Neurocomputing 177 (2016) 215–227

Contents lists available at ScienceDirect

Neurocomputing

journal homepage: www.elsevier.com/locate/neucom

Denoising magnetic resonance images using collaborative non-local means

Geng Chen [a,b], Pei Zhang [b], Yafeng Wu [a,**], Dinggang Shen [b,c,**], Pew-Thian Yap [b,*]

[a] Data Processing Center, Northwestern Polytechnical University, Xi'an, China
[b] Department of Radiology and Biomedical Research Imaging Center (BRIC), University of North Carolina at Chapel Hill, NC, USA
[c] Department of Brain and Cognitive Engineering, Korea University, Seoul, Republic of Korea

ARTICLE INFO

Article history:
Received 7 July 2015
Received in revised form
3 September 2015
Accepted 6 November 2015
Communicated by Shuiwang Ji
Available online 3 December 2015

Keywords:
Denoising
Non-local means
Block matching
Non-parametric regression

ABSTRACT

Noise artifacts in magnetic resonance (MR) images increase the complexity of image processing workflows and decrease the reliability of inferences drawn from the images. It is thus often desirable to remove such artifacts beforehand for more robust and effective quantitative analysis. It is important to preserve the integrity of relevant image information while removing noise in MR images. A variety of approaches have been developed for this purpose, and the non-local means (NLM) filter has been shown to be able to achieve state-of-the-art denoising performance. For effective denoising, NLM relies heavily on the existence of repeating structural patterns, which however might not always be present within a single image. This is especially true when one considers the fact that the human brain is complex and contains a lot of unique structures. In this paper we propose to leverage the repeating structures from *multiple* images to *collaboratively* denoise an image. The underlying assumption is that it is more likely to find repeating structures from multiple scans than from a single scan. Specifically, to denoise a target image, multiple images, which may be acquired from different subjects, are spatially aligned to the target image, and an NLM-like block matching is performed on these aligned images with the target image as the reference. This will significantly increase the number of matching structures and thus boost the denoising performance. Experiments on both synthetic and real data show that the proposed approach, collaborative non-local means (CNLM), outperforms the classic NLM and yields results with markedly improved structural details.

1. Introduction

Due to thermal noise, magnetic resonance (MR) images are susceptible to noise artifacts resulting from random fluctuation of the MR signal. Such artifacts cause uncertainty in signal measurements and unreliability in quantitative analyses performed using these images. It is thus critical to denoise these images to improve the robustness and reliability of subsequent analysis.

There are in general two kinds of approaches to noise removal in images. One is the hardware approach [1], which involves scanning an object of interest multiple times and averaging the resulting signals to increase signal-to-noise ratio (SNR). This approach is not always practical due to the long acquisition time. The other is the software approach [2–9], which uses computer algorithms to extract the true signals from noisy measurements. In this work, we focus on the second approach because it can be applied to existing data without requiring expensive equipment upgrades.

Among the large number of algorithms developed for noise removal, a frequently used approach is to attempt to recover the true intensity value of a voxel by averaging the intensity values of neighboring voxels [10]. A popular example is the Gaussian smoothing filter. However, this kind of local averaging technique will remove not only noise but also structural details such as anatomical boundaries. The loss of such details is undesirable due to their potential clinical diagnostic value, such as in characterizing small pathological changes in the brain. To deal with this issue, patch-based approaches have been shown to obtain considerable improvements. Especially notable patch-based methods are the non-local means (NLM) algorithm [11] and the block matching and 3D filtering (BM3D) algorithm [12]. Instead of relying on voxels that are spatially close to each other, the NLM filter averages across (potentially distant) voxels that capture similar structures and thus avoids blurring structural details. The assumption is that real

* Primary corresponding author.
** Secondary corresponding authors.
E-mail addresses: yfwu@nwpu.edu.cn (Y. Wu), dgshen@med.unc.edu (D. Shen), ptyap@med.unc.edu (P.-T. Yap).

http://dx.doi.org/10.1016/j.neucom.2015.11.031
0925-2312/

images often have many self-similar structures that are not necessarily spatially close to each other and that these repeating patterns may be used for effective noise removal. Similar to the NLM filter, the BM3D filter utilizes redundant information distributed throughout the whole image for effective denoising. It arranges similar patches into groups and then carries out denoising by shrinkage of the transform coefficients of the group of patches. We mainly focus on the NLM filter in this work because of its simplicity.

Although the NLM filter has been successfully applied to MR image denoising [2–4], it fails when self-repeating structures cannot be located. A natural solution to this problem is to extend the spatial search range to increase the chance of finding similar structures. However, this will increase computation time dramatically and there is no guarantee of success in finding similar structures. In fact, small weights given to a large number of dissimilar structures will often overwhelm the weights of a few true matching structures. To increase the number of matching structures, Prima and Commowick [3] proposed to capitalize on the bilateral symmetry of the human brain to double the chance of finding matching structures. This is achieved by using information from both ipsilateral and contralateral hemispheres. Despite the promising results, this approach only increases moderately the chance of finding matching structures because only information from one single image is used.

An alternative approach is to borrow information across multiple images for denoising [13]. For example, VBM3D [14] utilizes redundant information found within a frame as well as other frames to carry out efficient video denoising. Foi [15] combines VBM3D and a variance-stabilization approach for multi-image denoising [16]. Note however that all of the above methods require repeatedly acquired images of the same object and are hence not applicable to MRI denoising. Repeated acquisition in MRI increases scan times and is hence clinically prohibitive.

