This paper is © IOP and appeared reformatted in *Physics in Medicine and Biology*, 60(2):869-880, 2015
DOI: 10.1088/0031-9155/60/2/869

Low Dose CT Image Restoration Using a Database of Image Patches

Sungsoo Ha[1] **and Klaus Mueller**[1]

[1] Center of Visual Computing, Computer Science Department, Stony Brook University, Stony Brook, New York 11794-4400, and SUNY Korea, Incheon, Korea

E-mail: {sunha, mueller}@cs.sunysb.edu

Abstract. Reducing the radiation dose in CT imaging has become an active research topic and many solutions have been proposed to remove the significant noise and streak artifacts in the reconstructed images. Most of these methods operate within the domain of the image that is subject to restoration. This, however, poses limitations on the extent of filtering possible. We advocate to take into consideration the vast body of external knowledge that exists in the domain of already acquired medical CT images, since after all this is what radiologists do when they examine these low quality images. We can incorporate this knowledge by creating a database of prior scans, either of the same patient or a diverse corpus of different patients, to assist in the restoration process. Our paper follows up on our previous work that used a database of images. Using images, however, is challenging since it requires tedious and error prone registration and alignment. Our new method eliminates these problems by storing a diverse set of small image patches in conjunction with a localized similarity matching scheme. We also empirically show that it is sufficient to store these patches without anatomical tags since their statistics is sufficiently strong to yield good similarity matches from the database, and as direct effect, produce image restorations of high quality. A final experiment demonstrates that our global database approach can recover image features that are difficult to preserve with conventional denoising approaches.

1. Introduction

An important issue in the field of Computed Tomography (CT) today is how to reduce the radiation dose during CT examinations without compromising image quality (Brenner and Hall 2007). The images obtained from *low-dose CT*, which can be achieved either by reducing the X-ray flux, the number of measurements, or both, usually suffer from severe noise artifacts and reduced feature details – all of which impede image readability in diagnostic tasks. Low dose CT has become an active research topic in recent years (see for example, Chen et al. 2008, Sidky and Pan 2008, Jia et al. 2010, Ritschl et al. 2011, Xu et al. 2012) and a visual interface that allows users to evaluate the trade-offs between X-ray dose, image quality and reconstruction speed has been presented by Zheng et al. (2013).

Removal or reduction of the low-dose artifacts is typically achieved by *regularization* which is the introduction of additional constraints or knowledge to steer an ill-posed problem to a plausible solution. One form of regularization is to penalize large local variations via Total Variation Minimization (TVM) (Rudin et al. 1992). Another form is to reduce low-dose noise by averaging a set of pixels that are part of statistically similar neighborhoods, which is the strategy of the Non-Local Means (NLM) filter (Buades et al. 2005). TVM has seen frequent application in low-dose CT reconstruction and it has shown promising results. However, being an iterative optimization approach it bears significant computational cost which can limit its practical application. The NLM filter, on the other hand, has been less popular so far but a number of papers have appeared that demonstrate its good performance in the domain of low-dose CT (Li et al. 2008, Huang et al. 2011, Xu et al. 2013). NLM filtering is non-iterative and highly parallelizable and so can be performed quite expediently (Xu and Mueller 2009). The framework we present makes use of the NLM similarity constraint.

Regularization can be wrapped into an iterative reconstruction framework, such as POCS (Sidky and Pan 2008) or SIR (Zhang et al 2014), or it can follow a single Filtered Backprojection reconstruction as a one-time post processing step (Zheng at al. 2013). The latter leads to the highest reconstruction speed possible, simply by virtue of the number of required operations. We follow this latter paradigm in the work we present here. Future work will then look into the performance of our framework within an iterative reconstruction pipeline.

The NLM filter is a general filter used for a wide range of denoising tasks. It traditionally looks for pixels with statistically similar local neighborhoods in the noisy image itself. But these regions are likely equally contaminated by noise and so a satisfactory regularization may not be found. Xu and Mueller (2012) proposed a solution that expands the search into a set of prior CT images taken from the same patient but at regular dose conditions. This method first simulated the low-dose artifacts in the regular dose prior. It then used these simulated low-dose priors for the match, but employed the corresponding high-quality pixels in the original prior for the update. They showed that this artifact-based matching can significantly improve the quality of the match, leading to improved feature detail in the regularized image. Other researchers in the CT domain also utilized a prior scan of the same patient for high-quality updates (Ma et al 2011, Yu et al 2009) or they constrained a reconstruction by images of the same dynamic scan (Chen at al. 2008). A downside of any of these methods is that they cannot be used when a prior or dynamic scan of the patient is not available.

To meet this inherent shortcoming Xu et al. (2013) proposed the idea of extending the search space even further, namely, to a collection of images of *different* patients. This approach, in fact, is quite alike the psycho-physical processes that occur in medical professionals when viewing degraded imagery. Medical doctors also borrow from their extensive medical training and experience to see the "true patterns behind the noise". In the proposed method the medical knowledge was replaced by (1) a large and diverse database with thousands or more regular dose CT images, (2) a sophisticated matching strategy to locate priors containing structures similar to those in the noisy low-dose target, and (3) a procedure that aligned the matches with the noisy target via non-linear registration (Liu et al. 2011) to constrain the search to anatomically identical regions.

There are, however, two shortcomings with this method: (1) the registration can be error-prone, not only because of the noise artifacts, but also due to the likely presence of high variations between the possible targets and the images stored in the database, and (2) the need to store a large number of full-size images can lead to excessive hardware requirements, such as disk space and management. In order to overcome these challenges we recently presented some preliminary work (Ha et al 2013) that replaced the image database with a database of small 7×7 patches. This has several advantages: (1) it simplifies the restoration pipeline because it eliminates the need for registration, and (2) by storing only patches we can remove a good amount of redundancy and in its place store more diversity.

