General and Efficient Super-Resolution Method for Multi-slice MRI

D.H.J. Poot[1,2], V. Van Meir[2], and J. Sijbers[2]

[1] BIGR, Erasmus Medical Center, Rotterdam
[2] Visionlab, University of Antwerp, Antwerp

Abstract. In this paper, a method is developed that reconstructs a high resolution image from an arbitrary set of multi-slice 3D MR images with a high in-plane resolution and a low through-plane resolution. Such images are often recorded to increase the efficiency of the acquisition. With a model of the acquisition of MR images, which is improved compared to previous super-resolution methods for MR images, a large system with linear equations is obtained. With the conjugated gradient method and this linear system, a high resolution image is reconstructed from MR images of an object. Also, a new and efficient method to apply an affine transformation to multi-dimensional images is presented. This method is used to efficiently reconstruction the high resolution image from multi-slice MR images with arbitrary orientations of the slices.

Keywords: Super-resolution, multi-slice imaging, reconstruction, conjugated gradients, affine transformation, multi-dimensional imaging.

1 Introduction

This paper describes how a high resolution isotropic 3D image can be reconstructed from a series of multi-slice 3D MR images, in which the slice thickness is substantially larger than the in plane resolution. Since this method improves the resolution in all directions up to the in-plane resolution, it is a super-resolution method. The motivation of this work is that multi slice images are often acquired with anisotropic voxels. The slice thickness can be substantially higher than the in-plane resolution in order to increase the Signal to Noise Ratio (SNR). As will be shown, several of these multi-slice MR images, recorded with different slice orientations, can be combined into a single, high resolution 3D image with isotropic voxels. The advantage of multi slice images, compared to full 3D acquisitions is that it is possible to interleave the acquisition of slices. That is, while waiting for the (T_1) relaxation of the magnetization of a slice, (a part of) the k-space of various other slices can be excited and recorded. In general, when the repetition time (TR) is limited by the T_1 decay, it is possible to record multi slice images significantly faster than full 3D images with the same resolution [1].

Previously, several attempts have been made to improve the resolution of MR images. The methods of Peled et. al. [2] and Carmi et al. [3] try to improve the in-plane resolution. The validity of such methods was questioned by Scheffler [4],

since each in-plane shifted image acquires the same points in k-space, only introducing a linear phase shift in the k-space samples. This means that each of the shifted images does contain the same information, except for measurement noise, with no possibility to improve the resolution. A different method, presented by Greenspan et al. [5], improves the resolution in the slice direction, i.e. the direction in which the different slices of a multi-slice MR image are recorded. Several multi-slice MR images with different positions in the slice direction are combined. This method is not limited to the original resolution as the acquisition in the slice direction is not band limited. However, only MR images with identical orientation can be combined. This limitation is partially removed by the method of Shilling et al. [6], in which multi-slice MR images rotated around a common frequency encoding axis are combined. It allows the reconstruction of high resolution slices with iterative projection reconstruction algorithms. In their method they state the projection as a linear system and solve the high resolution image from the set of linear equations with iterative solvers that are also used in Computed Tomography (CT) reconstructions.

In our work, the method of Shilling et al. [6] is extended to allow for any orientation of the slices of the multi-slice MR images, i.e. the images do not need to be rotated around a common frequency encoding axis. Furthermore, the projection via matrix multiplication is reformulated as affine transformation, followed by a filter operation, which reduces the number of computations substantially. Finally, the method presented in this paper uses the Conjugated Gradient method to solve the large linear system in a low number of iterations.

2   Methods

2.1   Introduction

The acquisition of multiple MR images with the same contrast can be seen as acquiring multiple samples from the same object. For each slice of a multi-slice MR image, the MR acquisition records a part of the k-space of the object. With the discrete Fourier transform, a projection of the object intensities in each slice is reconstructed on a discrete grid. Since MR acquisitions record a limited part of the k-space, the intensity at a grid point does not depend exclusively on the intensity of the object at the location of the grid point, but a low-pass filter is applied to the excited slice of the object before sampling the intensities at the grid nodes. An alternative interpretation is that the intensity value of each voxel of the MR image is obtained by multiplying a properly shifted version of a (3D) sampling function with the object.

