A multichannel block-matching denoising algorithm for spectral photon-counting CT images

Adam P. Harrison and Ziyue Xu[a)]
Center for Infectious Disease Imaging, Department of Radiology and Imaging Sciences, Clinical Center, National Institutes of Health, 10 Center Dr, Bethesda, MD 20814, USA

Amir Pourmorteza and David A. Bluemke
Department of Radiology and Imaging Sciences, Clinical Center, National Institutes of Health, 10 Center Dr, Bethesda, MD 20814, USA

Daniel J. Mollura
Center for Infectious Disease Imaging, Department of Radiology and Imaging Sciences, Clinical Center, National Institutes of Health, 10 Center Dr, Bethesda, MD 20814, USA

(Received 5 December 2016; revised 5 March 2017; accepted for publication 8 March 2017; published 22 May 2017)

Purpose: We present a denoising algorithm designed for a whole-body prototype photon-counting computed tomography (PCCT) scanner with up to 4 energy thresholds and associated energy-binned images.
Methods: Spectral PCCT images can exhibit low signal to noise ratios (SNRs) due to the limited photon counts in each simultaneously-acquired energy bin. To help address this, our denoising method exploits the correlation and exact alignment between energy bins, adapting the highly-effective block-matching 3D (BM3D) denoising algorithm for PCCT. The original single-channel BM3D algorithm operates patch-by-patch. For each small patch in the image, a patch grouping action collects similar patches from the rest of the image, which are then collaboratively filtered together. The resulting performance hinges on accurate patch grouping. Our improved multi-channel version, called BM3D_PCCT, incorporates two improvements. First, BM3D_PCCT uses a more accurate shared patch grouping based on the image reconstructed from photons detected in all 4 energy bins. Second, BM3D_PCCT performs a cross-channel decorrelation, adding a further dimension to the collaborative filtering process. These two improvements produce a more effective algorithm for PCCT denoising.
Results: Preliminary results compare BM3D_PCCT against BM3D_Naive, which denoises each energy bin independently. Experiments use a three-contrast PCCT image of a canine abdomen. Within five regions of interest, selected from paraspinal muscle, liver, and visceral fat, BM3D_PCCT reduces the noise standard deviation by 65.0%, compared to 40.4% for BM3D_Naive. Attenuation values of the contrast agents in calibration vials also cluster much tighter to their respective lines of best fit. Mean angular differences (in degrees) for the original, BM3D_Naive, and BM3D_PCCT images, respectively, were 15.61, 7.34, and 4.45 (iodine); 12.17, 7.17, and 4.39 (galodinium); and 12.86, 6.33, and 3.96 (bismuth).
Conclusion: We outline a multi-channel denoising algorithm tailored for spectral PCCT images, demonstrating improved performance over an independent, yet state-of-the-art, single-channel approach. [https://doi.org/10.1002/mp.12225]

Key words: block-matching 3D, collaborative filtering, denoising, multi-channel, photon-counting computed tomography

1. INTRODUCTION

Spectral computed tomography (CT) can allow the differentiation of multiple contrast agents in the body, enabling analyses that are not possible with single-energy imaging. Photon-counting computed tomography (PCCT) represents an exciting spectral-CT development, with the ability to achieve higher energy resolution, decreased spectral overlap, and superior dose efficiency.[1]

One of PCCT's advantages is its multiple energy windows, which, in turn, allow greater numbers of materials to be decomposed than using dual-energy CT. For instance, Fig. 1 depicts attenuation responses of three contrast-enhancing materials. Also shown are four thresholds, which can be used to produce four image channels, such as those in Fig. 2. With such four image channels, all three of the materials in Fig. 1 could be decomposed.

However, assuming constant radiation dose, more energy windows leads to less photons per image, producing higher noise levels and jeopardizing image quality. Hence, effective denoising is often needed before further analyses, such as material decomposition. Motivated by this need, we present a denoising algorithm for a new human prototype PCCT scanner. Our algorithm takes advantage of the correlation and

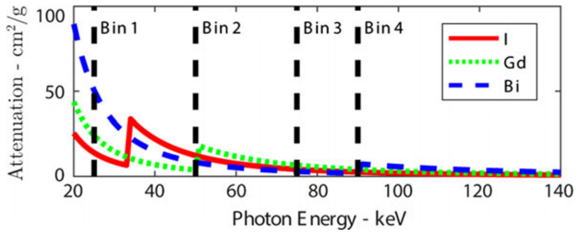

FIG. 1. Attenuation responses of three contrast-enhancing materials and the PCCT thresholds we used to generate four energy bins and their associated images. Attenuation responses obtained from Hubbell and Seltzer.[2] [Color figure can be viewed at wileyonlinelibrary.com]

