Efficient Low-Dose CT Denoising by Locally-Consistent Non-Local Means (LC-NLM)

Michael Green[1], Edith M. Marom[2], Nahum Kiryati[1],
Eli Konen[2], and Arnaldo Mayer[2(✉)]

[1] Department of Electrical Engineering, Tel-Aviv University, Tel-Aviv, Israel
green1@mail.tau.ac.il, nk@eng.tau.ac.il
[2] Diagnostic Imaging, Sheba Medical Center, Affiliated to the Sackler School of Medicine,
Tel-Aviv University, Tel-Aviv, Israel
{edith.marom,eli.konen,arnaldo.mayer}@sheba.health.gov.il

Abstract. The never-ending quest for lower radiation exposure is a major challenge to the image quality of advanced CT scans. Post-processing algorithms have been recently proposed to improve low-dose CT denoising after image reconstruction. In this work, a novel algorithm, termed the locally-consistent non-local means (LC-NLM), is proposed for this challenging task. By using a database of high-SNR CT patches to filter noisy pixels while locally enforcing spatial consistency, the proposed algorithm achieves both powerful denoising and preservation of fine image details. The LC-NLM is compared both quantitatively and qualitatively, for synthetic and real noise, to state-of-the-art published algorithms. The highest structural similarity index (SSIM) were achieved by LC-NLM in 8 out of 10 denoised chest CT volumes. Also, the visual appearance of the denoised images was clearly better for the proposed algorithm. The favorable comparison results, together with the computational efficiency of LC-NLM makes it a promising tool for low-dose CT denoising.

1 Introduction

In the last decade, low dose CT scan techniques have been successful at reducing radiation exposure of the patient by tens of percent, narrowing the gap with X-ray radiographs [1]. However, generating diagnostic grade images from very low SNR projections is too demanding for the classical filtered back-projection algorithm [2]. Iterative-reconstruction algorithms [3] have been successful at the task and were implemented by most commercial CT scan developers. These methods, however, are computationally very demanding, generally requiring significantly longer reconstruction times [1] and expensive computing power.

In recent years, post-processing algorithms have been proposed to denoise low-dose CT after the back-projection step [1, 4–8]. In [6], a previous high quality CT scan is registered to the low dose scan and a (non-linearly) filtered difference is computed. Eventually, a denoised scan is obtained by adding the filtered difference to the co-registered high quality scan. The method assumes the availability of an earlier scan and depends on registration accuracy.

In [1], a modified formulation of the block matching and 3-D filtering (BM3D) algorithm [9] is combined with a noise confidence region evaluation (NCRE) method that controls the update of the regularization parameters. Results are shown for a few phantom and human CT slices.

Non-local mean (NLM) filters [10] have also been proposed for low-dose CT denoising [4, 5, 7]. In [5], the weights of the NLM filters are computed between pixels in the noisy scan and an earlier, co-registered, high quality scan. Artificial streak noise is added to the high quality scan to improve weights accuracy. In [4], a dataset of 2-D patches (7×7) is extracted from high quality CT slices belonging to several patients. At each position in the noisy scan, the 40 Euclidean nearest-neighbors patches (hereafter denoted *filtering patches*) of the local 7×7 patch (hereafter denoted *query* patch) are found by exhaustive search in the patch dataset. The denoised value at each position is computed by NLM filtering of these 40 nearest-neighbor patches. While the method does not need previous co-registered scans, it requires the extraction of 40 nearest-neighbor patches for every pixel, at high computational cost. The method shows initial promising results, although it was validated on a very small set of 2-D images. It suggests that small CT patches originating from different scans for the same body area (chest in that case) contain similar patterns.

We further observe that in [4], local consistency in the filtered image is not explicitly enforced: the *filtering patches* are just required to be similar to the query patch, but not to neighboring patches which still have a significant overlap with the query patch. As a result, fine details may be filtered out in the denoised image.