2.2   Model of the MRI Acquisition

Let $\boldsymbol{o}$ $(n_o \times 1)$ be a vector containing the intensities of the continuous object o at the 3D grid points $\boldsymbol{x}_m$, $m \in \{1, \ldots, n_o\}$, where $\boldsymbol{x}_m$ is a coordinate in object space. Let $\boldsymbol{S}_j$ $(n_{S_j} \times 1)$ be the samples of the $j^{th} \in \{1, \ldots, N\}$ MR image S_j at the grid nodes $\boldsymbol{y}_l$, $l \in \{1, \ldots, n_{S_j}\}$, where $\boldsymbol{y}_l$ is a coordinate in the space of

the j^{th} MR image, and let the measurement noise be described by $\boldsymbol{e}_j$ ($n_{S_j} \times 1$). When the coordinates $\boldsymbol{x}$ and $\boldsymbol{y}$ are linked by a coordinate transform T_j, the acquisition of the MR images can be modeled with:

$$S = Xo + e, \tag{1}$$

where

$$\boldsymbol{S} = \begin{bmatrix} \boldsymbol{S}_1^T \dots \boldsymbol{S}_N^T \end{bmatrix}^T, \quad \boldsymbol{X} = \begin{bmatrix} \boldsymbol{X}_1^T \dots \boldsymbol{X}_N^T \end{bmatrix}^T, \quad \boldsymbol{e} = \begin{bmatrix} \boldsymbol{e}_1^T \dots \boldsymbol{e}_N^T \end{bmatrix}^T, \tag{2}$$

and

$$X_j(l, m) = w(T_j(\boldsymbol{x}_m) - \boldsymbol{y}_l), \tag{3}$$

where w is the sampling function. This sampling function is assumed to be separable in MR- image coordinates, $w(\boldsymbol{y}) = \prod_{i=1}^{3} w_i(y_i)$, which can, without loss of generality, be numbered 1,2, and 3 for the read, phase, and slice encoding directions, respectively. The sampling function w is (implicitly) defined by the MR image acquisition method. The in-plane sampling functions w_1 and w_2 are defined by the rectangular part of k-space that is sampled and thus are Dirichlet, or periodic sinc functions. For multi-slice MR images, w_3 depends on the slice selection excitation, which usually is either a (windowed) sinc or a Gaussian shaped RF pulse. In our experiments, with windowed sinc slice excitation, w_3 was modeled with a smoothed box function :

$$w_3(y) = \begin{cases} 1 & |y| \leq \frac{1}{3} \\ \frac{1}{2} - \frac{1}{2} \sin\left(3\pi(|y| - \frac{1}{2})\right) & \frac{1}{3} < |y| < \frac{2}{3} \\ 0 & \frac{2}{3} \leq |y| \end{cases} \tag{4}$$

Due to the aliasing caused by the slice selection function, it is easier to describe the acquisition in image space (Eq. (1)), rather than in the k-space.

Note that, in this model, the transformation $(T_j(\boldsymbol{x}))_2$ is not necessarily just a rotation or even only an affine transformation, but it might also contain other deformations. For example, when the images are recorded with Echo Planar Imaging (EPI), it might also describe displacements in the phase encoding direction due to inhomogeneities of the main magnetic field.

