SUPER-RESOLUTION OF 3D MRI CORRUPTED BY HEAVY NOISE WITH THE MEDIAN FILTER TRANSFORM

Karl Thurnhofer-Hemsi[1], *Ezequiel López-Rubio*[1], *Núria Roé-Vellvé*[2], *Lipika Deka*[3]

[1] Department of Computer Languages and Computer Science, University of Málaga
Biomedic Research Institute of Málaga (IBIMA)
[2] Centro de Investigaciones Médico-Sanitarias, General Foundation of the University of Málaga
[3] School of Computer Science and Informatics, De Monfort University

ABSTRACT

The acquisition of 3D MRIs is adversely affected by many degrading factors including low spatial resolution and noise. Image enhancement techniques are commonplace, but there are few proposals that address the increase of the spatial resolution and noise removal at the same time. An algorithm to address this vital need is proposed in this presented work. The proposal tiles the 3D image space into parallelepipeds, so that a median filter is applied in each parallelepiped. The results obtained from several such tilings are then combined by a subsequent median computation. The convergence properties of the proposed method are formally proved. Experimental results with both synthetic and real images demonstrate our approach outperforms its competitors for images with high noise levels. Moreover, it is demonstrated that our algorithm does not generate any hallucinations.

Index Terms— 3D magnetic resonance imaging, median filter, image denoising, single image super-resolution

1. INTRODUCTION

In Magnetic Resonance Imaging (MRI) data is affected by a number of degrading factors, notably the finite spatial resolution and noise. Lack of resolution may imply the loss of details that could be critical for medical diagnosis. Noise heavily affects image quality in some imaging modalities and, in general, it has a negative effect on the contrast and visibility of details that could contain vital information. Given that the acquisition times and protocols are limited for practical reasons, various postprocessing algorithms have been proposed to enhance resolution, and they are continuously being improved.

Interpolation techniques such linear or bicubic interpolation methods have been extensively used [1, 2], but the resulting images are typically blurred versions of the high resolution (HR) goal image. Moreover, these techniques may also introduce ringing and aliasing artifacts. Super-resolution (SR) techniques offer a better approach for resolution enhancement [3, 4, 5, 6]. Single-image SR employs a single low resolution (LR) image, which is typically an ill-posed inverse problem. Nearest neighbor search [7, 8] and sparse representation approaches [9, 10, 11] have been proposed. However, example-based SR highly depends on the database of LR and HR patch pairs. It is also difficult to apply these methods with arbitrary zoom factors, since the model needs to be retrained for each new zoom factor. Self-similarity based methods achieve SR by searching for similar patches in the same image, as done in [12], where the patches are also searched across different scales. However, self-similarity and example-based methods usually assume there is little noise in the images, or respond well only to low levels of noise [11]. Even though [10] has started to address this issue, it still works on a 2D-basis and requires either a low-noise single-image input or a low-noise database. A conventional approach to overcome noise-sensitivity is to apply a denoising algorithm before SR. This is not optimal, since denoising removes details and fine textures, and these effects can then be magnified in the SR step. An iterative combination of the two procedures (denoising and SR) is presented in [13], although it also works in 2D only. In order to avoid the difficulties associated with the use of an external database, there has been proposed to use an additional HR image of the same object instead [14, 15].

The proposed method allows enhancing resolution and removing noise at the same time for 3D images (not 2D slices) without using any additional images. The method is a modified version of the Median Filter Transform [16]. It is based on the computation of the median filter over randomly chosen, parallelepiped-shaped tilings of the 3D space, followed by the computation of the median of the results of these median filters. The algorithm does not make any assumptions about the noise type or the noise level, which facilitates its use for different imaging techniques. No assumptions are made either about the parameters of the imaging system, so the procedure is not affected by errors in the estimations of these parameters [17]. Moreover, arbitrary fractional scales are allowed without the help of an additional interpolation algorithm.

The remainder of this paper is as follows: Section 2 describes the proposed transform, named as MFT3D, and proves some important properties. Section 3 deals with the experiments on synthetic and real images. Finally, Section 4 is devoted to conclusions.