In the preliminary work just mentioned we added region tags to the stored patches. The region-tagged patches were extracted from relevant anatomical structures using a machine learning approach, and we subsequently called the database the *localized patch-database*. The same machine learning approach was then also used to map the patches to the appropriate anatomical regions in the noisy target. The region-tagging represented a conservative mechanism to explicitly guide the search for patches – to ensure that no patch from another anatomy was used in the regularization of a certain anatomical site in the target.

In this paper, we contrast the results obtained with this localized patch-database with a simpler one in which the region tags are not stored, call it the *global patch-database*. Such a global approach is strongly motivated by the success of the existing work that employed NLM filtering for low-dose CT restoration and also used patches originating from distant image regions (but in the same image). It provided good results because the pixel statistics encoded in the 7×7 patch has sufficiently strong discriminatory power to identify similar edge and contract configurations. The global patch database is desirable since it eliminates the tedious and possibly error-prone region tagging step and it saves us from the need to decide how differentiated the region tagging needs to be – for example, should we distinguish between portions of the heart or is a single region tag for the heart sufficient, and so on. Our paper presents evidence that a global patch database indeed might be sufficient.

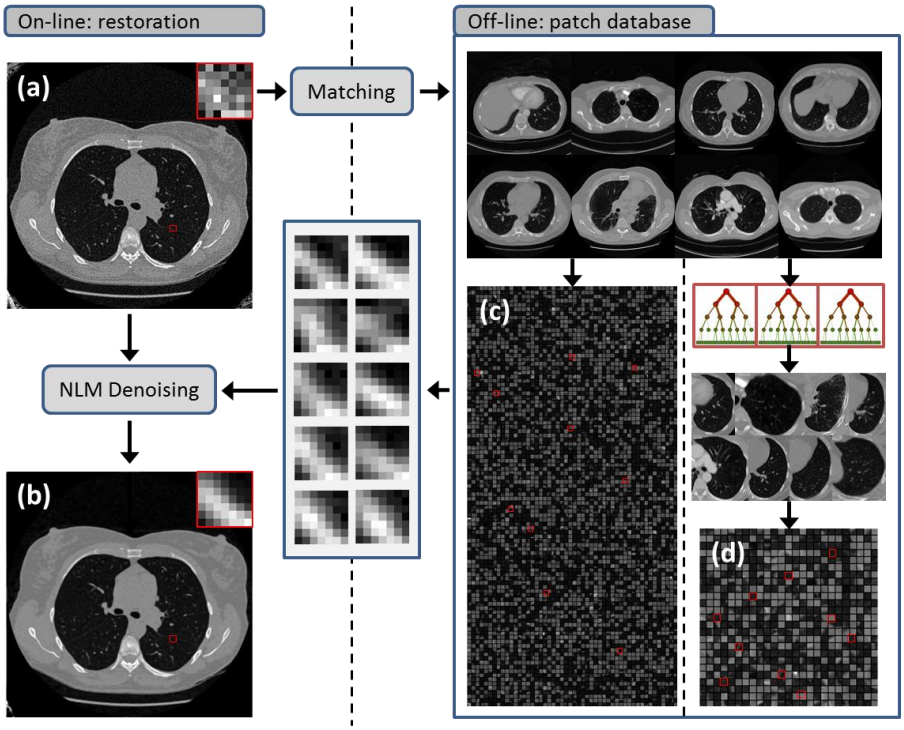

Figure 1. Our low-dose CT image restoration pipeline using the patch database: (a) is the low-dose CT image to be restored into (b) via NLM filtering using either (c) the global patch database or (d) the localized patch database. Both databases are generated from a collection of CT images (top), but the latter adds a region tag using random forest class segmentation.

2. Methods

Figure 1 illustrates the overall pipeline of our proposed patch-based low-dose CT image restoration process. It consists of an online and an offline phase. In the offline phase, the patch database is constructed using a collection of regular-dose CT images taken from the same or different patients. As mentioned, we have studied two types of patch databases – *global* (Figure 1c) and *localized* (Figure 1d) – according to the degree of database segmentation (see below). In the online phase a noisy low-dose CT (target) image is restored using the NLM filtering mechanism in conjunction with either the global or the localized patch database. Specifically, for each pixel p_x its 7×7 patch P_x centered at location x is first extracted and then matched with the regular-dose patches stored in the database to find N similar patches. Then, the NLM filter mechanism is applied to obtain the filtered pixel, p'_x according to this equation:

$$p'_x = \frac{\sum_{y \in S_N} \exp\left(-\sum_{t \in P} \frac{G_a(t)\left|p_{x+t} - p_{y+t}\right|^2}{h^2}\right) \cdot p_y}{\sum_{y \in S_N} \exp\left(-\sum_{t \in P} \frac{G_a(t)\left|p_{x+t} - p_{y+t}\right|^2}{h^2}\right)} \quad (1)$$

Here, S_N is the set of patches retrieved from the database which are similar to the target patch located at x, and y indexes the matched patches with values p_y at their centers. The patch similarity is measured by the Gaussian weighted L_2 distance between two patch vectors with t representing the index within a patch and G_a being a Gaussian kernel with standard deviation a. The factor h controls the strength of the filtering.