In this work, we propose a computationally efficient algorithm for the denoising of low-dose CT scans after image reconstruction. By using a database of high-SNR patches to filter the pixels while locally enforcing spatial consistency, the proposed algorithm achieves both powerful denoising and preservation of fine image details and structures. The remainder of this paper is organized as follows: The LC-NLM algorithm is detailed in Sect. 2 and validated on synthetic and real noisy data in Sect. 3. A discussion concludes the paper in Sect. 4.

2 Methods

The main steps of the proposed method, shown in Fig. 1 (left), are divided in two major parts: The creation of the patch dataset, which is performed only once, and the denoising of input low-dose (*LD*) scans. Both steps will be described in detail in the following subsections.

2.1 Patch Dataset Creation

The idea is to create a dataset of high SNR patches to approximate visually similar noisy patches in the *LD* scans. It is therefore necessary for the dataset to represent as much as possible the variability of the patches in order to increase the chance to find a good match. However, due to memory and performance limitations, the size of the patch dataset must be limited. Randomly sampling the full dataset is not a good option as most
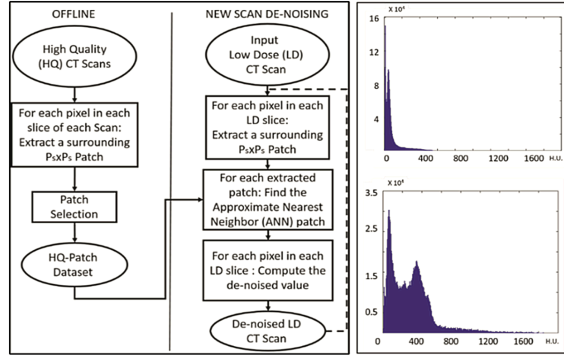
Fig. 1. (left) The main steps of the proposed method; (right) Patch variance histogram for the patch dataset: (top) for random sampling; (bottom) for roulettte wheel selection.

of the patches have fairly constant values as can be seen from the patch standard deviation (SD) histogram (Fig. 1, top-right). We propose a patch selection step in which patches containing more signal variation have a higher probability of being included in the dataset. For this purpose, a roulette wheel (RW) selection scheme [11] is implemented. RW samples uniformly a virtual disk along which patch indices are represented by variable length segments. The higher the patch SD, the longer the segment and accordingly the probability to select it. The resulting SD histogram for the selected patches (Fig. 1, bottom-right) is clearly more balanced than for the random selection case.

2.2 Denoising of Low Dose (*LD*) Scans

Given an input *LD* CT scan, a $P_s \times P_s$ patch centered at each pixel position is extracted in every slice. The approximate nearest neighbor (ANN) is then found for each patch using the state-of-the-art randomized kd-trees algorithm [12] implemented in the Fast
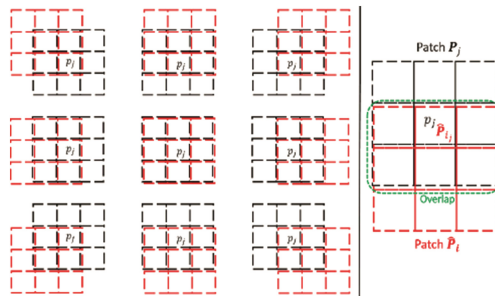

Fig. 2. (left) Pixel p_j can be viewed as the intersection of P_s^2 partially overlapping $P_s \times P_s$ patches (red). In the example $P_s = 3$; (right) patch $\hat{P}_i$ contributes to the denoised value $\hat{p}_j$ via its pixel $\hat{P}_{i_j}$, which the way noisy pixel p_j is perceived by high-SNR patch $\hat{P}_i$. (Color figure online)

Library for ANN (FLANN) [13] for the Euclidean norm. These patches will be desig-
nated as *high-SNR* patches. Any given pixel p_j, distant by at least $\dfrac{P_s}{2}$ pixels from the slice
borders, can be viewed as the intersection of P_s^2 overlapping $P_s \times P_s$ patches. An example
is shown in Fig. 2 (left) for $P_s = 3$.

