

Example-based Restoration of High-resolution Magnetic Resonance Image Acquisitions

Ender Konukoglu, Andre van der Kouwe, Mert R. Sabuncu, and Bruce Fischl

A. A. Martinos Center for Biomedical Imaging, MGH, Harvard Medical School, USA

Abstract. Increasing scan resolution in magnetic resonance imaging is possible with advances in acquisition technology. The increase in resolution, however, comes at the expense of severe image noise. The current approach is to acquire multiple images and average them to restore the lost quality. This approach is expensive as it requires a large number of acquisitions to achieve quality comparable to lower resolution images. We propose an image restoration method for reducing the number of required acquisitions. The method leverages a high-quality lower-resolution image of the same subject and a database of pairs of high-quality low/high-resolution images acquired from different individuals. Experimental results show that the proposed method decreases noise levels and improves contrast differences between fine-scale structures, yielding high signal-to-noise ratio (SNR) and contrast-to-noise ratio (CNR). Comparisons with the current standard method of averaging approach and state-of-the-art non-local means denoising demonstrate the method’s advantages.

1 Introduction

Resolution in medical imaging sets a fundamental limit on the scale of structures that can be visualized. Increasing resolution yields numerous benefits for both basic research and clinical applications. In magnetic resonance imaging (MRI), higher field strengths and array receive coils allow acquisition of increased resolution. This increase, however, comes at the expense of lower signal-to-noise ratio (SNR) and lower contrast-to-noise ratio (CNR). Scientists currently acquire multiple high-resolution (high-res) images of the same structure and average them to recover the SNR(CNR) lost by the increased resolution. This approach comes at an extremely steep price: doubling the resolution requires scanning 64 times as long to achieve comparable SNR for 3D encoded acquisitions. Obtaining high quality images with reduced scan time is a necessary step to make higher-resolution imaging feasible for clinical practice and available for a larger number of research studies. We believe that image restoration can provide a viable alternative to acquiring a large number of scans.

Restoration (denoising) is an active field of research in both computer vision and MRI literature. The most popular approaches that have been applied to MRI are based on Gaussian filtering[1], wavelet decompositions[14], anisotropic diffusion[9, 11], non-parametric estimation[2] and non-local means[8, 3, 5, 13]. While these are successful in denoising MRI, they do not take into account the specific

aspects of high-res acquisition. As a result, although they are able to increase the SNR, they are not effective in restoring the contrast between fine-scale anatomical details in the presence of severe noise.

High-res acquisitions in MRI are particularly difficult for restoration. The noise levels often severely distort the appearance of fine-scale structures, whose restoration based only on a small number of acquisitions is challenging (see Fig.1). Conversely, problem specific aspects related to high-res acquisition can help restoration. First, high-quality low-resolution (low-res) images (e.g. 1 mm^3) can be acquired rapidly. These provide coarse level prior information for restoration. However, low-res does not provide enough information for restoring fine scale structures. Second, similarity of anatomy across individuals can complement the short-comings of low-res acquisitions as prior information. Previously acquired high-quality low/high-res image pairs of different subjects provide empirical prior that can help link different resolutions and guide restoration.

This article presents a restoration method that aims to reduce the number of acquisitions required to obtain a high-quality high-res MRI. It integrates low-res acquisition and a training database of pairs of high-quality low/high-res images in a probabilistic formulation. This method shares similarities with dictionary-based methods for denoising, such as [7, 12]. However, the proposed method does not learn a dictionary to integrate the database into the restoration. Instead, it builds on a patch-based synthesis framework, which has been successfully used in super-resolution [15], image analogies [10] and synthesis [16]. Experimental results on five subjects demonstrate the capabilities of the method for achieving noise levels that would normally require more acquisitions. The proposed method improves SNR and CNR, revealing fine-scale structures. Comparisons with the state-of-the-art non-local means denoising algorithm illustrate the advantages of the proposed method for restoring high-res MRI.

