View metadata, citation and similar papers at core.ac.uk
brought to you by CORE
provided by Apollo

Journal of Magnetic Resonance 261 (2015) 27–37

Contents lists available at ScienceDirect

Journal of Magnetic Resonance

journal homepage: www.elsevier.com/locate/jmr

Quantitative mapping of chemical compositions with MRI using compressed sensing

Erik von Harbou [a], Hilary T. Fabich [b], Martin Benning [b,1], Alexander B. Tayler [b], Andrew J. Sederman [b], Lynn F. Gladden [b], Daniel J. Holland [b,*]

[a] Laboratory of Engineering Thermodynamics, University of Kaiserslautern, Erwin-Schrödinger-Straße 44, 67663 Kaiserslautern, Germany
[b] Department of Chemical Engineering and Biotechnology, University of Cambridge, Pembroke Street, Cambridge CB2 3RA, United Kingdom

ARTICLE INFO

Article history:
Received 20 July 2015
Revised 21 September 2015
Available online 19 October 2015

Keywords:
Concentration mapping
Compressed sensing
Fast acquisition
Quantitative MRI
Total variation
Bregman iteration

ABSTRACT

In this work, a magnetic resonance (MR) imaging method for accelerating the acquisition time of two dimensional concentration maps of different chemical species in mixtures by the use of compressed sensing (CS) is presented. Whilst 2D-concentration maps with a high spatial resolution are prohibitively time-consuming to acquire using full **k**-space sampling techniques, CS enables the reconstruction of quantitative concentration maps from sub-sampled **k**-space data. First, the method was tested by reconstructing simulated data. Then, the CS algorithm was used to reconstruct concentration maps of binary mixtures of 1,4-dioxane and cyclooctane in different samples with a field-of-view of 22 mm and a spatial resolution of 344 μm × 344 μm. Spiral based trajectories were used as sampling schemes. For the data acquisition, eight scans with slightly different trajectories were applied resulting in a total acquisition time of about 8 min. In contrast, a conventional chemical shift imaging experiment at the same resolution would require about 17 h. To get quantitative results, a careful weighting of the regularisation parameter (via the L-curve approach) or contrast-enhancing Bregman iterations are applied for the reconstruction of the concentration maps. Both approaches yield relative errors of the concentration map of less than 2 mol-% without any calibration prior to the measurement. The accuracy of the reconstructed concentration maps deteriorates when the reconstruction model is biased by systematic errors such as large inhomogeneities in the static magnetic field. The presented method is a powerful tool for the fast acquisition of concentration maps that can provide valuable information for the investigation of many phenomena in chemical engineering applications.
© 2015 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

1. Introduction

Maps of chemical compositions can provide valuable information for many applications, especially in chemical engineering. They can be used to gain a rigorous understanding of chemical processes and mass transfer phenomena occurring, for example, in catalyst beds, along interfaces, or in and near membranes. This understanding is important for a reliable design and scale-up of chemical processes. Taking samples and analysing them *ex situ* is often not feasible because the sampling disturbs the system and the effort is immense to obtain sufficient spatial resolution to resolve the processes. In this application, magnetic resonance imaging (MRI) offers great potential as it is a non-invasive, spatially resolved measurement technique able to probe optically opaque environments like reactors. In situ MRI has been successfully applied to study conversion and composition profiles or local reaction rates along fixed-bed reactors for various reactions using spatially resolved ^{1}H NMR-spectroscopy [1,2] and ^{13}C NMR-spectroscopy [3–5] also called chemical shift imaging (CSI). The acquisition time needed to obtain multidimensional, fully sampled concentration maps, however, may take several hours [3] which can be detrimental. First, the process has to be operated steadily for several hours so the consumption of chemicals is high which is costly and undesirable concerning the safety in laboratories. Second, transient phenomena that take place within minutes cannot be studied with this technique. This paper presents a method for accelerating the acquisition of spatially resolved concentration maps by the use of compressed sensing (CS).

CS enables the accurate reconstruction of an under-sampled signal by utilising the prior knowledge that the signal is compressible

* Corresponding author at: Department of Chemical and Process Engineering, University of Canterbury, Private Bag, 4800 Christchurch, New Zealand.
E-mail address: daniel.holland@canterbury.ac.nz (D.J. Holland).
[1] Current address: Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Wilberforce Road, CB3 0WA Cambridge, United Kingdom.

http://dx.doi.org/10.1016/j.jmr.2015.09.013
1090-7807/© 2015 The Authors. Published by Elsevier Inc.
This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

or sparse with respect to a specific representation [6,7]. As under-sampled signals can be used, CS provides a method of reducing the data acquisition times characteristic of many imaging techniques. CS has been successfully applied to reduce the acquisition time of MR images [8,9]. Holland et al. [10] and Tayler et al. [11] demonstrated the potential of CS by reconstructing velocity images in fixed-bed reactors and of multiphase flow, respectively from fast and under-sampled phase-encoded MR measurements. Furthermore, Holland et al. [12] and Kazimierczuk and Orekhov [13] applied CS for fast multidimensional NMR spectroscopy. Hu et al. [14] and Kampf et al. [15] used CS for the accurate reconstruction of three dimensional chemical shift imaging (CSI) of ^{13}C and ^{19}F markers, respectively from under-sampled data sets. When the chemical shift information of the observed chemical species is known and is incorporated into the model used for the reconstruction, images showing different species can be directly recovered with high resolution from the under-sampled signals by CS. Good results with a significant reduction of the scanning time compared to conventional methods have been achieved in medical applications with this method for imaging water and fat [16–18]. The focus of these works was to get a good separation of water and fat in the reconstructed images and not to obtain quantitative information on the composition.

In this work, we apply CS reconstruction to resolve spatially and quantitatively the compositions of different species in mixtures. This method enables the mapping of the composition directly as a function of space. Only the information about the chemical shift of the observed species are required for the reconstruction; there is no need for calibration prior to the analysis. This feature of the presented method is beneficial for many applications in chemical engineering where unstable intermediates are formed during the process that make a calibration impossible. To achieve a high accuracy of the concentration map, however, the parameters of the CS algorithm have to be correctly set. As mentioned above, CS exploits prior knowledge of the signal. This prior knowledge is integrated in the CS solver with a regulariser [6,7]. To get quantitative results, the systematic bias of the CS reconstruction has to be minimised, either by carefully weighting the regulariser or by applying contrast-enhancement approaches. Different generic approaches exist for the identification of good regularisation parameters. In the present work, two different approaches, the L-curve approach [19] and the Bregman iterations [20], are applied for the reconstruction of simulated data of a phantom sample and for the reconstruction of measured data from binary mixtures in different test samples. These results are used to assess the robustness of the approaches to yield concentration maps with a high accuracy. Finally, we present a discussion of the strengths and limitations of the method for the spatial quantification of chemical species.

2. Reconstruction using compressed sensing

2.1. Model equations

The measured **k**-space signal **S** at the echo time t is related to the concentration maps $\mathbf{x}_k$ of all species $k = 1, \ldots, M$ via the signal model [16]

$$\mathbf{S}(t) = \sum_{k=1}^{M}\left(\sum_{j=1}^{L_k} w_{k,j}\exp\left(2\pi i\,\delta_{k,j}t\right)\exp\left(-\frac{t+2\tau}{T_2^*}\right)\right.$$
$$\left.\times\int_{\Omega}\mathbf{x}_k(\mathbf{r})\exp\left(2\pi i\,\mathbf{k}(\mathbf{t})\cdot\mathbf{r}\right)\mathrm{d}\mathbf{r}\right)+\mathbf{v} \tag{1}$$

with

$$\mathbf{k}(t) = \frac{1}{2\pi}\int_0^t \gamma\mathbf{G}(t')\mathrm{d}t' \quad \text{and} \quad \Omega \subset \mathbb{R}^2. \tag{2}$$

In Eq. (1), **v** is the noise. $\delta_{k,j}$ denotes the relative chemical shift (related to the resonance frequency of the spectrometer) of the j-th group (peak) that belongs to species k. $w_{k,j}$ is the group weighting factor that exists for all groups $j = 1, \ldots, L_k$ of species k. It describes the mole of the nuclei (here ^{1}H: $n^{^1\mathrm{H}}$) in the j-th group per mole of species k (n_k^{species}), see Eq. (3). To get quantitative results from the measured **k**-space signal, the group weighting factors have to be set correctly.

$$w_{k,j} = \frac{n^{^1\mathrm{H}}}{n_{k,j}^{\text{group}}}\,\frac{n_{k,j}^{\text{group}}}{n_k^{\text{species}}}. \tag{3}$$

In Eq. (2), 2τ denotes the time from the excitation pulse to the centre of the echo and T_2^* denotes the apparent T_2-relaxation time. **G** is the vector of the magnetic field gradient that acts at the echo time t. Here, we subsample the **k**-space $S_{p,q}$ as $S(t)$. Eq. (1) can be abbreviated with linear operators, see Eq. (4). The explicit equations for the operators are given in Appendix A.

$$\mathbf{S} = \mathbf{CHS}\cdot\mathcal{F}_{\mathrm{u}}\cdot\mathbf{x}+\mathbf{v}. \tag{4}$$

In Eq. (4), **CHS** denotes the chemical shift operator, $\mathcal{F}_{\mathrm{u}}$ is the undersampled Fourier transform, and **x** is the concatenated matrix of all concentration maps $\mathbf{x}_k$ with $k = 1 \ldots M$. Eq. (4) can only be applied when spatial and temporal inhomogeneities in the $\mathbf{B}_0$-field are negligible.