Let NN_{p_j} be the set of P_s^2 high-SNR patches returned as nearest-neighbor of the P_s^2
overlapping patches. We propose to compute the denoised value $\hat{p}_j$ at pixel p_j as a func-
tion of the patches belonging to N_{p_j} (Eq. 1):

$$\hat{p}_j = F\left(\hat{\boldsymbol{P}}_1, \ldots, \hat{\boldsymbol{P}}_{p_s^2}\right) \ for \ \hat{\boldsymbol{P}}_i \in \left(NN_{p_j}\right) \tag{1}$$

where $F\left(\underline{X}\right):\mathcal{R}^{P_s^2} \to \mathcal{R}$ and $i = 1 \ldots P_s^2$. A simple choice of F may be the average of all
the pixels from the same location as p_j, leading to (Eq. 2):

$$\hat{p}_j = \frac{1}{P_s^2} \sum_{i=1}^{P_s^2} \hat{\boldsymbol{P}}_{i_j} \ for \ \hat{\boldsymbol{P}}_i \in \left(NN_{p_j}\right) \tag{2}$$

where $\hat{\boldsymbol{P}}_{i_j}$ stands for pixel value in patch $\hat{\boldsymbol{P}}_i$ at the location overlapping with pixel
p_j. In this approach, each overlapping patch $\hat{\boldsymbol{P}}_i$, (Fig. 2, right, in red) contributes to the
denoised value $\hat{p}_j$ via its pixel $\hat{\boldsymbol{P}}_{i_j}$, which is the way noisy pixel p_j is perceived by high-
SNR patch $\hat{\boldsymbol{P}}_i$.

A more powerful choice for F, inspired by the non-local means (NLM) filters [10],
is to weight the contributions $\hat{\boldsymbol{P}}_{i_j}$, $i = 1 \ldots P_s^2$ by a similarity measure between $\hat{\boldsymbol{P}}_i$ and the
original noisy patch centered in p_j, formalized by (Eq. 3):

$$\hat{p}_j = \frac{\sum_{i=1}^{P_s^2} \exp\left(-\dfrac{D\left(\boldsymbol{P}_j, \hat{\boldsymbol{P}}_i\right)}{h^2}\right)\hat{\boldsymbol{P}}_{i_j}}{\sum_{i=1}^{P_s^2} \exp\left(-\dfrac{D\left(\boldsymbol{P}_j, \hat{\boldsymbol{P}}_i\right)}{h^2}\right)}, \hat{\boldsymbol{P}}_i \in \left(NN_{p_j}\right) \tag{3}$$

where $h = P_s \cdot \gamma > 0$ is the filtering parameter [10] and γ a constant, $\boldsymbol{P}_j$ is the noisy patch
centered in p_j and $D(\boldsymbol{P}_j, \hat{\boldsymbol{P}}_i)$ is the average L_1 distance between the overlapping pixels of
$\boldsymbol{P}_j$ and $\hat{\boldsymbol{P}}_i$ (Fig. 2, right, in green). Thus, surrounding high-SNR $\hat{\boldsymbol{P}}_i$ patches with high
similarity to $\boldsymbol{P}_j$ at their overlap will contribute more to the denoised value of p_j. Conse-
quently, the proposed method explicitly promotes local-consistency in the denoised

image, resulting in a locally-consistent non-local means (LC-NLM) algorithm. Eventually, to further boost the denoising effect, the LC-NLM can be applied iteratively several times (Fig. 1, left, dashed line) as will be shown in the experiments.

3 Experiments

Quantitative validation on real data requires the availability of high-SNR and *LD* CT scans pairs acquired with exactly the same body position and voxel size for each subject. As these conditions are difficult to meet in practice, we have adopted the pragmatic approach proposed in [4]: *LD* scans were generated by adding artificial noise to real high-SNR CT scans. Zero-mean additive Gaussian noise ($\sigma = 2000$) was added to the sinogram re-computed from the high-SNR slices before returning to the image space.

