Proceedings of the 29th Annual International
Conference of the IEEE EMBS
Cité Internationale, Lyon, France
August 23-26, 2007.

Automated Analysis of ^{1}H Magnetic Resonance Metabolic Imaging Data as an Aid to Clinical Decision-Making in the Evaluation of Intracranial Lesions

Dikoma C. Shungu, Shuyan Du, Xiangling Mao, Linda A. Heier, Susan C. Pannullo, Paul Sajda

Abstract—**Proton magnetic resonance spectroscopic imaging (^{1}H MRSI) is a noninvasive metabolic imaging technique that has emerged as a potentially powerful tool for complementing structural magnetic resonance imaging (MRI) in the clinical evaluation of neurological disorders and diagnostic decision-making. However, the relative complexity of methods that are currently available for analyzing the derived multi-dimensional metabolic imaging data has slowed incorporation of the technique into routine clinical practice. This paper discusses this impediment to widespread clinical use of ^{1}H MRSI and then describes an automated data analysis approach that promises to facilitate use of the technique in the evaluation of intracranial lesions, with the potential to enhance the specificity of MRI and improve clinical decision-making.**

I. Introduction

Over two decades of technological advances have seen structural magnetic resonance imaging (MRI) with contrast-enhancement develop into the preferred diagnostic radiologic modality for noninvasive assessment of intracranial lesions and other diseased human soft tissues, producing exquisitely detailed anatomic images that have dramatically improved clinical evaluations and decision-making. However, contrast-enhanced MRI has a number of significant limitations: it often lacks the specificity to differentiate pathologic lesions (e.g., brain tumors) from normal post-operative or post-therapy changes, such as edema and radiation-induced necrosis, and it can miss pathologic changes that do not exhibit contrast-enhancement, such as some low-grade brain neoplasms [1]. As a result, there has been a great deal of interest in evaluating proton magnetic resonance spectroscopy (^{1}H MRS) – a closely related technique that permits noninvasive assessment of brain chemistry and can be performed in conjunction with MRI on virtually every clinical MR scanner, without the need to change hardware or move the patient – as a complement to contrast-enhanced MRI in the evaluation of intracranial lesions [2]-[5]. Numerous studies have shown ^{1}H MRS to be a robust technique that can measure the brain concentrations of several potential biomarkers of human disease with relative ease, either from a single volume element (voxel) [3],[5], or from multiple voxels in a single or multiple brain slices simultaneously [2],[4]. Metabolites of potential clinical interest that can be accessed by either of these ^{1}H MRS approaches include [3],[5] the putative neuronal density, viability and function marker, N-acetyl-aspartate (NAA), which is decreased in most brain lesions; the cell bioenergetics marker, total creatine (tCr), whose resonance includes a contribution from phosphocreatine (PCr) and remains unchanged under most *in vivo* conditions, except for those where there is extensive tissue damage or disrupted enzymatic homeostasis; the cell membrane biosynthesis and metabolism marker, total choline (tCho), which increases in processes that increase availability of free choline, such as cell membrane breakdown, and rapid proliferation; and the anaerobic energy metabolism marker, lactate, whose increases are often indicative of increased anaerobic glycolysis activity. Despite this clear potential, and after nearly 15 years of concerted effort in laboratories and medical centers worldwide to evaluate the complementary role of ^{1}H MRS in the clinical evaluation of brain disorders and decision-making [3],[5], the technique is still considered an 'investigational' tool that is not recommended in most cases for reimbursement by insurance companies. As a result, ^{1}H MRS has been slow to replicate the success of MRI in routine clinical practice.

In this paper, we first discuss how the limited spatial coverage and resolution of single-voxel ^{1}H MRS, and the relative complexity and high dimensionality of multi-voxel MRS (i.e., MRS imaging or ^{1}H MRSI) might be pitfalls, among others, that have contributed to the near permanent status of MRS as an 'investigational' technique. Then, we demonstrate a promising automated and robust ^{1}H MRSI data analysis approach, constrained non-negative matrix

Manuscript received April 23, 2007. This work was supported by an NSF CAREER Award (BES-0133804, PS), the DoD Multidisciplinary University Research Initiative (MURI) program administered by the Office of Naval Research (N00014-01-0625, PS), and Cornell University Faculty Development Funds (DCS).

D.C. Shungu, Department of Radiology, Weill Medical College of Cornell University, New York, NY 10021 USA (phone: 212-746-2481; fax: 212-746-6681; e-mail: dcs7001@med.cornell.edu).

S. Du, Department of Biomedical Engineering, Columbia University, New York, NY 10025; current address: Sanofi-Aventis, Bridgewater, NJ 08807 (email: shuyan.du@sanofi-aventis.com)

X. Mao, Department of Radiology, Weill Medical College of Cornell Univ., New York, NY 10021 USA (email: xim2004@med.cornell.edu).