2.1 The Localized Patch Database

To create the region tags for the patches in the localized patch database we first process the collection of regular-dose CT images with a localization algorithm to detect and identify their internal structures. In our

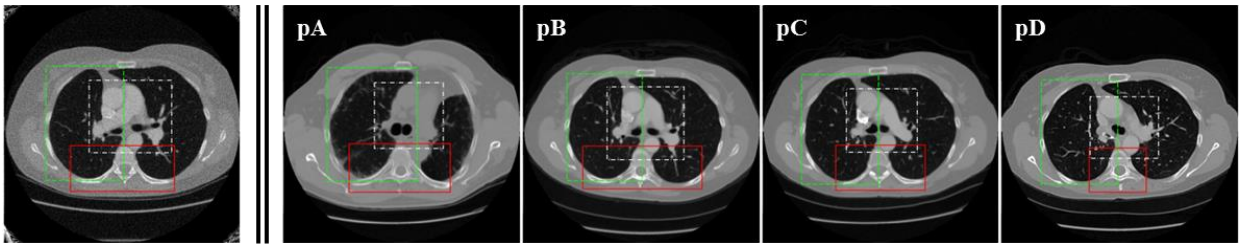

Figure 2. The lung dataset. The tagged anatomical regions are shown annotated by colored rectangular boxes (green: right lung, white: heart, red: spine region). The left-most image is the target low-dose CT image (reconstructed from 720 projections, 40 SNR Gaussian noise), while the other images are regular-dose CT images which we used to construct the global and localized patch database, respectively.

case, we are particularly interested in a localization algorithm that is applicable in a medical setting and for a variety of cases (patients). One of the state-of-the-art algorithms that meet this requirement is the regression forest method proposed by Criminisi et al. (2011). This method captures anatomical structures within rectangular bounding boxes aligned with the Cartesian coordinate system (see Figure 2). Its framework follows the intuitive observation that a set of distinct density clusters (the anatomical structures) can predict the position of another anatomical structure with high confidence. The localized patch database is then constructed by labeling the extracted patches with the corresponding localized structures. In the actual implementation, instead of labeling each patch with the detected structure(s), we bin the patches into sets according to the localized structures contained in them. Finally, the same algorithm is then also used to locate anatomical structures in the target images for the localized patch matching.

2.2 The Global Patch Database

The global patch database, on the other hand, is simply a set of image patches extracted from the collection of normal-dose CT images regardless of their anatomical structures. As mentioned, the existing work that used NLM filtering for low-dose CT restoration also did not make use of local information. This indicates that the feature vectors represented by the patches may have sufficiently strong discriminatory power to identify local structures implicitly. The assumption is further supported by the intuitive observation that each anatomical structure represented in CT imagery has a unique dynamic range and shape. For example, Ye et al. (2009) used local shape and intensity dispersion information for the detection of lung nodules in thoracic CT images. The intensity information has also often been used as a key parameter in segmentation tasks (Pham et al. 2000). And finally, the Hounsfield number (HN) ranges often overlap for pairs of anatomical regions (for example, the HN range of the heart is mostly contained in that of the liver) and therefore patches can be shared if the texture patterns of the different regions are sufficiently similar. In the following section, we provide some empirical evidence that all this seems indeed to be the case.

3. Comparing the localized patch database with the global patch database

To test our database-assisted low-dose CT restoration pipeline, we constructed a small case study with images obtained from the online human lung database (http://www.giveascan.org). We selected the CT images of five patients and matched their quality by re-generating them from 720 projections over 360° with a fan-beam geometry (fan angle = 20°). The 720 projections were the point at which there was no more quality improvement, as gauged by the RMS error with the original CT image.

We built three localized patch databases, one each for the right lung, the heart, and the spine regions using four of the five selected images. The anatomical structures of interest were marked with rectangular bounding boxes using the localization algorithm proposed by Criminisi et al. (2011). Each localized patch database was independent, storing only the patches for the anatomical region it was designed for. The global patch database, on the other hand, was constructed using the full set of patches extracted from the four images. This set contained patches not only from the tagged regions but also from surrounding regions not covered by the boxes shown in Figure 2. The global patch database was hence much larger.

Table 1. Characteristic of patch database and target low-dose CT image

Type	N[a]	Region	Size[c] (%)	Diversity
Localized	4	Heart	9	1.099
		Right Lung	16	1.102
		Spine	8	1.227
Global		All[b]	100	2.226
Target	1	Heart	11	0.134
		Right Lung	16	0.126
		Spine	9	0.140
		All[b]	100	0.345[d]

[a] The number of CT images (512 × 512, 262,144 patches) used for each patch set.
[b] The whole image area is covered.
[c] Normalized by the total number of patches used (1,048,576 and 262,144 for the patch database and target image, respectively).
[d] Only patches within the field of view (FOV) of the target image are considered (the noisy background region is excluded).

After constructing the various databases, we generated a low-dose CT target image by adding 40 SNR Gaussian noise to the sinogram of the 5th image before CT reconstruction. Figure 2 shows the CT image data used in our experiments, along with their anatomical annotations.

We used a patch size of 7×7 and ran the NLM filtering with the set of top 40 matches to the target. These top 40 matches were found based on the Gaussian weighted L_2-norm similarity metric. The two NLM parameters – h (the NLM smoothing parameter) and a (the Gaussian standard deviation) – were experimentally determined by the settings (h=600 and a=1) that gave the best restored image quality in terms of RMSE. We confirmed these finding by subsequent visual inspection. The parameter settings could be further optimized using the approach of Zheng et al. (2013). We chose 40 as the magnitude of the set of top matches since it yielded the best results. It is also close to the number of patches that fall into a typical NLM search window, as used for example by Xu et al. (2013). The 40 matches were identified via exhaustive search, but we accelerated this procedure on the GPU to make it computationally feasible. We currently do not utilize any space decomposition, dimension reduction, or clustering schemes to prune this search.

3.1 Patch Database Statistics

A database is considered *good* if it is of sufficient *size* (that is, it stores an appropriate number of patches) and of sufficient *diversity*. We measure size as the number of patches in a database, while diversity is estimated as the trace of the covariance matrix $\mathbf{\Sigma}$ of the database patches:

$$Diversity = \mathrm{Tr}(\mathbf{\Sigma}) \tag{2}$$

Table 1 shows the statistics of the patch databases we constructed – localized and global – as well as the patch-wise statistics of the target image we used in this study. The size is reported as the percentage amount of the full image. The global patch database covers the entire image and has a size of 100%, while the localized patch databases only cover small image areas, with the right lung being the largest at 10%.