2 Restoration Method

We model an MRI image, I , as a mapping from space to intensity values, i.e., $I : \Omega \rightarrow \mathbb{R}$, where $\Omega \subset \mathbb{N}^3$ is a discrete domain. A high-res MRI acquisition, \tilde{H}_m , is a noisy version of an ideal noise-free high-res image H . The current approach for restoring H from a set $\{\tilde{H}_m\}_{m=1}^M$ is the point-wise averaging $\mathbf{x} \in \Omega$: $\hat{H}(\mathbf{x}) = \sum_m \tilde{H}_m(\mathbf{x})/M$, which requires M to be as high as 7 or 8 to overcome the severe noise levels, for e.g. 0.5 mm resolution. Following, we present a restoration method that aims to reduce the required M .

Probabilistic Model: The inputs of the proposed method are: i) the high-res acquisitions, $\{\tilde{H}_m\}_{m=1}^M$; ii) the corresponding high-SNR low-res image, L , registered and up-sampled with tri-linear interpolation to the same grid as \tilde{H}_m 's; and iii) a *training* database of coupled high-SNR low/high-res images $\{(L_q, H_q)\}_{q=1}^Q$, previously acquired, from different subjects. The goal of the algorithm is to estimate the high-SNR high-res image H , denoted by \hat{H} .

The proposed method works on image patches. A patch of size $d \in \mathbb{N}$ in an image I at location \mathbf{x} is the set of intensities over the neighborhood voxels,

i.e., $I^d(\mathbf{x}) \triangleq \{I(\mathbf{y}) : \mathbf{y} \in W^d(\mathbf{x})\}$, where $W^d(\mathbf{x}) \triangleq \{\mathbf{y} : \|\mathbf{x} - \mathbf{y}\|_\infty \leq d\}$ is the neighborhood of \mathbf{x} . For instance, $W^1(\mathbf{x})$ is the set that includes \mathbf{x} and its 26 immediate neighbors, and $I^1(\mathbf{x})$ are the set of intensities within $W^1(\mathbf{x})$.

We estimate each patch in \hat{H} by maximizing the posterior probability:

$$\begin{aligned} \hat{H}^d(\mathbf{x}) &= \underset{H^d(\mathbf{x})}{\operatorname{argmax}} p(H^d(\mathbf{x})|L^d(\mathbf{x}), \{\tilde{H}_m^d(\mathbf{x})\}) = \underset{H^d(\mathbf{x})}{\operatorname{argmax}} p(H^d(\mathbf{x}), L^d(\mathbf{x}), \{\tilde{H}_m^d(\mathbf{x})\}) \\ &= \underset{H^d(\mathbf{x})}{\operatorname{argmax}} p(L^d(\mathbf{x}), H^d(\mathbf{x})) \prod_m p(\tilde{H}_m^d(\mathbf{x})|H^d(\mathbf{x})). \end{aligned} \quad (1)$$

To reach (1), we assumed (i) $p(\tilde{H}_m^d(\mathbf{x})|H^d(\mathbf{x}), L^d(\mathbf{x})) = p(\tilde{H}_m^d(\mathbf{x})|H^d(\mathbf{x}))$, i.e., in the presence of H , image L provides no additional information about each H_m ; and (ii) $p(\{\tilde{H}_m^d(\mathbf{x})\}|H^d(\mathbf{x})) = \prod_m p(\tilde{H}_m^d(\mathbf{x})|H^d(\mathbf{x}))$, i.e., each low-SNR image H_m is conditionally independent given H .

We further assume the following Gaussian noise model: $p(\tilde{H}_m^d(\mathbf{x})|H^d(\mathbf{x})) = \mathcal{N}(H^d(\mathbf{x}), \sigma_n \mathbf{I})$, where $\mathcal{N}(\cdot, \cdot)$ represents the normal distribution and \mathbf{I} is the identity matrix. The reasons for this choice are two-folds. First, empirically we observed that the noise distribution for each high-res acquisition can be well approximated with a Gaussian (see Fig. 1). Second, point-wise averaging is the solution of the model that ignores $p(L^d(\mathbf{x}), H^d(\mathbf{x}))$ in Equation 1. This second point links the proposed method to the current practice.

