

$$\|\mathbf{v}_s(\mathbf{i}) - \boldsymbol{\xi}(\mathbf{i})\| \leq L |W| \sqrt{nR} \|\mathbf{p}(\mathbf{i}) - \boldsymbol{\mu}_s\|. \quad (26)$$

Similarly,

$$|\eta(\mathbf{i}) - r_s(\mathbf{i})| \leq L |W| \|\mathbf{p}(\mathbf{i}) - \boldsymbol{\mu}_s\|. \quad (27)$$



Now, since  $\omega$  and  $\varphi$  are non-negative,  $\eta(\mathbf{i}) \geq \omega(\mathbf{0})\varphi(\mathbf{0})$ . Using the fact that  $1/|\eta(\mathbf{i})| \leq 1/\omega(\mathbf{0})\varphi(\mathbf{0})$  along with (25), (26) and (27), we obtain

$$\|\hat{\mathbf{g}}(\mathbf{i}) - \mathbf{g}(\mathbf{i})\| \leq C|W|\sqrt{n}L\|\mathbf{p}(\mathbf{i}) - \boldsymbol{\mu}_s\|, \quad (28)$$

where  $C = 2\sqrt{R}/\omega(\mathbf{0})\varphi(\mathbf{0})$ . On squaring (28) and summing over all pixels, we arrive at (19).

#### ACKNOWLEDGEMENTS

The authors thank the editor and the anonymous reviewers for their comments and suggestions. The authors would also thank the authors of [12], [14], [19], and [44] for distributing their code. The authors would also thank Jingxiang Yang for helping them debug the code of [45].