The diversity is more interesting. First, we notice that for each anatomical region the diversity is similar. Further, summing the individual diversity values for the three anatomical regions in the localized database yields a value that is only about 1.5 times greater than the diversity of the global database. This suggests that there is a considerable amount of redundancy and duplication in the localized databases (as is manifested by the overlapping regions) and it also demonstrates the compactness of the global database, at least when it comes to diversity. Another interesting observation is that even though the size of the global patch database is about 6 to 10 times larger than any of the localized patch databases, its diversity is only about twice of these. This suggests that the CT images involved in constructing the databases consist of a fairly limited number of local patterns and this confirms that a global patch database is likely all that is needed. We note that currently the global patch database has not been optimized for size via standard sampling and summarization methods. A large database size can lead to excessive delays in the matching phase as the application is scaled up. Optimizing the patch database in terms of diversity and size is subject of current research.

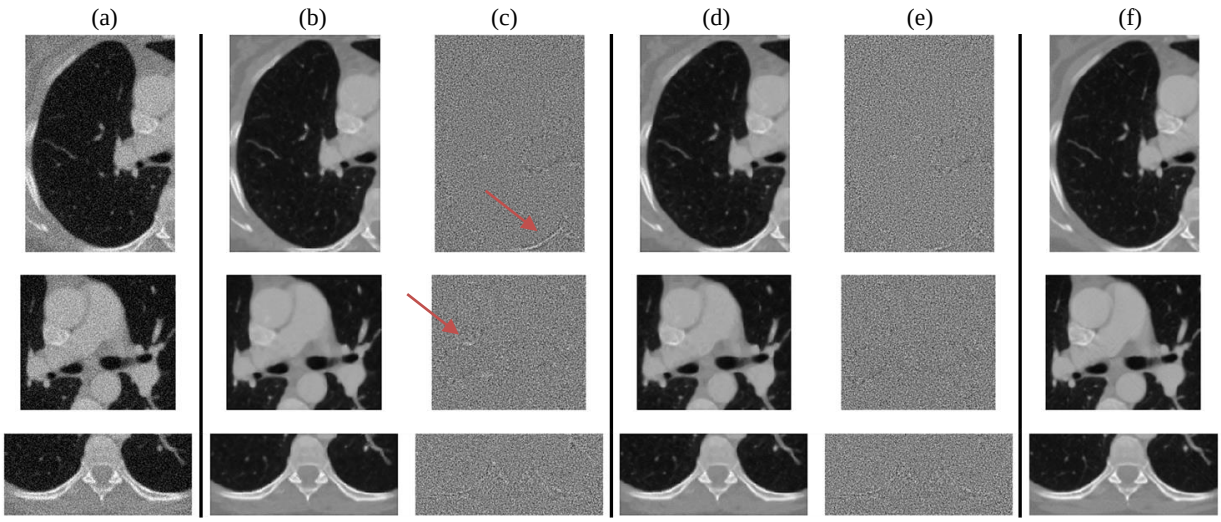

Figure 3. Qualitative comparisons of the restoration results. (a) low-dose (target) image, restored using (b) the localized database and (d) the global database; (c) difference image of local database restoration |(a)-(b)| and (e) difference image of global database restoration |(a)-(d)}; (f) regular dose version of the target.

We make similar observations also for the target image (see lower half of Table 1). In following, we will more closely study what image regions and features the patches are derived from and relate this to the restoration quality that can be obtained.

3.2 Restoration Quality

Figure 3 shows qualitative comparisons of the restored images for both the localized and the global patch databases. Note that we only restored pixels in the three anatomical regions we had trained the localization method to identify – right lung (top row), heart (center row), and spine (bottom row). Although the global patch database could have restored the other pixels as well, we constrained the restoration to these regions only. In Figure 3, column (a) shows the noisy target image, and columns (b) and (d) show the restoration results obtained with the local database and the global database, respectively. Columns (c) and (e) show the respective difference images: |target – restored|. Finally, column (f) shows the regular dose version of the noisy target image as a gold standard for comparison.

We observe that both databases succeed in removing the noise in the target image while preserving small detail to a similar extent. The difference images are fairly structure-less which means that the method was able to well separate the noise from the image features. However, it becomes also evident that the global database removes less structural information than the localized database. The difference image of the localized database shows several small features (indicated by arrows) which have been either completely removed from the restored image or are excessively smoothed. Conversely, these same features are still present or appear sharper in the image restored with the global database. This indicates that the larger patch variety provided by the global database can be an advantage, and it also indicates the need for a sufficiently feature-rich database. While our present database is admittedly relatively small, sizing it up to contain such richness is just a matter of (big) data gathering which is absolutely possible given access to such a trove of data.

3.3 Qualitative Match Location Analysis

For further insight, we have tracked the patches used for restoration in two aspects and for both databases. We have created what we call a *prior map* where we visualize for each target pixel the prior that contained the top ranked patch used for its restoration, and we have created a set of *reference maps*, one for each prior pA, pB, pC, and pD (see Figure 2), where we visualize the locations of these top ranked patches. Both maps can help in determining the role each prior played in the restoration process and they also show that the matching addressed appropriate regions in the priors.

