

HHS Public Access

Author manuscript

IEEE Trans Med Imaging. Author manuscript; available in PMC 2017 December 22.

Published in final edited form as: *IEEE Trans Med Imaging*. 2014 October ; 33(10): 1969–1985. doi:10.1109/TMI.2014.2329271.

MRI Upsampling Using Feature-Based Non-Local Means Approach

Kourosh Jafari-Khouzani

Athinoula A. Martinos Center for Biomedical Imaging, Department of Radiology, Massachusetts General Hospital and Harvard Medical School, Boston, MA, USA

Abstract

In magnetic resonance imaging (MRI), spatial resolution is limited by several factors such as acquisition time, short physiological phenomena, and organ motion. The acquired image usually has higher resolution in two dimensions (the acquisition plane) in comparison with the third dimension, resulting in highly anisotropic voxel size. Interpolation of these low resolution (LR) images using standard techniques, such as linear or spline interpolation, results in distorted edges in the planes perpendicular to the acquisition plane. This poses limitation on conducting quantitative analyses of LR images, particularly on their voxel-wise analysis and registration. We have proposed a new non-local means feature-based technique that uses structural information of a high resolution (HR) image with a different contrast and interpolates the LR image. In this approach, the similarity between voxels is estimated using a feature vector that characterizes the laminar pattern of the brain structures, resulting in a more accurate similarity measure in comparison with conventional patch-based approach. This technique can be applied to LR images with both anisotropic and isotropic voxel sizes. Experimental results conducted on brain MRI scans of patients with brain tumors, multiple sclerosis, epilepsy, as well as schizophrenic patients and normal controls show that the proposed method is more accurate, requires fewer computations, and thus is significantly faster than a previous state-of-the-art patch-based technique. We also show how the proposed method may be used to upsample regions of interest drawn on LR images.

Index Terms

Interpolation; upsampling; super-resolution; non-local means; magnetic resonance imaging

I. Introduction

Image spatial resolution is a limiting factor in quantitative analyses of medical images. In magnetic resonance imaging (MRI), the acquisition time, short physiological phenomena, and organ motion may limit the image resolution. For example, in dynamic contrast enhanced (DCE) and dynamic susceptibility contrast (DSC) MRI of brain, a tracer is injected into the vein during continuous rapid brain imaging. As the arrival of the bolus, its passage through the brain, and its wash-out occur in a short period of time, rapid imaging is required to achieve sufficient temporal resolution to capture the dynamics of the contrast. This poses limitations on the spatial resolution. In other cases, spatial resolution is limited by image acquisition time. For example, fluid attenuated inversion recovery (FLAIR) MRI

requires longer acquisition time compared to T1-weighted (T1W) MRI. Thus, spatial resolution may be compromised to achieve reasonable scan time and minimize the likelihood of motion artifact. This is usually done by the acquisition of relatively small number of 2D slices resulting in large slice thickness and spacing between slices. As a result, the images have lower resolution in the slice dimension compared to the other two dimensions (acquisition plane), resulting in highly anisotropic voxels (e.g., 1mm×1mm×6mm).

Quantitative analyses of these low resolution (LR) images are further distorted when these images are transformed into a different space. For instance, voxel-wise analysis usually requires all image volumes of one subject to be aligned and analyzed in one space. Thus, if slices of different image volumes are not already acquired exactly from the same location, they need to be coregistered, which involves interpolation. However, standard interpolation techniques are not accurate and result in distorted edges in the planes perpendicular to the acquisition plane. Fig. 1 shows the result of nearest neighbor (NN), linear, and spline interpolation techniques applied to an image with voxel size 1mm×1mm×6mm. The reason for an inaccurate interpolation is that the neighboring voxels, from which the intensity of an unknown voxel is estimated, may not have the same tissue type as the unknown voxel. This results in an incorrect estimation of the voxel intensity using standard interpolation techniques.

Several methods have been proposed to improve the interpolation of LR images. In one category of these methods, a single LR image from the subject is available for upsampling. Adjacent slices of the LR image may be registered using a non-rigid method to reconstruct the slices between them [1–6]. These registration-based methods best work when each slice is acquired from a thin slab of the organ as blurred edges between thick slices lead to poor coregistration. In [7], a family of adaptive 3D interpolation filters are designed and conditioned on different spatial contexts in order to reconstruct a slice based on four neighboring slices. This method, however, requires a training set. A diffusion-based technique has been used in [8] to reconstruct a slice using neighboring slices. In [9], a sparse super-resolution approach was proposed using overcomplete dictionaries. This method also requires a training set. Finally, in [10], a structure-preserving constraint is enforced by a patch-based reconstruction approach using a variant of non-local means (NLM) filter [11]. This approach is based on the self-similarity concept and this patch-based approach has been also applied successfully in [12] to enhance the resolution of 4D-CT images.

Our proposed method is related to another category of methods, in which a high resolution (HR) image with a different contrast is used to upsample the LR image [13–15]. Conventionally, a HR image, usually a T1W image with isotropic voxel size, is acquired to reveal the brain structure. A LR image may be interpolated with the help of this HR image. This is referred to as brain hallucination by Rousseau in [13] and later the method was elaborated in [14]. Rousseau uses patch-based image interpolation to upsample the LR image, which is based on NLM denoising [11, 16–18]. In NLM denoising, the intensity of a voxel is approximated by the linear combination of all intensities within the image [11]. Rousseau performs this approximation within a neighborhood of each voxel. In other words:

$$\mathbf{x}(\mathbf{v}) \approx \sum_{\mathbf{k}\in\Omega(\mathbf{v})} w(\mathbf{v},\mathbf{k}) \mathbf{x}(\mathbf{k})$$
 (1)

where **x** is the image intensity function and $\Omega(\mathbf{v})$ is a neighborhood of voxel **v**. The weights $w(\mathbf{v}, \mathbf{k})$ are calculated based on the similarity of 3D patches centered in voxels **v** and **k** in the HR image registered to the LR image. The same weights are used for interpolation of the LR image based on the assumption that the local patterns in HR and LR images are similar. This assumption may not be valid – for example when a lesion is clearly visible in the LR image but not in the HR image due to the difference between their contrasts. To resolve this issue, a combination of weights calculated from both HR and reconstructed images is used in an iterative algorithm in [14]. In each iteration, a correlation map showing the inconsistencies between HR and reconstructed images is generated to adjust the enforcement of the weights of the two images. The enforcement of the HR image pattern is lowered when the correlation is poor. In this paper, we refer to this method as the SRIP (super-resolution using intermodality priors) method.

In [15], a similar strategy is used with the difference that the similarity between two voxels of the HR image is based on the similarity of their intensities, as opposed to the similarity of patches centered in those voxels. However, to resolve HR and LR pattern inconsistencies mentioned above as well as any slight misalignment between HR and LR images, a self-similarity strategy is applied by taking into account the similarity between patches centered in two voxels in the reconstructed LR image. The weights of these two similarities (calculated from HR and LR images, respectively) are controlled by adjusting some parameters. In this paper, we refer to this method as the SSIP (self-similarity and image priors) method. SSIP has been shown to outperform SRIP in upsampling synthetic MRI from the BrainWeb database [15, 19].

To perform an accurate interpolation, it is critical to estimate the correct tissue type similarity for the LR image. A patch-based approach to measure the similarity between voxels may be suitable for images with directional textural patterns, which is not an accurate assumption for MR images. The brain gyri and sulci form laminar structures where voxels with similar tissue types form layers (as shown in Fig. 2). A patch-based approach is not rotation invariant and thus two close voxels of one tissue type placed at curved structures may be interpreted dissimilar. This is better shown in Figs. 3 (a)–(b) where the performance of a patch-based approach is illustrated by presenting the similarity map between one pixel (shown by a cross sign) and all other pixels in the image. The pixel is located at the boundary of white matter (WM) and gray matter (GM). As shown, the laminar shape of the brain structure is ignored by a patch-based approach, whereas our approach specifically incorporates this anatomy as a prior in the upsampling procedure. Not only is the assumption of laminar structures valid for normal brain structures, as will be shown by experiments, it also results in more accurate interpolation within lesions compared to a patch-based approach.