2. THE MFT3D MODEL

In this section the MFT3D is presented. First, the median of over non overlapping parallelepipeds (called *bins*) is computed. The output

This work is partially supported by the Ministry of Economy and Competitiveness of Spain (TIN2016-75097-P and PPIT.UMA.B1.2017). It is also partially supported by the Ministry of Science, Innovation and Universities of Spain (RTI2018-094645-B-I00), and by the Autonomous Government of Andalusia (Spain) under project UMA18-FEDERJA-084. All of them include funds from the European Regional Development Fund (ERDF). The authors thankfully acknowledge the computer resources, technical expertise and assistance provided by the SCBI (Supercomputing and Bioinformatics) center of the University of Málaga. They also gratefully acknowledge the support of NVIDIA Corporation with the donation of two Titan X GPUs. The authors acknowledge the funding from the Universidad de Málaga. Karl Thurnhofer-Hemsi is funded by a Ph.D. scholarship from the Spanish Ministry of Education, Culture and Sport under the FPU program (FPU15/06512).

HR voxels are not necessarily at the center of the bin which is used to compute its value. This procedure is repeated several times with different random template parallelepipeds to tile the space. Thus, a collection of different HR images are obtained and finally the median of each HR voxel over all the images is computed. In this way, the robustness properties of the median are preserved and the resulting HR image features a higher resolution. The MFT3D is defined in Subsection 2.1 and convergence is studied in Subsection 2.2.

2.1. Definition

To compute the MFT3D a median operation by using sample quantiles based on mid-distribution functions is implemented [18]. For discrete distributions, these sample quantiles behave more favorably than the classical sample quantiles, which have no asymptotic normality properties. The formulation and algorithm is specified with more details in [18]. Since we need the sample median, we only compute $\tilde{Q}(0.5)$ and we will call it the *mid-sample median*, where $\tilde{Q}$ is the sample quantile function.

Next, the MFT3D is defined. Let us consider an LR image $f(\mathbf{x})$ with voxels at coordinates $\mathbf{x} = (x_1, x_2, x_3) \in \mathbb{Z}^3$. The pixel coordinates in the HR image are $\mathbf{y} = (y_1, y_2, y_3) \in \mathbb{Z}^3$. The pixel at $\mathbf{x}$ in the LR image is associated to coordinates $\alpha\mathbf{x} = (\alpha x_1, \alpha x_2, \alpha x_3) \in \mathbb{R}^3$, $\alpha \in \mathbb{R}$, $\alpha > 1$, in the HR image, where α is the zoom factor. We define the MFT3D of f as

$$\hat{f}(\mathbf{y}) = \Theta\left(\{\psi(\mathbf{y}, \mathbf{A}_1, \mathbf{b}_1), ..., \psi(\mathbf{y}, \mathbf{A}_H, \mathbf{b}_H)\}\right) \tag{1}$$

$$\forall i \in \{1, ..., H\},\ \psi(\mathbf{y}, \mathbf{A}_i, \mathbf{b}_i) = \Theta\left(\zeta(\mathbf{y}, \mathbf{A}_i, \mathbf{b}_i)\right) \tag{2}$$

where H is a constant number of tilings, $\mathbf{A}_i$ are 3×3 matrices drawn at random from a suitable distribution $p(\mathbf{A})$, $\mathbf{b}_i$ are 3×1 vectors drawn at random from a suitable distribution $p(\mathbf{b})$, ζ is a set of voxel values of the LR image defined as

$$\zeta(\mathbf{y}, \mathbf{A}, \mathbf{b}) = \{f(\mathbf{x}) \mid \mathrm{round}(\mathbf{A}\alpha\mathbf{x} + \mathbf{b}) = \mathrm{round}(\mathbf{A}\mathbf{y} + \mathbf{b})\}, \tag{3}$$

and Θ stands for a median function which can be either the classical sample median or the mid-sample median. In both cases, the quantity $\psi(\mathbf{y}, \mathbf{A}, \mathbf{b})$ is seen as a random variable from which H samples are drawn in (1). The set $\zeta(\mathbf{y}, \mathbf{A}, \mathbf{b})$ contains all the LR voxel values which belong to the parallelepiped where the HR voxel $\mathbf{y}$ belongs, according to the tiling of the space defined by $\mathbf{A}$ and $\mathbf{b}$. The $\mathbf{A}$ matrix is obtained as the product of a rotation matrix $\mathbf{U}$ and a diagonal scaling matrix $\mathbf{\Lambda}$:

$$\mathbf{A} = \mathbf{U}\mathbf{\Lambda}, \tag{4}$$

where $\det(\mathbf{U}) = 1$, $\det(\mathbf{A}) = \det(\mathbf{\Lambda}) = \lambda_1\lambda_2\lambda_3$, and λ_i are the diagonal elements of $\mathbf{\Lambda}$. We define $BinSize$ as the length of the sides of the parallelepiped which defines the bin, measured in pixels in the LR image. Consequently, as the same size for all sides are used:

$$\lambda_1 = \lambda_2 = \lambda_3 = \frac{1}{BinSize}. \tag{5}$$

2.2. Convergence study

This section determines whether the MFT3D converges to any values when the number of random tilings $H \to \infty$, and it is also estimated how close the result is to the original noiseless HR image.