The key component of the proposed method is $p(L^d(\mathbf{x}), H^d(\mathbf{x}))$. There are multiple ways of defining this term. One could, for example, use a parametric form that models subsampling. Without anatomically informed priors, however, this approach would fail to model structures that are only visible at high-res. As a result, unless these structures are prominent in the noisy acquisitions, they cannot be restored. For brain MRI, an alternative approach is to leverage the anatomical similarities between individuals by using available training datasets. Here, we take this approach and use the training database $\{(L_q, H_q)\}_{q=1}^Q$ to estimate the joint distribution $p(L^d(\mathbf{x}), H^d(\mathbf{x}))$ using a non-parametric model:

$$p(L^d(\mathbf{x}), H^d(\mathbf{x})) = \frac{1}{Q|W^D(\mathbf{x})|} \sum_{q=1}^Q \sum_{\mathbf{y} \in W^D(\mathbf{x})} \mathbf{K}_\Sigma(L^d(\mathbf{x}), L_q^d(\mathbf{y})) \mathbf{K}_\Sigma(H^d(\mathbf{x}), H_q^d(\mathbf{y})) \quad (2)$$

where $W^D(\mathbf{x})$ is the D -neighborhood of \mathbf{x} , $|\cdot|$ denotes set cardinality, $\mathbf{K}_\Sigma(I, J) = \exp\{-\frac{1}{2}(I - J)^T \Sigma^{-1}(I - J)\} / \sqrt{2\pi \det(\Sigma)}$, $\Sigma(\mathbf{x}_1, \mathbf{x}_2) = \sigma^2 \exp\{-\|\mathbf{x}_1 - \mathbf{x}_2\|_2^2 / \alpha^2\}$ and $\det(\cdot)$ is the matrix determinant. Σ models the spatial correlation in the residuals and has two global parameters, σ and α . More refined parameterizations can also be used, for example by assigning locally varying parameters. In that case, however, the estimation of the parameters becomes more challenging. In Eq. 2, the summation over the voxel index \mathbf{y} allows us to consider patches at voxels other than \mathbf{x} . This enriches the training data used for voxel \mathbf{x} and models misalignments between the subjects.

Optimization: To solve Eq. 1 with the definition of Eq. 2, one can use numerical methods such as Expectation Maximization. This, however, becomes

computationally intractable because a separate iterative optimization needs to be run at each voxel \mathbf{x} , and the summation over all training patches can be expensive when the training dataset is large (e.g. when D is large). As a first order approximation we propose to use

$$p(L^d(\mathbf{x}), H^d(\mathbf{x})) \approx \max_{q, \mathbf{y} \in W^D(\mathbf{x})} \mathbf{K}_\Sigma(L^d(\mathbf{x}), L_q^d(\mathbf{y})) \mathbf{K}_\Sigma(H^d(\mathbf{x}), H_q^d(\mathbf{y})). \quad (3)$$

This type of approximation can be justified when the dimensionality of the problem is high and the training samples are sparse. In Eq. 3, in the same spirit as k-means clustering, the new image is associated with the closest training patch and the probability value is computed solely based on this association. The main advantage of adopting Eq. 3 is that it converts the problem given in Eq. 1 to

$$\operatorname{argmax}_{H^d(\mathbf{x})} \left\{ \max_{q, \mathbf{y} \in W^D(\mathbf{x})} \mathbf{K}_\Sigma(L^d(\mathbf{x}), L_q^d(\mathbf{y})) \mathbf{K}_\Sigma(H^d(\mathbf{x}), H_q^d(\mathbf{y})) \prod_m p(\tilde{H}_m^d(\mathbf{x}) | H^d(\mathbf{x})) \right\}, \quad (4)$$

which can be solved efficiently. We observe that for a fixed q and \mathbf{y} , the outer optimization of Eq. 4 yields a closed-form solution:

$$\hat{H}_{q, \mathbf{y}}^d(\mathbf{x}) = \left(\frac{M}{\sigma_n^2} \mathbf{I} + \Sigma^{-1} \right)^{-1} \left(\frac{1}{\sigma_n^2} \mathbf{I} \sum_{m=1}^M \tilde{H}_m^d(\mathbf{x}) + \Sigma^{-1} H_q^d(\mathbf{y}) \right). \quad (5)$$