In this paper, we use the structural information of a HR image to guide the interpolation of a LR image. We propose a novel similarity measure that takes the brain laminar structure into account by assuming that voxels that have equal distances to a strong edge are more likely to belong to the same tissue type. This approach results in higher interpolation accuracy compared to SSIP, a state-of-the-art approach [15], while also requiring fewer computations and being significantly faster. Experiments are conducted on MRI scans of patients with brain tumors, multiple sclerosis, and epilepsy, MRI dataset of schizophrenic patients and normal controls, as well as BrainWeb MRI with multiple sclerosis model [19]. We also test the performance of the algorithm in the presence of noise, intensity non-uniformity, and motion artifact in the HR image, as well as noise and susceptibility artifacts in the LR image. We will also show how the proposed algorithm may be used to upsample regions of interest (ROIs) manually drawn on LR images.

II. Method

A. Overview

In a standard interpolation technique, the continuous function (image intensity in our problem)–from which a limited number of samples is available–is locally modeled, and the value of the function in one point (voxel in our problem) is estimated by fitting the model to its neighboring points and calculating the value of the model at that point. For example, in linear interpolation, the intensities are assumed to change as a linear combination of the values at some neighboring points, similar to the approximation in equation (1). The weights $w(\mathbf{v}, \mathbf{k})$ should be proportional to the similarity between $\mathbf{x}(\mathbf{v})$ and $\mathbf{x}(\mathbf{k})$. Conventionally, $w(\mathbf{v}, \mathbf{k})$ correlates with the distance between points \mathbf{v} and \mathbf{k} . This is based on the assumption that points closer to each other are more likely to have similar values. This is not always true, particularly in images with highly anisotropic voxels. In the NLM approach, this distance is not the spatial distance between the two points. The similarity between two points may be calculated by a patch-based approach as:

$$w_{\mathbf{x}}(\mathbf{v}, \mathbf{k}) = \frac{1}{Z_{\mathbf{v}}} \exp\left(-\frac{\|P_{3D}\left(\mathbf{x}\left(\mathbf{v}\right)\right) - P_{3D}\left(\mathbf{x}\left(\mathbf{k}\right)\right)\|_{2}^{2}}{2N_{i}\beta\hat{\sigma}^{2}}\right)$$
(2)

where $P_{3D}(\mathbf{x}(\mathbf{v}))$ and $P_{3D}(\mathbf{x}(\mathbf{k}))$ are 3D patches of **x** centered in voxels **v** and **k**, respectively [16]. $Z_{\mathbf{v}}$ is a constant of normalization, N_i is the number of voxels of a 3D patch, $\hat{\sigma}^2$ is the standard deviation of the noise, and β is a constant. In this method, two points are more similar if they have similar neighborhoods. In this paper, we assume $w_{\mathbf{x}}(\mathbf{v}, \mathbf{k})$ to be:

$$w_{\mathbf{x}}\left(\mathbf{v},\mathbf{k}\right) = \frac{1}{N}P\left(\mathbf{v},\mathbf{k}\right) \quad (3)$$

where $P(\mathbf{v}, \mathbf{k})$ is the probability that voxels \mathbf{v} and \mathbf{k} have similar tissue types in terms of tissue properties specified by intensity function \mathbf{x} , and N is a normalization factor such that $\Sigma_{\mathbf{k}} w_{\mathbf{x}}(\mathbf{v}, \mathbf{k}) = 1$. This means that a voxel should be interpolated using neighboring voxels that are similar in tissue property to that voxel. The function $P(\mathbf{v}, \mathbf{k})$ may be estimated from

a HR image of the same subject. If the contrast of the two images is different, they may have different values of $w_{\mathbf{x}}(\mathbf{v}, \mathbf{k})$. Nonetheless, if the two images have some structural similarity, it is likely that the above estimation results in more accurate interpolation compared to a standard interpolation technique. To perform a more accurate interpolation, one needs to accurately estimate the values of $w_{\mathbf{x}}(\mathbf{v}, \mathbf{k})$ for the LR image.

Furthermore, a LR image may not be simply considered as a down-sampled version of a higher resolution image. Partial volume effect and geometric transformation are usually present in the LR images. We consider the following degradation model to take these into account:

$$y = Hx + n$$
 (4)

where \mathbf{y} is the LR image, \mathbf{x} is the original HR image of the same contrast, *H* is a matrix incorporating geometric transformation, partial volume effect, and sub-sampling, and \mathbf{n} is the observation noise. An approach to estimate the LR image \mathbf{x} is to minimize a least-square cost function such as:

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{\arg\min} \|\mathbf{y} - H\mathbf{x}\|^2 \quad (5)$$

As this problem does not have a unique solution, some prior information is needed to constrain the solutions. This prior information may be applied in the form of equation (1), for example by introducing a regularization term $R(\mathbf{x})$ and solving:

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{\arg\min} \|\mathbf{y} - H\mathbf{x}\|^{2} + \lambda R(\mathbf{x})$$
(6)

where $R(\mathbf{x})$ incorporates the constraint in (1). Since the assumption is that *H* is known, we would need to select an appropriate $R(\mathbf{x})$ and a minimization algorithm. We consider the following regularization term:

$$R(\mathbf{x}) = \sum_{\mathbf{v}} \left| \mathbf{x}(\mathbf{v}) - \sum_{\mathbf{k} \in \Omega(\mathbf{v})} w(\mathbf{v}, \mathbf{k}) \mathbf{x}(\mathbf{k}) \right|^{2}$$
(7)

This ensures the condition in equation (1) is met [20]. An overview of the proposed algorithm is presented in Fig. 4. An initial estimation of the upsampled image is generated by NN interpolation of the LR image. The HR image is then coregistered to our initial estimate, and the weights $w_{\mathbf{x}}(\mathbf{v}, \mathbf{k})$ are estimated by a method explained later. The optimal upsampled image is then generated by the minimization in equation (6) through use of parallel computing (multi-threading) to speed up the computations. Details of the laminar-based method will be presented next.

B. Similarity features for laminar structures

In this paper, to obtain the values of $w_{\mathbf{x}}(\mathbf{v}, \mathbf{k})$ we have made the following general assumption:

$$P(\mathbf{v}, \mathbf{k}) \propto \exp\left(-\|F(\mathbf{v}) - F(\mathbf{k})\|^2\right)$$
 (8)

where *F* is a feature vector representing tissue property of a voxel. Assuming that z represents the intensity of the HR image with a different contrast coregistered to the LR image, *F* could be set to z or a vector of z values in a neighborhood of the input voxel, e.g., a patch [11, 13–15]. However, such feature vector does not take the brain laminar structure into account and is not rotation-invariant. To capture the laminar structure, voxels having equal distances to the main edges should have similar *F* values. Equivalently, *F* contour maps should have isolines parallel to the main edges. The key point is to propagate the edge information perpendicular to the edge direction. The gradient operator can be used to extract edges. Fig. 5 (b) shows the contour map of a gradient operator applied to the sample image in Fig. 5 (a) (before applying the gradient operator, the image is slightly smoothed). The edge information, however, stays local in a gradient operator. This causes the information of sharp curves to be retained but the edge information is lost as one deviates from the edges.

To further propagate the edge information, we propose using a multi-resolution Gaussian filtering approach. In this approach, the image is filtered by a set of Gaussian kernels with different standard deviations. Figs. 5 (c)–(d) show the contour maps of the filtered image in Fig. 5 (a) using Gaussian kernels with two different standard deviations. As shown, by increasing the standard deviation, the edge information is further propagated but detailed information of smaller structures is lost. A combination of the above operators in the similarity vector F may be used in order to retain both high and low resolution information. We propose to use the following feature vector:

$$F_{\mathbf{z}}(\mathbf{v}) = \left(\mathbf{z}(\mathbf{v}), \left|\nabla \mathbf{z}(\mathbf{v})\right|, \mathbf{z} * \mathbf{g}_{1}|_{\mathbf{v}}, \mathbf{z} * \mathbf{g}_{2}|_{\mathbf{v}}, \dots, \mathbf{z} * \mathbf{g}_{k}|_{\mathbf{v}}\right)^{T}$$
(9)

where \forall is gradient operator, * is convolution operator, \mathbf{g}_i , *i*=1,2,...,*k* are Gaussian kernels, and *k* is the number of kernels. The required standard deviations of the Gaussian kernels depend on the desired voxel size for the upsampled image. Fig. 3 compares the similarity map of the proposed method for *k*=2 with that of a patch-based method. As shown, the proposed method has superior performance in finding voxels with similar tissue types in a laminar structure. Calculation of this similarity measure is faster than a patch-based similarity measure if the size of vector *F* is small as it requires fewer mathematical operations.