If $\psi(\mathbf{y}, \mathbf{A}, \mathbf{b})$ were a continuous random variable, then a convergence rate $1/H$ would be ensured [19]. But it is a discrete variable with many ties because the values of the LR voxels close to $\mathbf{y}$ are reused many times to obtain each $\psi(\mathbf{y}, \mathbf{A}_i, \mathbf{b}_i)$. Propositions 1 and 2 in [16] show results about the variance of the MFT3D for increasing finite values of H.

Let $\hat{f}_H$ denote the MFT3D computed with H tilings. Let us denote by $z_1, \ldots, z_N$ the possible values of the discrete random variable $\psi(\mathbf{y}, \mathbf{A}, \mathbf{b})$. Also, let us denote

$$q_i = P(\psi(\mathbf{y}, \mathbf{A}, \mathbf{b}) = z_1) \tag{6}$$

where $q_i > 0$, and $\sum_{i=1}^{N} q_i = 1$. If the following condition holds:

$$\exists j \in \{1, ..., N-1\}\ \sum_{i=1}^{j} q_i = \sum_{i=j+1}^{N} q_i = \frac{1}{2} \tag{7}$$

then both z_j and z_{j+1} are medians of $\psi(\mathbf{y}, \mathbf{A}, \mathbf{b})$. This is a degenerate case that is unlikely to occur in practice, since if q_i is chosen at random from any continuous distribution, then the probability that (7) holds is zero. If (7) does not hold, then there is a unique median $m = z_{j+1}$ with $j \in \{0, 1, ..., N-1\}$, where:

$$\left(\sum_{i=1}^{j} q_i < \frac{1}{2}\right) \text{ and } \left(\sum_{i=1}^{j+1} q_i > \frac{1}{2}\right) \tag{8}$$

This second case is the one to be expected in practice. Both cases are studied by Proposition 3 and 4 in [16], that prove that the MFT3D converges to the median of the distribution of the parallelepiped medians $\psi(\mathbf{y}, \mathbf{A}, \mathbf{b})$, which are the median of the values of the voxels inside the parallelepiped defined by $\mathbf{A}$ and $\mathbf{b}$. This median can be unique in the general case or two medians in the degenerate case.

Proposition 1. *The probabilities q_i are continuous functions of $\mathbf{y}$.*

Proof. In order to study the continuity of q_i as a function of $\mathbf{y}$ we must show that $|q_i(\mathbf{y} + \boldsymbol{\epsilon}) - q_i(\mathbf{y})| \to 0$ as $\|\boldsymbol{\epsilon}\| \to 0$. Let us consider that $\mathbf{y}$ and $\mathbf{y} + \boldsymbol{\epsilon}$ belong to different parallelepipeds if and only if $\mathrm{round}(\mathbf{A}\mathbf{y} + \mathbf{b}) \neq \mathrm{round}(\mathbf{A}(\mathbf{y} + \boldsymbol{\epsilon}) + \mathbf{b})$. Since the values of $\psi(\mathbf{y}, \mathbf{A}, \mathbf{b})$ and $\psi(\mathbf{y} + \boldsymbol{\epsilon}, \mathbf{A}, \mathbf{b})$ can only be different if the parallelepipeds are different, it follows that:

$$|q_i(\mathbf{y} + \boldsymbol{\epsilon}) - q_i(\mathbf{y})| \leq P\left(\mathrm{round}(\mathbf{A}\mathbf{y} + \mathbf{b}) \neq \mathrm{round}(\mathbf{A}(\mathbf{y} + \boldsymbol{\epsilon}) + \mathbf{b})\right) \tag{9}$$

From (4) and (5) we find that the norm of $\mathbf{A}$ is upper bounded:

$$\|\mathbf{A}\| = \|\mathbf{U}\mathbf{\Lambda}\| \leq \|\mathbf{U}\|\,\|\mathbf{\Lambda}\| = \|\mathbf{\Lambda}\| = \frac{1}{BinSize} \tag{10}$$

where $\|\cdot\|$ stands for the spectral norm of a matrix and $\|\mathbf{U}\| = 1$ because $\mathbf{U}$ is orthogonal, and the spectral norm of $\mathbf{\Lambda}$ is its largest element because $\mathbf{\Lambda}$ is a diagonal matrix with non-negative entries. Consequently $\|\mathbf{A}\boldsymbol{\epsilon}\|$ is also bounded:

$$\|\mathbf{A}\boldsymbol{\epsilon}\| \leq \|\mathbf{A}\|\,\|\boldsymbol{\epsilon}\| \leq \frac{1}{BinSize}\|\boldsymbol{\epsilon}\| \tag{11}$$

