

Minimum-Norm Reconstruction for Sensitivity-Encoded Magnetic Resonance Spectroscopic Imaging

Javier Sánchez-González,^{1*} Jeffrey Tsao,² Ulrike Dydak,² Manuel Desco,¹ Peter Boesiger,² and Klaas Paul Pruessmann²

In this work we propose minimum-norm reconstruction as a means to enhance the spatial response behavior in parallel spectroscopic MRI. By directly optimizing the shape of the spatial response function (SRF), the new method accounts for coil sensitivity variation across individual voxels and their side lobes. In this fashion, it mitigates the signal contamination and side-lobe aliasing, to which previous techniques are susceptible at low resolution. Although the computational burden is higher, minimum-norm reconstruction is shown to be feasible using an iterative algorithm. Benefits in terms of SRF shape and artifact suppression are demonstrated. Magn Reson Med 55: 287–295, 2006. © 2006 Wiley-Liss, Inc.

Key words: sensitivity encoding (SENSE); MR spectroscopic imaging (MRSI); minimum-norm reconstruction; strong reconstruction; iterative algorithm

In recent years, there has been significant progress in the development of parallel imaging techniques, such as SMASH (1), SENSE (2), SPACE-RIP (3), GRAPPA (4), and PILS (5). These techniques have permitted significant scan time reduction in both MRI and spectroscopic imaging (MRSI) (6,7). Although these techniques differ in implementation and underlying approximations, they are all based on the principle that the spatially varying sensitivities of the receiver coils complement the role of the magnetic field gradients in spatial encoding. As a result, it is feasible to reduce the sampling density in k -space without compromising the spatial resolution or the field of view (FOV).

A key departure of parallel imaging from conventional imaging is that the effects of the coil sensitivities are taken into account during reconstruction. This departure affects the properties of the reconstructed image, since the net encoding functions are no longer orthogonal. For example, the achievable signal-to-noise ratio (SNR) is spatially varying and, as the present work shows, the achievable resolution is also spatially varying.

Parallel imaging reconstruction can be performed very efficiently when k -space is sampled along a Cartesian grid. In this case, the effects of k -space undersampling are par-

ticularly easy to account for. In the image domain, they result in aliasing that occurs among small sets of equidistant voxels. Image reconstruction can then be achieved by individual unfolding of these aliased sets. This is the approach underlying the common SENSE method in the case of Cartesian k -space sampling (2). In the following it will be simply be referred as SENSE for easier reading.

Straightforward image-domain unfolding requires relatively little computation. It is important to note, however, that this advantage is the result of a mild approximation. SENSE strictly enforces the elimination of aliasing only in the voxel centers and was therefore dubbed *weak* reconstruction in Ref. (2). The potential downside of weak reconstruction is that appreciable residual aliasing may occur when coil sensitivities vary considerably over the extent of a voxel and its significant side lobes. This is typically not of concern for high-resolution imaging since coil sensitivities vary smoothly at the scale of common voxel sizes. However, it gradually becomes a problem when the scan resolution is reduced and becomes a serious issue at the very low resolutions that are typically used in MRSI.

In the present work, we propose an alternative reconstruction approach that overcomes the described restrictions. The basic idea is to optimize the spatial response function of reconstructed voxels as a whole rather than only at the voxel centers. This is achieved by formulating the encoding equation at an enhanced spatial resolution and taking its minimum-norm solution as the reconstructed image.

THEORY

We consider the most common type of MRSI where spatial encoding is performed by phase-encoding gradients, followed by readout along the time axis in the absence of gradient fields. With this approach, spectral and spatial reconstruction are independent, with spectral reconstruction amounting to Fourier transform of the whole data set along the time axis. To keep the following formal treatment as simple as possible we consider the situation after such an initial temporal Fourier transform. Out of the resulting k_x - k_y - ω data set we consider only a single k -space plane for some fixed value of ω , leaving us with the equivalent of a common MRI data set.

Using n_c independent receiver coils for data acquisition one obtains n_c such data sets, which shall be denoted as $m_{\gamma,\kappa}^{(\omega)}$, where γ counts the coils and κ is the sampled positions in k -space. The contributions to each data value are modeled as

$$m_{\gamma,\kappa}^{(\omega)} = \int \rho(\mathbf{r}, \omega) e_{\gamma,\kappa}(\mathbf{r}) d\mathbf{r} + \boldsymbol{\eta}_{\gamma,\kappa}^{(\omega)} \quad [1]$$

¹Medicina y Cirugía Experimental, Hospital General Universitario "Gregorio Marañón," Madrid, Spain.

²Institute for Biomedical Engineering, Swiss Federal Institute of Technology Zurich and University of Zurich, Zurich, Switzerland.

Grant sponsor: Strategic Excellence Projects Program, ETH Zurich; Grant number: SEP Life Science TH7/02-2; Grant sponsor: Spanish Ministerio de Sanidad; Grant number: Red Temática IM3 G03/185.

*Correspondence to: Javier Sánchez González, Laboratorio de Imagen, Medicina Experimental, Hospital General Universitario "Gregorio Marañón" Dr. Esquerdo 46, E-28007 Madrid, Spain. E-mail: jsanchez@mce.hggm.es

Received 14 April 2005; revised 19 July 2005; accepted 23 September 2005.