C. Algorithm

An iterative method is used for the minimization in (6). In this method, in iteration *t*+1, the image $\hat{\mathbf{x}}_{rec}^{t+1}$ is reconstructed by (1) from $\hat{\mathbf{x}}^t$ and the condition in (4) is enforced by:

$$\hat{\mathbf{x}}^{t+1} = \hat{\mathbf{x}}_{\text{rec}}^{t+1} - \text{NN} \left(H \hat{\mathbf{x}}_{\text{rec}}^{t+1} - \mathbf{y} \right)$$
(10)

where NN is the nearest neighbor interpolation operator. This constraint has been referred to as subsampling consistency [21]. To take into account the inconsistencies between the LR and HR weights, we apply a self-similarity approach based on the concepts from image denoising [16–18]. The values of $w_x(\mathbf{v}, \mathbf{k})$ are generated using both HR and LR images as follows:

$$w_{\mathbf{x}}\left(\mathbf{v},\mathbf{k}\right) = \frac{1}{N} \exp\left(-\alpha_{\mathbf{z}} \|F_{\mathbf{z}}\left(\mathbf{v}\right) - F_{\mathbf{z}}\left(\mathbf{k}\right)\|^{2} - \alpha_{\mathbf{x}} \|F_{\mathbf{x}}\left(\mathbf{v}\right) - F_{\mathbf{x}}\left(\mathbf{k}\right)\|^{2}\right)$$
(11)

where a_x and a_z equalize the contribution of the features of LR and HR images when the range of intensity values of F_z and F_x are different. Consider a situation in which a lesion is visible in the LR image but not in the HR image, e.g., F_z values are similar in a local neighborhood whereas F_x values are not. In this case, the LR image has dominant effect in equation (11) and the result of upsampling becomes similar to a self-similarity based denoising approach.

The most time-consuming part of the algorithm is the generation of $w_{\mathbf{x}}(\mathbf{v}, \mathbf{k})$ values. This makes the algorithm very slow if we recalculate the weights in each iteration. On the other hand, iteratively taking the LR image into account in the calculation of weights may result in suboptimal solutions because the structural information of the reconstructed LR image (which may not be accurate in the first place) may gradually dominate and reduce the accuracy by increasing iterations. To simultaneously expedite the process and avoid suboptimal solutions, we first calculate the weights $w_x(\mathbf{v}, \mathbf{k})$ using only the HR image (i.e., by setting $a_x = 0$). The LR image is then upsampled using these fixed weights in an iterative approach. Note that once $w_{\mathbf{x}}(\mathbf{y}, \mathbf{k})$ values are known, the minimization in (6) is done by the iterative approach, which is relatively fast. The upsampled LR image then becomes more similar to the correct result and thus can better contribute to the upsampling. In the next step, a non-zero $a_{\mathbf{x}}$ is used and the image is upsampled iteratively with new fixed weights $w_{\mathbf{x}}(\mathbf{v},$ **k**). Thus, the weights $w_{\mathbf{x}}(\mathbf{v}, \mathbf{k})$ are calculated twice in total. In addition, to save memory and time, we only keep the N_v highest $w_x(\mathbf{v}, \mathbf{k})$ values by assuming that there are N_v voxels with similar properties to the voxel of interest, and that the rest of voxels are irrelevant. N_v can be adjusted based on the neighborhood size and voxel size.

An overview of the algorithm is presented below. First a NN operator interpolates the image to generate an initial estimation of the upsampled image. This is equivalent to repeating each slice *n*-1 times when we would like to divide the slice thickness by *n*. We may want to insert enough number of slices to make the voxels of the resulting image near-isotropic. The HR image is then coregistered to the resulting image and the values of $w_x(\mathbf{v}, \mathbf{k})$ are estimated for each voxel.

1. Use NN operator to generate initial estimation $\mathbf{\hat{x}}^0$

- 2. Coregister the HR image to $\mathbf{\hat{x}}^0$ and call it \mathbf{z}
- **3.** Set $a_{\mathbf{x}} = 0$
- 4. Estimate the weighs $w_{\mathbf{x}}(\mathbf{v}, \mathbf{k})$ using \mathbf{z} and $\mathbf{\hat{x}}^t$
- 5. Repeat until $||\mathbf{\hat{x}}^{t+1} \mathbf{\hat{x}}^{t}||^2 < \varepsilon$
 - **a.** Reconstruct $\mathbf{\hat{x}}^{t+1}$ from $\mathbf{\hat{x}}^t$ using equation (1)
 - **b.** Correct $\mathbf{\hat{x}}^{t+1}$ for downsampling and partial volume effects by applying constraint (10)
- 6. Repeat steps 4–5 with a non-zero $a_{\mathbf{x}}$

III. Results

Experiments were conducted on artificially downsampled images and clinically acquired LR datasets. Assuming that partial volume effect can be modeled by averaging, downsampling by a factor of n in one direction was performed by averaging every n slices of the HR image **x**. This is equivalent to the application of operator H in (4). The downsampled image was then reconstructed by the HR image **z** and evaluated by two quality measures, namely Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM). PSNR is defined by:

$$\operatorname{PSNR}\left(\hat{\mathbf{x}}\right) = 10 \log_{10} \left(\frac{d^2}{\operatorname{MSE}\left(\hat{\mathbf{x}}\right)}\right) \quad (12)$$

where $\mathbf{\hat{x}}$ is the reconstructed image, $MSE(\mathbf{\hat{x}}) = 1/|\Omega| \cdot \sum_{\mathbf{v} \in \Omega} (\mathbf{x} (\mathbf{v}) - \mathbf{\hat{x}} (\mathbf{v}))^2$, Ω is the region covering the organ, i.e., the whole image excluding background, and *d* is the actual dynamic range of the image, i.e., $d = \max(\mathbf{x}) - \min(\mathbf{x})$. SSIM [22] is a measure more consistent with the human visual system. For images \mathbf{x} and $\mathbf{\hat{x}}$, SSIM is defined as:

$$\operatorname{SSIM}\left(\mathbf{x}, \hat{\mathbf{x}}\right) = \frac{\left(2\mu_{\mathbf{x}}\mu_{\hat{\mathbf{x}}}\right)\left(2\sigma_{\mathbf{x}\hat{\mathbf{x}}} + c_{2}\right)}{\left(\mu_{\mathbf{x}}^{2} + \mu_{\hat{\mathbf{x}}}^{2} + c_{1}\right)\left(\sigma_{\mathbf{x}}^{2} + \sigma_{\hat{\mathbf{x}}}^{2} + c_{2}\right)} \quad (13)$$

where $\mu_{\mathbf{x}}$ and $\mu_{\mathbf{\hat{x}}}$ are the mean value of images \mathbf{x} and $\mathbf{\hat{x}}$, $\sigma_{\mathbf{x}}$ and $\sigma_{\mathbf{\hat{x}}}$ are the standard deviation of images \mathbf{x} and $\mathbf{\hat{x}}$, $\sigma_{\mathbf{x}\mathbf{\hat{x}}}$ is the covariance of \mathbf{x} and $\mathbf{\hat{x}}$, $c_1 = (K_1d)^2$ and $c_2 = (K_2d)^2$ ($K_1 = 0.01$ and $K_2 = 0.03$). For image quality assessment, it is useful to apply SSIM index locally, rather than globally. To this end, a spatially varying quality map of the image is generated for images \mathbf{x} and $\mathbf{\hat{x}}$, and the average SSIM is calculated. SSIM ranges from -1 to 1 and the value 1 is only reachable in the case of two identical images. For clinical applications (where the original HR image \mathbf{x} was not available), the images were assessed visually.

Experimentally, N_v was set to 10 to maximize PSNR and minimize computation time and memory usage for one sample MR image (shown later in the experiments). The effect of parameter *a* has been previously evaluated [16]. We set *a* to $1/2\mu^2$ where μ is the mean

value of the absolute values of the image (which makes the algorithm work for images with negative values as well). This choice of *a* normalizes the intensity values and results in accurate outcomes independent of the range of intensities. We also used a 6-parameter affine coregistration based on normalized mutual information and with trilinear interpolation in Statistical Parameter Mapping software package (SPM8; http://www.fil.ion.ucl.ac.uk/spm/) for the image coregistration in step 2. Furthermore, with smaller values of *k* (the number of filters), the method becomes faster in each iteration although the total computation time depends on the number of iterations as well. In this paper, we set *k*=2 with Gaussian kernels with standard deviations of $2h/p_i$ and $5h/p_i$ (in voxel space) along dimension *i*, where $h=1/\left(4\sqrt{2\ln 2}\right)$ and p_i is the voxel size along dimension *i*. These parameters were determined empirically to provide accurate results. In order to expedite the calculations, we

used multi-threading implementation in both the calculation of the weights as well as reconstruction. All experiments were done on a quad-core Xeon 5472 3.0GHz CPU.