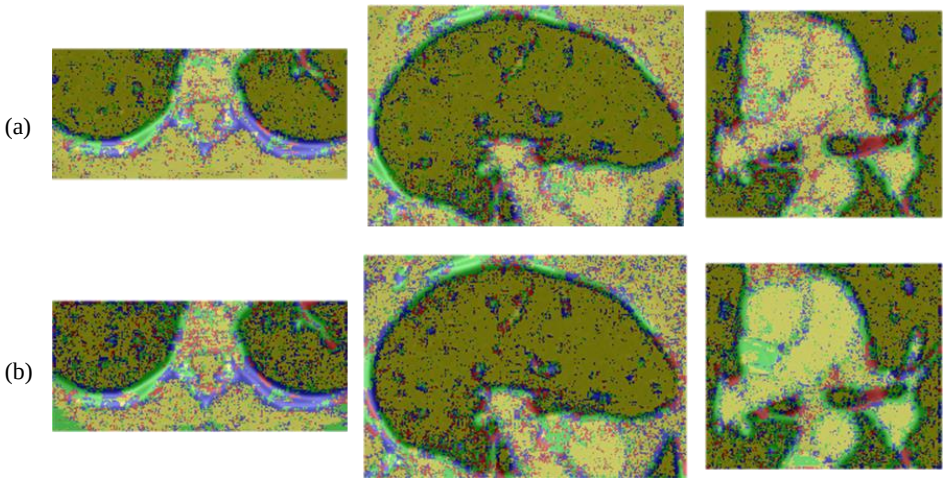

Figure 4. Prior maps for the three anatomical regions we studied (a) the global and (b) the localized patch database. Red, green, blue and yellow colors are used to represent the priors pA, pB, pC and pD, respectively (see Figure 1 for the priors).

In these maps, we only consider the top ranked prior patch for each pixel because in most cases: (1) it has the largest impact on restoration quality, (2) it is most similar to the corresponding gold standard patch, and (3) the next best patches are usually quite close to it.

Figure 4 shows the prior map for each of the three anatomical regions we studied in our test data. We observe that for both databases there is a tendency to borrow prior values from pD for the plain regions like tissues and air (the low-frequency regions), while the other priors (pA, pB, pC) are mostly used to restore detail information, such as bones, vessels and other boundaries of sub-structures (or high-frequency regions) in each anatomical structure. We, however, also notice that the global database leads to a somewhat nosier prior map possibly due to the greater variety of patches.

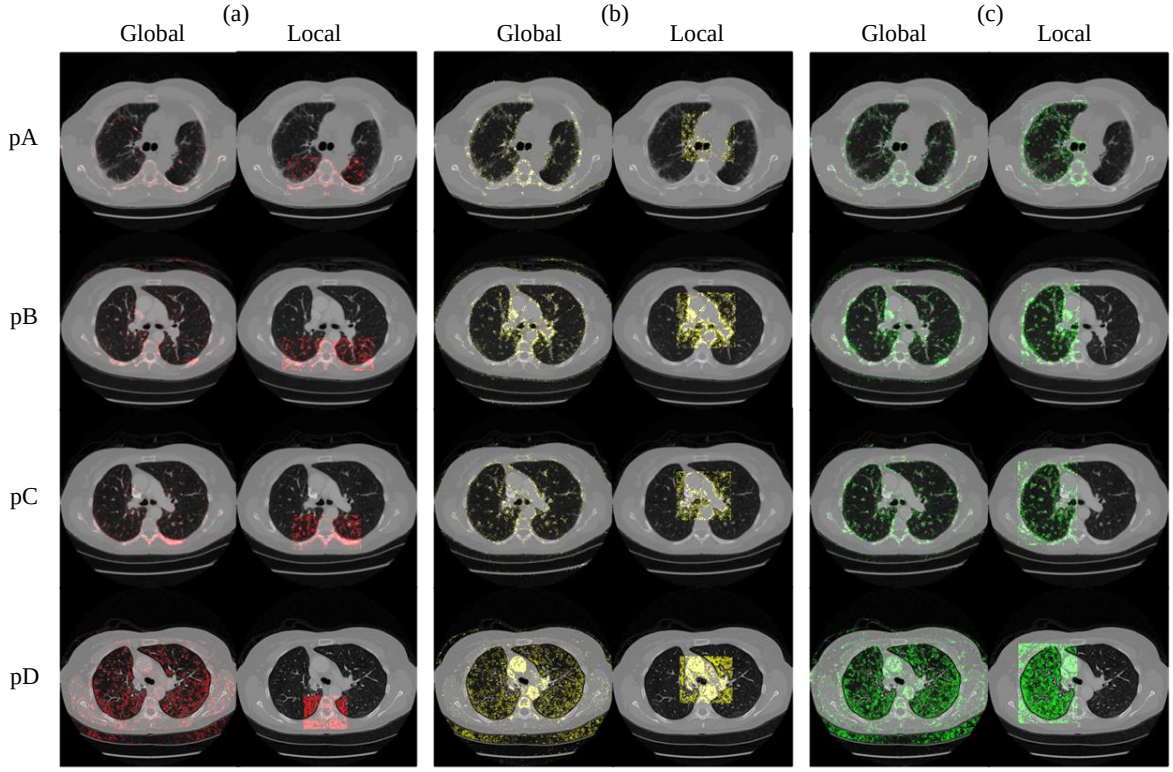

Figure 5. Reference maps for each prior (pA, bP, pC, and pD) for (a) spine (red), (b) heart (yellow) and (c) left lung (green) for the global database (left column) and localized database (right column).