DOI 10.1002/mrm.20758

Published online 11 January 2006 in Wiley InterScience (www.interscience.wiley.com).

where \mathbf{r} denotes 3D position and $\rho(\mathbf{r},\omega)$ represents the spatio-spectral signal density of the scanned object, which depends both on the object and on the sequence parameters. $e_{\gamma,\kappa}(\mathbf{r})$ denotes the net spatial encoding function composed of harmonic modulation by the phase-encoding gradients and the complex-valued spatial sensitivity $s_\gamma(\mathbf{r})$ of the γ th coil:

$$e_{\gamma,\kappa}(\mathbf{r}) = e^{ik_x \mathbf{r}} \cdot S_\gamma(\mathbf{r}). \quad [2]$$

Note that the encoding functions $e_{\gamma,\kappa}(\mathbf{r})$ are not restricted to any particular k -space sampling patterns. Therefore, the model covers the case of arbitrary sampling patterns in k -space. Note also that the encoding functions do not vary over time, so they are also the same for all ω . $\eta_{\gamma,\kappa}^{(\omega)}$ in Eq. [1] represents stochastic noise in the intermediate data. It has zero-mean Gaussian statistics and is uncorrelated between different sampling times. Due to these characteristics the noise statistics do not change under the initial Fourier transform. The second-order statistics of $\eta_{\gamma,\kappa}^{(\omega)}$ are hence given by

$$\tilde{\Psi}_{(\gamma,\kappa),(\gamma',\kappa')} = \langle \eta_{\gamma,\kappa}^{(\omega)} \overline{\eta_{\gamma',\kappa'}^{(\omega)}} \rangle = \Psi_{\gamma,\gamma'} \delta_{\kappa,\kappa'}, \quad [3]$$

where $\langle \cdot \rangle$ and the bar indicate averaging and complex conjugation, respectively. Ψ denotes the common noise covariance matrix of the coil array (2) and the Kronecker delta $\delta_{\kappa,\kappa'}$ reflects the fact that noise is uncorrelated between different k -space positions. As Eq. [3] reflects, noise correlation does occur among data that differ only in their coil indices if the coil array as such exhibits noise correlation. To simplify this situation, we further assume that the receiver channels are initially precombined such as to decorrelate and normalize their noise contents as described in Ref. (8). In the statistics literature this step is known as prewhitening. It is a simple and computationally cheap linear operation, creating a set of virtual data channels, which can then be treated exactly like physical ones. When using prewhitening it is important to perform the corresponding precombination not only to the raw data but also to the coil sensitivities that occur in Eq. [2]. The benefit of this operation is that the noise covariance matrix in Eq. [3] then becomes equal to identity and can hence be omitted. For the following it is generally assumed that prewhitening has been performed.

By discretizing the spatial coordinates, the signal part of Eq. [1] can be expressed more conveniently in matrix form as

$$\mathbf{m} = \mathbf{E}\boldsymbol{\rho}. \quad [4]$$

Here, all values of $m_{\gamma,\kappa}^{(\omega)}$ for the chosen ω have been assembled in the column vector \mathbf{m} of length $n_c n_\kappa$, n_κ denoting the number of sampled k -space positions. $\boldsymbol{\rho}$ is a column vector of length n_v , listing unknown image values for the finite set of pixel positions according to the discretization. \mathbf{E} is the $(n_c n_\kappa) \times n_v$ encoding matrix, listing the values of the encoding functions at the pixel positions. Note that the matrix representation is not limiting because the discretization can be made arbitrarily fine in order to achieve any desired level of accuracy.

Image reconstruction amounts to solving Eq. [4] for $\boldsymbol{\rho}$. For a sufficiently fine discretization, Eq. [4] is underdetermined. In this case it has an infinite number of solutions. In the absence of any additional information (9,10), a particularly attractive choice is the minimum-norm solution (11). As the name suggests, the minimum-norm solution has the smallest norm of all solutions satisfying Eq. [4]. It can be argued to be optimal since it recovers all signal that was actually encoded, while setting the remainder to zero and thus eliminating all inconsistent noise.

The key difference between the proposed reconstruction and SENSE is that for minimum-norm reconstruction the encoding matrix \mathbf{E} is discretized at a finer resolution than the nominal image resolution, which is determined by the extent of the sampled area in k -space. This finer discretization is key because it is the basis of controlling the spatial response between voxel centers.

The minimum-norm solution $\hat{\boldsymbol{\rho}}$ of Eq. [4] is

$$\hat{\boldsymbol{\rho}} = \mathbf{E}^\dagger \mathbf{m}, \quad [5]$$

where the dagger indicates the Moore–Penrose pseudoinverse. The task of inverting the generally rectangular matrix \mathbf{E} can be translated into a more amenable symmetric inversion by using two equivalent expansions of the pseudoinverse (12),

$$\hat{\boldsymbol{\rho}} = (\mathbf{E}^H \mathbf{E})^\dagger \mathbf{E}^H \mathbf{m} \quad [6]$$

$$\hat{\boldsymbol{\rho}} = \mathbf{E}^H (\mathbf{E} \mathbf{E}^H)^\dagger \mathbf{m}, \quad [7]$$

where the superscript H denotes the complex conjugate transpose. Equations [6] and [7] are equivalent mathematically. However, Eq. [7] is more manageable computationally, since the matrix inversion only involves a size of $(n_c n_\kappa) \times (n_c n_\kappa)$, which does not increase when the spatial discretization is refined. An implementation of Eq. [7] can be viewed as consisting of two steps:

$$\mathbf{c} = (\mathbf{E} \mathbf{E}^H)^\dagger \mathbf{m} \quad [8]$$

$$\hat{\boldsymbol{\rho}} = \mathbf{E}^H \mathbf{c}. \quad [9]$$

The first step (Eq. [8]) involves a base change from the data domain to a space spanned by the left singular vectors of \mathbf{E} , which are equal to the rows of the matrix \mathbf{U} in the singular value decomposition

$$\mathbf{E} = \mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^H, \quad [10]$$

yielding a set of coefficients \mathbf{c} . In the second step (Eq. [9]), the reconstructed image $\hat{\boldsymbol{\rho}}$ is generated as a linear combination of the complex conjugates of the encoding functions. Equation [9] illustrates that the achievable resolution depends directly upon the encoding functions. In particular, it shows that coil sensitivity gradients can enhance the resolution by adding to the spatial signal variation induced by the common B_z gradients. This effect is relevant whenever the coil sensitivities vary significantly across the nominal voxel size. It is hence most appreciable in low-resolution imaging and close to coil conductors

where the coil sensitivity varies the most. As a consequence, the spatial resolution of sensitivity-encoded MRI and MRSI can be somewhat nonuniform across the image.