In the following experiments, we evaluated performance of the proposed method in comparison with SSIP (also with multi-threading implementation), a state-of-the-art method. In order to compare the methods in different imaging conditions and in the presence of lesions with different sizes, we selected images with different contrasts and resolutions from three datasets of patients with brain tumors, multiple sclerosis, and epilepsy. The Institutional Review Board approved this study. We also tested the algorithm on publicly available MRI dataset of schizophrenic patients and normal controls provided by National Alliance for Medical Image Computing (NAMIC) as well as BrainWeb MRI. A fixed set of upsampling parameters was used for the proposed method in all experiments. These parameters have been optimized for a single image in the brain tumor dataset. We also evaluated the effect of upsampling parameters, slice thickness, noise, intensity nonuniformity, as well as presence of motion and susceptibility artifacts, and showed the application of the proposed method in upsampling of ROIs. We limited the number of iterations in SSIP method to 30 in all experiments to avoid convergence to a suboptimal solution. This number was selected empirically by a compromise between accuracy and computation time.

A. Comparison with previous work

i) Brain tumor—We selected 5 patients with brain tumors from our dataset of glioblastoma patients [23, 24]. Each patient had sagittal T2-weighted (T2W) and Gadolinium-enhanced MPRAGE images, both with 1mm isotropic voxel size (matrix size $256 \times 256 \times 176$). For each subject, the T2W image was downsampled axially to generate a LR image with the slice thickness of 6 mm, which is a common slice thickness for clinically acquired LR images. The downsampled image was then upsampled using the Gadolinium-enhanced MPRAGE image. We also upsampled the T2W image using NN, linear, cubic, and spline interpolation techniques, as well as SSIP.

Fig. 6 shows the results of upsampling using four different methods, namely NN interpolation, spline interpolation, SSIP, and our proposed method. For a better comparison, a smaller field of view of each image is displayed. As shown, the standard interpolation techniques result in distorted edges. However, SSIP and the proposed method better

reconstruct the edges. Fig. 7 shows a magnified section of two images presented in Fig. 6 to better illustrate the difference between the performances of SSIP and the proposed method. As shown, the proposed method has visually superior performance, particularly near the edges. This is mainly because of the better performance of the proposed feature-based similarity measure compared to the patch-based method, which was illustrated in Fig. 3.

Table I compares the PSNR and SSIM values of different upsampling methods applied to the above images. The table also presents the computation times. The proposed method was significantly faster and more accurate than SSIP. Average computation time for the proposed method was 6 minutes whereas the average computation time for SSIP was 84 minutes. Nevertheless, both SSIP and the proposed methods outperformed standard interpolation techniques.

Fig. 8 (a) presents the PSNR curve for the image presented in the first column of Fig. 6 (subject 1 in Table I). To better show the trend of changes in PSNR values of the SSIP method, we have presented the PSNR values for 60 iterations. The proposed method calculates the weights only twice, which correspond to the two exponential-like parts of the curve of the proposed method in Fig. 8. The proposed method outperformed SSIP even if the weights were extracted only from the HR image. As shown, SSIP has decreasing performance after seven iterations while the termination condition is not met. This may be due to the problem with iteratively taking the LR image into account explained in the Methods section. SSIP had similar decreasing performance in the rest of images in Fig. 6. Note that each iteration of the SSIP method takes significantly longer than each iteration of the proposed method as SSIP requires the weights to be calculated iteratively. In the case presented in Fig. 8 (a), maximum PSNR for SSIP method was achieved in seven iterations and took about 19 minutes, which is still longer than the 6 minutes using the proposed method.

To test whether the assumption of laminar structures of the MR images is also more accurate than the patch-based similarity assumption within a lesion, we calculated the values of PSNR and SSIM within the tumors. Table II reports the values. As shown, the proposed method has higher values of PSNR and SSIM in all subjects.

ii) Multiple sclerosis—We selected 5 patients with multiple sclerosis. Each patient had a sagittal FLAIR image with 0.5mm×0.5mm×1mm voxel size (matrix size 384×384×160) and a sagittal MPRAGE image with 0.9 mm isotropic voxel size (matrix size 256×256×192). The FLAIR image was downsampled sagittally to generate a LR image with the slice thickness of 6 mm (i.e., 0.5mm×0.5mm×6mm voxel size). Similar to the previous experiment, the downsampled FLAIR image was upsampled by MPRAGE image using the proposed method as well as SSIP and four standard interpolation techniques. Table III compares the PSNR and SSIM values and computation times of different upsampling methods. Both SSIP and the proposed method is significantly faster and is more accurate than SSIP. Fig. 9 compares a sample slice of the FLAIR image of subject 1 in Table 3 upsampled using SSIP and the proposed method. As shown, small lesions that are not visible in the images upsampled by NN and spline techniques are visible in the upsampled image using

the proposed method as well as SSIP method, even though the lesions in the MPRAGE and FLAIR images do not always have similar patterns of contrasts. Fig. 8 (b) presents the PSNR curve for subject 1 in Table 3 for SSIP and the proposed method.

iii) Epilepsy—We selected 5 patients with epilepsy. Each patient had a sagittal FLAIR image (matrix size $192 \times 192 \times 176$), a sagittal T2W image (matrix size $192 \times 192 \times 160$), and a sagittal MPRAGE image (matrix size 256×256×176) - all with 1 mm isotropic voxel size. In one experiment, the FLAIR image was downsampled coronally to generate a LR image with slice thickness of 6 mm and then was upsampled using the T2W image. In another experiment, the downsampled FLAIR image was upsampled using the MPRAGE image. Table IV compares the PSNR and SSIM values and computation times of different upsampling methods applied to the above images. Similar to the previous experiments, both SSIP and the proposed method outperformed standard interpolation techniques. However, the proposed method was significantly faster and more accurate than SSIP. In this dataset, we did not find a significant difference in using MPRAGE and T2W images for upsampling the FLAIR image based on PSNR. However, the images upsampled using T2W images had generally higher values of SSIM. Fig. 10 presents the upsampling results using NN interpolation, spline interpolation, SSIP, and the proposed method for subject 1 in Table IV. Although the MPRAGE image had better overall performance compared to T2W image in this subject, T2W image had better reconstruction in the substantia nigra and red nucleus region since it had better contrast in this region. Fig. 8 (c) presents the PSNR curves for this subject using SSIP and the proposed method upsampled by T2W and MPRAGE images.

iv) NAMIC Brain Multimodality dataset-NAMIC Brain Multimodality (http:// hdl.handle.net/1926/1687) is a publicly available MRI dataset of 10 schizophrenic patients and 10 normal controls. Each subject has T1W and T2W images with 1 mm isotropic resolution (matrix size $256 \times 256 \times 176$). In this experiment, we were interested in evaluating the effects of noise and intensity non-uniformity on upsampling. The T2W image was first downsampled axially to generate a LR image with the slice thickness of 9 mm. We added Rician noise with levels of 0% and 5% (relative to the brain mean intensity) to both LR and HR images. We also applied intensity non-uniformity to the HR image, modeled by a 3D Gaussian function. The downsampled T2W image was then upsampled by the T1W image using the proposed method as well as SSIP and four standard interpolation techniques. We tested the upsampling algorithm with and without noise and intensity non-uniformity, comprising 80 upsampling operations. In each experiment, similar levels of noise were applied to both LR and HR images. For a fair comparison, we denoised the LR images in all methods in the same way the LR images were denoised in SSIP method. Table V presents the mean and standard deviation values of PSNR and SSIM over all 20 subjects. The upsampled images were compared with noise-free images. As shown, the proposed method outperformed all other methods. It also outperformed SSIP in computation time. Furthermore, as expected, noise and intensity non-uniformity generally decreased both PSNR and SSIM values. However, the proposed method was less affected by noise and intensity non-uniformity compared to SSIP. Fig. 11 presents a sample upsampling result with %5 noise level in the LR and HR images and intensity non-uniformity applied to the HR image. Fig. 8 (d) presents the PSNR curves for this image.

v) BrainWeb—BrainWeb [19] (http://brainweb.bic.mni.mcgill.ca/brainweb/) is a simulated brain MRI database providing two anatomical models: normal and multiple sclerosis (MS). MRI scans have been simulated using T1W, T2W, and proton-density-weighted (PDW) sequences. We were interested in evaluating the performance of the algorithm in upsampling along all three dimensions. We downsampled the T2W and PDW images of multiple sclerosis model to 3 mm, 6 mm, and 9 mm isotropic resolutions and then upsampled by T1W image to 1 mm isotropic resolution. Table VI presents the values of PSNR and SSIM for standard interpolation methods as well as the proposed method. The proposed method outperformed all standard methods. Fig. 12 displays the results of upsampling the T2W and PD images with 9 mm isotropic resolution.