2.2. Solving strategy

The goal of the reconstruction is to find well resolved concentration maps **x** from the under-sampled **k**-space measurements **S** so that the signal model according to Eq. (4) is fulfilled. In CS, the reconstruction is obtained by solving a Tikhonov-type optimisation problem of the form (for details, see e.g. Benning et al. [21]):

$$\mathbf{x}_{\text{reconstructed}} = \arg\min_{\mathbf{x}}\left\{\frac{1}{2}\|\mathbf{S}-\mathbf{CHS}\cdot\mathcal{F}_{\mathrm{u}}\cdot\mathbf{x}\|_2^2+\sum_{k=1}^{M}\alpha_k J(\Psi\mathbf{x}_k)\right\}. \tag{5}$$

The first term in Eq. (5) is the fidelity term that models Eq. (1). Here $\|v\|_2 := \sqrt{\sum|v(i)|^2}$is the standard Euclidean 2-norm. The second term is the regularisation with $J(\Psi\mathbf{x}_k)$ as regularisation functional that enables the incorporation of prior information on the reconstruction. Ψ is a linear operator that transforms the concentration maps **x** to another domain where they are sparse. Thus, the solution of Eq. (5) yields concentration maps that have a sparse representation in the transform domain and that are, according to Eq. (1), consistent with the measured **k**-space data in the least squares sense. The parameter α_k is a positive regularisation parameter that weights the influence of the fidelity and the regularisation term. We found that quantitative reconstruction results are only obtained when the parameters $\alpha_1, \alpha_2, \ldots, \alpha_M$ are not chosen independently but based on the group weighting factors $w_{k,j}$ and a constant positive regularisation parameter α,

$$\alpha_k = \alpha\sum_{j=1}^{L_k} w_{k,j}. \tag{6}$$

The concentration maps of the test samples used in the present work to test the method contain sharp edges. Thus, a finite-difference approximation of the gradient operator is used as the sparsifying transform Ψ for all reconstructions carried out in this work. For a discrete, isotropic total variation the regularisation functional becomes $J(\Psi x_k) = \|\Psi x_k\|_{2,1} = \|\|\Psi x_k\|_2\|_1$. (details of the computation of this term are given in Appendix B) Depending on the features of the concentration maps, further sparsifying transforms, such as wavelet transforms, which are used for smooth changes in the concentration maps, or other one-norm-based

regularisers like the Total Generalised Variation can be chosen as well. A detailed discussion of different regularisers and their implementation is given, for example, by Benning et al. [21].

To solve Eq. (5), we used an inhouse-code written in MATLAB (product of MathWorks, Natick, USA) that is based on a scaled alternating direction method of multipliers (ADMM) [22]. The under-sampled Fourier transformation was performed using a non-uniform fast Fourier transform algorithm that had been developed by Fessler and Sutton [23] and that is accessible online as an open source toolbox [24]. The concentration map $\mathbf{x}_k$ obtained by solving Eq. (5) is given in arbitrary units (a.u.). To get the concentration map $\mathbf{x}_k^{(n)}$ in mole fractions, the concentration $x_k(p)$ of species k has to be scaled in each pixel p,

$$x_k^{(n)}(p) = \frac{x_k(p)}{\sum_{k=1}^{M} x_k(p)}. \tag{7}$$

If pixel p^{out} lies outside the sample where none of the species are present, the sum $\sum_{k=1}^{M} x_k(p^{out})$ of a well reconstructed concentration map approaches zero and here the concentration of each species $x_k^{(n)}(p^{out})$ is set to zero by default. In this work, the pixels that lie outside the sample were identified from the "best" of the Bregman iteration reconstructions. The same pixels were set to zero in all reconstructions. It would also be possible and potentially advantageous to identify these pixels from an independent experiment, see for example [10,25], however that was not done here.

The correct choice of the regularisation parameter α is not trivial. In this regard, different approaches have been described in literature. For example, Holland et al. [10] used simulated data to determine the regularisation parameter that yielded the best reconstruction results for a given signal to noise ratio (SNR). The drawback to this approach is that it is always necessary to simulate data very similar to the system under investigation. Hansen [19] suggested the L-curve as a more generic approach to choose a regularisation parameter. The L-curve plots the regularisation term – l_1-norm, here: $\sum_{k=1}^{M}\sum_{j=1}^{L_k} w_{k,j}\|\Psi\mathbf{x}_k\|_{2,1}$ – versus the norm of the fidelity term – l_2-norm, here: $\|\mathbf{S} - \mathbf{CHS}\cdot\mathcal{F}_u\cdot\mathbf{x}\|_2$. An example of an L-curve is depicted in Fig. 1. The L-curve starts at low values of the fidelity term and high values of the regularisation term. In other words, the reconstruction fits the measurements precisely but the image likely contains a lot of noise or artefact as it is not well regularised. As the value of α increases, the l_1-norm of the regularisation term decreases. Initially large changes in the l_1-norm are associated with only small changes in the data fidelity term, thus the curve is steep. At some value of α, further increases in α result in small decreases in the l_1 term and large increases in the l_2 term, thus the curve becomes flat. The resulting curve looks approximately "L"-shaped. The point at which the curve turns from a sharp decrease to a flat line is known as the corner of the L-curve, and is indicated by the arrow on Fig. 1. This L-curve represents the range of possible solutions that provide a compromise between the two-norm of the fidelity and the one-norm of the regularization and as such is often considered as a Pareto frontier [26]. The best regularisation parameter corresponds to the reconstruction result that appears on the L-curve in that corner (or a little bit to the right) [19]. Thus, by varying the regularisation parameter in a broad range and plotting the L-curve, a selection criterion for an optimal regularisation parameter is provided.

Benning et al. [21] applied a different approach called Bregman iterations to obtain quantitative phase reconstruction from velocity-encoded MRI measurements. For Bregman iteration, the regulariser is replaced by its Bregman distance in order to create an iterative procedure that refines the solution the further one iterates. For this approach, the regularisation parameter α is set to a value that strongly overweights the regularisation term (cf. Eq. (5)) and the following iterative procedure is carried out [27]:

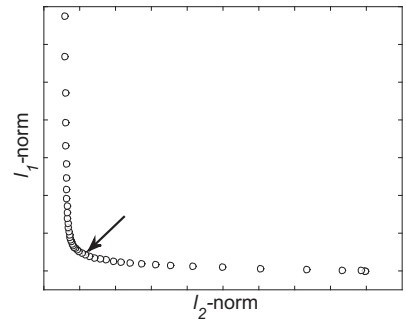

Fig. 1. An example of an L-curve. By varying the regularisation parameter α, the regularisation term can be plotted versus the norm of the fidelity term. The arrow indicates the corner of the L-curve which corresponds to the optimal regularisation parameter α.

$$\mathbf{x}^m = \arg\min_{\mathbf{x}} \left\{ \frac{1}{2}\left\|\mathbf{S}^{m-1} - \mathbf{CHS}\cdot\mathcal{F}_u\cdot\mathbf{x}\right\|_2^2 + \alpha\sum_{k=1}^{M}\sum_{j=1}^{L_k} w_{k,j}\|\Psi\mathbf{x}_k\|_{2,1} \right\} \tag{8a}$$

$$\mathbf{S}^m = \mathbf{S}^{m-1} + (\mathbf{S} - \mathbf{CHS}\cdot\mathcal{F}_u\cdot\mathbf{x}^m) \quad \text{with} \quad \mathbf{S}^0 = \mathbf{S} \tag{8b}$$

The iteration given in Eqs. (8a) and (8b) is repeated until a stop criterion is satisfied. Benning et al. [21] and Yin et al. [27] demonstrated that Morozov's discrepancy principle [28], given in Eq. (9), yielded satisfactory reconstruction results in combination with the Bregman iteration.

$$\|\mathbf{S} - \mathbf{CHS}\cdot\mathcal{F}_u\cdot\mathbf{x}\|_2 \leqslant \sigma\sqrt{N_{\text{samples}}} \tag{9}$$

In Eq. (9), σ denotes the standard deviation of the noise and N_{samples} is the number of samples. Thus, the right hand side of Eq. (9) refers to the noise level. Morozov's discrepancy principle states that the error between the sub-sampled Fourier transform of the reconstruction and the measured **k**-space data differ by less than the normally distributed noise (which has mean zero and standard deviation σ). As long as this deviation is larger, data and reconstruction will differ by more than just noise. This stopping criterion is also applicable for selecting the optimum α value using the L-curve approach. It has the advantage that it is mathematically well-defined compared to the selection criterion "in the corner of the L-curve".