In general, Eq. [8] can be solved in two ways. The first is to solve $(\mathbf{E} \mathbf{E}^H) \mathbf{c} = \mathbf{m}$ directly for the specific \mathbf{m} of interest. The second way is to determine the pseudoinverse $(\mathbf{E} \mathbf{E}^H)^\dagger$ explicitly and then multiply it by \mathbf{m} . The first approach is generally faster. However, the second approach has the advantage that once the pseudoinverse is calculated, it can be applied again to any \mathbf{m} . For MRSI where the number of points along the spectroscopic axis is large, the computational load becomes either comparable or even smaller with the second approach. The latter has the additional advantage that the noise characteristics of the reconstructed image can be determined efficiently from the pseudoinverse, as described at the end of this section. Therefore, the second approach is adopted in this work. The pseudoinverse $(\mathbf{E} \mathbf{E}^H)^\dagger$ is calculated by applying an eigendecomposition to $(\mathbf{E} \mathbf{E}^H)$. This decomposition is useful because it yields the eigenvalues as a byproduct, enabling direct control over the amount of noise amplification during matrix inversion (13).

In general, the inversion of $(\mathbf{E} \mathbf{E}^H)$ should be regularized (as described under Methods) in order to limit noise amplification. The final image for the chosen frequency ω is generated from the vector \mathbf{c} (Eq. [9]), either directly by multiplying with \mathbf{E}^H or procedurally as described in Ref. (8), which reduces the computation effort and eliminates the need to store \mathbf{E}^H explicitly.

Point Spread Function and Spatial Response Function

The properties of a linear reconstruction scheme such as the minimum-norm reconstruction can be characterized by the point spread function (PSF) and the spatial response function (SRF). The PSF specifies the reconstructed image for a point source at a given position, whereas the SRF specifies the spatial weighting of signals contributing to a given pixel value. The PSFs and SRFs are the columns and rows, respectively, of the net depiction matrix \mathbf{D} , which is obtained by concatenating the reconstruction and encoding matrices:

$$\mathbf{D} = \mathbf{E}^H (\mathbf{E} \mathbf{E}^H)^\dagger \mathbf{E}. \quad [11]$$

Equation [11] reveals several properties of the PSF and SRF. First, each pixel has an individual PSF and SRF, corresponding to its individual column and row in the matrix \mathbf{D} . Thus, the shape of the reconstructed voxel will generally vary with position, which, among others, can reflect nonuniform spatial resolution as mentioned before. Second, for a given voxel position, the corresponding PSF and the SRF are complex conjugates of one another, since \mathbf{D} is Hermitian according to Eq. [11]. This is important because the SRF is often more difficult to calculate directly than the PSF. Given an efficient reconstruction algorithm, the PSF can be calculated by reconstructing simulated data from a point source. With this Hermitian relationship, the corresponding SRF can then be readily obtained by complex conjugation.

As illustrated by Eq. [11], all SRFs are linear combinations of the encoding functions, which form the rows of

the encoding matrix. It is interesting to note that these specific linear combinations have an optimality property that is closely related to the minimum-norm criterion. Without regularization, the SRF of each voxel is the least-squares optimal approximation of a Dirac distribution placed at the voxel position. In the limit of very fine discretization, the minimum-norm approach hence converges to the strong reconstruction approach proposed in Ref. (2). In this case, the matrix $(\mathbf{E} \mathbf{E}^H)$ approximates the matrix of mutual scalar products of the continuous encoding functions, which is inverted and then likewise left-multiplied by \mathbf{E}^H in the strong approach. In this respect, discretization along a refined grid represents the practical bridging step that connects the computation-efficient but less stringent weak approach and the more stringent but less practical strong approach.

Noise Map Determination

The noise in the reconstructed image is given by

$$\boldsymbol{\eta}_\rho = \mathbf{E}^\dagger \boldsymbol{\eta}. \quad [12]$$

Thus, the noise variance of the reconstructed pixel values is given by the diagonal elements of

$$\langle \boldsymbol{\eta}_\rho, \boldsymbol{\eta}_\rho^H \rangle = \mathbf{E}^\dagger \langle \boldsymbol{\eta}, \boldsymbol{\eta}^H \rangle \mathbf{E}^H = \mathbf{E}^\dagger \mathbf{E}^H, \quad [13]$$

where the second equality holds due to the initial prewhitening. Based on the singular value decomposition of \mathbf{E} , Eq. [13] can be rewritten as

$$\langle \boldsymbol{\eta}_\rho, \boldsymbol{\eta}_\rho^H \rangle = \mathbf{V} \boldsymbol{\Sigma}^{-2} \mathbf{V}^H \quad [14]$$

The square root of the diagonal entries of $\langle \boldsymbol{\eta}_\rho, \boldsymbol{\eta}_\rho^H \rangle$ yields the SD of noise in each reconstructed voxel. This is referred to as the noise map.

METHODS

Calculation of $(\mathbf{E} \mathbf{E}^H)^\dagger$ for Minimum-Norm Reconstruction

The eigendecomposition of the $(\mathbf{E} \mathbf{E}^H)^\dagger$ matrix was computed using the Lanczos method (14). The Lanczos method calculates a tridiagonal decomposition of $\mathbf{E} \mathbf{E}^H$ in a finite number of steps. The resulting tridiagonal matrix was then eigendecomposed with the QL method (15).

An important advantage of the Lanczos method is that it preserves and exploits the symmetry of the matrix. It is computationally efficient, since it relaxes the accuracy for eigenvectors with small eigenvalues. In MRSI experiments, where small eigenvalues are discarded by regularization, it is unnecessary to estimate the complete set of eigenvectors, thus enabling the computational burden to be reduced. For this reason, we use the number of significant eigenvectors determined to provide a measure of the efficiency of the Lanczos method. Additionally, like many efficient linear solvers (15), it only calculates the image of a specific vector under $\mathbf{E} \mathbf{E}^H$, but does not require an explicit representation of the matrix itself. Applying the effects of $\mathbf{E} \mathbf{E}^H$ procedurally is in fact more efficient than multiplying a precalculated $\mathbf{E} \mathbf{E}^H$ matrix explicitly (8).

After the eigendecomposition, the regularized pseudo-inverse ($\mathbf{E} \mathbf{E}^H$)[†] is calculated by inverting the eigenvalues. To avoid excessive noise amplification, for eigenvalues below a condition number (CN) threshold the inverse was set to zero. The CN is the ratio between the maximum and the minimum eigenvalues included in the pseudoinverse estimation.