Table 2. Comparisons of top first ranked prior patches used for restoration

Region	Same patch	Different patch					
		Same prior			Different prior		
		Total	HIT	SSIM[b]	Total	HIT	SSIM
Heart	0	19,696 (69.70%)	11,989 (60.87%)[a]	0.89 (0.06)	8,564 (30.30%)	517 (6.04%)	0.85 (0.09)
Right Lung	0	31,206 (74.62%)	15,702 (50.32%)	0.91 (0.05)	10,614 (25.38%)	639 (6.02%)	0.85 (0.09)
Spine	0	14,238 (63.64%)	5,860 (41.16%)	0.82 (0.10)	8,136 (36.36%)	269 (3.31%)	0.82 (0.11)

[a] Percentage is computed over corresponding *Total* number. [b] Mean of SSIM with variance in parentheses

Table 3. Similarity of top first ranked prior patches in SSIM index and L_2-norm

	Heart	Right Lung	Spine
SSIM	0. 93	0.94	0.93
L_2-norm (%)	95.6	95.3	95.3

The reference maps in Figure 5 show the locations of the top ranked patches in the four priors, pA, pB, pC, and pD, for both the global and the localized database. These maps confirm what we have already seen in Figure 4 – the prior pD is used most pre-dominantly to restore the low-frequency regions of the target, most likely because it has a similar dynamic range. Conversely, the remaining three priors may have a similar density change ratio (or gradient magnitude) around the boundaries of the internal structures and are therefore more often used to restore the boundary pixels of the target. We also observe that although the retrieved patches spread more widely in the global database, as opposed to the region-confined distribution in the localized database, the image features they access are similar. This visually confirms that the global database is able to return valid patches for image restoration, but without having the explicit structural knowledge the localized database has.

3.4 Quantitative Match Location Analysis

We have also conducted a more quantitative study of the patches retrieved with the global and localized databases. Specifically, we have compared the top ranked prior patches with regards to: (1) type, (2) origin and match with the intended target region, and (3) similarity with patches retrieved from the localized database. Table 2 shows a summary of our results, to be discussed in the following.

Consistency of patch: We first studied whether the top ranked patches obtained from the global and localized patch databases were the same or different. For this test, we relaxed the meaning of *same patch* to patches that originated from the same prior than the reference patch and were within a radius of 13 pixels from the reference patch location. We found that even with this relaxed criterion none of the top ranked patches retrieved from the global database fit this measure of equality with respect to the localized database. This finding splits Table 2 into two main sections: *same patch* and *different patch*.

Consistency of prior: The second test we ran was to determine whether the top ranked patches extracted from the global database came from the same prior in the localized database. Here we found that this was true for 63-75% of the test patches, depending on anatomy. This breaks the second section of Table 2 into two subsections: *same prior* and *different prior*.

HIT test: The *HIT* test checked for all patches retrieved from the global database how many of these fell within the bounding box of the corresponding region. We found that this was true on average for about 50% of the top ranked patches in the same prior category and for about 5% patches in a different prior category.

3.5 Appearance Analysis

The previous tests checked for location similarity only. But as suggested before, the likelihood to find a similar patch elsewhere in an image is quite high in the domain of CT images. Our ultimate goal is restore a noisy image into one that matches the perceived appearance of a human observer. So it is appropriate to

compare two candidate patches with a suitable perceptual metric, such as the Structural Similarity (SSIM) index defined by Wang et al. (2004). We compute the SSIM index of two patches, P_x and P_y, as follows:

$$SSIM(P_x, P_y) = \frac{(2\mu_x\mu_y + c_1)(2\sigma_{xy} + c_2)}{(\mu_x^2 + \mu_y^2 + c_1)(\sigma_x^2 + \sigma_y^2 + c_2)} \quad (3)$$

Here, μ and σ^2 is the average and the variance of a patch, respectively, and c_1 and c_2 are constant values to stabilize the division for weak denominators. We set the constant values as suggested by Wang et al. (2004) to 6.5 (c_1) and 58.5 (c_2) for CT images normalized to [0 255]. Here SSIM index values of 0.8 and greater are considered acceptable, values exceeding 0.9 are considered excellent, while a value of 1.0 is only reachable when P_x and P_y are identical.

Inter-database test: The first appearance test we conducted was to compare the corresponding patches retrieved from the global and localized databases. The results are reported in the right-most columns of the *same prior* and *different prior* sections in Table 2 where we list the mean SSIM index over all patches in each category along with the standard deviation of each such distribution. The results confirm that most of the patches found from the global database have acceptable to excellent structural similarity with the corresponding patches retrieved from the localized database despite the fact that they originate from different locations either in the same or in a different prior.

Gold standard test: The final and possibly ultimate test is to compare the patches retrieved from the global database with the regular dose version of the target image – the gold standard (Figure 2f). A strong similarity of these pairs would suggest a successful restoration to be taking place since these patches form the basis of this process. For this purpose, we compared all top ranked patches with their corresponding gold standard patches and measured their similarity using the SSIM index and the L_2-norm. Table 3 presents the results of these measurements. We observe that the mean SSIM index is an excellent 0.93 for each of the three regions and the L_2-norm is over 95%, which are very comforting results.

4. Comparing our global database method with other popular image denoising methods

We conclude our paper with a study that compares our global database approach with other denoising methods that have become popular in low dose CT reconstruction – specifically NLM and TVM. To make our study comparable, the NLM and TVM filters were also applied as a post-processing step following a filtered back-projection reconstruction of the same data. For this purpose we selected a new patient from the online human lung database and followed the low-dose image generation procedure described in Section 3. No changes were applied to the database – we used the same global patch database that was employed throughout this study. All parameters of the NLM, TVM, and proposed filters were carefully adjusted to yield the best restoration quality. Our results are presented in Figure 6, windowed to two different grey level ranges (row 1 and 2). The first row (W1) presents the [-1000, 1000] Hounsfield window and allows an assessment of the overall restoration quality, while the second window (W2) is narrower [-1000, 00] and is specifically tuned for the lung region. The third row shows close-up views of the fissures of the lung for the images in row 2.