This reduces the problem in Eq. 4 to solving:

$$\operatorname{argmax}_{q, \mathbf{y}} \mathbf{K}_\Sigma(L^d(\mathbf{x}), L_q^d(\mathbf{y})) \mathbf{K}_\Sigma(\hat{H}_{q, \mathbf{y}}^d(\mathbf{x}), H_q^d(\mathbf{y})) \prod_m p(\tilde{H}_m^d(\mathbf{x}) | \hat{H}_{q, \mathbf{y}}^d(\mathbf{x})).$$

Combining the last two terms and Eq. 5, we can rewrite this as:

$$q^*, \mathbf{y}^* = \operatorname{argmax}_{q, \mathbf{y}} \mathbf{K}_\Sigma(L^d(\mathbf{x}), L_q^d(\mathbf{y})) \mathbf{K}_{\Sigma + \mathbf{I}\sigma_n^2/M} \left(\frac{1}{M} \sum_m \tilde{H}_m^d(\mathbf{x}), H_q^d(\mathbf{y}) \right) \quad (6)$$

and the final estimate is given as $\hat{H}^d(\mathbf{x}) = \hat{H}_{q^*, \mathbf{y}^*}^d(\mathbf{x})$.

Equation 6 can be solved using the powerful patch-matching procedure borrowing ideas from patch-based segmentation systems [6]. Following the brain-specific strategy as given in [6] we linearly align the new subject data L , $\{\tilde{H}_m\}$ with each training subject $\{(L_q, H_q)\}$ via affine registration, and perform exhaustive search over a restricted spatial neighborhood, $W^D(\mathbf{x})$. The cost of searching over $W^D(\mathbf{x})$ can be reduced by employing a multi-resolution grid pyramid.

The presented method restores $\hat{H}^d(\mathbf{x})$ for each \mathbf{x} independently. $\hat{H}^d(\mathbf{x})$ contains the intensity estimates for \mathbf{x} and all its neighbors in $W^d(\mathbf{x})$. We compute the final estimate $\hat{H}(\mathbf{x})$ by averaging the estimates from all the patches containing \mathbf{x} . The interactions between neighboring voxels could alternatively be modeled as a prior distribution over H . This, however, would remove the possibility of solving each $H^d(\mathbf{x})$ independently, which enables parallelization.

Variation on the Model: A variation of the presented model is to remove the

dependence on $L^d(\mathbf{x})$ and only model $p(H^d(\mathbf{x}), \{\tilde{H}_m^d(\mathbf{x})\})$. This case corresponds to only using the high-res acquisitions and the high-res images in the database to restore H . In this case, most of the derivations follow suit and the restoration process reduces to solving $q^*, \mathbf{y}^* = \operatorname{argmax}_{q, \mathbf{y}} \mathbf{K}_{\Sigma + \mathbf{I}\sigma_n^2/M} \left(\frac{1}{M} \sum_m H_m^d(\mathbf{x}), H_q^d(\mathbf{y}) \right)$, where the restored image patch is computed as $\hat{H}^d(\mathbf{x}) = \hat{H}_{q^*, \mathbf{y}^*}^d(\mathbf{x})$.

Setting the Parameters: The probabilistic model has five free parameters, σ_n , σ , α , d and D . The first three we set using heuristic strategies. We first assume the noise variance σ_n is constant across subjects, i.e., the noise properties remain similar across images. Thus we can directly estimate σ_n on the training dataset, where each H_q is associated with multiple low-quality acquisitions $\{\tilde{H}_{q,m}\}$. We estimate σ_n as the square root of the mean square difference between H_q and $\{\tilde{H}_{q,m}\}$ on the training dataset. σ and α values define the influence domain of each kernel in the non-parametric distribution in Eq. 2. For a given subject, we estimate these parameters based on the high-SNR low-res images $L^d(\mathbf{x})$ and $\{L_q^d(\mathbf{x})\}$. We first compute the following empirical covariance matrix $S = \frac{1}{N} \sum_{\mathbf{x}_n} \left(L^d(\mathbf{x}_n) - L_{q_n^*}^d(\mathbf{y}_n^*) \right) \left(L^d(\mathbf{x}_n) - L_{q_n^*}^d(\mathbf{y}_n^*) \right)^T$, where $(q_n^*, \mathbf{y}_n^*) = \operatorname{argmax}_{q, \mathbf{y}} \|L^d(\mathbf{x}) - L_q^d(\mathbf{y})\|_2$, using N randomly selected voxels $\mathbf{x}_n \in \Omega$. We then determine σ and α by minimizing the square difference between $\Sigma(\cdot, \cdot)$ and the sample covariance S . d and D are set empirically.