The eigendecomposition of $\mathbf{E} \mathbf{E}^H$ yields the matrices \mathbf{U} and $\mathbf{\Sigma}$ from Eq. [13]. To determine the noise map via Eq. [14], the i th right singular vector is determined as follows:

$$\mathbf{V}_i = \mathbf{E}^H \mathbf{U}_i \mathbf{\Sigma}_i^{-1}. \quad [15]$$

SENSE Reconstruction

For comparison, data were reconstructed with SENSE as well. In contrast to the procedure reported in Ref.(6), the low-resolution aliased images were first zero-padded in k -space to the same voxel size as the sensitivity maps before SENSE reconstruction. This zero-padding does not affect the reconstruction, but allows better observation of any potential artifacts. In fact, if this reconstruction is resampled at the lower resolution, it is identical to the reconstructed image from the unpadded data.

The sensitivity maps used to generate the data were also used for reconstruction. Portions of the sensitivity maps outside the object borders were set to zero. In some simulations, the sensitivity maps for both reconstructions were only set to zero beyond 10 voxels outside the object borders to test the effects of sensitivity extrapolation, previously proposed by Dydak et al. (6).

For the simulations and MRI studies, the accuracy of the reconstruction was measured as the root-mean-square (RMS) difference between the original and the reconstructed low-resolution image. For MRSI experiments, we calculated the RMS difference between the fully sampled image and the undersampled image after reconstruction.

Simulation

Simulation was used to test the performance of minimum-norm and SENSE reconstruction without noise contamination. A six-element phased-array coil was assumed, with complex-valued sensitivities calculated according to Biot–Savart’s law at a spatial resolution of 256×256 . Low-resolution data for a single frequency ω were generated from the numerical phantom by multiplying the phantom with the complex-valued coil sensitivities, and the central 32×32 k -space points were obtained as the fully sampled low-resolution data. The k -space data were then subsampled by a factor of 2 along k_x and k_y to simulate a net fourfold SENSE acceleration.

Spatial Response Function

The pixelwise SRF was calculated following Eq. [11], with a CN of 10. Noise maps were also estimated for minimum-norm reconstruction, following Eq. [14], and for SENSE reconstructions.

Synthetic Image

The numerical phantom used to test the algorithm consisted of ellipses with different intensities inside a circle.

The main circle contains a bright rim on one side, in order to test potential problems with bright signals close to the border, as occurring with residual fat signals in typical neurologic SI applications.

Noisy synthetic images were generated by adding noise with zero-mean complex-valued Gaussian noise to the k -space data of each coil without correlation between coils. For each coil the noise SD σ was set to 0.1% of the signal level at the k -space center.

Experimental Verification

The reconstruction was also tested with real data from MRI and MRSI experiments. In all cases, slightly extrapolated sensitivity maps were used in image reconstruction as previously proposed in the literature (16).

Magnetic Resonance Imaging

Regular MR imaging was performed in a phantom using a Philips Intera 1.5 T scanner (Philips Medical Systems, Best, The Netherlands). The image was acquired with a gradient echo sequence with a FOV of 210×210 mm² and a matrix size of 256×256 . The resolution was then reduced, selecting a 32×32 window in the k -space center. Sensitivities were estimated from the high-resolution (256×256) image from each coil.

The high signal-to-noise ratio of MRI studies, compared with MRSI experiments, allowed us to better study the behavior of the minimum-norm reconstruction for different condition number thresholds (CN = 10,100, and 1000).

Spectroscopic Imaging

MRSI data were acquired from a phantom on a Philips Intera 3 T scanner (Philips Medical Systems). As described in Ref. (6), the phantom consisted of an elliptical tank with three glass spheres mounted inside. The tank contained water with 5 mM of *N*-acetylaspartate (NAA) and 5 mM of lactate. The spheres (see Fig. 4) contained water doped as follows: 10 mM NAA (left), 5 mM NAA + 5 mM lactate (middle), and 10 mM o lactate (right). Spectra were obtained from a single slice with sampling of a disk within the central 32×32 k -space region. The FOV was 230×230 mm², the slice thickness 20 mm, the spectral bandwidth 2000 Hz, TE 144 ms, and TR 1700 ms. Water suppression was achieved with a chemical-shift-selective (CHESS) pulse centered on the water frequency. For SENSE, the data were again decimated twofold along k_x and k_y to simulate a net $4\times$ acceleration. Sensitivity and B_0 maps were determined from separate reference images obtained with a gradient echo sequence. To reduce noise, we applied a 10-Hz Lorentzian apodization in the time domain. The residual water signal was reduced by multiplying with a sine function centered on the water frequency along the spectral domain. After the spectroscopic image reconstruction, the spectrum of each spatial position was shifted according to the B_0 maps. Finally, metabolite images were formed by integrating the modulus over the spectral frequency ranges pertaining to NAA (2.07–1.97 ppm) and lactate (1.35–1.25 ppm).

In vivo spectroscopic image were obtained in a similar fashion (TR = 1500 ms, TE = 144 ms, bandwidth =