The L-curve and Bregman iterations were applied in the present work to reconstruct quantitatively maps of the composition of simulated data and data from real measurements of test samples. By comparing the obtained maps of the composition with the expected values, the performance of both approaches and the applicability of the selection criterion (in the corner of the L-curve) and the stopping criterion (Morozov's discrepancy principle) is assessed.

3. Experiments

All experiments were performed on a Bruker AV-400 spectrometer (Rheinstetten, Germany) operating at a ^{1}H resonance frequency of 400.25 MHz with a vertical 9.4 T superconducting magnet. The spectrometer was equipped with a 25 mm diameter birdcage radio-frequency coil and with a shielded and water cooled gradient system producing a maximum gradient strength of 1.46 T/m in the x, y, and z directions.

3.1. Sampling scheme and acquisition parameter

The concentration maps were obtained with a slice selective two-dimensional spin echo pulse sequence using a 90° hard pulse and a 180° gaussian shaped soft pulse. Spiral trajectories were chosen to subsample **k**-space. As demonstrated by Tayler et al. [11], spiral trajectories present a suitable sampling scheme for CS. In

the present work, however, two different spirals strung together into a single trajectory were used to ensure the centre of the echo was formed near the centre of **k**-space. The spirals were constructed using an algorithm that had been developed by Lustig et al. [29] and that is accessible online as open source toolbox [30]. Two different basic types of trajectories (type A and type B) were designed. Type A starts at the centre of **k**-space, spirals out, comes back straight, goes through the centre of **k**-space, out again, and finally, it spirals back to the centre of **k**-space. Trajectory B is simply the opposite. Starting at the centre of **k**-space, it goes straight out, spirals back to the centre of **k**-space, spirals out again, and comes straight back to the centre of **k**-space. The entire sampling scheme employed for the concentration map consists of eight trajectories (four of type A and four of type B). Each trajectory is turned about the centre of **k**-space by a different angle so that a good coverage of **k**-space is achieved. To increase the randomness of the sampling scheme, the spirals are distorted with a sinusoidal oscillation at a higher frequency and lower amplitude than the main spiral trajectory. The direction of the oscillation was chosen such that it was perpendicular to the direction of the main spiral at all times. The amplitude and frequency were different for each spiral. Furthermore, different numbers of points were added at the beginning and/or removed at the end of each of the eight trajectories. In this way, all eight trajectories have the same length (number of points) but the centre of the echo is formed at different locations that are distributed around the centre of **k**-space. A better resolution of the chemical shift information is obtained in experiments using this type of trajectory compared to experiments with trajectories where the echo is always formed in the centre of **k**-space. The trajectories were then further processed using the algorithms of Lustig et al. [29]. By adding and removing points within the trajectories the algorithm ensures that the trajectories yield the desired field-of-view (FOV) and that they do not exceed the maximum gradient strength and slew rate achievable by the hardware. Other sampling schemes based on Lissajou curves or lemniscates were tested as well but gave significantly worse results compared to the results obtained with the spiral based trajectories.

For a good reconstruction of the concentration maps, the trajectories generated by the gradients during the acquisition have to be known very precisely. For that reason, the trajectories were measured using the technique of Duyn et al. [31]. To reduce errors in the phase measurement associated with inhomogeneities in the $\mathbf{B}_0$-field, the technique was slightly modified and a volume selective excitation was used as suggested by Tayler et al. [32].

In the present work, a sampling scheme designed as described above was employed with 8×551 complex data points and a dwell time of 2.5 μs. The sampling scheme is depicted in Fig. 2. The concentration maps were obtained with a field-of-view of 22 mm × 22 mm and a resolution of 344 μm × 344 μm for a slice thickness of 0.5 mm. The repetition time of the experiment was approximately 15 s and a 4 step phase-cycle was used, giving a total acquisition time of about 8 min.

The sampling scheme is obtained by integration (cumulative summation) of the acquire data as describe by Duyn et al. [31]. Thus the small measurement errors add up so that points at the end of the sampling scheme are subjected to larger errors than points earlier in the sampling scheme. For that reason, the reproducibility in the measurements at the same values of k_x and k_y is better in the part of the sampling scheme shown in Fig. 2 (b) at $k_x \approx 7.5\ \mathrm{cm}^{-1}$ and $k_y \approx -6\ \mathrm{cm}^{-1}$ when compared to the part of the sampling scheme at $k_x \approx 8\ \mathrm{cm}^{-1}$ and $k_y \approx -5\ \mathrm{cm}^{-1}$; the latter points being acquired much later during the acquisition.

3.2. Generation of simulated data

To generate simulated data, first, phantom concentration maps of a binary mixture of species A and B were created. In the present work, the concentrations were set to a constant value ($x_A^{(n)} = 0.667$ mol/mol, $x_B^{(n)} = 0.333$ mol/mol). The relative chemical shifts and group weighting factors of species A were $\delta_A = (800\ \mathrm{Hz}, 200\ \mathrm{Hz}, -400\ \mathrm{Hz})$ and $\mathbf{w}_A = (3/8, 1/4, 1/8)$, respectively. The relative chemical shift and weighting factor of species B was $\delta_B = 0$ Hz and $w_B = 1/4$, respectively. The image size was set to 64×64 pixels. The phantom concentration map of species A is shown in Fig. 3.

By means of Eq. (4), simulated data were generated for these phantom concentration maps. The noise $\mathbf{v}$ was Gaussian distributed and the noise level was set according to the experimental noise level determined by repeated measurements.

3.3. Preparation of test samples

To test the reconstruction method experimentally, two different test samples of about 5 ml were prepared in vials (inner diameter: 19 mm). Test sample A was a binary homogeneous mixture of cyclooctane and 1,4-dioxane ($x_{\mathrm{dioxane}} = 0.761$ mol/mol). For the preparation of test sample B, a small vial (inner diameter: 11 mm) was inserted into the large vial. Both vials were filled with binary homogeneous mixtures of cyclooctane and 1,4-dioxane with different compositions. The concentration of 1,4-dioxane in the small, inner vial was $x_{\mathrm{dioxane}}^{\mathrm{inner}} = 0.666$ mol/mol and in the large, outer vial $x_{\mathrm{dioxane}}^{\mathrm{outer}} = 0.415$ mol/mol. Additionally, a Teflon tube

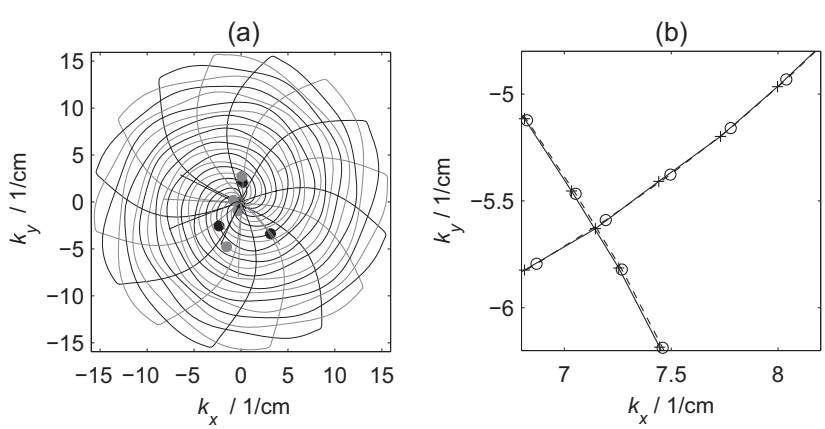

Fig. 2. (a) Spiral based sampling scheme with 8×551 data points. The bullets indicate the location of the centre of the spin echo. (b) Zoomed view of a comparison of repeated measurements (+/∘). The measured points in k-space are indicated by the symbols; the lines are linear interpolations to guide the eye.