A dataset was built from 10 high-SNR real full chest scans (voxel size $= 1 \times 1 \times 2.5 \, \text{mm}^3$) and the 10 corresponding synthetic *LD* scans. The patch size, P_s, and the γ constant, were set to 5 and 18, respectively, for all the experiments. A leave-one-out cross validation methodology was implemented: 10 training-test sets were generated, keeping each time another single *LD* scan for testing and using the remaining 9 high-SNR scans to build the high-SNR patch dataset (training set). The training set was generated using the roulette wheel methodology (see Sect. 2) to sample 2 % of the total number of patches in the 9 training scans while preserving inter-patch variability. The structural similarity index (SSIM) [14] was used to quantify the similarity between the denoised *LD* scan and the corresponding high-SNR scan used as ground truth. For a pair of *LD* and high-SNR patches, denoted P_1 and P_2, respectively, the SSIM is given by (Eq. 4):

$$SSIM(P_1, P_2) = \frac{(2\mu_1\mu_2 + c_1)(2\sigma_{12} + c_2)}{(\mu_1^2 + \mu_2^2 + c_1)(\sigma_1^2 + \sigma_2^2 + c_2)} \tag{4}$$

where, μ_1, μ_2 and σ_1, σ_2 are the mean and standard deviation, respectively, for P_1 *and* P_2, and σ_{12} their covariance. We set the constants $c_1 = 6.5$, $c_2 = 58.5$ and quantized the pixels to [0 255], as suggested in [4]. The LC-NLM was compared to three state-of-the-art denoising algorithms: (1) The original NLM [10] as implemented in [1] with default parameters, except for *max_dist*, that was set to 3. (2) Our implementation of the patch-database-NLM (PDB-NLM) [4] with parameters taken from [4]. (3) BM3D [9] with code and parameters from [1, 15], respectively.

In Table 1, the average *SSIM* is given for the LC-NLM (1-4 iterations), NLM and BM3D algorithms for the 10 scans datasets. Due to its very high running time, the PDB-NLM algorithm is compared to LC-NLM for a single representative slice per scan in separate Table 2. In this case, leave-one-out is performed on a total dataset of 10 slice pairs *(LD* and high-SNR), instead of 10 full scans. In 8 out of 10 cases, the best *SSIM* score (bold in Tables 1 and 2) is obtained for the proposed LC-NLM algorithm. In two cases alone (#2 & #8), BM3D reached a higher score than LC-NLM. For LC-NLM, the *SSIM* generally improved for 3–4 iterations before decreasing.

Table 1. Average *SSIM* for 10 full scans: LC-NLM vs. other methods (best score in bold)

case #	LC-NLM, it#1	LC-NLM, it#2	LC-NLM, it#3	LC-NLM, it#4	NLM [10]	BM3D [9]
1	0.614	**0.633**	0.629	0.620	0.463	0.623
2	0.549	0.600	0.619	0.624	0.513	**0.639**
3	0.563	0.601	**0.614**	0.608	0.478	0.610
4	0.571	0.610	**0.627**	0.620	0.490	0.625
5	0.595	**0.605**	0.595	0.581	0.419	0.584
6	0.609	0.62	**0.625**	0.615	0.447	0.618
7	0.589	**0.611**	0.607	0.598	0.453	0.606
8	0.608	0.636	0.639	0.634	0.494	**0.643**
9	0.600	**0.618**	0.614	0.605	0.454	0.607
10	0.589	**0.594**	0.581	0.567	0.410	0.579

Table 2. LC-NLM vs PDB-NLM (PDB-NLM) [4] for the 10 slices dataset (best score in bold)

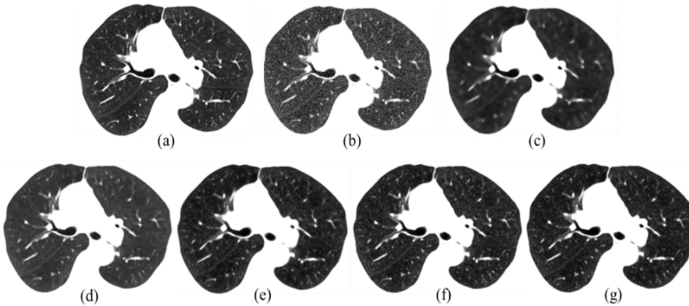| slice #	LC-NLM it#1	LC-NLM it#2	LC-NLM it#3	LC-NLM it#4	PDB-NLM [4]
1	0.571	0.616	0.631	**0.635**	0.620
2	0.616	0.646	**0.650**	0.64	0.638
3	0.565	0.614	0.631	**0.636**	0.620
4	0.587	0.632	0.648	**0.653**	0.637
5	0.639	0.670	**0.677**	0.675	0.667
6	0.639	0.667	**0.669**	0.66	0.661
7	0.591	**0.611**	0.609	0.601	0.600
8	0.587	0.634	0.649	**0.652**	0.638
9	0.527	0.624	**0.657**	0.648	0.645
10	0.572	0.621	0.637	**0.641**	0.628