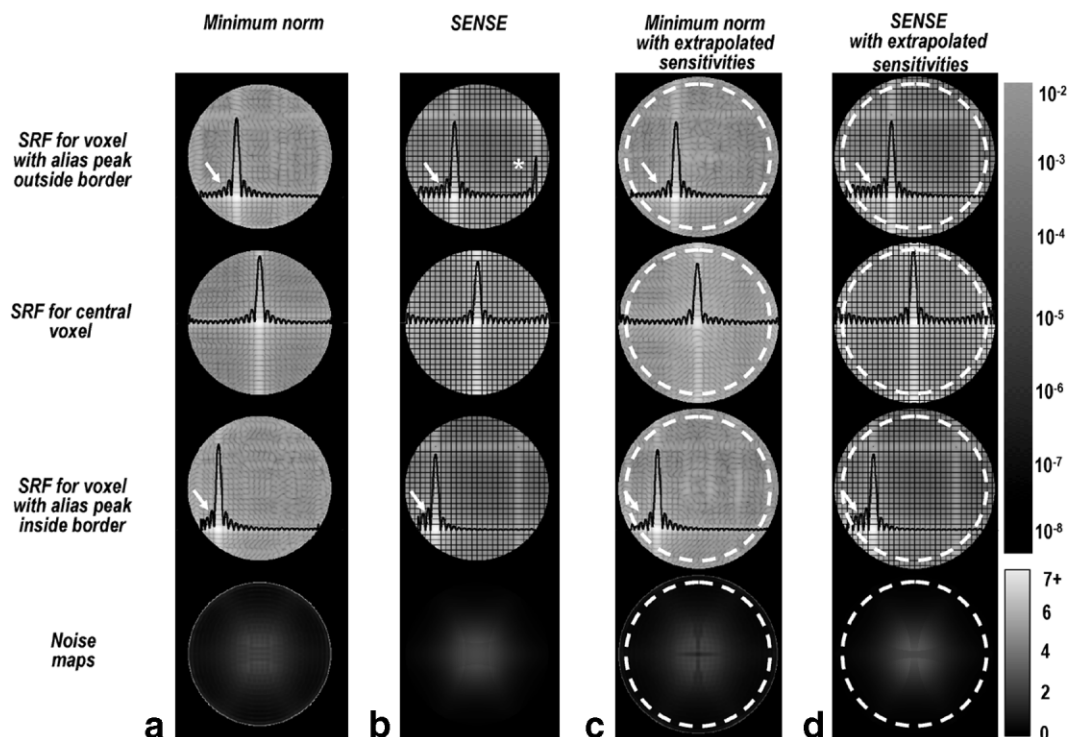


Figure 1. Spatial response function for three voxel locations (three first rows) and noise maps (bottom row). For the SRFs, one-dimensional profiles through the center are overlaid to illustrate the details. The columns from left to right correspond to (a) minimum-norm, (b) SENSE, (c) minimum-norm with sensitivity extrapolation, and (d) SENSE with sensitivity extrapolation. The extrapolation range is indicated by white dashed lines in (c) and (d). Sensitivity extrapolation overcomes residual aliasing for SENSE, while it overcomes noise increase close to the object border for minimum-norm reconstruction.

2250 Hz), using sampling of a circle within the central 24×24 k -space region (FOV = 240×240 mm²). As with the phantom data, water suppression was achieved by a CHES sequence and SENSE data were twofold decimated along the k_x and k_y directions.

In the spectroscopic dimension we applied the same processing as in the in vitro experiments. No postprocessing was applied for lipid suppression in the final reconstructed image. Images were generated by modulus integration of the frequency band for the NAA (2.07–1.97 ppm) and lipid (1.35–1.25 ppm) peaks.

RESULTS

Simulation

Spatial Response Function

Figure 1 shows the SRFs of three pixels for minimum-norm and SENSE reconstructions, without (a, b) and with (c, d) sensitivity extrapolation. The first voxel (top row in Fig. 1) was chosen such that two of its aliasing partners lie just outside the object. The second voxel (second row in Fig. 1) was the central voxel in the image with no direct aliasing partners and the third voxel (third row in Fig. 1) represented the generic case of full fourfold aliasing inside the object.

Without sensitivity extrapolation, SENSE may lead to aliasing artifacts if one or more of the alias peaks lies outside the object border (see asterisk in Fig. 1b, top row). This problem was studied previously in Refs. (6,17). It

occurs when an alias peak falls closely beneath the object borders and is hence not considered in SENSE reconstruction. In such a situation the shoulders and side lobes of the neglected aliasing peak may still cause appreciable artifact. One previously proposed solution is to extrapolate the sensitivity maps slightly beyond the object borders (6). In contrast, minimum-norm reconstruction takes into consideration the entire SRF, so this problem of unsuppressed alias side lobes is avoided (Fig. 1a, top row). The side lobes close to the SRF center are also lower, with a more symmetric distribution compared to SENSE (see arrows in Figs. 1a and b, first and third rows). These improvements are achieved at the expense of increased noise, as reflected in the noise maps (bottom row in Fig. 1). In particular, minimum-norm reconstruction suffers from significantly increased noise close to the object border. This kind of noise increase is known from the literature on finite-support extrapolation (10,18). It is caused by the sensitivity maps being set to zero outside the object borders, which implicitly enforces the constraint that all signals must originate from within the object borders only. Noise components that are inconsistent with this constraint result in error particularly near the borders (10,18). A simple solution to this problem is to extrapolate the sensitivity maps slightly. Most of the noise fringe then falls outside the object.

Figures 1c and d show the reconstruction results with sensitivity extrapolation. For SENSE, the residual aliasing is now eliminated since the alias peak is now located in

Table 1
Standard Deviation of Noise for Voxels Inside the Object in Fig. 1

	Minimum-norm	SENSE	Minimum-norm with sensitivity extrapolation	SENSE with sensitivity extrapolation
Average noise	1.81	1.19	1.31	1.23
Maximum noise	16.16	4.50	4.50	5.41

the extrapolated region (Fig. 1d, top row) (6). However, the SRF remains asymmetric about its main lobe (see arrows). For minimum-norm reconstruction, the SRF shape remains similar to before. However, the peripheral rim of strong noise is now at the border of the extrapolated region and hence outside the object (Fig. 1c, bottom row).

For the second and third voxels (second and third rows in Fig. 1), the SRFs are approximately the same for both methods, with and without sensitivity extrapolation. Nevertheless, the SRF for minimum-norm reconstruction has smaller side lobes for the central voxel and is more symmetric for the generic case of fourfold aliasing.

Table 1 shows the SD of noise for both methods with and without extrapolation. In all cases, the statistics were performed for voxels inside the object borders only. The results show that the minimum-norm reconstruction without sensitivity extrapolation suffers a threefold increase in maximum noise compared with reconstruction. However, this problem is overcome with extrapolated sensitivity maps, as the strong noise amplification then occurs beyond the object borders.