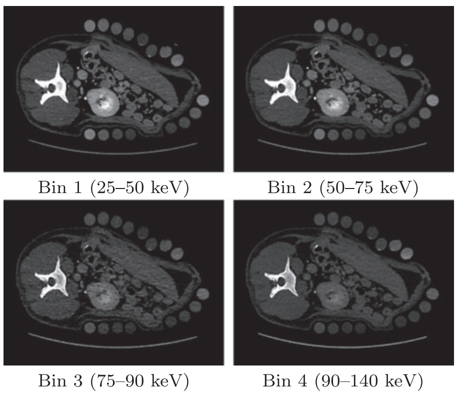

FIG. 2. Four PCCT images of a canine abdomen, imaged with iodine, galodinium, and bismuth contrast agents. Images were generated based on the four energy thresholds depicted in Fig. 1.

perfect alignment between each simultaneously acquired PCCT image, using the image created from all photons to better denoise each individual energy bin. In addition, our algorithm decorrelates attenuation values across PCCT channels, further improving performance.

We test our algorithm on PCCT images of a canine abdomen, qualitatively and quantitatively demonstrating the benefits of our algorithm over a general-purpose, yet otherwise state-of-the-art, approach. We illustrate improvements specific to spectral CT imaging, such as more tightly grouped contrast-agent attenuations.

2. METHODS

To denoise PCCT images, we adapt the block-matching 3D (BM3D) denoising algorithm,[3] which has garnered impressive results when applied to natural images and other modalities. The BM3D algorithm's success is based on its collaborative filtering approach, which we tailor specifically for PCCT images based on two improvements: a shared patch grouping and interchannel decorrelation. We call the resulting algorithm BM3D_PCCT.

2.A. BM3D algorithm

The BM3D algorithm operates patch-by-patch, performing two fundamental operations for each patch location: *patch grouping* and *collaborative filtering*. Focusing on patch grouping first, for each image patch, x_i, centered at pixel i, similar patches are collected based on a distance metric, e.g., the Euclidean distance:

$$d(x_i, x_j) = \frac{\|x_i - x_j\|_2^2}{N^2}, \quad (1)$$

where j denotes the pixel location of a candidate patch, and N represents the patch width and height. In practice, the importance of accurate patch grouping means the BM3D algorithm usually executes (1) on filtered versions of the patches, e.g., on the discrete-cosine transform (DCT) coefficients after hard thresholding. Grouped patches must also meet a minimum similarity. Regardless, for each patch a set of most similar patches are collected, resulting in a 3D stack.

With the 3D stacks collected, each are used as input for the next operation—collaborative filtering. The high correlation, both spatially within each patch and also across the patches, makes it possible to gain a highly sparse representation of the 3D stack using an appropriate linear transform. To accomplish this, BM3D executes a 3D Fourier-like transform, e.g., a combination of DCT and wavelet transforms, upon the 3D stack. This produces many low-magnitude coefficients, which can be set to zero based on a threshold. An inverse transform returns the patches to the spatial domain. After collaborative filtering is performed for each 3D stack, a final aggregation step merges together any overlapping patches.

These steps produce a *basic estimate* of the denoised image. With this in hand, the patch-grouping and collaborative filtering operations are repeated once more to produce a *final estimate*. In this second stage, the original images are denoised again, but the basic estimate is used to compute a more accurate patch grouping. In addition, during the collaborative filtering step, the basic estimate provides a rough estimate of the true coefficient values of each 3D stack after the 3D Fourier-like transform. These coefficient estimates are then used to perform empirical Wiener filtering,[4] replacing the hard thresholding used in the basic estimate stage. This second stage comes with its own transformation, threshold, and distance-metric specifications.

For both stages of the algorithm other hyper-parameters include the size of the patches, their overlap, and the size and range of the patch grouping. Dabov et al. published configurations that have proven effective,[3] which we follow. Interested readers are referred there for more details.

BM3D was originally designed to denoise single-channel images. The simplest way to apply BM3D to PCCT is to denoise each channel independently. We call such an approach BM3D_Naive. However, such an approach fails to take advantage of the correlation and alignment between PCCT channels, hampering performance. We discuss two improvements below, which are designed to overcome this limitation.