We observe that all three methods can restore the bone structures well, but the noise was not overly significant there. Other structures, such as tissues, vessels, and tiny (but important) features, show greater differences. Here both NLM and TVM cause blurring – TVM more than NLM – and therefore lose small details in their attempt to remove the noise. They also reduce the sharpness of the edges, such as those of the bony structures. Our global patch database method on the other hand performs significantly better in all of these respects. For example, only our database method can preserve the fissures of the Lungs (annotated with a red arrow), which is an important anatomic landmark in the interpretation of chest CT scans (Hayashi et al. 2001). The restoration of this intricate detail is only possible because our patch database contained a similar local pattern, extracted automatically in the database building process from the regular dose lung images of different patients. The implicit filtering guidance by a template of this feature is the crucial difference that distinguishes our database from the standard NLM and others filtering methods which have no access to such deep anatomical knowledge.

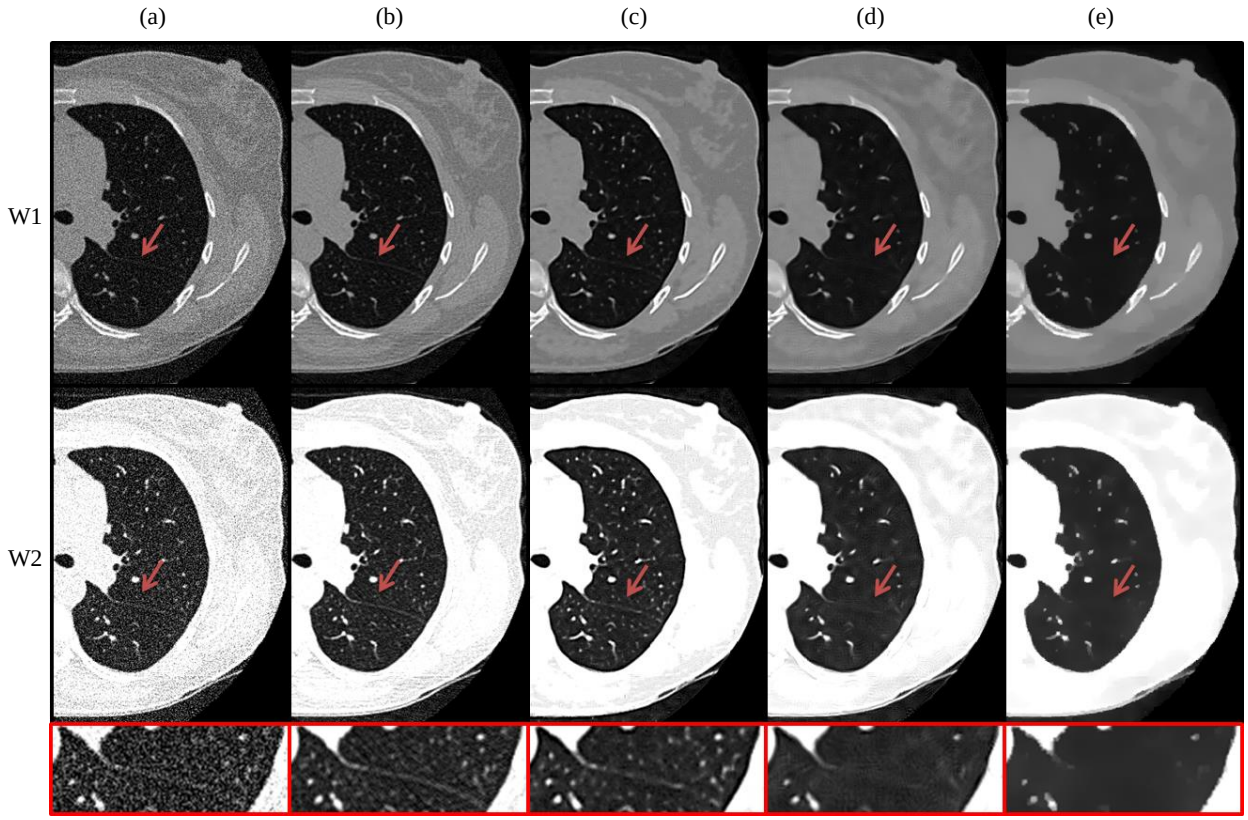

Figure 6: Comparing our method with other popular denoising schemes for a new lung dataset visualized with two Hounsfield windows (row 1: wide [-1000 1000], row 2: lung and [-1000 0]) and row 3: close-up views of the fissure of the Lung in row 2). (a) the low-dose CT image subject to denoising, (b) gold standard image, (c) NLM with our global patch database, (d) standard NLM, (e) TVM.

5. Conclusions

Using a database of priors in conjunction with image filtering for the restoration of noisy low-dose CT images can recover image features that are difficult to preserve with conventional denoising approaches. In this work, we have described a new regularization method that employs a database of 7×7 high-quality CT image patches to remove the noise artifacts that typically appear in low-dose CT images obtained with reduced X-ray flux. The database can be private and only store CT image patches from the patient's prior regular-dose scan, or it can be public and store CT image patches of many regular-dose scans of a diverse patient population. Our method uses similarity matching to retrieve, for each noisy pixel in the target image, a set of suitable patches from this database and then applies them for NLM-type similarity-weighted filtering. Using patches, as opposed to images, simplifies the overall complexity of the restoration pipeline and eliminates the tedious and error-prone image registration that plagued previous methods of this kind. Further, by storing a database of image patches we can remove a great amount of redundancy and in its place store more diversity.

In our studies we have shown that it is sufficient to store the patches without anatomical tags since their statistics is sufficiently strong to yield good similarity matches from the database. Matches of high quality can be obtained this way, which then enable image restorations of equally high quality. Since no classification and segmentation of images is needed our scheme is quite easy to implement.