3 Experiments

We tested the proposed restoration method on a dataset of five subjects. For each subject, seven high-res T1w images at a resolution of $(500\mu m)^3$ and one low-res T1w image at the resolution of 1 mm^3 were acquired on a 3T Siemens Trio scanner with a 3D encoded MPRAGE using a 32-channel receive coil. The protocol took 7 minutes to acquire each high-res image and 3.5 minutes to acquire the low-res images. Fig. 1-(a-c) show, for a subject, the low-res image, a high-res acquisition and the average high-res image, respectively. Fig. 1-d plots the histogram of the difference between each high-res acquisition shown in (c) and the average of seven shown in (b), “the noise distribution”, along with the best fitting Gaussian distribution in red.

We performed leave-one-out experiments, where for each test case the images of the remaining subjects were used to construct the corresponding training database. We used two different ways to construct the training database high-res images: i) by averaging the seven high-res acquisitions (set1) and ii) by first averaging then denoising the average image by the non-local means (NLM) algorithm as proposed in [5] (set2). For each test case, we performed seven restorations, where for each restoration we assumed a different number, M , of high-res low-SNR scans. The restoration quality is quantified using two measures, SNR and CNR, which are computed based on the claustrum: a fine scale structure that is only visible in the high-res. Two regions-of-interests were drawn on the average of seven high-res acquisitions, one within the claustrum and another within the

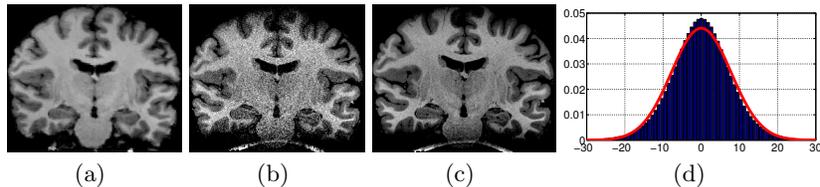


Fig. 1. (a) T1w image at 1 mm^3 resolution (b) T1w image at $(500 \text{ } \mu\text{m})^3$ acquisition (c) Average of seven $(500 \text{ } \mu\text{m})^3$ images (d) Histogram of the difference between (b) and (c), and overlaid is the best fit Gaussian distribution to the histogram.

external capsule, which borders the claustrum (see Fig. 3-a). SNR was computed as the ratio of the average intensity value to the standard deviation within the claustrum ROI. CNR was computed as the absolute difference of mean intensities of the two ROIs divided by the combined standard deviation.

We tested three variants of the proposed model: (i) “model HL” uses both L and \hat{H}_m for the test images and set1 type training database, (ii) “model H” uses only \hat{H}_m and set1 type database, and (iii) “model HLD” uses both L and \hat{H}_m and set2 type database. We used a patch size of $d = 1$ and ran patch-matching in a multi-resolution pyramid of three levels at resolutions $(2 \text{ mm})^3$, 1 mm^3 and $(500 \text{ } \mu\text{m})^3$ with a search neighborhood of $D = 8 \text{ mm}$, 4 mm , 2 mm , respectively. We compared the restoration quality of the proposed method with four benchmark methods: (i) the point-wise averaging, (ii) block-wise NLM [5] (NLMBO), (iii) oracle-based DCT filter [13] (ODCT) and (iv) pre-filtered rotationally invariant NLM [13] (PRINLM) all applied to the point-wise averaging result¹. The noise level estimations for the latter three methods were performed using the method proposed in [4] and all the parameters were set as suggested in [13].