#### REFERENCES

- [1] P. Perona and J. Malik, "Scale-space and edge detection using anisotropic diffusion," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 12, no. 7, pp. 629–639, Jul. 1990.
- [2] V. Aurich and J. Weule, "Non-linear Gaussian filters performing edge preserving diffusion," in *Proc. Mustererkennung*, 1995, pp. 538–545.
- [3] S. M. Smith and J. M. Brady, "SUSAN—A new approach to low level image processing," *Int. J. Comput. Vis.*, vol. 23, no. 1, pp. 45–78, 1997.
- [4] C. Tomasi and R. Manduchi, "Bilateral filtering for gray and color images," in *Proc. IEEE Int. Conf. Comput. Vis.*, Jan. 1998, pp. 839–846.
- [5] D. Comaniciu and P. Meer, "Mean shift analysis and applications," in *Proc. IEEE Int. Conf. Comput. Vis.*, vol. 2, Sep. 1999, pp. 1197–1203.
- [6] Z. Farbman, D. Lischinski, and R. Szeliski, "Edge-preserving decompositions for multi-scale tone and detail manipulation," *ACM Trans. Graph.*, vol. 27, no. 3, p. 67, Aug. 2008.
- [7] E. S. L. Gastal and M. M. Oliveira, "Domain transform for edge-aware image and video processing," *ACM Trans. Graph.*, vol. 30, no. 4, pp. 69:1–69:12, 2011.
- [8] K. He, J. Sun, and X. Tang, "Guided image filtering," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 6, pp. 1397–1409, Jun. 2013.
- [9] A. Buades, B. Coll, and J.-M. Morel, "A non-local algorithm for image denoising," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, vol. 2, Jun. 2005, pp. 60–65.
- [10] S. Paris, P. Kornprobst, J. Tumblin, and F. Durand, "Bilateral filtering: Theory and applications," *Found. Trends Comput. Graph. Vis.*, vol. 4, no. 1, pp. 1–73, 2009. [Online]. Available: <http://www.nowpublishers.com/article/Details/CGV-020>
- [11] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising with block-matching and 3D filtering," *Proc. SPIE*, vol. 6064, no. 30, p. 606414, Feb. 2006.
- [12] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, "Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising," *IEEE Trans. Image Process.*, vol. 26, no. 7, pp. 3142–3155, Jul. 2017.
- [13] J. Darbon, A. Cunha, T. F. Chan, S. Osher, and G. J. Jensen, "Fast nonlocal filtering applied to electron cryomicroscopy," in *Proc. IEEE Int. Symp. Biomed. Imaging, From Nano to Macro*, May 2008, pp. 1331–1334.
- [14] E. S. L. Gastal and M. M. Oliveira, "Adaptive manifolds for real-time high-dimensional filtering," *ACM Trans. Graph.*, vol. 31, no. 4, pp. 33:1–33:13, 2012.
- [15] J. Wang, Y. Guo, Y. Ying, Y. Liu, and Q. Peng, "Fast non-local algorithm for image denoising," in *Proc. IEEE Int. Conf. Image Process.*, Oct. 2006, pp. 1429–1432.
- [16] J. Chen, S. Paris, and F. Durand, "Real-time edge-aware image processing with the bilateral grid," *ACM Trans. Graph.*, vol. 26, no. 3, p. 103, Jul. 2007.
- [17] A. Adams, J. Baek, and M. A. Davis, "Fast high-dimensional filtering using the permutohedral lattice," *Comput. Graph. Forum*, vol. 29, no. 2, pp. 753–762, 2010.
- [18] A. Adams, N. Gelfand, J. Dolson, and M. Levoy, "Gaussian KD-trees for fast high-dimensional filtering," *ACM Trans. Graph.*, vol. 28, no. 3, p. 21, 2009.
- [19] M. G. Mozerov and J. Van De Weijer, "Global color sparseness and a local statistics prior for fast bilateral filtering," *IEEE Trans. Image Process.*, vol. 24, no. 12, pp. 5842–5853, Dec. 2015.
- [20] S. Paris and F. Durand, "A fast approximation of the bilateral filter using a signal processing approach," in *Proc. Eur. Conf. Comput. Vis.*, 2006, pp. 568–580.
- [21] T. Tasdizen, "Principal neighborhood dictionaries for nonlocal means image denoising," *IEEE Trans. Image Process.*, vol. 18, no. 12, pp. 2649–2660, Dec. 2009.
- [22] K. N. Chaudhury, D. Sage, and M. Unser, "Fast  $O(1)$  bilateral filtering using trigonometric range kernels," *IEEE Trans. Image Process.*, vol. 20, no. 12, pp. 3376–3382, Dec. 2011.
- [23] F. Durand and J. Dorsey, "Fast bilateral filtering for the display of high-dynamic-range images," *ACM Trans. Graph.*, vol. 21, no. 3, pp. 257–266, 2002.
- [24] F. Porikli, "Constant time  $O(1)$  bilateral filtering," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2008, pp. 1–8.
- [25] K. Sugimoto and S.-I. Kamata, "Compressive bilateral filtering," *IEEE Trans. Image Process.*, vol. 24, no. 11, pp. 3357–3369, Nov. 2015.
- [26] Q. Yang, K.-H. Tan, and N. Ahuja, "Real-time  $O(1)$  bilateral filtering," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2009, pp. 557–564.
- [27] Q. Yang, N. Ahuja, and K.-H. Tan, "Constant time median and bilateral filtering," *Int. J. Comput. Vis.*, vol. 112, no. 3, pp. 307–318, 2015.
- [28] K. N. Chaudhury and S. D. Dabhade, "Fast and provably accurate bilateral filtering," *IEEE Trans. Image Process.*, vol. 25, no. 6, pp. 2519–2528, Jun. 2016.
- [29] S. Ghosh and K. N. Chaudhury, "Fast bilateral filtering of vector-valued images," in *Proc. IEEE Int. Conf. Image Process.*, Sep. 2016, pp. 1823–1827.
- [30] S. Ghosh and K. N. Chaudhury, "Color bilateral filtering using stratified Fourier sampling," in *Proc. IEEE Global Conf. Signal Inf. Process.*, Nov. 2018.
- [31] K. N. Chaudhury, "Acceleration of the shiftable  $O(1)$  algorithm for bilateral filtering and nonlocal means," *IEEE Trans. Image Process.*, vol. 22, no. 4, pp. 1291–1300, Apr. 2013.
- [32] P. Nair and K. N. Chaudhury, "Fast high-dimensional filtering using clustering," in *Proc. IEEE Int. Conf. Image Process.*, Sep. 2017, pp. 240–244.
- [33] S. Zhuo, X. Zhang, X. Miao, and T. Sim, "Enhancing low light images using near infrared flash images," in *Proc. IEEE Int. Conf. Image Process.*, Sep. 2010, pp. 2537–2540.
- [34] Q. Yuan, L. Zhang, and H. Shen, "Hyperspectral image denoising employing a spectral-spatial adaptive total variation model," *IEEE Trans. Geosci. Remote Sens.*, vol. 50, no. 10, pp. 3660–3677, Oct. 2012.
- [35] M. A. Westenberg and T. Ertl, "Denoising 2-D vector fields by vector wavelet thresholding," *J. WSCG*, vol. 13, no. 1, pp. 33–40, 2005.
- [36] Q. Yang, "Hardware-efficient bilateral filtering for stereo matching," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 36, no. 5, pp. 1026–1032, May 2014.
- [37] D. H. Foster, K. Amano, S. M. C. Nascimento, and M. J. Foster, "Frequency of metamerism in natural scenes," *J. Opt. Soc. Amer. A*, vol. 23, no. 12, pp. 2359–2372, 2006.
- [38] P.-N. Tan, M. Steinbach, and V. Kumar, *Introduction to Data Mining*. White Plains, NY, USA: Longman, 2005.
- [39] R. Deriche, "Recursively implementing the Gaussian and its derivatives," INRIA, Rocquencourt, France, Res. Rep. RR-1893, 1993.
- [40] I. T. Young and L. J. Van Vliet, "Recursive implementation of the Gaussian filter," *Signal Process.*, vol. 44, no. 2, pp. 139–151, 1995.
- [41] *Hyperspectral Image Database*. Accessed: Aug. 1, 2018. [Online]. Available: <http://lesun.weebly.com/hyperspectral-data-set.html>
- [42] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.
- [43] X.-W. Yang and X.-B. Lin, "Brightening and denoising lowlight images," in *Proc. Int. Conf. Inf., Commun. Signal Process.*, Dec. 2015, pp. 1–5.
- [44] F. Fan, Y. Ma, C. Li, X. Mei, J. Huang, and J. Ma, "Hyperspectral image denoising with superpixel segmentation and low-rank representation," *Inf. Sci.*, vols. 397–398, pp. 48–68, Aug. 2017.
- [45] Y.-Q. Zhao and J. Yang, "Hyperspectral image denoising via sparse representation and low-rank constraint," *IEEE Trans. Geosci. Remote Sens.*, vol. 53, no. 1, pp. 296–308, Jan. 2015.

**Authors'**, photographs and biographies not available at the time of publication.