On the other hand:

$$\|(\mathbf{A}(\mathbf{y} + \boldsymbol{\epsilon}) + \mathbf{b}) - (\mathbf{A}\mathbf{y} + \mathbf{b})\| = \|\mathbf{A}\boldsymbol{\epsilon}\| \tag{12}$$

Then it turns out that $\|(\mathbf{A}(\mathbf{y} + \boldsymbol{\epsilon}) + \mathbf{b}) - (\mathbf{A}\mathbf{y} + \mathbf{b})\| \to 0$ as $\|\boldsymbol{\epsilon}\| \to 0$. Since $\mathbf{b}$ is uniformly distributed on $[0, 1]^3$, the fractional part of $\mathbf{A}\mathbf{y} + \mathbf{b}$ is also uniformly distributed on $[0, 1]^3$ irrespective of the values of $\mathbf{A}$ and $\mathbf{y}$. Consequently,

$$P\left(\mathrm{round}(\mathbf{A}\mathbf{y} + \mathbf{b}) \neq \mathrm{round}(\mathbf{A}(\mathbf{y} + \boldsymbol{\epsilon}) + \mathbf{b})\right) \leq 8\,\|(\mathbf{A}(\mathbf{y} + \boldsymbol{\epsilon}) + \mathbf{b}) - (\mathbf{A}\mathbf{y} + \mathbf{b})\| \tag{13}$$

so that the probability in (9) tends to 0 as $\|\boldsymbol{\epsilon}\| \to 0$, which implies that q_i is a continuous function of $\mathbf{y}$. □

From this result we can deduce that the functions of the form $\sum_{i=1}^{j} q_i$ are continuous, since they are sum of continuous functions. The MFT3D defines regions so that inside a region it always converges to the same value as $H \to \infty$. This can be derived from (8) and Proposition 4 in [16]. It means that this method is valid for subvoxel value estimation, because as $H \to \infty$, a subvoxel and all its immediate neighbors lie in the same region and converge to the same value. Lastly, the following proposition proves the convergence in distribution of the MFT3D when the mid-sample median is employed.

Proposition 2. *Let $\hat{f}_H$ be the MFT3D computed for an image $\mathbf{y}$ using the mid-sample median, and consider the distribution of the discrete random variable $\psi(\mathbf{y}, \mathbf{A}, \mathbf{b})$, from which we have H independent observations. Let $z_1 < \cdots < z_N$ be the support and $q_1, \ldots, q_N$ the corresponding probabilities of the distribution. Define $\pi_k = \sum_{i=1}^{k-1} q_i + q_k/2$. Then,*

$$\sqrt{H}(\hat{f}_H - \tilde{f})\frac{1}{2}\frac{q_{k+1}+q_{k+2}}{z_{k+2}-z_{k+1}} \to N(0,\sigma^2) \quad \textit{in distribution as } H \to \infty, \tag{14}$$

$$\sigma^2 = \frac{1-\{1-(\lambda-1)^2\}q_{k+1}-(1-\lambda^2)q_{k+2}}{4}, \tag{15}$$

where $0 < \lambda < 1$ and for some $k \in \{0, \ldots, N-2\}$,

$$0.5 = \lambda\pi_{k+1} + (1-\lambda)\pi_{k+2}, \tag{16}$$

$$\tilde{f} = \lambda z_{k+1} + (1-\lambda)z_{k+2} \tag{17}$$

Proof. It follows immediately from the case 3 of Theorem 2 in [18], where we take $p = 0.5$. The other cases of Theorem 2 in [18] are extreme cases that have a negligible probability of occurrence for real medical images. □

3. EXPERIMENTS AND RESULTS

In this section, the results of our SR experiments are presented. Since there are not many developed SR algorithms for three dimensional MR images, and most of them do not have available code, we have used MR denoising algorithms as the first step, and then we resample the denoised image to the desired resolution using a 3D cubic spline interpolation as implemented in *Matlab*. The following denoising methods were used to compare the proposed method: NLM3D [20, 21, 22], WSM [23], ODCT3D [24], PRI-NLM3D [24], BD4M [25]. In addition, a 3D SR algorithm for MRIs is tested: NLMU [26].