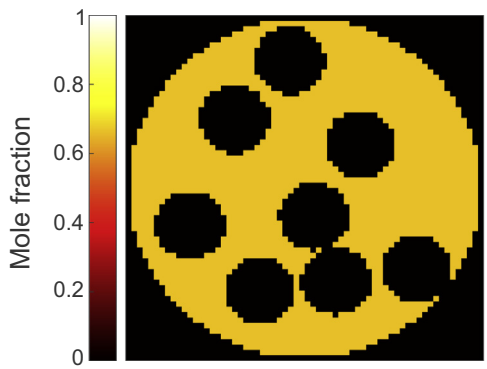

Fig. 3. Phantom concentration map of species A with $x_A^{(n)} = 0.667$ mol/mol. Image size: 64 × 64 pixels.

(inner diameter 4 mm) was inserted into the large, outer vial to test the resolution of the reconstruction method. As the Teflon tube was open at the ends, the composition of the liquid inside the Teflon tube was the same as the composition of the liquid contained in the large, outer vial.

The chemical shift of interest in these reconstructions is the chemical shift relative to the resonant frequency of the acquisition. This chemical shift is measured relative to one of the frequencies in the sample, and not relative to a standard reference species, such as tetramethylsilane. In this case, we set the resonant frequency of the instrument to the frequency of the peak for 1,4-dioxane. The relative chemical shift for the binary mixture of 1,4-dioxane and cyclooctane was then determined by the acquisition of a standard ^{1}H-spectrum of the sample. The results are summarised in Table 1.

4. Results and discussion

4.1. Reconstruction of simulated data

Fig. 4(a) shows the relative error of the reconstructed concentration map of species A compared to the concentration map that was input to the simulation as shown in Fig. 3. In Fig. 4(b), a comparison of the reconstructed concentration profile compared to the input concentration profile in the middle of the sample is depicted. The reconstruction was carried out with a regularisation parameter $\alpha = 0.004$ determined by the L-curve approach (for details, see below). The results demonstrate that both a good spatial resolution is achieved in the reconstructed image and that the reconstruction yields almost perfect quantitative results. Only at the corners and edges larger deviations occur in the reconstructed concentration map. The reason for this behaviour is a systematic error that is introduced because of the discretisation of the gradient operator Ψ with finite differences (cf. Eq. (5)). Other discretisation approaches exist with a lower systematic error [33]. The optimisation of the gradient operator, however, is not in the scope of this paper. The reconstruction results using 19 Bregman iterations give similar results (not shown here).

Table 1
Measured relative chemical shift and weighting factor (mole of ^{1}H per mole of the species).

Species	Group	Rel. chem. shift (Hz)	Weighting
1,4-Dioxane	CH_2	0	8
Cyclooctane	CH_2	−814	16

4.1.1. Optimisation of reconstruction

As mentioned above, both the L-curve approach and Bregman iterations were applied to reconstruct the simulated data of the phantom concentration maps. The reconstruction results are summarised in Fig. 5 for the L-curve approach and in Fig. 6 for the Bregman iterations. Figs. 5(a/b) and 6(a/b) show the reconstructed mean mole fraction of species A and the standard deviation of the mole fraction of species A (related to the true value) as a function of the regularisation parameter and number of Bregman iterations, respectively. Fig. 5(c) shows the L-curve and Fig. 6(c) shows the l_2-norm as a function of the number of Bregman iteration. As expected, the reconstruction results depend strongly on the chosen regularisation parameter α and on the number of Bregman iterations m_{Bregman}, respectively. However, when the stop or selection criteria discussed above are applied, a parameter (α and m_{Bregman}, respectively) can be found for both approaches that yield a concentration map which represents an almost perfect reconstruction of the phantom concentration map.

The regularisation parameter that corresponds to the corner in the L-curve (cf. Fig. 5(c)) is about $\alpha = 0.004$. This regularisation parameter was used for the reconstruction of the concentration map shown in Fig. 4. Morozov's discrepancy principle (cf. Eq. (9)) is also applicable for the L-curve approach as the noise level intersects with the corner of the L-curve. For this value of the regularisation parameter ($\alpha = 0.004$), the relative error in the reconstructed mean mole fraction of species A is 0.05% and the standard deviation of the reconstructed mole fractions (related to

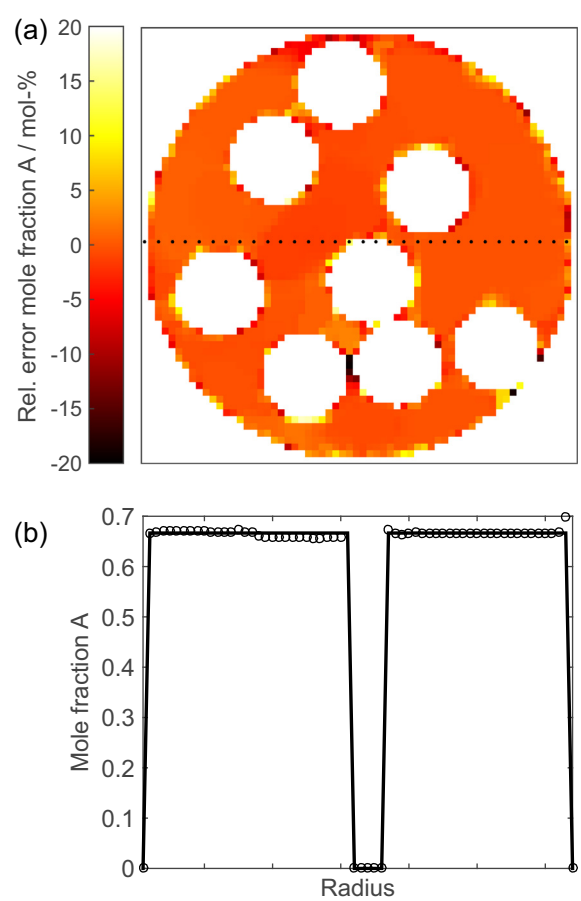

Fig. 4. (a) Relative error of the reconstructed concentration map compared to the set concentration map of species A in the phantom. (b) Radial concentration profile of species A through the sample at the location indicated with a dotted line in the part (a) of the figure. The reconstruction was carried out with $\alpha = 0.004$ (L-curve apporach). ∘ reconstructed mole fraction, – set mole fraction.

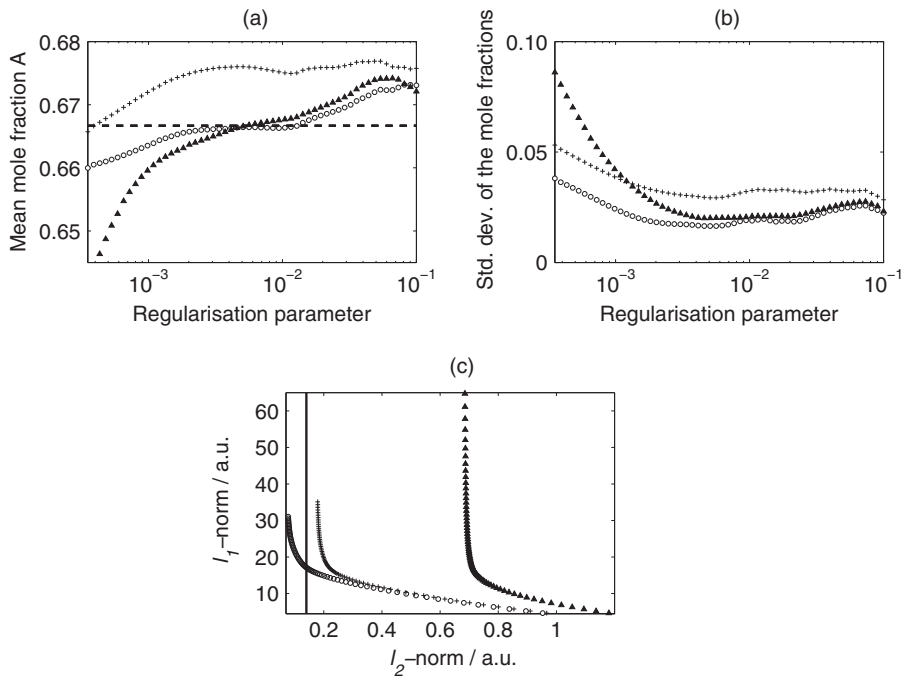

Fig. 5. (a) Mean mole fraction of species A in the reconstructed phantom concentration map, (b) the standard deviation of the reconstructed mole fractions of species A as a function of the regularisation parameter and (c) illustration of the "L-curve" approach whereby the regularisation term (l_1-norm) is shown as a function of the fidelity term (l_2-norm). The optimum regularisation parameter would correspond to the value required to obtain a result at the corner of the l-shaped curves shown. ∘ without systematic error, + systematic error in the chemical shift, ▲ systematic error in the sampling scheme, – – true mole fraction, — noise level.