2.B. Using a shared patch grouping

An issue in applying BM3D_Naive to PCCT is that the limited photon counts in each channel produces relatively high noise levels, increasing errors in the distance metrics of (1) and, in turn, the entire denoising process. To address this, we use an improvement inspired by color-based BM3D.[5] In this BM3D variant, the luminance channel is used to calculate the patch grouping for it and the other two chrominance channels, based on the assumption that the former is less susceptible to noise than the latter.

An equivalent low-noise channel would also benefit PCCT denoising. Fortunately, this is directly available. In short, as all PCCT images are perfectly aligned, an image created from all the photons in each channel can be easily created. Because this image exhibits a higher signal-to-noise ratio than any individual channel, it provides a more accurate source for calculating patch similarity, which is crucial for BM3D accuracy.[3]

Thus, to denoise PCCT images, we execute a shared patch grouping, which calculates patch similarity for all channels using the following expression:

$$d(x_i, x_j) = \frac{\|z_i - z_j\|_2^2}{N^2}, \tag{2}$$

where $z_{(.)}$ denotes patches from the image created from all the photons, or alternatively linearly transformed versions, as discussed above.

2.C. Decorrelating image channels

BM3D's effectiveness stems from using collaborative filtering to obtain sparse representations of 3D stacks of patches. The spatial correlations within and across each image patch make this possible. Yet, when working with spectral PCCT images, there is also very high interchannel correlation, which should also be exploited.

Hence, our second improvement to BM3D_Naive decorrelates each PCCT channel using an across-channel transform. We choose the DCT transform based on its simplicity, unitary properties, and success in image and audio compression. We also tried the popular decorrelation stretch algorithm,[6] but found DCT to provide superior results. The decorrelated channels are then each denoised using BM3D with the shared patch grouping, producing even greater levels of sparsity for collaborative filtering.

2.D. BM3D_PCCT

The additions of shared patch grouping and interchannel decorrelation are two improvements designed for more effective PCCT denoising. We call the resulting algorithm BM3D_PCCT. Figure 3 depicts a flowchart of the algorithm's operation. The use of an additional all-photon image and the inclusion of DCT interchannel decorrelation differentiates our method from color BM3D.[5] Apart from the BM3D hyperparamters,[3] which we leave unchanged, the final algorithm requires only one parameter specifying the denoising strength based on an estimate of the noise standard deviation. This can be determined through calibration, by using noise estimation techniques, or through empirical testing. For this work, we rely on the latter approach.

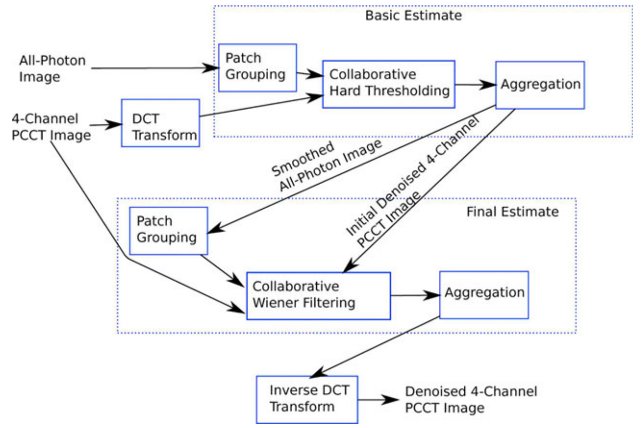

FIG. 3. Flowchart of the BM3D_PCCT algorithm's operation. The algorithm is broken into two main stages: basic and final estimation. Each stage is composed of the patch grouping, collaborative filtering, and aggregation operations. Unlike the original BM3D algorithm, BM3D_PCCT relies on the all-photon image for patch grouping and decorrelates the four PCCT channels using a DCT transform. [Color figure can be viewed at wileyonlinelibrary.com]

2.E. Animal model

We study one adult male mongrel canine (30 kg) in this Institutional Animal Care and Use Committee-approved study (DRD 14-05). The animal model was part of a larger study and was not created solely for the purpose of this experiment. General anesthesia was induced by intramuscular injection of a combination of ketamine (5.5 mg/kg), acepromazine (0.05–0.11 mg/kg), and torbugesic (0.2 mg/kg) with atropine (0.04 mg/kg) as a separate injection and intravenous meloxicam (0.2 mg/kg) for analgesia. After endotracheal intubation, anesthesia was maintained via inhalation with isoflurane (1–2%) with oxygen at a total flow rate of 2 L/min during imaging. Automatic mechanical ventilation with a tidal volume of 15 mL/kg and a breathing frequency of 15 breaths/min was performed with continuous recording of a pulse oximeter and a four-lead electrocardiogram.