Dimensions of the original images were trimmed to be congruent with the zoom factor. LR images were obtained by downsampling the original HR images using two steps [27]: first, the noisy HR image was filtered using a 3-D Gaussian smoothing kernel with unit standard deviation. Then, noisy LR image of the required size is obtained by cubic spline interpolation of the noisy HR image.

Furthermore, all experiments with our method were carried out using $H = 150$ median filters. Integer zoom factors $\alpha \in \{2, 3, 4\}$ and fractional zoom factors $\alpha \in \{2.5, 3.5\}$ have been considered. None of the competitors except NLMU use the zoom factor because they are denoising methods, while the super-resolution step is done by the bicubic interpolation algorithm.

Two quality measures were used to evaluate the proposed approach: Mean Squared Error (MSE), where a lower value is better, and Structural Similarity Index (SSIM), where higher is better. In addition, a qualitative evaluation is performed by examination of the residual images, i.e. the difference between the input noisy image and the restored image obtained by each method. The reported experiments have been carried out in *Matlab* on a 64-bit PC with an eight-core Intel i7 3.60GHz CPU and 32 GB RAM.

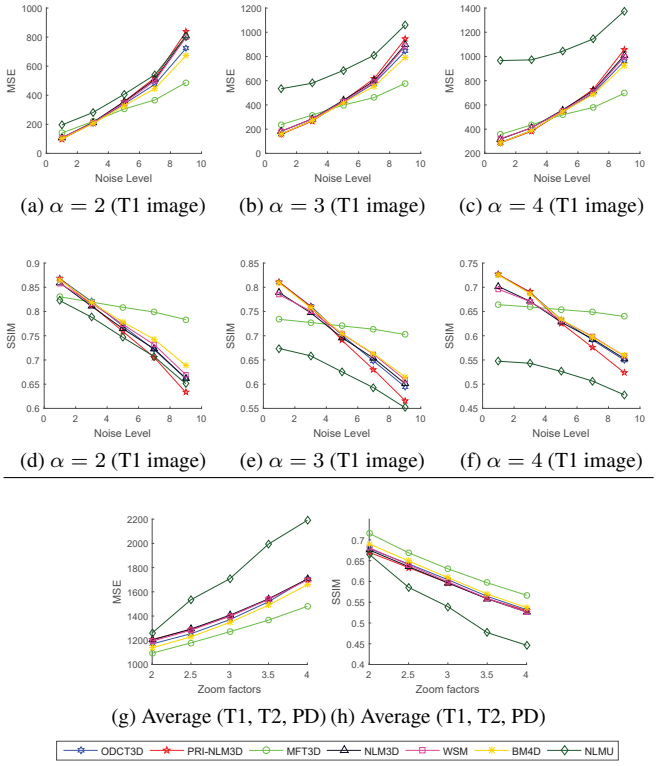

(a) $\alpha = 2$ (T1 image) (b) $\alpha = 3$ (T1 image) (c) $\alpha = 4$ (T1 image)

(d) $\alpha = 2$ (T1 image) (e) $\alpha = 3$ (T1 image) (f) $\alpha = 4$ (T1 image)

(g) Average (T1, T2, PD) (h) Average (T1, T2, PD)

Fig. 1. Quantitative results varying zoom factor and noise levels experiments using the Brainweb database.

3.1. Synthetic data

The freely available open access BrainWeb benchmark database was used to carry on the synthetic experiments [28]. T1, T2 and PD weighted volumes of $181 \times 217 \times 181$ voxels and voxel resolution of $1 \times 1 \times 1 mm^3$ were used, with zero noise and zero non-uniformity. Five Rician noise levels have been tested: $1\%, 3\%, 5\%, 7\%, 9\%$. The noisy images have also been obtained from the Brainweb database.

Quantitative results for the tested integer zoom factors, and noise levels are reported in Figure 1(a-f). As seen, the MFT3D method obtain the best results when moderate and heavy noise is present in the image ($\geq 5\%$). This happens both at the voxel level (MSE) and at the local structure level (SSIM). For lower noise levels our method produces slightly larger errors than the competitors. As expected, all the NLM based filters have a similar behavior and almost the same values, they suffer from its lack of adaptation to the properties of Rician noise. The pure SR competitor, NLMU, obtains the worst results. In Figure 1(g-h) is summarized in two plots the results of the images extracted from the Brainweb dataset. For each zoom factor $\alpha \in \{2, 2.5, 3, 3.5, 4\}$, the average over all the images and noise levels has been computed. MFT3D method clearly overcomes the competitors for all the zoom factors analyzed, both for MSE and SSIM measures, which confirms the better performance of MFT3D.