2.3   Reconstruction of the Object

When the MR images $\boldsymbol{S}$ are acquired, the reconstruction of the object intensities $\boldsymbol{o}$ can be stated as a regularized least squares problem:

$$\hat{\boldsymbol{o}} = \arg\min_{\boldsymbol{o}}(\boldsymbol{X}\boldsymbol{o} - \boldsymbol{S})^T(\boldsymbol{X}\boldsymbol{o} - \boldsymbol{S}) + \lambda \boldsymbol{o}^T \boldsymbol{K} \boldsymbol{o}, \tag{5}$$

where $\boldsymbol{K}$ specifies the regularization term, which will be explained below, and λ is a scalar constant that scales the regularization. In general, the solution of this regularized least squares problem is given by:

$$\hat{\boldsymbol{o}} = (\boldsymbol{X}^T \boldsymbol{X} + \lambda \boldsymbol{K})^{-1} \boldsymbol{X}^T \boldsymbol{S} \tag{6}$$

2.4   Regularization

Regularization is a standard technique to solve under-determined or badly conditioned problems [7]. The current problem is badly conditioned due to the high resolution of the grid on which the object intensities are reconstructed. This grid should contain all spatial frequencies present in the MR images. Thus, it will most likely also contain spatial frequencies that are (almost) not sampled by any of the MR images. As there is no information about these (high) spatial frequencies, it is best to force the amplitude of them towards zero by adding the power in the (high) frequencies to the minimization criterium. In this work, this is achieved with the square of the discrete second derivative of the reconstructed o:

$$\boldsymbol{o}^T \boldsymbol{K} \boldsymbol{o} = \left(\frac{\partial^2 \boldsymbol{o}}{\partial x_1{}^2}\right)^2 + \left(\frac{\partial^2 \boldsymbol{o}}{\partial x_2{}^2}\right)^2 + \left(\frac{\partial^2 \boldsymbol{o}}{\partial x_3{}^2}\right)^2 . \tag{7}$$

The strength of the regularization is controlled by the variable λ in Eq. (6). Increasing λ increases the bias in the solution, while reducing the variance. A good value for λ is a value that (approximately) minimizes the Root Mean Square Error (RMSE).

For realistic image dimensions, the matrices in the general solution given in Eq. (6) are too large to actually store, even as sparse matrices. Furthermore, the solution of the linear system with QR or LU decomposition would consume way to much computation time for any realistic size of images. Therefore, the conjugated gradients method [8,9], which is an efficient iterative method to solve linear systems, was used to obtain the solution.

2.5   Affine Transformation

It is possible to evaluate the matrix vector multiplications that are required for the conjugated gradient method explicitly. However, this will require a large amount of computation time. When the transform T_j is an affine transform, or a subset of an affine transform, such as a combination of translation, scaling, and rotation, the acquisition of the MR images can be reformulated as an affine transform of the object o, followed by subsequent filter operations with the three orthogonal sampling functions w_i.

An affine transform of the initial coordinate $\boldsymbol{x}$ to the final coordinate $\boldsymbol{y}$ is given by

$$\boldsymbol{y} = \boldsymbol{T} \boldsymbol{x} + C, \tag{8}$$

where $\boldsymbol{T}$ $(n \times n)$ specifies the affine transformation matrix in n dimensions, and $\boldsymbol{C}$ $(n \times 1)$ specifies the translation part of the transformation. Then, the affinely transformed image $\tilde{o}$ of a continuous image o is given by

$$\tilde{o}(\boldsymbol{x}) = o(\boldsymbol{T}\boldsymbol{x} + \boldsymbol{C}). \tag{9}$$

This transform of a continuous image specifies the ideal transform that we try to approximate with the discretely sampled images. Different affine transformation methods exist. The most common methods interpolate the source image,

with linear interpolation or (approximate) sinc interpolation. However, for multi-dimensional images, a better transformation method is given below.

In order to properly evaluate the different transformation methods, both the spatial and the frequency domain should be studied. This is easy since an affine transform of an image also causes an affine transformation and a frequency dependent phase shift in the frequency domain:

$$\tilde{\mathcal{O}}(\boldsymbol{f}) = \mathcal{O}(\boldsymbol{f}\boldsymbol{T}^{-1})e^{i\boldsymbol{f}\boldsymbol{C}}, \tag{10}$$