2.F. PCCT system and protocol

We test our algorithm using images obtained from a new whole-body prototype PCCT scanner, which is based on a second-generation dual-source CT scanner (SOMATOM Definition Flash, Siemens Healthcare, Forchheim, Germany) where a cadmium telluride-based photon-counting detector replaces one of the conventional energy-integrating detectors. The scanner can produce up to four energy-binned images between 20 and 90 keV at 1 keV increments. Readers are encouraged to consult Kappler et al.[7] for more detailed information.

Imaging was performed after intravenous administration of gadolinium-based contrast (40–60 mL, 3 mL/sec, Dotarem, Guerbet) followed after 3–4 min by iodine-based contrast (20 mL, 3 mL/sec, Isovue 370, Bracco). Images were acquired at the level of the left renal pelvis to visualize the contrast enhancement and excretion in the kidney for both contrast agents. As indicated in Fig. 4, calibration test tubes

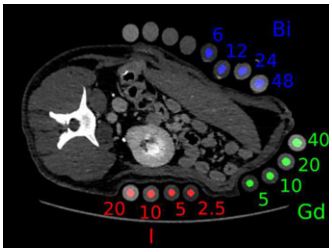

FIG. 4. Abdominal PCCT scan of the subject, with the calibration test tubes, regions of interest used, and their concentrations marked. All units are in mM. The rendered image corresponds to the first energy channel. [Color figure can be viewed at wileyonlinelibrary.com]

of iodine (I), gadolinium (GD), and bismuth (Bi) agents under different concentrations were also included. The PCCT scanner was configured to 140 kVp tube voltage and a 300 mAs tube current with a rotation time of 0.5 s. Energy thresholds correspond to those of Fig. 1, i.e., 25, 50, 75, and 90 keV. Previous phantom-based experiments demonstrated that these thresholds allow for good differentiation of I-, Gd-, and Bi-based contrast agents. Reconstruction was performed using a quantitative soft tissue kernel (D30f), slice thickness of 1 mm, and an increment of 1 mm. Based on these settings, the quality of these PCCT images has been shown to be comparable, if not better, than those produced from conventional energy-integrated detector scanners.[8] Figure 2 depicts the four image channels produced under this setup.

3. RESULTS

Results compare the BM3D_PCCT algorithm against non-denoised images and also against BM3D_Naive. Figure 5 depicts all energy bins, zoomed in on the subject's left kidney, before and after denoising. As can be seen, BM3D_Naive manages to suppress much of the visually apparent noise. However, certain streaks are left, indicated by the arrows. In contrast, these streaks are better corrected by BM3D_PCCT. Importantly, BM3D_PCCT does so while keeping differences in contrast apparent. These qualitative results are supported by improved quantitative measures. For instance, within four regions of interest, drawn from paraspinal muscle, liver, and visceral fat tissue, BM3D_Naive reduced the noise standard deviation by 40.4%, whereas BM3D_PCCT did so by 65.0%. Experiments also compared the pixel response across the four I calibration vials, providing a rough indication of spatial resolution. As Fig. 6 illustrates, spatial resolution seems to remain stable after application of BM3D_Naive and BM3D_PCCT.

In addition to these general measures, it is also important to examine metrics related to spectral imaging, particularly the attenuation responses of contrast-enhancing agents under the concentrations indicated by Fig. 4. Figure 7 plots a scatter plot of the responses. As can be seen, contrast-agent attenuation responses clustered much tighter to their respective lines of best fit after application of BM3D_PCCT compared to BM3D_Naive.