Fig. 3. A sample slice for case #2: (a) Original image; (b) synthetic low-dose image; (c) NLM [10]; (d) BM3D [9]; (e) PDB-NLM [4]; (f) iter.2 of LC-NLM; (g) iter.3 of LC-NLM.

A sample slice for case #2 is shown in Fig. 3. The original high-SNR slice (a) is shown alongside the corresponding synthetic *LD* image (b), the NLM, BM3D, PDB-NLM results (c–e), and the LC-NLM results for 2 and 3 iterations (f–g). BM3D performed significantly better than NLM, but over-smoothed the fine lung texture while PDB-NLM and LC-NLM better preserved fine details and resulted in more realistic and contrasted images (Table 3).

Table 3. Average running time for the denoising of a CT slice by the compared algorithms.

LC-NLM (1 iter.)	NLM [10]	BM3D [9]	PDB-NLM [4]
~4 s	~7 s	~16 s	~1800 s

In Fig. 4, a pathological free-breathing porcine slice (1 mm thick), acquired at both (a) normal (120 kVp–250 mAs) and (b) low-dose (80 kVp–50 mAs) is shown after denoising the low-dose acquisition by: (c) NLM [10], (d) BM3D [9], (e) PDB-NLM [4], (f) LC-NLM-2iterations, and (g) LC-NLM-3iterations. In the highlighted area (red), a 3 mm part-solid nodule (yellow box, a–g) was further magnified (g). The solid part of the nodule is clearly visible (thick arrows) only in the normal-dose slice (a) and the LC-NLM outputs (f–g). Conversely, a barely distinguishable smeared spot is observed (thin arrow) for BM3D [9] (d), while the nodule has completely disappeared with NLM [10] (c) and PDB-NLM[4] (e), and is undistinguishable from noise in the low-dose slice (b). The dose length product (DLP) was 670.2 and 37.7 mGy·cm for normal and low-dose porcine scans, respectively, reflecting a dose reduction of about 94 % between them.
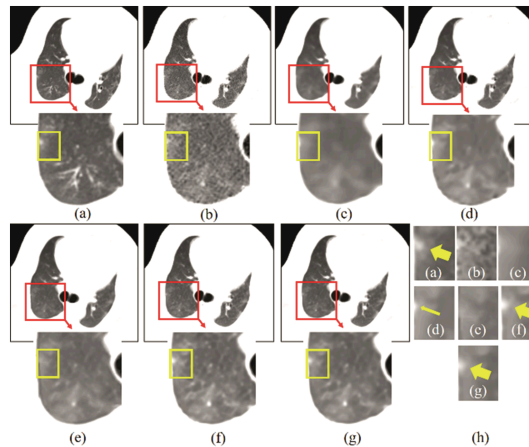

Fig. 4. Porcine CT scan slice acquired at (a) 120 kVp and 250 mAs, and (b) 80 kVp and 50 mAs. Low dose slice denoised by: (c) NLM [10]; (d) BM3D [9]; (e) PDB-NLM [4]; (f) LC-NLM-2iter; (g) LC-NLM-3iter. The highlighted lung area (red) is zoomed-in. (h) Zoomed view of a 3 mm part-solid nodule (yellow boxes) for images a–g. The solid part of the nodule is only clearly visible (thick arrows) in the normal-dose slice (a) and the LC-NLM outputs (f–g). A barely distinguishable smeared spot is observed (thin arrow) for BM3D [9] (d), while it has completely disappeared with NLM [10] (c) and PDB-NLM [4] (e), and is undistinguishable from noise in the low-dose slice (b). (Color figure online)