To solve the above problem we propose to harness repeating structures from MR images of different individuals to boost image denoising performance. This is a generalization of the classic NLM filter. The underlying idea is that although human brains differ from each other, they all have many common structures that may be used for effective denoising. To increase the probability of finding matching structures, we spatially align a group of images, called co-denoising images, to a target noisy image and use them to improve denoising. NLM-like block matching is performed to locate matching blocks, not only in the target image itself, but also in the co-denoising images, significantly increasing the number of matching blocks. Such technique has been applied in multi-atlas segmentation [17], but its application in image denoising has not been investigated. Extensive experiments on both simulated and real datasets show that the proposed approach, called collaborative non-local means (CNLM), yields results with markedly improved structural details when compared with the classic NLM filter.

The rest of the paper is organized as follows. In Section 2, we will describe the proposed method. In Section 3, we will then describe the datasets used for evaluation. In Section 4, we will demonstrate the effectiveness of the proposed algorithm for both synthetic and real data. In Section 5, we will provide additional discussion and conclude the paper.

2. Method

2.1. Non-local means filter

We first introduce the classic NLM filter. Let $\mathrm{NL}(u)(x_i)$ be a restored value of a given voxel at location $x_i \in \mathbb{R}^3$. It can be computed as the weighted average of all voxels within a search volume $V(x_i)$, i.e.

$$\mathrm{NL}(u)(x_i) = \sum_{x_j \in V(x_i)} w(x_i, x_j) u(x_j),$$

where $V(x_i)$ is a cubic volume centered at x_i, $u(x_j)$ is the intensity value of the voxel at x_j and $w(x_i, x_j)$ is the weight. The size of $V(x_i)$ is $(2M+1)^3$, where M is a search radius. For structural matching, we define a smaller local cubic neighborhood $N(x_i)$ around x_i. The size of $N(x_i)$ is $(2d+1)^3$, where d is a neighborhood radius. Let $u(N(x_i))$ be a vector which represents the intensity values of all voxels within $N(x_i)$, then $w(x_i, x_j)$ may be defined as a Gaussian function of the Euclidean distance between vectors $u(N(x_i))$ and $u(N(x_j))$ by

$$w(x_i, x_j) = \frac{1}{Z_i} \exp\left\{ -\frac{\| u(N(x_i)) - u(N(x_j)) \|_2^2}{h_i^2} \right\}, \tag{1}$$

where h_i controls the attenuation of the exponential function and Z_i is a normalization constant to ensure that $w(x_i, x_j)$ sums up to one.

If h_i is too large, all voxels tend to have a same weight, leading to a strong smoothing effect. If h_i is too small, only a few very similar voxels will be involved in denoising, and the difference between the denoised image and the original image will be subtle. Coupé et al. [2] suggested to set $h_i = \sqrt{2\beta \hat{\sigma}_i^2 |N(x_i)|}$, where $\hat{\sigma}_i$ is an estimate of the standard deviation of the noise at voxel x_i, β is a constant and is set to 1 [2] and $|N(x_i)|$ is the size of $N(x_i)$. The weight $w(x_i, x_j)$ is required to satisfy $0 \le w(x_i, x_j) \le 1$ and $\sum_{x_i \in V(x_i)} w(x_i, x_j) = 1$. If x_i and x_j are the same, the weight is too large. Hence, $w(x_i, x_i)$ is set according to $w(x_i, x_i) = \max(w(x_i, x_j)), \forall i \neq j$.

2.2. Collaborative NLM denoising

NLM relies on recurring image information. But when the number of similar structures is small, particularly in regions that contain a corner or an edge, one encounters the *rare patch effect* [18–20]. This phenomenon leads to degradation of fine details and often manifests as halos around object boundaries. In the following, we will reformulate NLM to work with images scanned from different subjects to overcome the problem of insufficient structural recurrence. Unlike [2,5,4,3,21], which are focused on locating similar structures within an image, our approach, called collaborative non-local means (CNLM), will allow leveraging of common structures in different scans to improve denoising performance.

Suppose that we have a target noisy image and a group of co-denoising images with indices denoted as set S. Let $V_k(x_i)$ be the search volume centered at x_i in image $k \in S$ and $\hat{w}_k(x_i, x_j)$ be an unnormalized weight, then the CNLM compute the restored value of the voxel at x_i as

$$\mathrm{NL}(u)(x_i) = \frac{\sum_{k \in S} \sum_{x_j \in V_k(x_i)} \hat{w}_k(x_i, x_j) u(x_j)}{\sum_{k \in S} \sum_{x_j \in V_k(x_i)} \hat{w}_k(x_i, x_j)}.$$

If we let

$$\mathrm{NL}_k(u)(x_i) = \frac{\sum_{x_j \in V_k(x_i)} \hat{w}_k(x_i, x_j) u(x_j)}{\sum_{x_j \in V_k(x_i)} \hat{w}_k(x_i, x_j)},$$

and

$$Z_{i_k} = \sum_{x_j \in V_k(x_i)} \hat{w}_k(x_i, x_j),$$

then we have

$$\mathrm{NL}(u)(x_i) = \frac{\sum_{k \in S} Z_{i_k} \mathrm{NL}_k(u)(x_i)}{\sum_{k \in S} Z_{i_k}}.$$

Hence, the restored value given by CNLM is just a weighted