This improvement is quantified in Table I, which tabulates mean angular errors of each material vs. the lines of best fit depicted in Fig. 7. As the table demonstrates, even a naive application of BM3D is able to reduce the angular errors by wide margins. Importantly, BM3D_PCCT reduces the mean angular difference by even greater margins margins, i.e., by 2.89, 2.78, and 2.37 degrees over BM3D_Naive for I, Gd, and Bi, respectively, demonstrating the impact of our modifications. In addition, we also tabulate results when BM3D_PCCT is applied without the DCT step, which, when compared to BM3D_Naive, helps reveal the benefits of using the shared patch grouping. As well, by comparing the former to BM3D_PCCT, the impact of applying the DCT step is also

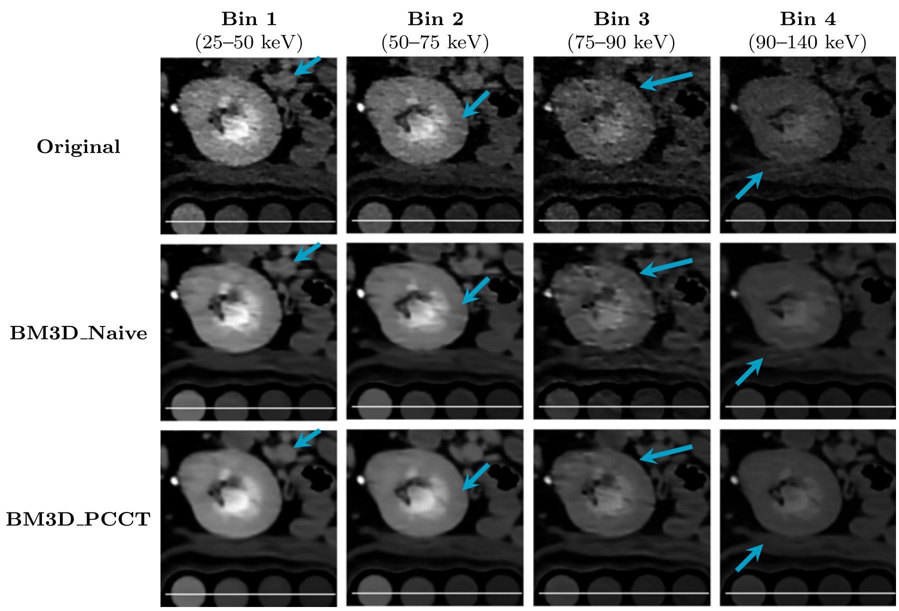

FIG. 5. PCCT images of the subject's left kidney, before and after denoising, corresponding to the four different energy bins. Arrows point to example areas where improvement can be discerned between BM3D_Naive and BM3D_PCCT. The white lines rendered across the I calibration vials illustrate the location of the pixel responses graphed in Fig. 6. [Color figure can be viewed at wileyonlinelibrary.com]

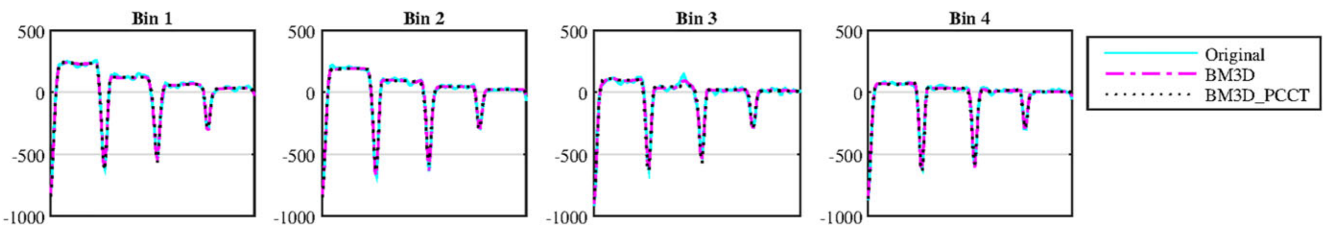

FIG. 6. Pixel responses across the I calibration vials corresponding to the white lines rendered in Fig. 5's images. Best viewed in color. [Color figure can be viewed at wileyonlinelibrary.com]

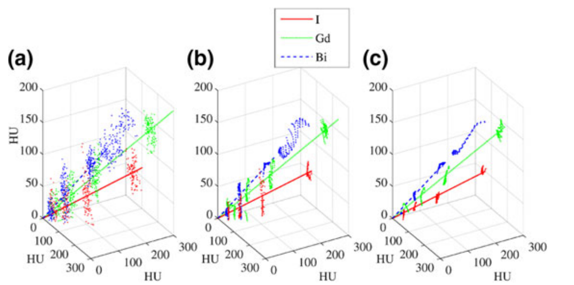

FIG. 7. Contrast-agent pixel responses from the canine abdomen PCCT images. (a)–(c) depicts scatter plots of the original, BM3D_Naive, and BM3D_PCCT images, respectively. Contrast-agent pixel responses are drawn from the first three energy. Lines of best fit are calculated from all four energy bins. Best viewed in color. [Color figure can be viewed at wileyonlinelibrary.com]