In the case without sensitivity extrapolation, the Lanczos algorithm required only 961 of the 1536 eigenvalues/eigenvectors of the $\mathbf{E} \mathbf{E}^H$ matrix to reach a numerical convergence limit of 10^{-5} relative to the maximum eigenvalue. In the case of sensitivity extrapolation, 1121 of 1536 eigenvalues/eigenvectors were calculated to reach the same accuracy.

Synthetic Image

The numerical phantom is shown in Fig. 2a. The top and bottom rows of Fig. 2 show the reconstruction results

without and with noise added to the data, respectively. In the noiseless simulation (top row), errors in the reconstructed images reflect inaccuracies in signal localization only, whereas the data in the bottom row reflect genuine noise as well.

Without sensitivity extrapolation, SENSE exhibits residual aliasing, as mentioned previously (6). It also exhibits significant ringing artifacts, as expected from the larger side lobes of the SRF (Fig. 1). In comparison, the minimum-norm reconstruction does not exhibit appreciable aliasing. Also, without sensitivity extrapolation, the reconstructed image is seen to have improved resolution close to the object border. This resolution improvement is expected from finite-support extrapolation (10,18) and is directly linked to the enhanced noise close to the borders.

With sensitivity extrapolation, the residual aliasing in SENSE is reduced considerably, but remains noticeable from the object's bright rim (Fig. 2e). This is expected from the SRF (Fig. 1d), since the SRF for the central voxel exhibits slightly elevated "wings" toward the periphery, indicating that the central voxel receives increased contributions from those regions, compared to minimum-norm reconstruction. In addition, it can be seen that the sensitivity extrapolation hardly reduces ringing with SENSE. Some ringing occurs also in the minimum-norm reconstruction when extrapolating the sensitivity maps. However, even with extrapolation, the minimum-norm reconstruction provides improved reconstruction, suffering less ringing and avoiding side-lobe aliasing between the bright rim and the center.

Quantitatively, Table 2 shows the root-mean-squares error for both reconstruction methods with and without

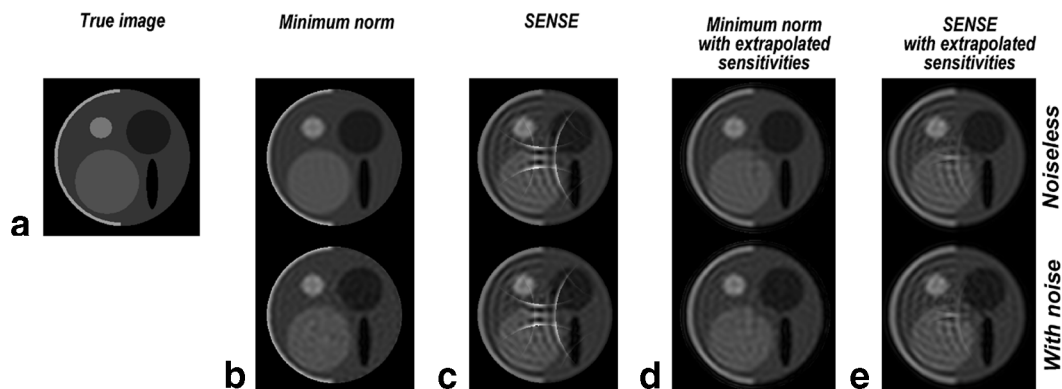


Figure 2. Reconstruction results using minimum-norm (b, d) and SENSE (c, e) reconstruction without (b, c) and with (d, e) sensitivity extrapolation. The true image is shown in high-resolution for comparison (a). The top row shows the noiseless simulation results, so artifacts result from inexact signal localization only. The bottom row shows the simulation with noise, so image errors result from both inexact signal localization and genuine noise.

Table 2
Root-Mean-Square (RMS) Reconstruction Error in Simulation Results Shown in Fig. 2

	Minimum-norm	SENSE	Minimum-norm with sensitivity extrapolation	SENSE with sensitivity extrapolation
RMS error	0.043	0.108	0.064	0.099
Maximum RMS error	0.48	1.12	0.56	0.63

sensitivity extrapolation. In all cases, the statistics were performed for voxels inside the object borders only. These results show that on average the minimum-norm reconstruction outperforms regular Cartesian SENSE both without and with sensitivity extrapolation.

Experimental Verification

Nonspectroscopic Imaging

Figure 3 shows the original high-resolution image, the SENSE reconstruction, and the results after minimum-norm reconstruction for three different CN thresholds (10, 100, and 1000) using extrapolated sensitivity maps. For this experiment, the numerical convergence limit was held tighter (error less than 10^{-8} of the maximum eigenvalue), so the Lanczos method calculated the complete encoding subspace (1536 of 1536 eigenvalues/eigenvectors).

SENSE reconstruction shows some aliasing contamination from the bright rim of the phantom in the center of the image (Fig. 3e). On the contrary, minimum-norm images do not show any aliasing artifacts at any CN (Figs. 3b, c, and d). As the CN increases, the image resolution increases at the periphery and the ringing artifacts are reduced. The improvement in image resolution is related both to complementary encoding from the steep sensitivity gradients and to the prior information inherent to the finite support of the sensitivity maps. It is important to note, however, that the resolution benefit comes at the expense of noise amplification.

Table 3 shows the RMS difference between the original high-resolution image and the subsampled low-resolution images (middle row). According to these results, the minimum-norm images are of higher fidelity than both the SENSE and the fully sampled low-resolution image, particularly for high CN threshold. The benefit stems mainly from increasing effective resolution in parts of the image as the CN threshold increases. These beneficial effects also cause increasing deviations between the minimum-norm images and the low-resolution image, as data are partially

recovered from the borders in k -space surrounding the sampled central region in the low-resolution image.

Spectroscopic Imaging

Figure 4 shows the results of the in vitro experiment, displaying metabolite maps for NAA and Lac. It shows a fully acquired image for reference (Fig. 4a), a minimum-norm image with CN = 10 (Fig. 4b), and a SENSE image (Fig. 4d). Figures 4c and e show the differences between the fully sampled image and the minimum-norm and SENSE images, respectively. Both minimum-norm and SENSE reconstructions were similar to the fully sampled images, but with some noise increase as expected. The two reconstruction methods yielded nearly identical results, presumably because noise obscures the differences, as expected from Fig. 2. Also, this particular phantom did not show strong metabolite contrast close to its surface, which is most critical according to the numerical study.