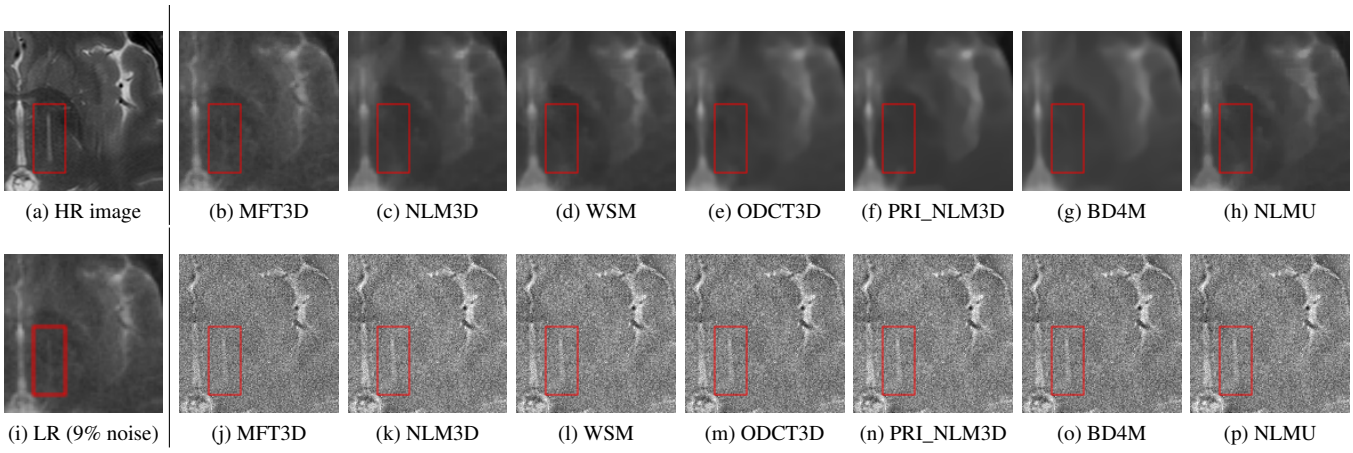

(a) HR image (b) MFT3D (c) NLM3D (d) WSM (e) ODCT3D (f) PRI_NLM3D (g) BD4M (h) NLMU

(i) LR (9% noise) (j) MFT3D (k) NLM3D (l) WSM (m) ODCT3D (n) PRI_NLM3D (o) BD4M (p) NLMU

Fig. 2. Qualitative results for T2 weighted real image from CIMES, $\alpha = 2$, axial view.

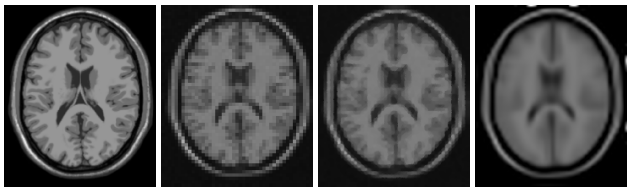

(a) HR (b) LR (c) MFT3D (d) BD4M

Fig. 3. Example of presence of hallucinations in restored images for NLM based methods (BD4M), $\alpha = 4$.

An axial slice of T1 Brainweb image is presented in Figure 3. It can be noted that white spots appear in the superior and right part of the restored image. These spots must not appear since neither the noiseless image nor the LR noisy image evince any imperfections outside the brain. Therefore, they are hallucinations. NLM learns from the data, which could be from any similar but far apart region of the image, and may introduce features from other regions of the brain in the wrong place. However, MFT3D only takes into account the voxels which are neighboring to the estimated one so it does not suffer this kind of imperfections.

3.2. Real data

A real image provided by CIMES, a medical research center of the General Foundation of the University of Málaga (FGUMA), was tested in order to study the algorithm's performance. T1-weighted and T2-weighted MR images were provided where the MRI studies were acquired on a Philips 3T Intera MRI scanner. The image size is $256 \times 256 \times 190$, with pixel spacing $0.93 \times 0.93 \times 1.0 mm^3$. Only the qualitative results for these images are shown since the original images already have noise. The noise present in the real images is too low, so as proposed in [29], it has been augmented by simulated 9% Rician noise to achieve a better discrimination between the different compared methods.