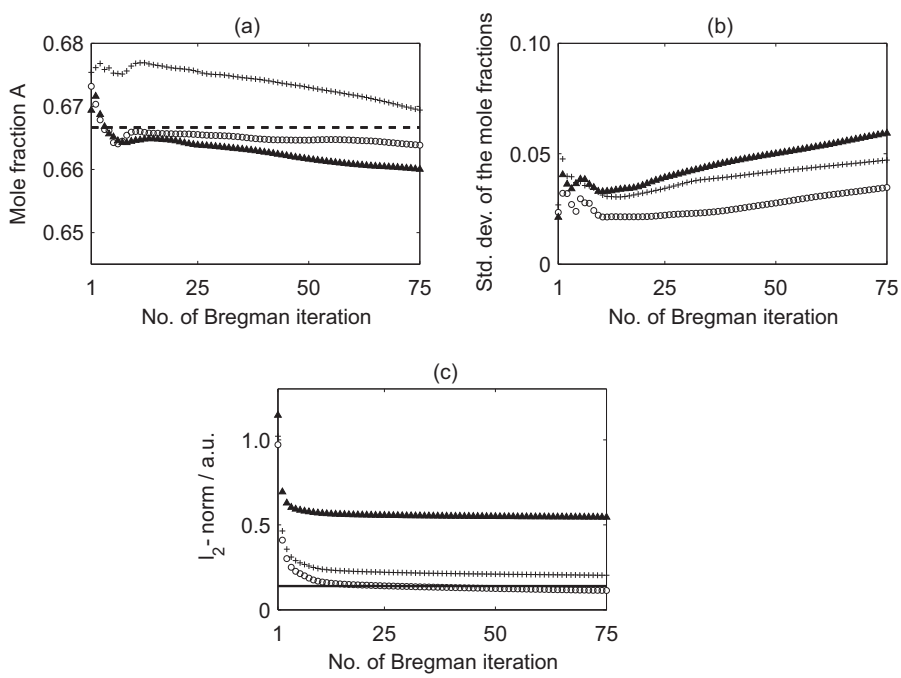

Fig. 6. (a) Mean mole fraction of species A in the reconstructed phantom concentration map, (b) and the standard deviation of the reconstructed mole fractions of species A, and (c) the fidelity term (l_2-norm) as a function of the number of Bregman iterations. ∘ without systematic error, + systematic error in the chemical shift, ▲ systematic error in the sampling scheme, – – true mole fraction, — noise level.

the set mole fraction) exhibits a minimum with a value of 4×10^{-4} mol/mol. The same is true for the approach using Bregman iterations. After 19 iterations, the l_2-norm intersects the noise level and the stop criterion given in Eq. (9) is fulfilled. Here, the relative error in the reconstructed mean mole fraction of species A and the standard deviation of the reconstructed mole fractions show again a minimum (cf. Fig. 6(a/b)). The relative error in the mean mole fraction is 0.3% and the standard deviation of the reconstructed mole fractions is 5×10^{-4} mol/mol showing that an almost perfect reconstruction of the concentration map is achieved.

With respect to the robustness of the two approaches, it is important to evaluate the sensitivity of the reconstruction results on the chosen regularisation parameter and number of Bregman

iterations, respectively. As can be seen in Figs. 5 and 6, the mean mole fraction and the standard deviation are almost constant for a range near the regularisation parameter or number of Bregman iterations that are chosen according to the selection criterion. Hence, the quality of the reconstruction results is not very sensitive to the choice of the regularisation parameter or the number of Bregman iterations, as long as this choice is in a range near the optimal values. This low sensitivity of the reconstruction results on the regularisation parameter near the optimal regularisation parameter is a very important feature of the L-curve approach, since the corner in the L-curve shown in Fig. 5 is not sharp and thus its location is not exactly defined.

Additionally, the optimal range of regularisation parameters or the number of Bregman iterations can be quite well estimated by evaluating the reconstructed images. If the regularisation parameter is chosen too high or the number of Bregman iterations are too few, the image is oversmoothed and the spatial resolution deteriorates significantly caused by the overweighted TV operator. On the other hand, if the regularisation parameter is chosen too small or the number of Bregman iterations are too many, the fidelity term is overweighted and the resulting image looks pixelated. Thus, as expected, when an image with good spatial resolution is obtained (neither oversmoothed nor pixelated) the quantitative information, i.e. the concentration map, is correctly recovered.

4.1.2. Sensitivity to systematic errors

To test the sensitivity of the reconstruction, the simulated data was reconstructed with systematic errors introduced to the model. First, the concentration map was reconstructed with a biased relative chemical shift of species A and B ($\delta_A^{biased} = \delta_A - 20\,\text{Hz}$, $\delta_B^{biased} = \delta_B + 30\,\text{Hz}$). The line widths of the NMR samples studied were typically about 100 Hz, therefore the combined shift of 50 Hz in the estimated frequency corresponds to a worst case estimate of the expected error in the chemical shift. Second, the concentration map was reconstructed assuming that there were errors in the measured trajectory map. In order to simulate error in the trajectory map, the reconstructions were performed using an effective sampling scheme given by $\mathbf{k}^{eff} = \mathbf{k} + \mathbf{v}$ with $\mathbf{v}$ as Gaussian distributed noise. The standard deviation of the noise was determined by repeated measurements of the sampling trajectories as shown in Fig. 2(b). Also the influence of the apparent T_2-relaxation time on the reconstruction results was examined. Since the time to acquire data along a sampling trajectory is only 1.4 ms (i.e. short compared with the T_2-relaxation time of the samples considered in this work), the term $\exp\left(-\frac{t+2\tau}{T_2}\right)$ can be neglected and the chosen value of the apparent T_2-relaxation time has no effect on the reconstruction results. The reconstruction results biased by systematic errors are included both for the L-curve approach and for the Bregman iteration approach in Figs. 5 and 6, respectively.

An error in the chemical shift has a large impact on the reconstructed mole fractions (cf. Figs. 5(a) and 6(a)) whilst the spatial resolution is nearly unaffected. To demonstrate the effect of a systematic error in the chemical shift on the spatial resolution, Fig. 7 shows the sum of the unscaled concentration maps of species A and B when the reconstruction is carried out without systematic errors (Fig. 7(a)) and with a systematic error in the chemical shift (Fig. 7(b)). In both figures the spatial resolution is good and the sum of the concentration maps of species A and species B is almost the same. The systematic error of the chemical shift results in a small amount of signal being incorrectly assigned outside the sample and changes the ratio of species A to species B. These changes cause the concentration of species A and B to be estimated incorrectly (cf. Fig. 5(a)), though the effect is not too severe (1%); the spatial resolution of the image is almost unaffected by a systematic error in the chemical shift.

A systematic error in the sampling scheme causes the reconstruction results to deteriorate compared to the reconstruction without systematic error. The systematic error in the k-space trajectory causes a large shift in the l_2-norm, and hence the L-curve (cf. Figs. 5(c) and 6(c)). The shift of the L-curve does not effect the shape of the L-curve (cf. Fig. 5) and the optimum regularisation parameter α is still located in the corner of the L-curve. The value of α that corresponds to the corner of the L-curve is about 0.007 (the optimal value of α in the corner of the L-curve obtained without systematic error is 0.004). Thus the change of the optimal value of α caused by an introduction of a systematic error is only minor (the overall variation of the value of α along the L-curve is from 1×10^{-1} to 5×10^{-5}). By contrast, Morozov's discrepancy principle, cf. Eq. (9), is not applicable now since the l_2-norm and the noise level do not intersect. However, the L-curve approach can be adapted for use with Bregman iterations. A plot of the l_2-norm versus the number of Bregman iterations shows a corner, as with the L-curve. Here, the optimal range for the number of Bregman iterations is located a little bit to the right of that corner (cf. Fig. 6(c)). The optimal range for the number of Bregman iterations can also be identified by evaluating the reconstructed images. To the left of the optimal range in (i.e. $m_{Bregman} < 10$), where the l_2-norm has a steep slope (cf. Fig. 6 (c)), the reconstructed images are oversmoothed and to the right of the optimal range (i.e. $m_{Bregman} > 30$), where the l_2-norm reaches a constant level, the reconstructed images are pixelated. Within the range ($10 < m_{Bregman} < 30$), little change is seen between images.