In all the above experiments, the LR and HR images had different contrasts (e.g., LR T2W and HR T1W images). However, the method works on images with similar contrasts as well, and indeed due to their contrast similarity, it should result in higher PSNR. In practice, it is uncommon to have LR and HR images from the same contrast. Nevertheless, this may find application in upsampling dynamic images given a static HR image with similar contrast. In one experiment, we upsampled the BrainWeb T2W and PD images with 6 mm isotropic resolution using the T2W and PD images with 1 mm isotropic resolution and obtained PSNR (SSIM) values of 21.78 (0.893) and 21.99 (0.927), respectively, compared to the values of 19.59 (0.807) and 20.01 (0.853) in Table VI.

B. Effects of parameters and slice thickness

For the image presented in the first column of Fig. 6 (subject 1 in Table I), we evaluated the effects of the number of voxels N_{ν} , neighborhood size, number of Gaussian kernels and their standard deviations, parameter e, and slice thickness. The results are presented in Fig. 13. Choosing a small N_{ν} makes the method faster and requires less memory to store the weight array. As shown in Fig. 13 (a), the best performance is achieved with $N_{\nu}=10$, although PSNR is at most reduced only by 0.2 using other values of N_{ν} . The optimal value of N_{ν} may, however, change with the desired resolution for the upsampled image. The use of only 10 voxels for reconstruction makes the method less memory demanding. Both SRIP and SSIP use a patch of $7 \times 7 \times 7$ containing 343 voxels, which requires a large amount of memory to store the weight array for the entire image. In SSIP, however, the weights do not need to be stored in memory as they are used only once. Still, the number of calculations is very high for both SRIP and SSIP methods.

The effect of neighborhood size on PSNR is shown in Fig. 13 (b) for $N_v=10$. As shown, when increasing the neighborhood size, the PSNR ratio saturates, although the computation time increases as well. We chose to use a neighborhood size of 7 mm (i.e., a cube of size $7 \times 7 \times 7$ around the voxel of interest in the image with $1 \text{mm} \times 1 \text{mm} \times 1 \text{mm}$ resolution). Similar to N_v , the optimal neighborhood size may change with the desired resolution for the upsampled image.

To evaluate the effect of the number of Gaussian kernels and their standard deviations, we evaluated six filters with standard deviations of $\sigma_n = (2n + 1)h$ and filter size $(2n + 1)\times(2n + 1)\times(2n + 1)$, where n = 1, 2, ..., 6 and $h = 1/(4\sqrt{2\ln 2})$. In one experiment, single Gaussian

kernels (*k*=1) with standard deviations of σ_n , n = 1,2,...,6 were used to evaluate the effect of standard deviation on the upsampling accuracy. The results are presented in Fig. 13 (c). As shown, PSNR changed with changing the standard deviation. However, the change in PSNR was below 0.2 for the specified standard deviations. In another experiment, sets of Gaussian kernels with standard deviations of $\{\sigma_1\}$, $\{\sigma_1, \sigma_2\}$, $\{\sigma_1, \sigma_2, \sigma_3\}$, $\{\sigma_1, \sigma_2, \sigma_3, \sigma_4\}$, $\{\sigma_1, \sigma_2, \sigma_3, \sigma_4, \sigma_5\}$, and $\{\sigma_1, \sigma_2, \sigma_3, \sigma_4, \sigma_5, \sigma_6\}$ were used to evaluate the effect of the number of kernels (i.e., *k*). The results are presented in Fig. 13 (d). As shown, the combination of different Gaussian kernels changed the PSNR but, similar to the previous experiment, the change was below 0.2. In this experiment, the use of *k*=6 kernels doubled the computation time in comparison with *k*=2 kernels. However, in general, the total computation time will depend on both the number of iterations and the duration of each iteration. Thus, a more complex and time consuming calculation of the weights does not necessarily result in a slower algorithm as fewer iterations may be required to reach a solution.

In another experiment on subject 1 from the brain tumor dataset, we changed the threshold e to alter the start point of the first weight update and also repeated step 6 of the algorithm twice, resulting in three weight updates. The PSNR curves with two weight updates are presented in Fig. 13 (e) (PSNR after the third weight update was not shown for better visualization). The final PSNR slightly decreased with increasing e, and the computation time increased due to an additional weight update. However, the third weight update did not notably change the PSNR (less than 0.08 change) and actually slightly decreased PSNR for e = 0.05 and e = 0.03. Thus two weight updates were sufficient to achieve reasonable PSNR.

Fig. 13 (f) shows how performance of the proposed method changes with different slice thicknesses in comparison with SSIP and spline interpolation techniques. As expected, PSNR decreases with increased slice thickness. The difference between the proposed method and spline technique increases at larger slice thicknesses. The same is true for the SSIP method, though it has lower PSNR compared to the proposed method.

C. Effects of motion and susceptibility artifacts

We were interested in the performance of our method in the presence of two common artifacts, motion artifact in the HR image and susceptibility artifact causing signal drop and geometric distortion in the LR image. In one experiment (Fig. 14), we selected a HR MPRAGE image with 1mm×1mm×1mm resolution with motion artifact and used it to upsample a LR (downsampled) T2W image with a 1mm×1mm×6mm voxel size. As shown, the result is still visually acceptable, though the structural information in the HR image was limited. It also provided higher values of PSNR and SSIM (PSNR=27.93, SSIM=0.773) compared to NN interpolation (PSNR=26.88, SSIM=0.735), linear (PSNR=27.41, SSIM=0.738), cubic (PSNR=27.72, SSIM=0.755), spline (PSNR=27.77, SSIM=0.757), and SSIP (PSNR=26.25, SSIM=0.710) techniques.

In another experiment, we upsampled spin echo (SE) and gradient echo (GE) baseline DSC images (median of 4D images over time) using a HR T2W image (1mm isotropic). The DSC images, and in particular GE DSC, had susceptibility artifact appearing as signal drop and geometric distortion. The resolution of baseline DSC images was 1.2mm×1.2mm×6.5mm, which was upsampled to 1.2mm×1.2mm×1.08mm. The results obtained by NN

interpolation, SSIP method, and the proposed method are presented in Fig. 15. The upsampled SE DSC images using SSIP and the proposed method have visually more details than the ones upsampled using NN method. While the upsampled GE DSC images have more details compared to the upsampled ones using NN method, they have some artifacts near the cortex edges due to higher geometric distortion in this image (resulting in inaccurate coregistration), particularly near the temporal lobe. On the other hand, there is inter-slice intensity variation appearing as dark and bright slices in the LR image, which is also present in the upsampled images using both SSIP and the proposed method.

D. Clinical application

We modeled the partial volume effect with simple averaging. In practice, this assumption may not be valid. Fig. 15 presented the results for upsampling of real clinical images (i.e., not artificially downsampled). In another experiment, we used the proposed algorithm to upsample a FLAIR image with 0.43mm×0.43mm×6mm resolution using a MPRAGE image as well as a T2W image, both with 1mm×1mm×1mm resolution. The upsampled images have 0.43mm×0.43mm×0.43mm×1mm resolution and are displayed in Fig. 16. As shown, the proposed method has a visually superior performance compared to the standard interpolation methods. In particular, some details in the GM and WM boundaries within the cortex are now visible in the upsampled image. Note that the FLAIR image had higher in-plane resolution (0.43mm×0.43mm) compared with the MPRAGE and T2W images (1mm×1mm). Nevertheless, the upsampling algorithm had an acceptable performance. As shown, the choice of HR images had little impact on the final outcome. However, a comparison of the two images' details reveals that each upsampled image has its own advantages and disadvantages.

E. ROI upsampling

The proposed method can be used to upsample ROIs manually drawn on LR images. Note that due to the large number of slices, manual drawing of ROIs in a HR image is a tedious task. On the other hand, a reliable automated tool to outline the regions of interest may not exist. Thus, in practice, the outlines may have to be manually drawn on the LR images. The proposed method could be used to upsample ROIs using a HR image with a similar contrast. The algorithm to upsample ROIs is similar to the one introduced above with two differences: (i) partial volume condition (i.e., (10)) is not enforced as the ROI values are all constant, (ii) the weights are calculated only once using the HR image (i.e., step 6 of the algorithm is eliminated). The resulting image has a range of values which needs to be thresholded at 0.5 (if ROI values are 1) in order to get the final upsampled ROI.