Our approach yields good results, specially in terms of sharpness. MFT3D can recover brain structures that the other methods loose with large zoom factors. This can be seen in the red rectangles in Figure 2, where there are very definite surcus and gyrus in the brain of the original images but they can not be appreciated in the competitors' results. It can be seen that the restored images obtained by the other methods are oversmoothed, and MFT3D is the only one that can remove the aliasing while keeping the small features to obtain a HR image similar to the original one, even when large zoom factors are employed. If we compare the residual images, the competitors manifest accrued losses of brain structures, i.e. anatomical features are visible in the residual image for the competitors. All methods remove noise properly, but NLM based methods remove more in the background, which is useless. MFT3D provides a good balance between noise removal and detail preservation.

In which respects to the CPU time, on average the competing methods need around 2 minutes at most and MFT3D method spends at least 1 hour to compute the restored image, depending on the size factor employed for the SR. The use of a GPU would improve between 25 and 50 times the processing time, since each of the H tilings of the MFT3D can be computed in parallel.

4. CONCLUSIONS

A new method for both superresolution and noise removal applied to 3D magnetic resonance images is presented. It is based on the application of median filters in two steps. Since the median is used for denoising, the noise distribution in the image is not relevant, and could be completely unknown. Besides, the convergence of the method has been demonstrated. The proposed method yields better quantitative results than other state-of-the-art methods when 5% or more Rician noise is present in the images, irrespective of the integer or fractional scale factor employed in the super-resolution process. MSE and SSIM results confirm the better quality of the HR images obtained by our proposal. It has been shown that MFT3D also achieves good qualitative results when heavy noise is present in the image. The noise is removed while the finer details of the image are preserved. In contrast to this, the competing methods suffer from oversmoothing and relevant anatomical features are lost. This can incur in errors since a lesion can disappear due to its similarity to noisy pixels. Moreover, it does not generate hallucinations in the HR image, the presence of which could be could be largely detrimental for medical diagnosis. Further works will evaluate the effect of different super-resolution methods in combination with noise removal algorithms. Besides, the inclusion of the median filter transform might be beneficial if super-resolution deep networks could integrate them.

5. REFERENCES

[1] Thomas M. Lehmann, Claudia Gönner, and Klaus Spitzer, "Survey: Interpolation methods in medical image processing," *IEEE Transactions on Medical Imaging*, vol. 18, pp. 1049–1075, 1999.

[2] P. Thevenaz, T. Blu, and M. Unser, "Interpolation revisited [medical images application]," *IEEE Transactions on Medical Imaging*, vol. 19, no. 7, pp. 739–758, July 2000.

[3] Jing Tian and Kai-Kuang Ma, "A survey on super-resolution imaging," *Signal, Image and Video Processing*, vol. 5, no. 3, pp. 329–342, 2011.

[4] Eric Van Reeth, Ivan W. K. Tham, Cher Heng Tan, and Chueh Loo Poh, "Super-resolution in magnetic resonance imaging: A review," *Concepts in Magnetic Resonance Part A*, vol. 40A, no. 6, pp. 306–325, 2012.

[5] Jianchao Yang and Thomas Huang, "Image super-resolution: Historical overview and future challenges," in *Super-resoluton imaging*. 2010, CRC Press.

[6] Chih-Yuan Yang, Chao Ma, and Ming-Hsuan Yang, "Single-image super-resolution: A benchmark," in *Proceedings of European Conference on Computer Vision*, 2014.

[7] Hong Chang, Dit-Yan Yeung, and Yimin Xiong, "Super-resolution through neighbor embedding.," in *CVPR (1)*, 2004, pp. 275–282.

[8] William T. Freeman, Thouis R. Jones, and Egon C Pasztor, "Example-based super-resolution," *IEEE Comput. Graph. Appl.*, vol. 22, no. 2, pp. 56–65, Mar. 2002.

[9] Weisheng Dong, Lei Zhang, Guangming Shi, and Xiaolin Wu, "Image deblurring and super-resolution by adaptive sparse domain selection and adaptive regularization," *Trans. Img. Proc.*, vol. 20, no. 7, pp. 1838–1857, July 2011.

[10] Dinh Hoan Trinh, Marie Luong, Françoise Dibos, Jean-Marie Rocchisani, Canh Duong Pham, and Truong Q. Nguyen, "Novel example-based method for super-resolution and denoising of medical images," *IEEE Trans. Image Processing*, vol. 23, no. 4, pp. 1882–1895, 2014.