Fig. 2 plots the results with respect to M (and acquisition times corresponding to each M in parenthesis). The bars correspond to average values obtained over five test cases and the errorbars are the standard errors. In terms of SNR, model HLD achieved the highest values for all M . The highest CNR values were obtained by HLD for $M < 4$ and NLMBO for $M \geq 3$. The high CNR and SNR values achieved by the models HL and HLD demonstrate that the proposed approach can drastically reduce noise and improve contrast simultaneously. The differences between models H and HL show that the integration of the low-res image is advantageous. Considering acquisition times, the proposed methods, in particular model HLD, provides substantial benefits for low M .

Fig. 3 displays visual results from two subjects: (a) slices of high-res images and the ROIs (blue = claustrum, red = external capsule) and (b) restoration results of NLMBO, HL and HLD for $M = 1, 3$. These images demonstrate that depending on the noise level NLM might not be enough. The proposed models HL and HLD are able to restore the image even in high noise levels. e.g. $M = 1$.

¹ The implementations for (ii-iv) are from <http://personales.upv.es/jmanjon/denoising/prinlm.html>.

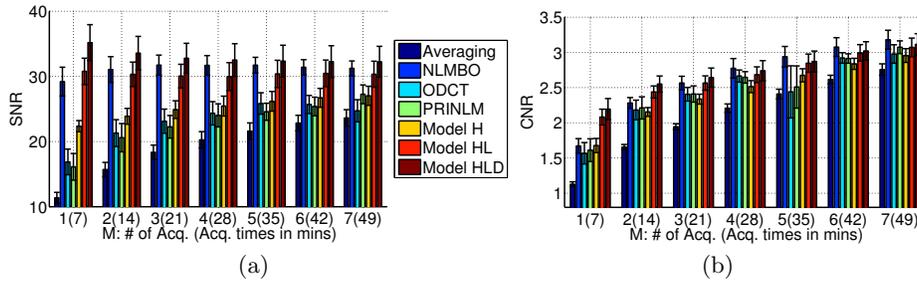


Fig. 2. Quantitative restoration results vs. number of acquisitions and acquisition times in minutes: (a) SNR at the ROI drawn on the claustrum, see Fig. 3, (b) CNR computed between the claustrum and the external capsule. The bars are mean statistics over 5 subjects and errorbars are standard errors.

4 Conclusions

In this work we proposed a restoration method for improving the quality of high-res MRI acquisitions. Our experiments demonstrate that the method is able to reduce the severe noise levels and improve the contrast between neighboring fine-scale structures. The method achieves this by leveraging low-res images and a training database, which provides an empirical prior on the appearance of structures at different resolutions. The preliminary results are promising and suggest the possibility of reducing the number of acquisitions needed to obtain high-quality high-res MRI at higher-field strengths.

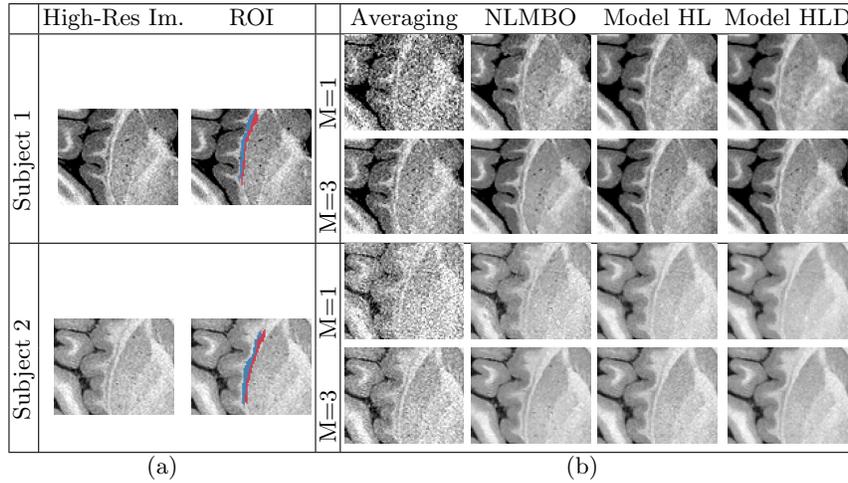


Fig. 3. Visual results: (a-left) Two different subjects' high-res images zoomed around the claustrum (left). (a-right) ROIs (blue = claustrum, red = external capsule). (b) Restored images: (left to right) averaging, NLMBO and the proposed models.