Fig. 17 shows the result of ROI upsampling on two subjects with brain tumors. In one subject, the tumor ROI was drawn on post-contrast T1W image with 0.43mm×0.43mm×6mm resolution and then upsampled using a HR post-contrast MPRAGE image (1mm×1mm×1mm resolution). In another subject, the edema ROI was drawn on FLAIR image with 0.43mm×0.43mm×6mm resolution and then upsampled using a HR T2W image (1mm×1mm×1mm resolution). The results are overlaid on the HR images.

Note that segmentation algorithms that work based on the intensity of region may not provide good results here as the region may have varying intensity. However, the proposed upsampling works even if the outlined regions present varying intensity.

IV. Discussion

A new method was proposed to upsample LR images. The proposed method outperformed SSIP, a state-of-the-art upsampling method, in PSNR, SSIM, and computation time. The entire upsampling takes about 6 minutes for a $256 \times 256 \times 176$ image on a quad-core Xeon 5472 3.0GHz CPU. The upsampling parameters were optimized based on the observations in a sample image and were fixed in all other experiments. This showed the method had a robust performance when image resolution or contrast changed. The method successfully worked on the LR images, both with anisotropic and isotropic voxel sizes. In addition, we showed how the method could be applied to upsample ROIs.

We used features based on image intensity, gradient, and multi-resolution Gaussian filtering. Other feature vectors may be proposed that can better estimate the similarity of the voxels on a LR image with laminar structures. However, the proposed features are simple and produce reasonably accurate results in a reasonable time. Although the assumption of having laminar structures is valid in healthy brain tissues, the results showed the method had superior performance in the abnormal tissues as well. The proposed method had decreased but still acceptable performance in the presence of noise, intensity non-uniformity and motion artifact in the HR image, and susceptibility artifact in the LR image. It was also observed that intensity artifacts in the LR image will remain in the upsampled image.

For the majority of subjects, the SSIP method had decreased PSNR with increasing iterations and thus we limited the number of iterations to 30, resulting in an average of 84 minutes of computation time for patients with brain tumors. Our algorithm is faster than the SSIP algorithm because it only requires two updates to the weights instead of each iteration, and this calculation is the most time consuming part. One may want to keep SSIP weights constant to speed up the algorithm, but since the weight data, which is a large amount of data, would need to be saved in the memory, its access may be very slow, resulting in an undesirable decrease in speed.

It is required for the HR image to have some structural similarity to the LR image. In the presence of multiple HR images, we may choose the one that is closest in contrast to the LR image. However, the proposed method may also be extended in future work to combine several HR images to generate more accurate weights. The weights may be generated by multiplying the $P(\mathbf{v}, \mathbf{k})$ values of the two HR images.

In the presence of large geometric distortion in the LR image due to susceptibility artifact, the performance of the proposed method decreases. This is because the calculated weights using the HR image do not correspond to the correct voxels in the LR image. To reduce this effect, we may use a non-rigid registration technique to register the HR image to the LR image [25]. The registration needs to be smooth enough for the HR image to avoid following the blocky edges of the LR image interpolated by the NN algorithm (step 2 of the proposed

algorithm). After upsampling, the inverse registration may be applied to correct for geometric distortion of the upsampled image. This, however, does not correct the signal drop due to susceptibility artifact. We also showed the performance of the proposed method on the baseline image of LR dynamic MRI scans. The method may also be applied to each time point of the dynamic image but must be fast to reduce the total reconstruction time.

Another application of the proposed method is in the coregistration of LR images. The LR images may be upsampled by a HR image separately and then coregistered to each other. This may have application for the accurate correction of motion artifacts in dynamic MRI scans.

Acknowledgments

This work was supported in part by the National Institutes of Health under Advanced Multimodal Neuroimaging Training Program Grant R90-DA023427, Grant 5U01CA154601-03, Grant 2P41EB015896-16, Grant R01CA129371, and Grant K24CA125440A.

We would like to thank Tracy Batchelor for providing the MRI data of glioblastoma patients, Caterina Mainero and Sindhuja Tirumalai Govindarajan for providing the MRI data of multiple sclerosis patients, Steven M. Stufflebeam and Naoaki Tanaka for providing the MRI data of epilepsy patients, and Elizabeth Gerstner for providing ROIs in the ROI upsampling experiments. We would like to thank José V. Manjón and colleagues for providing their upsampling codes and helping with using the codes (SSIP method) and Alireza Akhondi-Asl, Ali Gholipour, and Jayashree Kalpathy-Cramer for their useful comments.

References

- Frakes DH, Dasi LP, Pekkan K, Kitajima HD, Sundareswaran K, Yoganathan AP, Smith MJ. A new method for registration-based medical image interpolation. IEEE Trans Med Imag. Mar; 2008 27(3): 370–7.
- Penney GP, Schnabel JA, Rueckert D, Viergever MA, Niessen WJ. Registration-based interpolation. IEEE Trans Med Imag. Jul; 2004 23(7):922–926.
- Ólafsdóttir H, Pedersen H, Hansen MS, Lyksborg M, Hansen MF, Darkner S, Larsen R. Registration-based interpolation applied to cardiac MRI. SPIE Medical Imaging. 2010;762336-762336–8.
- Laursen, LF., Ólafsdóttir, H., Bærentzen, JA., Hansen, MS., Ersbøll, BK. Registration-based interpolation real-time volume visualization. Proceedings of the 28th Spring Conference on Computer Graphics; 2012. p. 15-21.
- 5. Leng J, Xu G, Zhang Y. Medical image interpolation based on multi-resolution registration. Computers & Mathematics with Applications. Aug; 2013 66(1):1–18.
- Ólafsdóttir, H., Pedersen, H., Hansen, MS., Larsson, H., Larsen, R. Improving image registration by correspondence interpolation. Biomedical Imaging: From Nano to Macro, 2011 IEEE International Symposium on; 2011. p. 1524-1527.
- Alipour, S., Wu, X., Shirani, S. Context-based interpolation of 3-D medical images. Conference proceedings: Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Conference; 2010. p. 4112-5.
- Neubert A, Salvado O, Acosta O, Bourgeat P, Fripp J. Constrained reverse diffusion for thick slice interpolation of 3D volumetric MRI images. Computerized medical imaging and graphics : the official journal of the Computerized Medical Imaging Society. Mar.2012 36:130–8. [PubMed: 21920702]
- 9. Rueda A, Malpica N, Romero E. Single-image super-resolution of brain MR images using overcomplete dictionaries. Medical image analysis. Jan.2013 17:113–32. [PubMed: 23102924]
- Manjon JV, Coupe P, Buades A, Fonov V, Louis Collins D, Robles M. Non-local MRI upsampling. Medical image analysis. Dec.2010 14:784–92. [PubMed: 20566298]