L.A. Heier, Department of Radiology, Weill Medical College of Cornell Univ., New York, NY 10021 USA (email: laheier@med.cornell.edu)

S.C. Pannullo, Department of Neurological Surgery, Weill Medical College of Cornell University, New York, NY 10021 USA (email: scp2002@med.cornell.edu)

P. Sajda, Department of Biomedical Engineering, Columbia University, New York, NY 10025 (email: ps629@columbia.edu)

factorization (cNMF) [6]-[8], that can enhance decision-making and promote routine clinical use of the technique by allowing rapid recovery, simplification, and presentation of the complex spectral and spatial information content of ^{1}H MRSI data sets in a manner that would facilitate interpretation by clinicians [7].

II. Single- vs Multi-voxel ^{1}H MRS

A patient who presents with a hyper-intense or contrast-enhancing lesion on structural brain MRI might be eligible to undergo a ^{1}H MRS study. If the examination is for a post-operative or post-therapy brain tumor, the referring physician might request that the lesion be examined for signs of recurrent tumor or radiation-induced necrosis, which might be difficult to differentiate based on MRI alone. This information would clearly inform the clinician's decision regarding the case. With respect to the ^{1}H MRS study to perform, the most forward option would be to conduct a single-voxel study, which has a number of notable advantages: (i) the method is available and is automated - from acquisition to processing and presentation – on nearly every clinical MR system; and (ii) it is fast, requiring only a few minutes of scanner time. However, as Fig. 1 illustrates, the limited spatial resolution of single-voxel ^{1}H MRS is a pitfall that can adversely affect the clinical utility of the technique in the evaluation of post-operative and post-therapy lesions, as in such cases, regions of interest tend to be confined to the edge or rim of the lesions and would be completely missed. In Fig. 1, single-voxel MRS (Fig. 1A) and MRSI (Fig. 1C) data are shown for a patient with a previously resected right frontal glioma that appeared hyper-intense on MRI (Fig. 1B). A single-voxel spectrum recorded from the center of the lesion showed complete absence of metabolite resonances (Fig. 1A), suggesting the presence of a fluid-filled, tumor-free, resection cavity. In fact, this example is an actual clinical case in which the MRI technologist, erroneously attributing the absence of metabolites to instrument malfunction, repeated the single-voxel scans five times with slightly different voxel sizes, but obtained identical results each time. The reason for the failure of single-voxel scans to detect any metabolite signals from the center of the lesion became apparent on close examination of ^{1}H MRSI data subsequently acquired with a substantially higher spatial resolution: the residual tumor was confined to a thin rim along the medial edge of the lesion (Fig. 1C), which was too small to contribute appreciable signal intensity to the overall spectrum recorded from the large region of interest sampled by single-voxel MRS. Clearly, single-voxel ^{1}H MRS is of limited clinical value for such studies, and use of ^{1}H MRSI would be indicated. However, the latter approach also has limitations, which are briefly discussed in the next section.

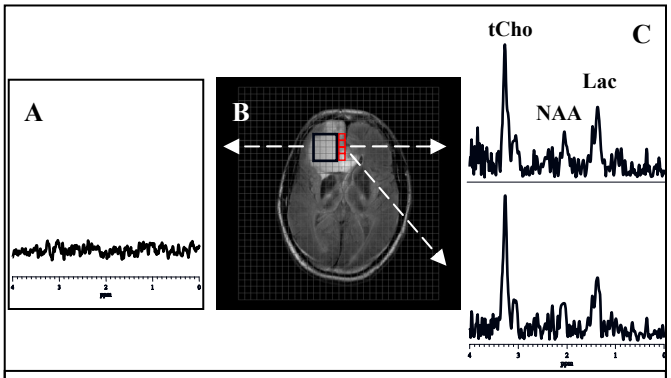

Fig. 1: [A] Single-voxel ^{1}H MRS spectrum recorded from the single large box in the center of the hyper-intense lesion shown on the anatomic MR image in [B]. Note complete absence of metabolite resonances in [A]. [C] Spectra from two of the voxels depicted in red recorded from the medial edge of the tumor using ^{1}H MRSI clearly show the presence of increased total choline (tCho), decreased N-acetyl-aspartate (NAA), and elevated lactate (Lac), consistent with the presence of a residual tumor.