Quantitatively, Table 4 shows the RMS reconstruction difference for both methods with sensitivity extrapolation and the mean signal intensity from fully sampled and undersampled images. Again, the statistics were performed for voxels inside the object borders only. They confirm the visual assessment that minimum-norm and conventional reconstruction performed nearly equally well in this configuration.

The upper row of Fig. 5 shows NAA maps obtained in vivo with minimum-norm (CN = 10) and SENSE reconstruction, as well as the respective differences from a fully sampled image (Fig. 5, bottom row). Although fat suppression was incomplete, these results illustrate the advantage of minimum-norm reconstruction with respect to adverse side-lobe effects of the SRF. The SENSE image shows some residual ringing (see arrows), which is obviated with the minimum-norm approach.

Table 5 shows the RMS error between the subsampled images and the fully sampled image. Minimum-norm reconstruction performs slightly better than SENSE, mainly due to some residual aliasing with the latter.




Figure 3. (a) Original image acquired with a spatial resolution of 256×256 . (b) Original low-resolution fully sampled image. (c, d, and e) Images obtained from undersampled data with minimum-norm reconstruction at different condition number thresholds (10, 100, and 1000, respectively). (f) Image obtained with SENSE.

Table 3
RMS Errors in Experimental Images shown in Fig. 3

	Fully sampled low-resolution image	Min.-norm, extrapolated sensitivities, CN = 10	Min.-norm, extrapolated sensitivities, CN = 100	Min.-norm, extrapolated sensitivities, CN = 1000	SENSE with sensitivity extrapolation
RMS deviation from high-resolution reference	55.98	54.24	47.38	41.91	59.51
RMS deviation from low-resolution reference	0	16.43	28.40	33.49	18.49

Note. Middle row, RMS deviation from high-resolution reference image. Bottom row, RMS deviation from fully sampled low-resolution image.

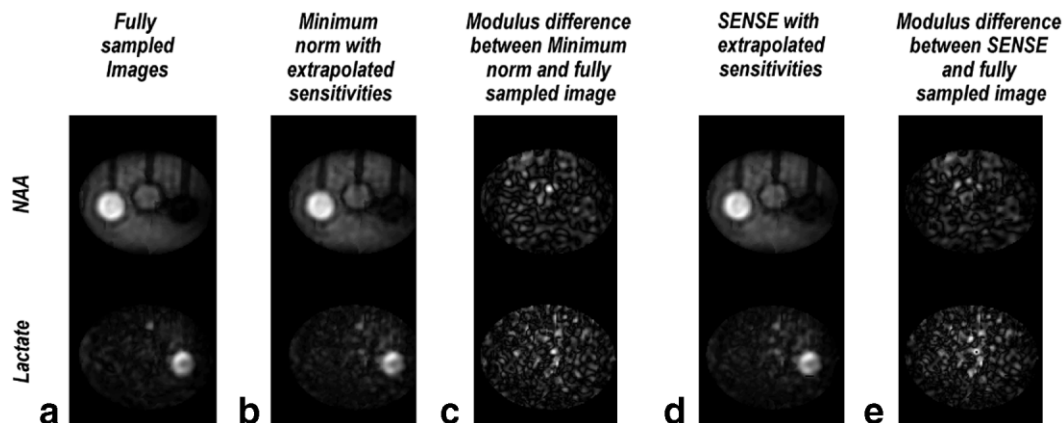


Figure 4. Metabolite images for NAA (top row) and lactate (bottom row) from fully sampled data (a) and from fourfold undersampled data, using minimum-norm reconstruction with CN = 10 (b) and SENSE reconstruction with sensitivity extrapolation (d). c and e show the modulus difference between these results and fully sampled images.

To obtain the singular value decomposition of the $\mathbf{E} \mathbf{E}^H$ matrix for the phantom and in vivo experiments, the Lanczos algorithm required 863 of 1170 and 530 of 666 eigenvectors, respectively, to achieve a numerical convergence limit of 10^{-5} .

DISCUSSION

SENSE achieves its computational efficiency by resolving the signal aliasing only at the voxel centers. This treatment is sufficient for high-resolution imaging, but it may lead to residual aliasing when the acquisition resolution is reduced. We show that minimum-norm reconstruction can overcome this problem by controlling the SRF with higher than nominal resolution. In addition to reducing residual aliasing, the more exact treatment of the SRF also results in more symmetric SRFs and lower side lobes, reducing ring-

ing artifacts. As revealed by the SRF, differences between minimum-norm and SENSE reconstructions are more pronounced for signals that originate close to the object borders. Therefore, minimum-norm reconstruction may be particularly useful for brain MRSI without PRESS or outer volume suppression where bright residual fat signals may originate from the scalp region.

For both minimum-norm and SENSE reconstruction, sensitivity extrapolation was found to improve image quality, but for different reasons. In SENSE, extrapolation is needed to avoid residual aliasing (6). In minimum-norm reconstruction, extrapolation helps avoiding the noise increase close to the border caused by finite-support effects (9,10).

The improved SRF of minimum-norm reconstruction is achieved at the expense of potential noise increase as reflected by the condition number. Therefore, the improve-

Table 4
Errors and Mean Signal Amplitudes in in Vitro Metabolite Images shown in Fig. 4

	Fully sampled low-resolution image	Minimum-norm with extrapolated sensitivities (CN = 10)	SENSE with extrapolated sensitivities
NAA RMS error	0	0.028	0.025
Mean NAA signal	22.10	22.68	22.63
Lactate RMS error	0	0.010	0.010
Mean lactate signal	1.98	2.24	2.21