- Buades A, Coll B, Morel J-M. A review of image denoising algorithms, with a new one. Multiscale Modeling & Simulation. 2005; 4:490–530.
- Zhang Y, Wu G, Yap P, Feng Q, Lian J, Chen W, Shen D. Hierarchical Patch-Based Sparse Representation-A New Approach for Resolution Enhancement of 4D-CT Lung Data. IEEE Trans Med Imag. Nov.2012 31(11)
- 13. Rousseau, F. Springer. Computer Vision-ECCV 2008. 2008. Brain hallucination; p. 497-508.
- Rousseau F. A non-local approach for image super-resolution using intermodality priors. Medical image analysis. Aug.2010 14:594–605. [PubMed: 20580893]
- Manjon JV, Coupe P, Buades A, Collins DL, Robles M. MRI superresolution using self-similarity and image priors. International journal of biomedical imaging. 2010; 2010:425891. [PubMed: 21197094]
- Coupe P, Yger P, Prima S, Hellier P, Kervrann C, Barillot C. An optimized blockwise nonlocal means denoising filter for 3-D magnetic resonance images. IEEE Trans Med Imag. Apr; 2008 27(4):425–41.
- Manjon JV, Carbonell-Caballero J, Lull JJ, Garcia-Marti G, Marti-Bonmati L, Robles M. MRI denoising using non-local means. Medical image analysis. Aug.2008 12:514–23. [PubMed: 18381247]
- Protter M, Elad M, Takeda H, Milanfar P. Generalizing the nonlocal-means to super-resolution reconstruction. IEEE Trans Image Process. Jan; 2009 18(1):36–51. [PubMed: 19095517]
- Cocosco, CA., Kollokian, V., Kwan, RK-S., Evans, AC. BrainWeb: Online interface to a 3D MRI simulated brain database. NeuroImage, Proceedings of 3-rd International Conference on Functional Mapping of the Human Brain; Copenhagen. 1997.
- 20. Mignotte M. A non-local regularization strategy for image deconvolution. Pattern Recognition Letters. 2008; 29:2206–2212.
- Banerjee, J., Jawahar, C. Super-resolution of text images using edge-directed tangent field. Document Analysis Systems, 2008. DAS'08. The Eighth IAPR International Workshop on; 2008. p. 76-83.
- Wang Z, Bovik AC, Sheikh HR, Simoncelli EP. Image quality assessment: from error visibility to structural similarity. IEEE Trans Image Process. Apr; 2004 13(4):600–612. [PubMed: 15376593]
- 23. Batchelor TT, Sorensen AG, di Tomaso E, Zhang WT, Duda DG, Cohen KS, Kozak KR, Cahill DP, Chen PJ, Zhu M, Ancukiewicz M, Mrugala MM, Plotkin S, Drappatz J, Louis DN, Ivy P, Scadden DT, Benner T, Loeffler JS, Wen PY, Jain RK. AZD2171, a pan-VEGF receptor tyrosine kinase inhibitor, normalizes tumor vasculature and alleviates edema in glioblastoma patients. Cancer cell. Jan.2007 11:83–95. [PubMed: 17222792]
- 24. Batchelor TT, Duda DG, di Tomaso E, Ancukiewicz M, Plotkin SR, Gerstner E, Eichler AF, Drappatz J, Hochberg FH, Benner T, Louis DN, Cohen KS, Chea H, Exarhopoulos A, Loeffler JS, Moses MA, Ivy P, Sorensen AG, Wen PY, Jain RK. Phase II study of cediranib, an oral panvascular endothelial growth factor receptor tyrosine kinase inhibitor, in patients with recurrent glioblastoma. Journal of clinical oncology : official journal of the American Society of Clinical Oncology. Jun 10.2010 28:2817–23. [PubMed: 20458050]
- Heinrich, MP., Jenkinson, M., Brady, M., Schnabel, JA. Robust Super-Resolution Reconstruction with Multi-Modal Registration and Guidance. Proceedings of the Medical Image Understanding and Analysis; Swansea, UK. 2012. p. 81-86.



Fig. 1.

A HR T2W image with 1mm×1mm×1mm resolution was downsampled to 1mm×1mm×6 mm by averaging 6 consecutive slices to simulate partial volume effect and then interpolated using (a) NN, (b) linear, and (c) cubic techniques and compared with (d) the original HR T2W image.



Fig. 2.

Laminar structure of cortex, courtesy of David H. Hubel (http://hubel.med.harvard.edu/). As shown, the cortex is made of layers of cells that form a laminar structure.



Fig. 3.

Comparison of the performances of similarity measurements by similarity maps. Higher values show more similarity. (a) original image with a point of interest shown by a cross sign, (b) similarity map showing similarity of image pixels with the point of interest using a 3×3 patch as proposed in [13], (c) similarity map using our proposed method.



Fig. 4.

Block diagram of the proposed upsampling method. The weights $w(\mathbf{v}, \mathbf{k})$ are estimated from both HR and LR images during minimization of $||\mathbf{y} - H\mathbf{x}||^2 + \lambda R(\mathbf{x})$.



Fig. 5.

Performance of the proposed features on a sample image with laminar structure (a), showing contour lines for gradient operator (b), and the same image after Gaussian filtering (c–d) with two different standard deviations with (d) having larger standard deviation compared to (c).



Fig. 6.

Upsampling of T2W images of five subjects from the brain tumor dataset using NN interpolation (row 1), spline interpolation (row 2), SSIP (row 3), and the proposed method (row 4), along with the original T2W image (ground truth, row 5) and HR MPRAGE image (row 6).



Fig. 7.

Magnified sections of two upsampled images in Fig 6 using SSIP method (first column), and the proposed method (second column), in comparison with the ground truth (third column).



Fig. 8.

Comparison of the PSNR values of SSIP and the proposed method. PSNR curves of upsampling the (a) T2W image of subject 1 of brain tumor patients, (b) FLAIR image of subject 1 of multiple sclerosis patients, (c) FLAIR image of subject 1 of epilepsy patients, and (d) T2W image of one subject of NAMIC Brain Multimodality dataset. Note that iterations of the two methods do not have similar computation times. To better show the trend of changes in PSNR values of SSIP method, we have presented the PSNR values for 60 iterations. However, the values presented in the tables were calculated based on 30 iterations. The SSIP upsampling in (d) stopped after 27 iterations.



Fig. 9.

Upsampling of FLAIR image of subject 1 from the multiple sclerosis dataset using HR MPRAGE image. First row (left to right): upsampled FLAIR image using NN, spline, and SSIP techniques. Second row (left to right): upsampled image using the proposed method, original HR FLAIR image (ground truth), and MPRAGE image that was use to upsample the LR FLAIR image.



Fig. 10.

Upsampling of FLAIR image of subject 1 from the epilepsy dataset using different methods. Row 1: NN interpolation, SSIP using T2W image and SSIP using MPRAGE image. Row 2: spline interpolation, the proposed method using T2W image, and the proposed method using MPRAGE image. Row 3: original FLAIR image (ground truth), HR T2W image, and HR MPRAGE image. The HR T2W and HR MPRAGE images have different performances in substantia nigra and red nucleus region since the contrast is higher in that region in T2W image and thus has better performance.



Fig. 11.

Upsampling of a sample T2W image from NAMIC Brain Multimodality dataset. The LR T2W image had 5% noise. The HR T1W image had intensity non-uniformity and 5% noise. First row (from left to right): upsampling using NN, spline, and SSIP methods. Second row (from left to right): upsampling using the proposed method, original HR T2W image (ground truth), and HR T1W image used for upsampling.



Fig. 12.

Upsampling of BrainWeb T2W and PD images with 9 mm isotropic resolution to 1 mm isotropic resolution using HR T1W image. First row (from left to right): upsampled T2W images using NN, spline, and the proposed method, the original HR T2W image (ground truth), and HR T1W image. Second row (from left to right): upsampled PD images using NN, spline, and the proposed method, the original HR PD image (ground truth), and HR T1W image.



Fig. 13.

Effect of (a) the number of selected voxels N_{ν} , (b) neighborhood size, (c) standard deviation of Gaussian kernels, (d) number of Gaussian kernels, (e) threshold ε , and (f) slice thickness on upsampling the LR T2W image of subject 1 from the brain tumor dataset.



Fig. 14.

Effect of motion artifact of the HR image on upsampling. First row (left to right): T2W image upsampled using NN (PSNR=26.88, SSIM=0.735), spline (PSNR=27.77, SSIM=0.757), and SSIP (PSNR=26.25, SSIM=0.710) techniques. Second row (left to right): T2W image upsampled using the proposed method (PSNR=27.93, SSIM=0.773), original T2W image (ground truth), and the HR MPRAGE image with motion artifact.



Fig. 15.

Effect of miscoregistration due to field inhomogeneity effect in the DSC image. First row: baseline SE DSC images after upsampling using NN method. Second row: baseline SE DSC images after upsampling using SSIP, Third row: baseline SE DSC images after upsampling using the proposed method. Fourth row: baseline GE DSC images after upsampling using NN method. Fifth row: baseline GE DSC images after upsampling using SSIP. Sixth row: baseline GE DSC images after upsampling using the proposed method. Both SE and GE images were upsampled using HR T2W image.



Fig. 16.

Upsampling FLAIR image with highly anisotropic voxel size (0.43mm×0.43mm×6mm) using (a) NN interpolation, (b) spline interpolation, (c) the proposed method using HR MPRAGE image, and (d) the proposed method using HR T2W image. Both HR images had 1mm×1mm voxel size. The upsampled image has 0.43mm×0.43mm×1mm voxel size.



Fig. 17.

Upsampling ROIs using the proposed method. (a) Sagittal view of a HR MPRAGE postcontrast image showing enhancing tumor coregistered to a LR T1W postcotrast image. The ROIs drawn on the LR image are overlaid on the MPRAGE image. (b) The result of 3D upsampling of the LR T1W ROI. (c) Sagittal view of a HR T2W image showing edema coregistered to a LR FLAIR image. The ROIs drawn on the LR image are overlaid on the T2W image are overlaid of 3D upsampling of the LR T1W and FLAIR images are not shown.