where $\mathcal{O} = \mathcal{F}(o)$, $\tilde{\mathcal{O}} = \mathcal{F}(\tilde{o})$ are the Fourier transformed o and $\tilde{o}$, and $\boldsymbol{f}$ $(1 \times n)$ is a position in the frequency domain. Obviously, the spatial transformations induced by a transformation method should be as specified by the affine transform. However, this is the case for all methods that are discussed and displayed in Fig. 1, so it cannot be used for discrimination. Also, the magnitudes of the spatial frequencies should not be modified by the transformation method. Linear interpolation fails in this respect. An stronger discriminating feature is aliasing. The energy in spatial frequencies of the source image that are aliased to different (thus: wrong) spatial frequencies should be minimal. The aliasing in multi dimensional $(n > 1)$ images might arise from three subtly different sources. First, spatial frequencies which are below the Nyquist frequency in the continuous source image o, might be transformed to above the Nyquist frequency in the destination image $\tilde{o}$. Secondly, spatial frequencies which are above the Nyquist frequency in the source image o, might be transformed to below the Nyquist frequency in the destination image $\tilde{o}$. Thirdly, spatial frequencies that are above the Nyquist frequency in both the continuous source and destination images might alias to below the Nyquist frequency of the destination image due to the sampling of this image.

The new transformation method proposed in this paper applies the affine transform T_f efficiently by splitting it into a **S**et of **SH**ear transformations (SSH). The shear transformations $\tilde{\boldsymbol{T}}_j$, $j \in \{1, \ldots, 2n\}$ satisfy $\boldsymbol{T}_f = \prod_{j=1}^{2n} \tilde{\boldsymbol{T}}_j$. Each $\tilde{\boldsymbol{T}}_j$ differs from the identity matrix only in row d_j, so the image is only deformed along the d_j^{th} main axis, allowing efficient interpolation with a 1D low pass filter. An optimal set of these $\tilde{\boldsymbol{T}}_j$ can be constructed with:

- Initialize all $\tilde{\boldsymbol{T}}_j$ to the $(n+1) \times (n+1)$ identity matrix.
- By computing the total cost of the transformations defined in the following steps, search for the optimal permutation p of the numbers $\{1, \ldots, n\}$.
- Initialize the remaining transform to the inverse transformation: $\boldsymbol{R}_0 = \boldsymbol{T}_f^{-1}$
- for $j = 1 \ldots n$:

$$(\tilde{\boldsymbol{T}}_j)_{p_j, p_{1\ldots j}} = \left(\left((\boldsymbol{R}_{j-1})_{p_{1\ldots j}, p_{1\ldots j}} \right)^{-1} \right)_{j, 1 \ldots j} u_j \tag{11}$$

$$\boldsymbol{R}_j = \tilde{\boldsymbol{T}}_j \boldsymbol{R}_{j-1}, \tag{12}$$

- for $j = n+1 \ldots 2n$

$$(\tilde{\boldsymbol{T}}_j)_{p_{j-n}, 1\ldots n+1} = \left(\boldsymbol{R}_{j-1}^{-1} \right)_{p_{j-n}, 1\ldots n+1} \tag{13}$$

$$\boldsymbol{R}_j = \tilde{\boldsymbol{T}}_j \boldsymbol{R}_{j-1}. \tag{14}$$
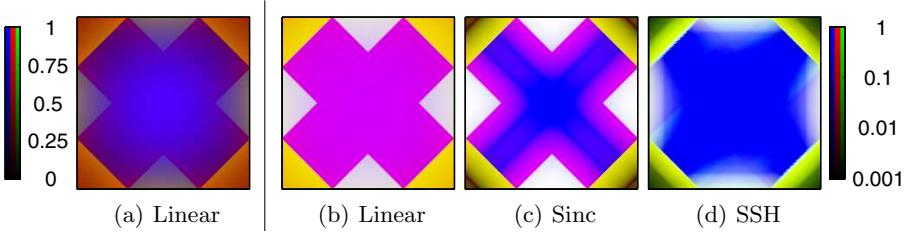
Fig. 1. Elementary frequency cell (zero spatial frequency at center, Nyquist frequency at edges) after rotating over 45 degrees with several methods. Colors: non constant Blue indicates spectral distortions. white: First type of aliasing, Yellow: Second type of aliasing, Purple: Third type of aliasing. Ideal transformation is Blue octagon and black triangles in the corners. Subfigure (a) and (b) display the result of the linear interpolator, scaled differently. Of the methods considered, linear interpolation is the only that has significant spectral distortions. (c) Sinc interpolation with a 11×11 neighborhood. (d) Rotation when the transform is split in a set of shear transforms. The length of the filters of (d), which is 15, was chosen such that the number of computations was (approximately) equal to (c).