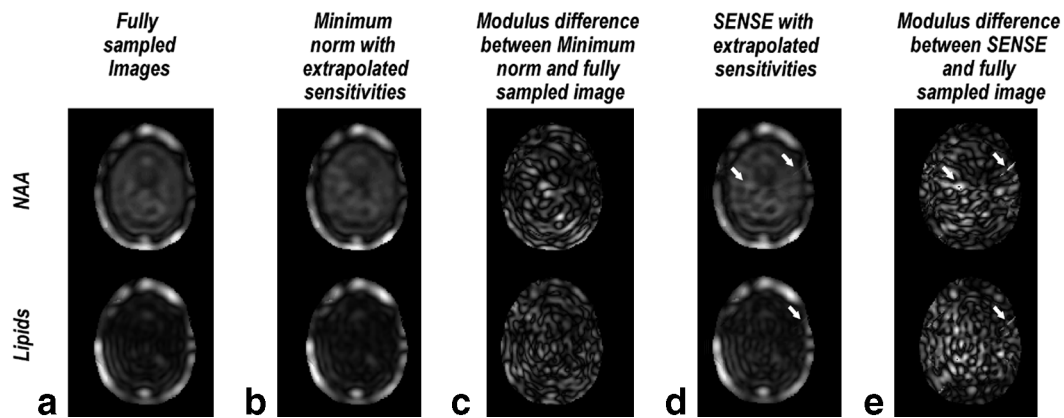


Figure 5. Metabolite images for NAA (top row) and lipids (bottom row) from fully sampled data (a) and from fourfold undersampled data, using minimum-norm reconstruction with $CN = 10$ (b) and SENSE reconstruction with sensitivity extrapolation (d). c and e show the respective modulus differences relative to fully sampled images.

Table 5
RMS Errors in in Vivo Metabolite Images Shown in Fig. 5

	Minimum-norm with extrapolated sensitivities ($CN = 10$)	SENSE with extrapolated sensitivities
NAA RMS error	0.030	0.033
Lipid RMS error	0.099	0.129

ment in image quality that can actually be realized depends on the signal-to-noise ratio of the raw MR data. For MRSI experiments with intrinsically low SNR the condition number threshold must be kept relatively low, thus limiting the image quality benefit. In practice, noise in the data may obscure the differences between minimum-norm and SENSE reconstructions. Nevertheless, it is important to note that the noise only makes the artifacts more difficult to spot, especially in low-resolution images typical of MRSI.

CONCLUSION

We propose the use of minimum-norm reconstruction for SENSE MRSI. The reconstruction procedure solves problems associated with low-resolution SENSE by considering the SRF as a whole rather than only in the voxel centers. The improved SRF quality comes at the expense of potential conditioning problems, which lead to noise amplification, particularly at the object borders. However, this issue can be mitigated by sensitivity extrapolation and regularization. Minimum-norm reconstruction requires significantly more computation compared to common Cartesian SENSE. In this work, we applied the Lanczos method to solve the underlying matrix inversion problem. The proposed reconstruction method is applicable to arbitrary k -space sampling patterns, including variable density sampling patterns.

ACKNOWLEDGMENTS

The authors are grateful for the continuing support of Philips Medical Systems.

REFERENCES

- Sodickson DK, Manning WJ. Simultaneous acquisition of spatial harmonics (SMASH): fast imaging with radiofrequency coil arrays. *Magn Reson Med* 1997;38:591–603.
- Pruessmann KP, Weiger M, Scheidegger MB, Boesiger P. SENSE: sensitivity encoding for fast MRI. *Magn Reson Med* 1999;42:952–962.
- Kyriakos WE, Panych LP, Kacher DF, Westin CF, Bao SM, Mulkern RV, Jolesz FA. Sensitivity profiles from an array of coils for encoding and reconstruction in parallel (SPACE RIP). *Magn Reson Med* 2000;44:301–308.
- Griswold MA, Jakob PM, Heidemann RM, Nittka M, Jellus V, Wang J, Kiefer B, Haase A. Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magn Reson Med* 2002;47:1202–1210.
- Griswold MA, Jakob PM, Nittka M, Goldfarb JW, Haase A. Partially parallel imaging with localized sensitivities (PILS). *Magn Reson Med* 2000;44:602–609.
- Dydak U, Weiger M, Pruessmann KP, Meier D, Boesiger P. Sensitivity-encoded spectroscopic imaging. *Magn Reson Med* 2001;46:713–722.
- Dydak U, Pruessmann KP, Weiger M, Tsao J, Meier D, Boesiger P. Parallel spectroscopic imaging with spin-echo trains. *Magn Reson Med* 2003;50:196–200.
- Pruessmann KP, Weiger M, Bornert P, Boesiger P. Advances in sensitivity encoding with arbitrary k -space trajectories. *Magn Reson Med* 2001;46:638–651.
- Tsao J, Behnia B, Webb AG. Unifying linear prior-information-driven methods for accelerated image acquisition. *Magn Reson Med* 2001;46:652–660.
- Tsao J. Extension of finite-support extrapolation using the generalized series model for MR spectroscopic imaging. *IEEE Trans Med Imaging* 2001;20:1178–1183.
- Sánchez J, Tsao J, Dydak U, Desco M, Boesiger P, Pruessmann K. Minimum norm reconstruction for parallel spectroscopic imaging. In: *Proceedings of the 11th Annual Meeting of ISMRM, Toronto, Canada, 2003*. p 111129.
- Albert A. In: Bellman R, editor. *Regression and the Moore–Penrose pseudoinverse*. New York: Academic Press; 1972. 180 p.
- Hansen PC. *Regularization tools: A Matlab package for analysis and solution of discrete ill-posed problems*. Numer Algor 1994;6:1–35.
- Lanczos C. An iterative method for the solution of the eigenvalue problem of linear differential and integral operators. *J Res Nat Bur Stand* 1950;45:255–282.
- Golub GH, Van Loan CF. *Matrix computations*. Baltimore: Johns Hopkins University Press; 1996.
- Dydak U, Pruessmann KP, Weiger M, Boesiger P. The spatial response function in SENSE-SI. In: *Proceedings of the 8th Annual Meeting of ISMRM, Denver, CO, USA, 2000*. p 81848.
- Zhao X, Prost RW, Li Z, Li SJ. Reduction of artifacts by optimization of the sensitivity map in sensitivity-encoded spectroscopic imaging. *Magn Reson Med* 2005;53:30–34.
- Plevritis SK, Macovski A. A spectral extrapolation of spatially bounded images. *IEEE Trans Med Imaging* 1995;14:487–497.