4   Conclusions

We have presented a computationally efficient algorithm for the denoising of *low dose* (LD) CT scans. The LC-NLM algorithm uses a database of high-SNR *filtering* patches to denoise low-dose scans. By explicitly enforcing similarity between the *filtering* patches and the spatial neighbors of the *query* patch, the LC-NLM better preserves fine structures in the denoised image. The algorithm compared favorably to previous state-of-the-art methods both qualitatively and quantitatively. While PDB-NLM produced better visual results than the other published algorithms (NLM and BM3D), it required prohibitive running times and, also in contrast to LC-NLM, was unable to preserve fine pathological details such as a tiny 3 mm nodule. Moreover, PDB-NLM [4] performs exhaustive search of 40-nearest neighbors. An efficient approximate-NN (ANN) implementation of [4] would strongly affect quality as patch similarity degrades rapidly with the ANN's neighbor order. This problem is avoided with LC-NLM which uses only 1st order ANNs, thus a much better approximation than the 40th. The encouraging results and the computational efficiency of LC-NLM make it a promising tool for significant dose reduction in lung cancer screening of populations at risk without compromising sensitivity. In future research, the patch dataset will be extended to include pathological cases and the method will be further validated on large datasets of real LD-scans in the framework of prospective studies.

References

1. Hashemi, S., Paul, N.S., Beheshti, S., Cobbold, R.S.: Adaptively tuned iterative low dose ct image denoising. Comput. Math. Methods Med. **2015** (2015)
2. Pontana, F., Duhamel, A., Pagniez, J., Flohr, T., Faivre, J.-B., Hachulla, A.-L., Remy, J., Remy-Jardin, M.: Chest computed tomography using iterative reconstruction vs filtered back projection (Part 2): image quality of low-dose CT examinations in 80 patients. Eur. Radiol. **21**(3), 636–643 (2011)
3. Beister, M., Kolditz, D., Kalender, W.A.: Iterative reconstruction methods in X-ray CT. Physica Med. **28**(2), 94–108 (2012)
4. Ha, S., Mueller, K.: Low dose CT image restoration using a database of image patches. Phys. Med. Biol. **60**(2), 869 (2015)
5. Xu, W., Mueller, K.: Efficient low-dose CT artifact mitigation using an artifact-matched prior scan. Med. Phys. **39**(8), 4748–4760 (2012)
6. Yu, H., Zhao, S., Hoffman, E.A., Wang, G.: Ultra-low dose lung CT perfusion regularized by a previous scan. Acad. Radiol. **16**(3), 363–373 (2009)
7. Zhang, H., Ma, J., Wang, J., Liu, Y., Lu, H., Liang, Z.: Statistical image reconstruction for low-dose CT using nonlocal means-based regularization. Comput. Med. Imaging Graph. **38**(6), 423–435 (2014)
8. Ai, D., Yang, J., Fan, J., Cong, W., Wang, Y.: Adaptive tensor-based principal component analysis for low-dose CT image denoising (2015)
9. Dabov, K., Foi, A., Katkovnik, V., Egiazarian, K.: Image denoising by sparse 3-D transform-domain collaborative filtering. IEEE Trans. Image Process. **16**(8), 2080–2095 (2007)

10. Buades, A., Coll, B., Morel, J.M.: A non-local algorithm for image denoising. In: IEEE Computer Society Conference on Computer Vision and Pattern Recognition, CVPR 2005. IEEE (2005)
11. Goldberg, D.E.: Genetic Algorithms in Search Optimization and Machine Learning, vol. 412. Addison-wesley, Reading Menlo Park (1989)
12. Silpa-Anan, C., Hartley, R.: Optimised KD-trees for fast image descriptor matching.In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2008. IEEE (2008)
13. Muja,M., Lowe, D.G.: Fast approximate nearest neighbors with automatic algorithm configuration. In: VISAPP, no. 1 (2009)
14. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. IEEE Trans. Image Process. **13**(4), 600–612 (2004)
15. Lebrun, M.: An analysis and implementation of the BM3D image denoising method. Image Process. Line **2**, 175–213 (2012)