4.1.3. Conclusions from simulations

The simulations demonstrate the potential of this Compressed Sensing based technique to reconstruct concentration maps accurately from significantly less data than would be required to obtain a full chemical shift image. The correct weighting of the fidelity term and the regularisation term is important for a good reconstruction result both concerning the spatial resolution and the quantitative information (concentration). Both approaches used in the present work facilitate the identification of an optimal range for the weighting that yield a good reconstruction result for the concentration maps. These two approaches are still applicable when the model used for the reconstruction is biased by systematic errors. An error in the relative chemical shift has a large effect on the reconstructed concentration map but only a minor effect on the l_2-norm of the fidelity term. In contrast, small deviations of the sampling schemes cause a large shift of the l_2-norm of the fidelity term but only minor shift of the reconstructed concentration.

4.2. Reconstruction of measured data

Fig. 8 shows the reconstructed concentration map of dioxane for the experimental test sample A (a binary homogeneous mixture of cyclooctane and 1,4-dioxane) that was obtained using Bregman iterations in combination with the selection criterion discussed above. The results are similar when the L-curve approach is applied and they are not shown here. The relative error of the reconstructed mean mole fraction of dioxane is 1.3% and the spatial deviations are low. This result demonstrates that the composition of samples can be spatially resolved with a high accuracy with the presented method.

As described above, the total acquisition time to obtain a concentration map was approximately 8 min. The recycle delay of 15 s before each acquisition had the main contribution to the total acquisition time along a sampling trajectory. The recycle delay was chosen to allow sufficient relaxation ($5 \times T_1$) such that quantitative

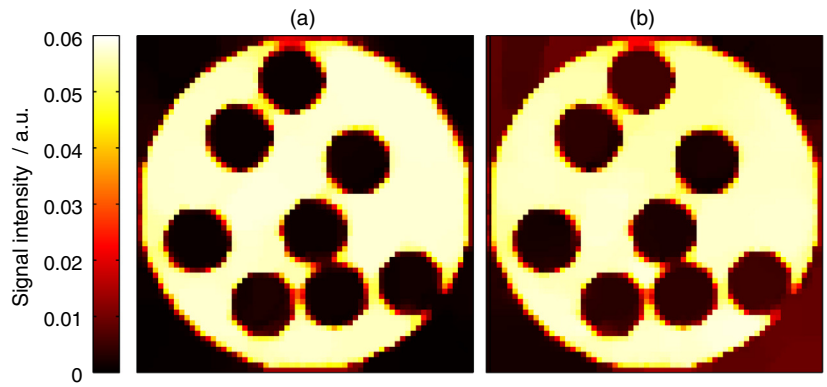

Fig. 7. Sum of the unscaled concentration maps of species A and B obtained by reconstruction of the simulated data: (a) without systematic error, (b) with systematic error in the chemical shift (b). The reconstruction was carried out with $\alpha = 0.004$ (L-curve approach). Image size: 64×64 pixels.

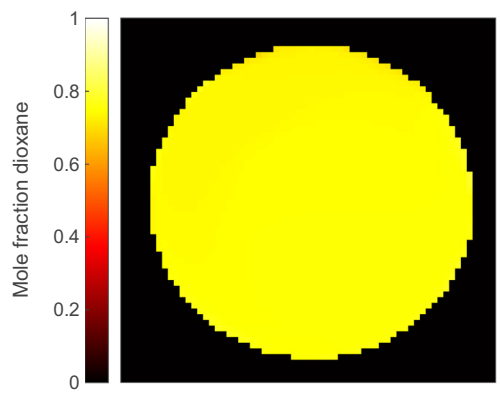

Fig. 8. Reconstructed concentration map of dioxane for test sample A using Bregman iteration. The resolution of the image was 344 μm × 344 μm. The concentration of dioxane in the sample was 0.761 mol/mol.

measurements were obtained. In some cases other approaches might be available to reduce the total acquisition time further. For example, a shorter phase cycle may be used or if the T_1-relaxation time of all species is approximately constant, a shorter recycle delay can be chosen.

To reconstruct the data acquired for the experimental test sample A, both the L-curve approach and Bregman iterations were applied. The reconstruction results are summarised in Fig. 9 for the L-curve approach and in Fig. 10 for the Bregman iteration. Figs. 9(a/b) and 10(a/b) show the reconstructed mean mole fraction of dioxane and the standard deviation of the mole fraction of dioxane (related to the expected mole fraction of dioxane) as a function of the regularisation parameter and number of Bregman iterations, respectively. Fig. 9(c) shows the L-curve and Fig. 10(c) the l_2-norm as a function of the number of Bregman iterations. The plots are very similar to the plots shown in Figs. 5 and 6 that were obtained for the simulated data. Also here, a range of regularisation parameters are given by the corner in the L-curve that yield good reconstruction results both concerning the spatial resolution and the accuracy of the concentration (cf. Fig. 9). For a regularisation parameter of about $\alpha = 0.012$, which corresponds to a result located in the corner of the L-curve, the relative error of the mean mole fraction of dioxane is 1.3% and the standard deviation shows a minimum.

The concentration maps that were reconstructed using Bregman iterations are also well resolved and a parameter set ($12 < m_{\text{Bregman}} < 30$) exists that yields a relative error in the mean fraction of dioxane of 1.4% and that has a minimum in the standard deviation of the mole fractions. The l_2-norm of the reconstruction result, however, is significantly larger than the estimated noise level (cf. Fig. 10) and hence the stopping criterion (Morozov's discrepancy principle) as defined in Eq. (9) is not applicable. As demonstrated in Section 4.1, the deviations in the reconstructed concentration and the large values for the l_2-norm of the fidelity term are attributable to systematic errors in the model. To verify this, the relative chemical shift of cyclooctane used in the model to reconstruct the concentration maps was decreased by 5%, which corresponds to 41 Hz. The results are included for both approaches in Figs. 9 and 10. The reconstructed concentration map when using a chemical shift that was 41 Hz lower than that measured relative chemical shift matches the known concentration more closely than the original chemical shift value. This change in chemical shift is attributed to errors in the shim of the sample making it difficult to identify the true chemical shift accurately. Interestingly, this small change of the chemical shift has no significant influence on the l_2-norm as shown in Figs. 9(c) and 10(c). However, when the sampling trajectory is slightly disturbed by noise, the whole curve of the l_2-norm is shifted significantly but there is almost no effect on the reconstructed concentration. In this case, the noise in the sampling trajectory was estimated from repeated measurements (see Fig. 2(b)). These results indicate that the quality of the reconstruction would likely be improved by more accurate measurement of the k-space trajectory.

Nevertheless, the Bregman iteration approach can still be used even though the stopping criterion is not applicable. As discussed above, the optimal range for the number of Bregman iterations can be identified both by evaluating the plot of the l_2-norm versus the number of Bregman iterations (the optimal range is here located a little bit to the right of the corner) or by evaluating the reconstructed images. The concentration map that has the best spatial resolution is also the concentration map that yields the best agreement with the expected concentration. As mentioned above, the reconstructed image is oversmoothed if the number of Bregman iterations is chosen too low and it becomes pixelated if it is too high.

Fig. 11 shows the reconstructed concentration map of dioxane for test sample B that was obtained using Bregman iterations in combination with the selection criterion discussed above. The results are similar when the L-curve approach is applied and they are not shown here. Table 2 lists a comparison of the reconstructed mean mole fraction with the expected mole fraction of dioxane. The results show that the concentration in the inner vial is well recovered (relative error less than 1%). In the outer vial, however, the error is larger (about 11%). Two reasons are presented for the larger error of the reconstructed concentration in the outer vial

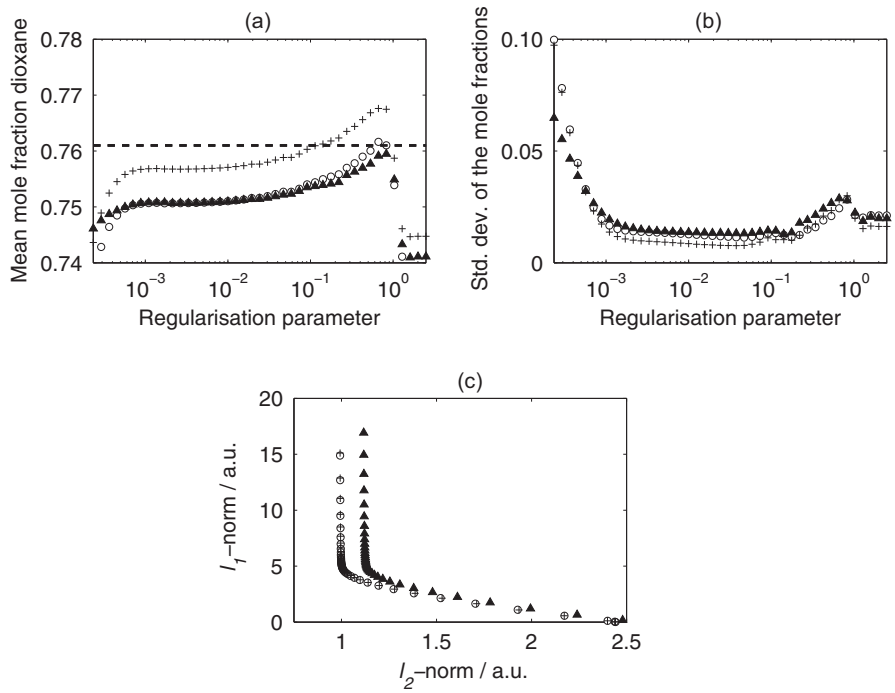

Fig. 9. (a) Mean mole fraction of dioxane in the reconstructed concentration map of test sample A and (b) the standard deviation of the reconstructed mole fractions of dioxane as a function of the regularisation parameter. (c) The regularisation term (l_1-norm) as a function of the fidelity term (l_2-norm) (c). ∘ reconstruction with measured relative chemical shift, + relative chemical shift of cyclooctane decreased by 5%, ▲ systematic error in the sampling scheme, – – expected mole fraction of dioxane in test sample A.