 – Find the lowest up-sampling factors u_j such that no distortion or aliasing is present in the final image and derive the cutoff frequency of the low-pass filter from the maximum frequency that needs to be transferred to the output image.

See Fig. 1 for a comparison of this SSH method to linear interpolation and a windowed sinc interpolation, when rotating a 2D image over 45 degrees. Note that, for 2D images, with the same computational complexity, the quality of the SSH method is substantially better than windowed sinc interpolation. When transforming 3D instead of 2D images, this method requires only approximately 50% more computations per voxel, which compares favorably to sinc interpolation where the number of computations is multiplied by the number of samples in the windowed sinc function.

3   Experiments

Several datasets of a bird were recorded with a Bruker small animal scanner. The resulting MR images were $192 \times 192 \times 32$ with voxel dimensions $0.125mm \times 0.125mm \times 0.75mm$. MR images in $N = 36$ different orientations were recorded and the reconstructed volume spanned the whole head in a volume of $21mm \times 21mm \times 22mm$ with isotropic voxel dimensions $0.1mm \times 0.1mm \times 0.1mm$.

As preprocessing step, the MR images were aligned by computing the required translations from the projections of the MR images to the object space,

$$\boldsymbol{b}_j = \boldsymbol{X}_j^T \boldsymbol{S}_j. \tag{15}$$

Since the images are assumed to have the same contrast, the optimal translation between two images was computed with a mean square difference measure,

$$\hat{\boldsymbol{\Delta}}_{j,k} = \underset{\boldsymbol{\Delta}_{j,k} \in \mathcal{R}^3}{\arg\min} \sum_{\boldsymbol{x}} \left(b_j(\boldsymbol{x}) - b_k(\boldsymbol{x} + \boldsymbol{\Delta}_{j,k}) \right)^2 . \tag{16}$$

The position of each image j was adjusted by $1/N \sum_{k=1}^{N} \hat{\boldsymbol{\Delta}}_{j,k}$.

After the alignment of the images, the conjugated gradient method with affine transformation with the SSH method was used to approximately solve Eq. (6). In our implementation each iteration of the conjugated gradient method took approximately 5m:35s of CPU time on one core of a Intel Core 2 Quad CPU @ 3.0 GHz, with 8GB of RAM.

4   Results

One of the original MR images and the high resolution reconstructed image are shown in Fig. 2. The 3D MR image is displayed in the read-phase, read-slice, and phase-slice directions, which, for this image, coincide with the coronal, sagittal and transversal directions in which the high resolution reconstructed image is shown. Note the substantially lower resolution of the MR image in the
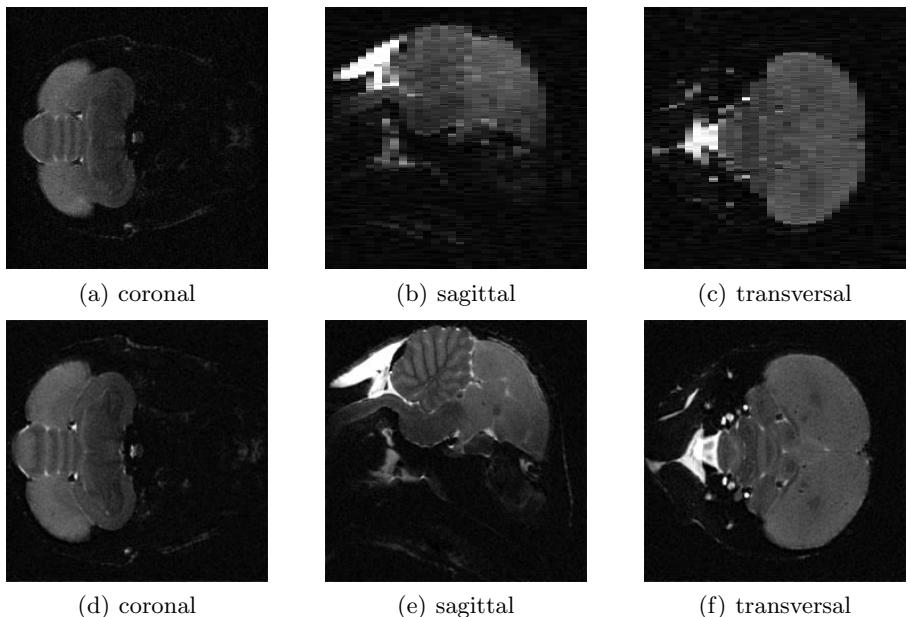