Acknowledgements

Support for this research was provided in part by the NIH (P41-RR14075, U24 RR021382, R01EB006758, AG022381, 5R01AG008122-22, RC1 AT005728-01, R01 NS052585-01, 1R21NS072652-01, 1R01NS070963, 1S10RR023401, 1S10RR019307, 1S10RR023043, R21 MH096559, R01 HD071664, R33 DA026104) and The Autism & Dyslexia Project funded by the Ellison Medical Foundation, and by the NIH Blueprint for Neuroscience Research (5U01-MH093765), part of the multi-institutional Human Connectome Project.

References

1. Ashburner, J., Friston, K.J.: Voxel-based morphometry: the methods. *Neuroimage* 11(6), 805–821 (2000)
2. Awate, S.P., Whitaker, R.T.: Feature-preserving mri denoising: A nonparametric empirical bayes approach. *IEEE TMI* 26(9), 1242–1255 (2007)
3. Buades, A., Coll, B., Morel, J.M.: A non-local algorithm for image denoising. In: *CVPR*. vol. 2, pp. 60–65. IEEE (2005)
4. Coupé, P., Manjón, J.V., Gedamu, E., Arnold, D.L., Robles, M., Collins, D.L., et al.: Robust rician noise estimation for mr images. *Med Image Anal* 14(4), 483–93 (2010)
5. Coupé, P., et al.: An optimized blockwise nonlocal means denoising filter for 3-d magnetic resonance images. *IEEE TMI* 27(4), 425–441 (2008)
6. Coupé, P., et al.: Patch-based segmentation using expert priors: Application to hippocampus and ventricle segmentation. *Neuroimage* 54(2), 940–954 (2011)
7. Elad, M., Aharon, M.: Image denoising via sparse and redundant representations over learned dictionaries. *IEEE TIP* 15(12), 3736–3745 (2006)
8. Fischl, B., Schwartz, E.L.: Adaptive nonlocal filtering: a fast alternative to anisotropic diffusion for image enhancement. *IEEE TPAMI* 21(1), 42–48 (1999)
9. Gerig, G., Kubler, O., Kikinis, R., Jolesz, F.A.: Nonlinear anisotropic filtering of mri data. *IEEE TMI* 11(2), 221–232 (1992)
10. Hertzmann, A., Jacobs, C.E., Oliver, N., Curless, B., Salesin, D.H.: Image analogies. In: *SIGGRAPH*. pp. 327–340. ACM (2001)
11. Krissian, K., Aja-Fernández, S.: Noise-driven anisotropic diffusion filtering of mri. *IEEE TIP* 18(10), 2265–2274 (2009)
12. Mairal, J., Bach, F., Ponce, J., Sapiro, G., Zisserman, A.: Non-local sparse models for image restoration. In: *CVPR*. pp. 2272–2279. IEEE (2009)
13. Manjón, J.V., Coupé, P., Buades, A., Louis Collins, D., Robles, M.: New methods for mri denoising based on sparseness and self-similarity. *Med Image Anal* 16, 18–27 (2012)
14. Pizurica, A., Philips, W., Lemahieu, I., Acheroy, M.: A versatile wavelet domain noise filtration technique for medical imaging. *IEEE TMI* 22(3), 323–331 (2003)
15. Rousseau, F.: Brain hallucination. In: Forsyth, D., Torr, P., Zisserman, A. (eds.) *ECCV 2008*. LNCS, vol. 5302, pp. 497–508. Springer Berlin Heidelberg (2008)
16. Roy, S., Carass, A., Prince, J.: A compressed sensing approach for mr tissue contrast synthesis. In: Székely, G., Hahn, H. (eds.) *IPMI 2011*. LNCS, vol. 6801, pp. 371–383. Springer Berlin Heidelberg (2011)