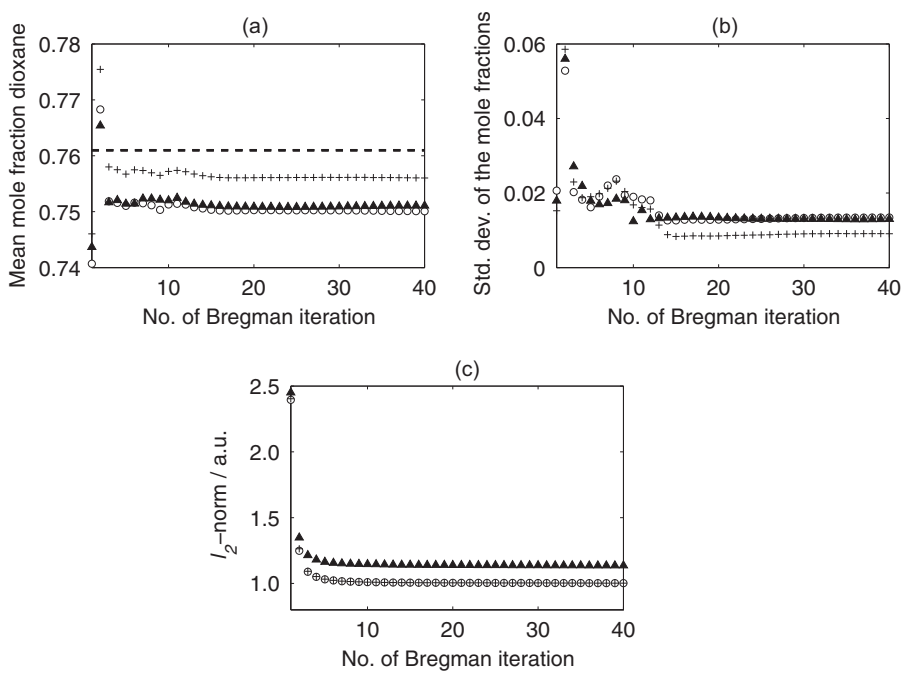

Fig. 10. (a) Mean mole fraction of dioxane in the reconstructed concentration map of test sample A, (b) the standard deviation of the reconstructed mole fractions of dioxane, and (c) the fidelity term (l_2-norm) as a function of the number of Bregman iterations. ∘ reconstruction with measured relative chemical shift, + relative chemical shift of cyclooctane decreased by 5%, ▲ systematic error in the sampling scheme, – – expected mole fraction of dioxane in test sample A.

compared to the inner vial. First, sharp corners are present in the outer vial, and as shown in Fig. 4 and by Benning et al. [21], sharp corners and confined spaces are challenging for the reconstruction algorithm and thus they are often not correctly recovered even though no systematic error is present in the model. Second, spatial inhomogeneities are present in the $\mathbf{B}_0$-field. These inhomogeneities are interpreted by the model as a relative chemical shift and thus they have a similar effect on the reconstruction result as an error in the chemical shift, namely the accuracy of the quantitative information, i.e. the concentration, deteriorates whilst the spatial resolution remains good. Thus, these inhomogeneities cause differences in the reconstructed concentration between the upper part of the image and the lower part of the image where there should be none.

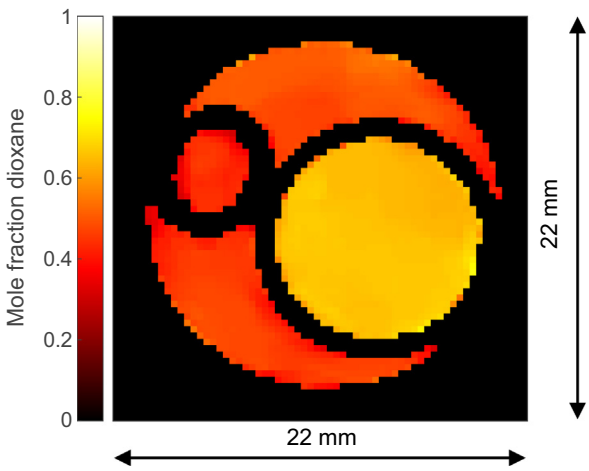

Fig. 11. Reconstructed concentration map of dioxane for test sample B using Bregman iteration. The resolution of the image was 344 μm × 344 μm.

Table 2
Comparison of the mean mole fraction in the reconstructed concentration map with the expected mole fraction of dioxane in test sample B. The error of the expected value is estimated based on the accuracy of the used scale.

Location	Mean mole fraction dioxane	
	Expected	Measured
Inner vial	0.666 ± 0.001	0.66
Outer vial	0.415 ± 0.002	0.46

It may be possible to improve the reconstruction shown in Fig. 11 by incorporating the spatial inhomogeneities of the magnetic field in the model. In this case, the concentration map and the map of the field inhomogeneities have to be reconstructed from the measured signal data by the solution algorithm. This additional reconstruction of the map of the field inhomogeneities is numerically expensive as the order of the chemical shift operator and the undersampled Fourier transform in Eq. (4) have to be changed and thus the Fourier transform has to be performed not only M times (M – number of species, here 2) as it is done in Eq. (4) but N_t times (N_t – number of data points in each trajectory, here 551). Furthermore, when the map of the field inhomogeneities is included in the model, the equation becomes non-linear which makes the solution algorithm more challenging compared to the algorithm used in the present work. As the scope of this paper is to show the principle of quantitative concentration mapping with MRI using compressed sensing, we refrain from a detailed discussion on solution algorithms for non-linear equations which include the reconstruction of the map of the field inhomogeneities. More details on that topic are given for example by Doneva et al. [16].

5. Conclusion

A fast MR imaging method is presented that enables the composition of mixtures of chemical species to be resolved. The method enables acquisition of quantitative maps of the chemical composition in as little as 8 min, when a full chemical shift image at the same resolution would require 17 h. The method is fast because it is based on a compressed sensing algorithm that uses prior-knowledge to obtain the concentration image from under-sampled data. Further reductions in acquisition time may be possible through optimisation of the pulse sequence. No calibration is necessary prior to the analysis in order to get quantitative information with an accuracy of ±2 mol-%. Therefore, the method is valuable for many applications, e.g. in chemical engineering where unstable intermediates may form during the process and hence prohibit a calibration or in medical sciences and biology.

The prior-knowledge that is necessary for the reconstruction of information from under-sampled data is incorporated in the algorithm via a regularisation term. In this work, a spatial finite differences ("total variation") based regularisation is used as images are piece-wise constant. For other systems, regularisers such as Total Generalised Variation or wavelets may be preferable. Regardless of the form of the regularisation function, the weighting has to be carefully chosen in order to get both a good spatial resolution and a high quantitative accuracy in the concentration map. In the present work, the L-curve approach and Bregman iterations, and different selection and stop criteria were used to find an optimal weight for the regularisation term. The two approaches and the selection and stop critera were tested by reconstructing both simulated data from a phantom concentration map and measured data from different samples of binary mixtures.

The mathematically well-defined stopping criterion that is based on Morozov's discrepancy principle is not applicable to the experimental data owing to systematic errors in the model, mainly deviations in the measured sampling trajectories. Nevertheless, the selection criterion that is based on a graphical evaluation of the reconstruction results enables well resolved concentration maps to be obtained using both the L-curve and the Bregman iteration approaches. Furthermore, the fact that the optimal parameters for the regularisation are based on a selection criterion that requires a graphical evaluation of the reconstruction results is not disadvantageous for the quality of the results because the reconstruction result is insensitive to the choice of these parameters in an interval near the optimal parameters. Thus, both the L-curve and the Bregman iteration are generic and robust approaches to achieve quantitative results.