(a) coronal   (b) sagittal   (c) transversal

(d) coronal   (e) sagittal   (f) transversal

Fig. 2. (a), (b), and (c) show one of the 36 original MR images. (d), (e), and (f) show the high resolution reconstructed image.

slice direction, especially when compared to the high resolution reconstructed image. Also, observe the factor 2 reduction of the noise standard deviation. Furthermore, note the improved detail in the coronal views, which is due to the reduced blurring in the slice direction.

The conjugated gradient method converged quickly. After 15 iterations the update to o was approximately 0.001 of the magnitude of o, which is below the noise magnitude, indicating that 15 iterations are sufficient. Note that this number of iterations needed for convergence is substantially lower than the approximately 1000 iterations required by the methods in [6]. This is due to the high rate of convergence of the Conjugated Gradient method.

5   Conclusion

In this paper, a method was developed by which a high resolution isotropic image can be reconstructed from a set of anisotropic multi-slice MR images, recorded with different slice orientations. In contrast to previous reconstruction methods, this new method does not constrain the slice orientations. The reconstruction method uses an improved model of the MR acquisition, but does not require any prior knowledge about the imaged object. The high resolution image of the object is accurately reconstructed by the conjugated gradient method in a small number of iterations, substantially less than previous methods. The experiments show that the quality of the reconstructed isotropic image is substantially better, both in resolution and SNR, than any of the original MR images.

References

1. Zimmerman, R.A., Gibby, W.A., Carmody, R.F.: Neuroimaging: clinical and physical principles. Springer, New York (2000)
2. Peled, S., Yeshurun, Y.: Superresolution in MRI; application to human white fibre vizualization by diffusion tensor imaging. Magn. Reson Med. 45, 29–35 (2001)
3. Carmi, E., Liu, S., Alon, N., Fiat, A., Fiat, D.: Resolution enhancement in MRI. Magn. Reson Imaging 24, 133–154 (2006)
4. Scheffler, K.: Superresolution in MRI? Magn. Reson Med. 48, 408 (2002)
5. Greenspan, H., Oz, G., Kiryati, N., Peled, S.: MRI inter-slice reconstruction using super-resolution. Magn. Reson Imaging 20, 437–446 (2002)
6. Shilling, R.Z., Robbie, T.Q., Bailloeul, T., Mewes, K., Mersereau, R.M., Brummer, M.E.: A super-resolution framework for 3-D high-resolution and high-contrast imaging using 2-D multislice MRI. IEEE T. Med. Imaging 28(5), 633–644 (2009)
7. Engl, H.W., Hanke, M., Neubauer, A.: Regularization of inverse problems. Kluwer Academic Publishers, Dordrecht (2000)
8. Hestenes, M.R., Stiefel, E.: Methods of conjugate gradients for solving linear systems. J. Res. Nat. Bur. Stand 49(6), 409–436 (1952)
9. Shewchuk, J.R.: An introduction to the conjugate gradient method without the agonizing pain. Technical report, Carnegie Mellon University, Pittsburgh, PA, USA (1994)