TABLE I. Tabulated circular mean of the angles, in degrees, between pixels of the contrasting agents and their lines of best fit, with circular standard deviation in brackets. Note that circular standard deviation is defined differently than its linear counterpart. We use a popular bounded version that falls within $[0, \sqrt{2}]$.[9]

Mean angular errors	Original	BM3D_ Naive	BM3D_ PCCT No DCT	BM3D_PCCT
I:	15.61 (0.26)	7.34 (0.11)	5.86 (0.09)	4.45 (0.07)
Gd:	12.17 (0.20)	7.17 (0.11)	6.22 (0.09)	4.39 (0.06)
Bi:	12.86 (0.23)	6.33 (0.11)	5.03 (0.08)	3.96 (0.07)

measured. As can be seen, both enhancements are crucial toward the overall performance of BM3D_PCCT. Importantly, as Table II illustrates, the direction of the lines of best fit remain very near the original attenuation responses. Angular deviations from the original line of best fit for BM3D_Naive were 0.23, 0.13, and 0.27 degrees for iodine, gadolinium, and bismuth, respectively. Whereas for BM3D_PCCT angular deviations were 0.30, 0.22, and 0.54 degrees for iodine, gadolinium, and bismuth, respectively.

We also measured the mean pixel response and standard deviation of the high-concentration calibration vials across the four energy channels. As Fig. 8 demonstrates, the mean response remains stable after application of BM3D_Naive and BM3D_PCCT. In addition, the standard deviation decreases going from the original images to BM3D_Naive and finally to BM3D_PCCT. These results show promise for subsequent spectral image processing, e.g., material decomposition.

TABLE II. Tabulated directions of the lines of best fit, normalized by the first energy bin.

Material	Original	BM3D_Naive	BM3D_PCCT
I:	(1 0.81 0.42 0.27)	(1 0.81 0.42 0.28)	(1 0.80 0.42 0.27)
Gd:	(1 0.99 0.65 0.51)	(1 1.00 0.65 0.52)	(1 0.99 0.64 0.51)
Bi:	(1 0.80 0.76 0.94)	(1 0.81 0.76 0.95)	(1 0.81 0.75 0.96)

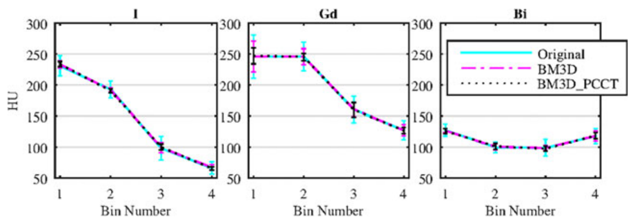

FIG. 8. Attenuation responses of the high-concentration vials of the contrast agents. Mean attenuations of the original, BM3D, and BM3D_PCCT images are graphed together for each contrast agent. Errors bars depict standard deviation. Pixel responses are drawn from the highest concentration vials for I and Gd. Because of settling effects, we use the vial with second highest concentration for Bi. Best viewed in color. [Color figure can be viewed at wileyonlinelibrary.com]

4. DISCUSSION

As our results demonstrate, BM3D_PCCT outperformed a naive application of the BM3D algorithm, making it a promising postprocessing denoising method. This is particularly valuable in the frequent case where raw data are unavailable. Further validation should include comparison to reconstruction-based[10,11] and postprocessing[12] approaches. In particular, the merits of Manhard et al.'s joint bilateral filter approach[12] vs. a BM3D-based one should be quantified, as the former also relies on an all-photon image for guidance.

To expand upon these preliminary findings, future work includes investigating other decorrelation methods, extending BM3D_PCCT to 3D volumes, and using more efficient searching for the patch grouping operations. As well, the hyper-parameter space of the BM3D algorithm should be rigorously explored to determine the optimal PCCT configuration. Another focus is determining the energy thresholds that best balance noise and contrast levels post-BM3D_PCCT. As well, future work should also consider how to best manage the possibly different noise levels between energy channels, which may require computing a per-channel denoising strength. Finally, the impact of using BM3D_PCCT on material decomposition should be directly characterized. Further phantom and animal studies will help pursue these directions of inquiry.