Having demonstrated the power and validity of our method in principle, future work aims to scale up our system and construct a comprehensive global patch databases with clinical data. This database would eventually cover different sections of anatomy, such as chest, neck, head, thorax, abdomen, brain and so on and also different CT scanners, scan protocols, and reconstruction parameters. Our present study

already included a diversity of scan configurations, albeit only of the lung. Once the database would get larger and extend to more diverse anatomical regions, we might re-introduce some tagging to cull the search space. But in that case the tagging would be at a higher level than in our current local patch database, say to distinguish brain from abdomen. Finding the best strategy will be subject of further study. But in any event, to make such an effort practical, an efficient patch extraction framework will be needed to deal with this large body of data. For this purpose, we will apply techniques from big data management and accelerate them on GPUs.

Acknowledgments
This research was partially supported by NSF grant IIS-1117132 and the MSIP (Ministry of Science, ICT and Future Planning), Korea, under the "IT Consilience Creative Program (ITCCP)" (NIPA-2013-H0203-13-1001) supervised by NIPA (National IT Industry Promotion Agency). We also thank Medtronic, Inc. for their continued support.

References

Brenner D and Hall E 2007 Computed tomography-an increasing source of radiation exposure *New England Journal of Medicine* **357** 2277–2284

Buades A, Coll B and Morel J 2005 A review of image denoising algorithms, with a new one *Multiscale Modeling & Simulation*. **4** 490-530

Chen G, Tang J and Leng S 2008 Prior image constrained compressed sensing (PICCS): A method to accurately reconstruct dynamic CT images from highly undersampled projection data sets *Medical Physics* **35** 660–663

Criminisi A, Shotton J, Robertson D and Konukoglu E 2011 Regression forests for efficient anatomy detection and localization in CT studies *Medical Computer Vision. Recognition Techniques and Applications in Medical Imaging* 106-117

Ha S and Mueller K 2013 Low Dose CT Image Restoration Using a Localized Patch Database 2013 *Nuclear Science Symp. Conf. Record.*

Hayashi K, Aziz A, Ashizawa K, Hayashi H, Nagaoki K and Otsuji H 2001 Radiographic and CT Appearances of the Major Fissures 1 *Radiographics* **21**(4) 861-874

Huang J, Ma J, Liu N, Zhang H, Bian Z, Feng Y, Feng Q, and Chen W 2011 Sparse angular CT reconstruction using non-local means based iterative-correction POCS *Computers in Biology and Medicine* **41** 195–205

Jia X, Lou Y, Li R, Song W, Jiang S 2010 GPU-based fast cone beam CT reconstruction from undersampled and noisy projection data via total variation *Medical Physics* **37** 3441–3447

Li Z, Yu L, Trzasko J, Fletcher J, McCollough C and Manduca A 2012 Adaptive nonlocal means filtering based on local noise level for CT denoising *Proc. SPIE* 8313 83131H

Liu C, Yuen J and Torralba A 2011 SIFT flow: Dense correspondence across scenes and its applications *Pattern Analysis and Machine Intelligence* **33** 978-994

Ma J, Huang J, Feng Q, Zhang H, Lu H, Liang Z and Chen W 2011 Low-dose computed tomography image restoration using previous normal-dose scan *Medical Physics*. **38** 5713-5732

Ye X, Lin X, Dehmeshki J, Slabaugh G and Beddoe G 2009 Shape-Based Computer-Aided Detection of Lung Nodules in Thoracic CT Images *Biomedical Engineering, IEEE Transactions* **56** 1810-1820

Pham D L, Xu C and Prince J L 2000 Current methods in medical image segmentation 1 *Annual review of biomedical engineering* **2** 315-337

Ritschl L, Bergner F, Fleischmann C, and Kachelrieß M 2011 Improved total variation-based CT image reconstruction applied to clinical data *Physics in Medicine and Biology* **56** 1545-1561

Rudin L, Osher S and Fatemi E 1992 Nonlinear Total Variation based noise removal algorithms *Physica D* **60** 259–268

Sidky E and Pan X 2008 Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization *Physics in Medicine and Biology* **53** 4777–4807

Wang Z, Bovik A C, Sheikh H R and Simoncelli E P 2004 Image quality assessment: from error visibility to structural similarity *Image Processing, IEEE Transactions* **13** 600-612

Xu Q, Yu H, Mou X, Zheng L, Hsieh J and Wang G 2012 Low-dose x-ray CT reconstruction via dictionary learning *IEEE Trans. on Medical Imaging* **31** 1682–1697

Xu W and Mueller K 2009 A performance-driven study of regularization methods for GPU-accelerated iterative CT *2nd High Performance Image Reconstruction Workshop* (Beijing, China) pp 20–23

Xu W and Mueller K 2012 Efficient low-dose CT artifact mitigation using an artifact-matched prior scan *Medical Physics*. **39** 4748-4760

Xu W, Ha S and Mueller K 2013 Database-assisted low-dose CT image restoration *Medical Physics* **40** 031109

Xu W, Ha S, Zheng Z and Mueller K 2013 A comparative study of neighborhood filters for artifact reduction in iterative low-dose CT *Fully 3D Image Reconstruction in Radiology and Nuclear Medicine*

Yu H, Zhao S, Hoffman E and Wang G 2009 Ultra-low dose lung CT perfusion regularized by a previous scan *Academic Radiology* **16** 363-373

Zhang H, Han H, Wang J, Ma J, Liu Y, Moore W and Liang Z 2014 Deriving adaptive MRF coefficients from previous normal-dose CT scan for low-dose image reconstruction via penalized weighted least-squares minimization *Medical Physics* **41** 041916

Zheng Z, Papenhausen E and Mueller K 2013 DQS Advisor: A visual interface and knowledge-based system to balance dose, quality, and reconstruction speed in iterative CT reconstruction with application to NLM-regularization *Physics in Medicine and Biology* **58** 7857-73