[11] Jianchao Yang, John Wright, Thomas S. Huang, and Yi Ma, "Image super-resolution via sparse representation," *Trans. Img. Proc.*, vol. 19, no. 11, pp. 2861–2873, Nov. 2010.

[12] Daniel Glasner, Shai Bagon, and Michal Irani, "Super-resolution from a single image," in *ICCV*, 2009.

[13] A. Singh, F. Porikli, and N. Ahuja, "Super-resolving noisy images," in *2014 IEEE Conference on Computer Vision and Pattern Recognition*, June 2014, pp. 2846–2853.

[14] José V. Manjón, Pierrick Coupé, Antonio Buades, D. Louis Collins, and Montserrat Robles, "MRI superresolution using self-similarity and image priors," *Journal of Biomedical Imaging*, vol. 2010, pp. 17:1–17:11, Jan. 2010.

[15] François Rousseau, "Brain hallucination," *Computer Vision–ECCV 2008*, pp. 497–508, 2008.

[16] Ezequiel López-Rubio, "Superresolution from a single noisy image by the median filter transform," *SIAM J. Imaging Sciences*, vol. 9, pp. 82–115, 2016.

[17] Sina Farsiu, Dirk Robinson, Michael Elad, and Peyman Milanfar, "Advances and challenges in super-resolution," *International Journal of Imaging Systems and Technology*, vol. 14, no. 2, pp. 47–57, 2004.

[18] Yanyuan Ma, Marc G. Genton, and Emanuel Parzen, "Asymptotic properties of sample quantiles of discrete distributions," *Annals of the Institute of Statistical Mathematics*, vol. 63, no. 2, pp. 227–243, 2011.

[19] Herbert A. David and H. N. Nagaraja, *Order statistics*, Wiley, Hoboken, N.J., 3rd ed. edition, 2003.

[20] Pierrick Coupé, Pierre Yger, Sylvain Prima, Pierre Hellier, Charles Kervrann, and Christian Barillot, "An optimized blockwise nonlocal means denoising filter for 3-D magnetic resonance images," *IEEE Transactions on Medical Imaging*, vol. 27, no. 4, pp. 425–441, 2008.

[21] Hong Liu, Cihui Yang, Ning Pan, Enmin Song, and Richard Green, "Denoising 3D MR images by the enhanced non-local means filter for Rician noise," *Magnetic Resonance Imaging*, vol. 28, no. 10, pp. 1485–1496, 2010.

[22] J.V. Manjón, P. Coupé, L. Martí-Bonmatí, D.L. Collins, and M. Robles, "Adaptive non-local means denoising of MR images with spatially varying noise levels," *Journal of Magnetic Resonance Imaging*, vol. 31, no. 1, pp. 192–203, 2010.

[23] Pierrick Coupé, Pierre Hellier, Sylvain Prima, Charles Kervrann, and Christian Barillot, "3d wavelet subbands mixing for image denoising," *Journal of Biomedical Imaging*, vol. 2008, pp. 1:1–1:11, Jan. 2008.

[24] José V. Manjón, Pierrick Coupé, Antonio Buades, D. Louis Collins, and Montserrat Robles, "New methods for MRI denoising based on sparseness and self-similarity.," *Medical Image Analysis*, vol. 16, no. 1, pp. 18–27, 2012.

[25] Matteo Maggioni, Vladimir Katkovnik, Karen Egiazarian, and Alessandro Foi, "Nonlocal transform-domain filter for volumetric data denoising and reconstruction," *IEEE Transactions on Image Processing*, vol. 22, no. 1, pp. 119–133, 2013.

[26] José V. Manjón, Pierrick Coupé, Antonio Buades, Vladimir Fonov, D. Louis Collins, and Montserrat Robles, "Non-local MRI upsampling.," *Medical Image Analysis*, vol. 14, no. 6, pp. 784–792, 2010.

[27] Weisheng Dong, Lei Zhang, Rastislav Lukac, and Guangming Shi, "Sparse representation based image interpolation with nonlocal autoregressive modeling," *IEEE Trans. Image Processing*, vol. 22, no. 4, pp. 1382–1394, 2013.

[28] R.K.-S. Kwan, A.C. Evans, and G.B. Pike, "BrainWeb: Simulated brain database," `http://mouldy.bic.mni.mcgill.ca/brainweb/`, 2010.

[29] C. Shyam Anand and Jyotinder S. Sahambi, "Wavelet domain non-linear filtering for MRI denoising," *Magnetic Resonance Imaging*, vol. 28, no. 6, pp. 842 – 861, 2010.