To conclude, the presented method is a powerful tool for the fast acquisition of concentration maps. These concentration maps can provide valuable information for the investigation of many phenomena in chemical engineering applications.

Acknowledgments

Financial support was provided by Microsoft Research Cambridge, Cambridge, and the EPSRC (EP/K039318/1 and EP/K008218/1). Erik von Harbou was the recipient of a scholarship from the German Academic Exchange Service (DAAD).

Appendix A. Description of the linear operators

The transformation of the M concentration maps $\mathbf{x}_k$ into the signal $\mathbf{S}$ shown in Eq. (1) can be abbreviated by linear operators, see Eq. (4). First, the concentration maps are subjected to a Fourier transformation,

$$\mathbf{S}_k^* = \mathcal{F}_\mathrm{u}\{\mathbf{x}_k\} \quad \text{with}: \ \mathbf{x}_k \in \mathbb{R}^{N\times N}, \quad \mathbf{S}_k^* \in \mathbb{C}^{[N_\mathrm{samples}\times 1]} \text{ and } k = 1,\ldots,M$$

$\mathcal{F}_\mathrm{u}$ is the discrete non-uniform Fourier transform operator that is described in detail by Fessler and Sutton [23]. $N \times N$ is the size (number of pixels) of the concentration map $\mathbf{x}_k$. N_samples is the number of samples.

To get the signal $\mathbf{S}$, the chemical shift operator $\mathbf{CHS}$ is applied to the Fourier transformed concentration map $\mathbf{S}_k^*$:

$$\mathbf{S} = \mathbf{CHS} \cdot \begin{pmatrix} \mathbf{S}_1^* \\ \vdots \\ \mathbf{S}_M^* \end{pmatrix} \quad \text{with}: \ \mathbf{S} \in \mathbb{C}^{N_\mathrm{samples}\times 1}$$

The chemical shift operator $\mathbf{CHS}$ is a matrix:

$$\mathbf{CHS} = \left[\mathrm{diag}\left(CHS_1(t_1), \ldots, CHS_1(t_{N_{\text{samples}}})\right), \ldots, \mathrm{diag}(CHS_M(t_1), \ldots, CHS_M(t_{N_{\text{samples}}})\right)\right] \text{ with : } \mathbf{CHS} \in \mathbb{C}^{N_{\text{samples}} \times N_{\text{samples}} \cdot M}$$

Here, the operator $CHS_k(t_j)$ is defined as

$$CHS_k(t_l) = \sum_{j=1}^{L_k} w_{k,j} \exp\left(2\pi i\, \delta_{k,j} t_l\right) \exp\left(-\frac{t_l + 2\tau}{T_2^*}\right).$$

Appendix B. Description of the total variation regularisation

Total variation regularisation is the 1-norm penalty on a discrete finite difference approximation of the two-dimensional gradient ∇ [21]. The two-dimensional gradient is defined as

$$\nabla_1 x_k(i,j) = \begin{cases} x_k(i+1,j) - x_k(i,j) & \text{if } i < n_1 \\ 0 & \text{if } i = n_1 \end{cases}$$

$$\nabla_2 x_k(i,j) = \begin{cases} x_k(i,j+1) - x_k(i,j) & \text{if } j < n_2 \\ 0 & \text{if } j = n_2 \end{cases}$$

for $i = 1, \ldots, n_1$ and $j = 1, \ldots, n_2$.

Thus the discrete total variation functional is given by

$$J(\Psi x_k) = \|\Psi x_k\|_{2,1} = \|\nabla x_k\|_{2,1} = \sum_{i,j} \sqrt{|\nabla_1 x_k(i,j)|^2 + |\nabla_2 x_k(i,j)|^2}.$$

References

[1] E. Yuen, A. Sederman, L. Gladden, Appl. Catal. A: Gen. 232 (2002) 29–38.
[2] I. Koptyug, A. Lysova, A. Kulikov, V. Kirillov, V. Parmon, R. Sagdeev, Appl. Catal. A: Gen. 267 (2004) 143–148.
[3] B.S. Akpa, M.D. Mantle, A.J. Sederman, L.F. Gladden, Chem. Commun. (2005) 2741–2743.
[4] A.J. Sederman, M.D. Mantle, C.P. Dunckley, Z. Huang, L.F. Gladden, Catal. Lett. 103 (2005) 1–8.
[5] L.F. Gladden, F.J. Abegão, C.P. Dunckley, D.J. Holland, M.H. Sankey, A.J. Sederman, Catal. Today 155 (2010) 157–163.
[6] D. Donoho, IEEE Trans. Inf. Theory 52 (2006) 1289–1306.
[7] E. Candes, J. Romberg, T. Tao, IEEE Trans. Inf. Theory 52 (2006) 489–509.
[8] M. Lustig, D. Donoho, J.M. Pauly, Magn. Reson. Med. 58 (2007) 1182–1195.
[9] R. Otazo, D. Kim, L. Axel, D.K. Sodickson, Magn. Reson. Med. 64 (2010) 767–776.
[10] D. Holland, D. Malioutov, A. Blake, A. Sederman, L. Gladden, J. Magn. Reson. 203 (2010) 236–246.
[11] A.B. Tayler, D.J. Holland, A.J. Sederman, L.F. Gladden, Phys. Rev. Lett. 108 (2012) 264505.
[12] D.J. Holland, M.J. Bostock, L.F. Gladden, D. Nietlispach, Angew. Chem. Int. Ed. 50 (2011) 6548–6551.
[13] K. Kazimierczuk, V.Y. Orekhov, Angew. Chem. Int. Ed. 50 (2011) 5556–5559.
[14] S. Hu, M. Lustig, A.P. Chen, J. Crane, A. Kerr, D.A. Kelley, R. Hurd, J. Kurhanewicz, S.J. Nelson, J.M. Pauly, D.B. Vigneron, J. Magn. Reson. 192 (2008) 258–264.
[15] T. Kampf, A. Fischer, T. Basse-Lüsebrink, G. Ladewig, F. Breuer, G. Stoll, P. Jakob, W. Bauer, J. Magn. Reson. 207 (2010) 262–273.
[16] M. Doneva, P. Börnert, H. Eggers, A. Mertins, J. Pauly, M. Lustig, Magn. Reson. Med. 64 (2010) 1749–1759.
[17] S.D. Sharma, H.H. Hu, K.S. Nayak, Magn. Reson. Med. 67 (2012) 650–659.
[18] S.D. Sharma, H.H. Hu, K.S. Nayak, Magn. Reson. Med. 69 (2013) 456–466.
[19] P. Hansen, SIAM Rev. 34 (1992) 561–580.
[20] S. Osher, M. Burger, D. Goldfarb, J. Xu, W. Yin, Multiscale Model. Simul. 4 (2005) 460–489.
[21] M. Benning, L. Gladden, D. Holland, C.-B. Schönlieb, T. Valkonen, J. Magn. Reson. 238 (2014) 26–43.
[22] D. Gabay, in: M. Fortin, R. Glowinski (Eds.), Augmented Lagrangian Methods: Applications to the Numerical Solution of Boundary-Value Problems, vol. 15, Elsevier, 1983, pp. 299–331.
[23] J. Fessler, B. Sutton, IEEE Trans. Signal Process. 51 (2003) 560–574.
[24] J.A. Fessler, Nonuniform FFT (NUFFT) Matlab Toolbox <http://web.eecs.umich.edu/fessler/irt/fessler.tgz>, October 2012.
[25] T.Z. Teisseyre, J.L. Paulsen, V.S. Bajaj, N.W. Halpern-Manners, A. Pines, J. Magn. Reson. 216 (2012) 13–20.
[26] E. van den Berg, M. Friedlander, SIAM J. Sci. Comput. 31 (2008) 890–912.
[27] W. Yin, S. Osher, D. Goldfarb, J. Darbon, SIAM J. Imag. Sci. 1 (2008) 143–168.
[28] V.A. Morozov , Soviet Math. Dokldy 7 (1966) 414–417.
[29] M. Lustig, S.-J. Kim, J. Pauly, IEEE Trans. Med. Imag. 27 (2008) 866–873.
[30] M. Lustig, Time Optimal Gradient Design <http://www.eecs.berkeley.edu/mlustig/Software.html>, October 2012.
[31] J.H. Duyn, Y. Yang, J.A. Frank, J.W. van der Veen, J. Magn. Reson. 132 (1998) 150–153.
[32] A.B. Tayler, D.J. Holland, A.J. Sederman, L.F. Gladden, J. Magn. Reson. 211 (2011) 1–10.
[33] A. Chambolle, S. Levine, B. Lucier, SIAM J. Imag. Sci. 4 (2011) 277–299.