5. CONCLUSIONS

Compared to the single-spectrum CT modality, PCCT provides a unique and powerful multispectral image of attenuation responses. Nonetheless, assuming the same dosage, each PCCT energy bin has a lower photon count compared to standard CT. Therefore, effective denoising algorithms will play a crucial role for this emerging modality. For instance, when performing tissue decomposition, denoising can help produce improved maps of contrast-agent concentration.

We present a denoising algorithm tailored for PCCT images called BM3D_PCCT, which is an adaptation of the popular BM3D algorithm for natural images. By using information in the least noisy all-photon image to denoise individual energy bins and by decorrelating PCCT image channels, our algorithm is designed to gain superior denoising performance.

When tested on a three-contrast PCCT image, our method demonstrates improved denoising performance and better separation between different materials compared to a naive application of the state-of-the-art BM3D algorithm. Aside from the improved performance, this work also highlights the importance of using algorithms tailored for the multichannel and highly correlated nature of spectral PCCT images.

ACKNOWLEDGMENTS

This research is supported by the Intramural Research Program of the National Institutes of Health.

CONFLICTS OF INTEREST

David A. Bluemke receives research support in the form of a Cooperative Research and Development Agreement with Siemens Healthcare.

[a]Author to whom correspondence should be addressed. Electronic mail: ziyue.xu@gmail.com.

REFERENCES

1. Taguchia K, Iwanczyk JS. Vision 20/20: single photon counting x-ray detectors in medical imaging. *Med Phys*. 2013;40:100901.
2. Hubbell J, Seltzer S. Tables of x-ray mass attenuation coefficients and mass energy-absorption coefficients from 1 keV to 20 MeV for elements Z = 1 to 92 and 48 additional substances of dosimetric interest, Tech. Rep. NISTIR 5632 (National Institute of Standards and Technology; 1996.
3. Dabov K, Foi A, Katkovnik V, Egiazarian K. Image denoising by sparse 3d transform-domain collaborative filtering. *IEEE Trans Image Process*. 2007;16:2080.
4. Ghael SP, Sayeed AM, Baraniuk RG. Improved wavelet denoising via empirical wiener filtering. In: Ghael S, Sayeed AM, Baraniuk RG, eds. *SPIE Technical Conference on Wavelet Applications in Signal Processing, San Diego, CA*. Proceedings of the SPIE. Vol. 3169. Bellingham, WA: SPIE; 1997:389–399.
5. Dabov K, Foi A, Katkovnik V, Egiazarian K. Color image denoising via sparse 3d collaborative filtering with grouping constraint in luminance-chrominance space. In: *2007 IEEE International Conference on Image Processing*, *San Antonio, TX*. Vol. 1, New York, NY: IEEE; 2007: I -313–I -316.
6. Soha JM, Schwartz AA. Multispectral histogram normalization contrast enhancement. In: *Proceedings of the 5th Canadian Symposium on Remote Sensing*. Ottawa, Canada: Canadian Aeronautics and Space Institute; 1978: 86–93.
7. Kappler S, Henning A, Kreisler B, Schoeck F, Stierstorfer K, Flohr T. Photon counting ct at elevated x-ray tube currents: contrast stability, image noise and multi-energy performance. In: Whiting BR, Hoeschen C, eds. *Medical Imaging 2014: Physics of Medical Imaging*, Vol. 9033, Bellingham, WA: SPIE; 2014: 90331C–90331C-8.
8. Pourmorteza A, Symons R, Sandfort V, et al. Abdominal imaging with contrast-enhanced photon-counting ct: first human experience. *Radiology*. 2016;279:239.
9. Berens P. Circstat: A matlab toolbox for circular statistics. *Journal of Statistical Software* 2009;31:1.
10. Rigie DS, Rivire PJL. Joint reconstruction of multi-channel, spectral ct data via constrained total nuclear variation minimization. *Phys Med Biol*. 2015;60:1741.
11. Yu Z, Leng S, Li Z, McCollough CH. Spectral prior image constrained compressed sensing (spectral piccs) for photon-counting computed tomography. *Phys Med Biol*. 2016;61:6707.
12. Manhart M, Fahrig R, Hornegger J, Drfler A, Maier A. Guided noise reduction for spectral ct with energy-selective photon counting detectors. in Proceedings of the Third CT Meeting; 2014.