2
Ē
ō
~
Ma
Manu
Manus
Manuscri

compared with four standard interpolation techniques as well as SSIP. The computation times have been reported for the proposed method and SSIP. Comparison of PSNR and SSIM values of different upsampling methods applied to T2W images of brain tumor patients. The proposed method is

			PSN	R (SSIM)			Comput	ation Time (minutes)
Subject Number	NN	Linear	Cubic	Spline	SSIP	Proposed Method	SSIP	Proposed Method
1	24.80 (0.706)	25.25 (0.706)	25.56 (0.726)	25.62 (0.728)	26.42 (0.783)	28.10 (0.838)	8.68	5.6
7	25.60 (0.706)	26.15 (0.710)	26.42 (0.727)	26.45 (0.728)	26.89 (0.759)	27.69 (0.810)	83.6	6.5
ę	24.13 (0.702)	24.61 (0.705)	24.95 (0.725)	25.00 (0.727)	25.70 (0.771)	27.63 (0.835)	88.3	5.8
4	25.84 (0.709)	26.17 (0.708)	26.43 (0.725)	26.46 (0.727)	26.78 (0.745)	28.15 (0.807)	80.2	6.4
5	24.93 (0.712)	25.35 (0.712)	25.64 (0.729)	25.69 (0.731)	26.29 (0.758)	27.89 (0.817)	76.4	6.5

TABLE II

Comparison of PSNR and SSIM values of SSIP and the proposed method within the tumors of brain tumor patients.

Subject Number	1	2	3	4	5
SSIP	30.83 (0.783)	33.95 (0.759)	26.90 (0.771)	26.45 (0.745)	23.59 (0.758)
Proposed Method	32.32 (0.901)	34.89 (0.911)	29.11 (0.877)	28.27 (0.850)	25.55 (0.817)

TABLE III

Comparison of PSNR and SSIM values of different upsampling methods applied to FLAIR images of multiple sclerosis patients. The proposed method is compared with four standard interpolation techniques as well as SSIP. The computation times have been reported for the proposed method and SSIP.

			NSA	R (SSIM)			Comput	ation Time (minutes)
Subject Number	NN	Linear	Cubic	Spline	SSIP	Proposed Method	SSIP	Proposed Method
1	26.31 (0.685)	27.16 (0.683)	27.49 (0.706)	27.50 (0.709)	28.53 (0.788)	29.50 (0.839)	113.0	43.1
2	25.39 (0.692)	25.89 (0.686)	26.16 (0.710)	26.63 (0.712)	27.64 (0.785)	28.48 (0.852)	108.0	45.9
С	24.58 (0.681)	24.98 (0.671)	25.27 (0.696)	25.50 (0.698)	26.96 (0.786)	27.55 (0.846)	110.9	45.2
4	24.45 (0.661)	24.83 (0.647)	25.15 (0.673)	25.27 (0.675)	26.78 (0.765)	27.08 (0.849)	133.9	26.3
5	25.35 (0.694)	26.06 (0.692)	26.47 (0.718)	26.57 (0.722)	27.76 (0.778)	28.01 (0.841)	128.9	44.2

Author Manuscript

TABLE IV

Comparison of PSNR and SSIM values of different upsampling methods applied to FLAIR images of epilepsy patients. The proposed method is compared with four standard interpolation techniques as well as SSIP. The computation times have been reported for the proposed method and SSIP. Both HR MPRAGE and T2W images have been used to upsample the LR FLAIR images.

				ALLC I						monundi		(cm
Subject Number	NN	Linear	Cubic	Spline	SSIP using MPRAGE	SSIP using T2W	Proposed Method using MPRAGE	Proposed Method using T2W	SSIP using MPRAGE	SSIP using T2W	Proposed Method using MPRAGE	Proposed Method using T2W
-	23.60 (0.683)	24.01 (0.680)	24.24 (0.701)	25.12 (0.704)	25.63 (0.783)	25.88 (0.794)	27.27 (0.817)	26.62 (0.831)	30.1	26.7	5.8	5.6
2	24.23 (0.691)	24.82 (0.693)	25.14 (0.720)	25.37 (0.719)	26.25 (0.792)	26.98 (0.820)	27.40 (0.829)	27.97 (0.857)	36.2	33.9	5.6	4.8
3	24.59 (0.701)	25.18 (0.705)	25.50 (0.727)	26.04 (0.732)	26.98 (0.805)	27.35 (0.816)	28.39 (0.840)	27.95 (0.853)	33.2	29.9	4.5	7.1
4	22.72 (0.691)	23.40 (0.693)	23.78 (0.717)	24.07 (0.721)	25.44 (0.806)	26.44 (0.831)	26.61 (0.839)	27.41 (0.864)	36.8	34.8	5.1	4.9
5	23.16 (0.628)	23.80 (0.630)	24.21 (0.654)	24.28 (0.658)	25.93 (0.760)	26.18 (0.796)	26.67 (0.795)	26.40 (0.835)	36.5	36.7	6.0	5.4

Comparison of PSNR and SSIM values of different upsampling methods applied to NAMIC Brain Multimodality images. The algorithm has been tested for T2W images with 9 mm slice thickness and 1 mm interpolation techniques as well as SSIP. The computation times have been reported for the proposed method and SSIP. Values of mean and standard deviations are reported as Mean±StD over all 20 subjects. in-plane resolution and with different intensity non-uniformity and noise levels. In each experiment, the HR and LR images had similar levels of noise. The proposed method is compared with four standard

Intensity Non-	love T code			PSNR ((WISS)			Computa (min	tion Time utes)
Uniformity	Noise Level	NN	Linear	Cubic	Spline	SSIP	Proposed Method	SSIP	Proposed Method
Ĩ	0 %	$24.54\pm2.07(0.662\pm0.043)$	$24.95\pm 2.15(0.671\pm 0.046)$	$25.26\pm 2.18(0.685\pm 0.045)$	$25.29\pm 2.19(0.680\pm 0.046)$	$28.24\pm2.06(0.834\pm0.023)$	$28.76\pm 1.98(0.841\pm 0.025)$	85.6±2.2	8.0 ± 0.8
ON	5 %	$24.46\pm 2.07(0.638\pm 0.046)$	$24.91\pm2.15(0.648\pm0.049)$	$25.17\pm2.18(0.661\pm0.048)$	$25.19\pm 2.19(0.657\pm 0.049)$	27.43±2.07(0.790±0.029)	$28.43\pm 2.05(0.809\pm 0.029)$	85.0±4.8	12.5±0.9
Vec	% 0	I	I	I	I	27.34±2.08(0.798±0.028)	$28.29\pm1.93(0.822\pm0.027)$	85.2±2.2	7.9±0.9
162	5 %	I	I	I	Ι	$26.71\pm2.11(0.754\pm0.035)$	$28.06{\pm}2.07(0.796{\pm}0.031)$	62.2 ± 4.3	12.5 ± 0.8

TABLE VI

BrainWeb database with epilepsy model. Upsampling was performed along all three directions. The computation times have been reported for the Comparison of PSNR and SSIM values of different standard upsampling methods and the proposed method applied to T2W and PDW images of proposed method.

5				PSNR (SSIV	(J		Computation Time (minutes
sequence	Kesoluuon	NN	Linear	Cubic	Spline	Proposed Method	Proposed Method
	3 mm isotropic	17.74 (0.717)	19.03 (0.727)	19.78 (0.775)	20.08 (0.789)	24.25 (0.940)	15.9
T2W	6 mm isotropic	14.87 (0.442)	15.93 (0.442)	16.39 (0.476)	16.58 (0.478)	19.59 (0.807)	7.0
	9 mm isotropic	13.80 (0.332)	14.82 (0.334)	15.17 (0.343)	15.36 (0.337)	17.14 (0.659)	7.5
	3 mm isotropic	18.08 (0.777)	20.78 (0.805)	21.82 (0.846)	22.33 (0.856)	22.39 (0.949)	5.7
PDW	6 mm isotropic	14.43 (0.531)	16.89 (0.562)	17.68 (0.592)	18.07 (0.595)	20.01 (0.853)	6.2
	9 mm isotropic	13.04 (0.426)	15.13 (0.450)	15.75 (0.458)	16.11 (0.454)	17.62 (0.744)	9.1