III. ^{1}H MRSI Data Acquisition and Analysis Methods

Although, as demonstrated in the preceding example of a resected brain tumor case, ^{1}H MRSI might offer greater reliability and diagnostic power than single-voxel MRS in the evaluation of certain types of intracranial lesions, implementation of the technique in routine clinical practice still requires a great deal of expertise due to the relative complexity of the associated data acquisition and analysis methods. It can, in fact, be argued that the clinical success of MRI is owed in large measure to the high degree of automation that has been achieved in nearly every aspect of the technique. Currently, MRI data can be routinely and efficiently acquired, processed, and presented for interpretation even before the patient has been removed from the scanner table. This suggests that in order for ^{1}H MRSI to become widely accepted by radiologists and other clinicians as a viable diagnostic tool, it might be necessary to adopt approaches similar to those used for clinical MRI.

Recent developments in data acquisition software and scanner hardware have led to significant advances in ^{1}H MRSI data acquisition speed and automation. However, these advances have resulted in dramatic increases in the sizes of the recorded ^{1}H MRSI data that have become challenging for existing data analysis approaches. There is thus a need for alternative data analysis methods that can quickly recover the important spectral features within each large ^{1}H MRSI data set - which may consist of thousands of spectra – and, then present these to clinicians in a manner that allows rapid and reliable assessment of the underlying diagnostic information.

Since each ^{1}H MRSI data set consists of a mixture of spatial and spectral information, a principled approach for analysis of such data would simultaneously consider the spatial and spectral patterns of the entire data set, taking advantage of the intrinsic relationships among these patterns to recover automatically the underlying diagnostic information. In the remainder of this paper, we describe and demonstrate our recent development and implementation of constrained non-negative matrix factorization (cNMF) [6]-

[8], an automated ^{1}H MRSI data analysis method, based on the non-negative matrix factorization (NMF) algorithm of Lee and Seung [9], that incorporates these desirable features.

IV. Automated ^{1}H MRSI Data Analysis by cNMF

A. The Basic cNMF Concepts

We had previously described cNMF [6]-[8] as an automated ^{1}H MRSI data analysis method that decomposes the recorded raw data into two non-negative matrices representing, on one hand, the underlying tissue-specific spectral patterns and, on the other hand, the spatial distribution of these patterns to rapidly recover meaningful diagnostic and biochemical information contained in each data set. The details of NMF and cNMF have been described previously [6]-[9]. Both approaches essentially implement a conic encoder [10], where the recovered sources envelop the observed data with minimal error. In cNMF, there is a further constraint, namely, that even if some spectral resonances may have negative values, the underlying sources and concentrations are forced to lie only in the positive hyper-quadrant. Briefly and formally, we express the observed MRSI data, $\mathbf{X}$, as

$$\mathbf{X} = \mathbf{AS} + \mathbf{N} \qquad [1]$$

where the columns in $\mathbf{A}$ represent the concentration of metabolites in the constituent tissue within each voxel and the rows in $\mathbf{S}$ represent the corresponding spectral resonances. $\mathbf{N}$ represents additive noise. The concentration matrix $\mathbf{A}$ has M columns (one for each spectral pattern) and N rows (one for each voxel). $\mathbf{X}$ and $\mathbf{S}$ have L columns (one for each resonance). Since one can interpret $\mathbf{A}$ as tissue concentration, the matrix can be constrained to be non-negative. In addition, since the constituent spectra $\mathbf{S}$ represent amplitudes of resonances, in theory the smallest resonance amplitude is zero, corresponding to the absence of resonance at a given frequency (where we ignore cases of negative peaks such as in J-modulation). The factorization of Eq. [1] is, therefore, constrained by,

$$\mathbf{X} \geq \mathbf{0}, \mathbf{S} \geq \mathbf{0} \qquad [2]$$

This non-negative constraint is the key to the ability of cNMF to blindly recover biochemically meaningful as well as tissue-specific diagnostic information. Tables 1 and 2, respectively, provide summaries of the basic cNMF algorithm and its extension with iterative data selection for artifact removal as described in detail elsewhere [6]-[8].

B. Demonstration of cNMF in the Automated Analysis of a Brain tumor Case

To demonstrate the ability of cNMF to recover diagnostically-specific biochemical information, we present the results of applying the method to automatically analyze ^{1}H MRSI data in a patient with a World Health Organization (WHO) grade III choroid plexus carcinoma (Figs. 2 and 3). However, we should first note that, whereas in conventional ^{1}H MRS metabolite images the amplitudes or concentrations of *individual resonances* within a spectrum can vary from voxel to voxel across the brain, the variations in pixel intensities in the spatial distribution of cNMF-recovered images correspond only to the variations in the spatial concentration of a specific "spectrum" recovered by cNMF. Consequently, we have used the term "spectral pattern" rather than "spectrum" to the designate cNMF-recovered spectral information.

Table 1: Procedure for cNMF without data selection

1	Initialize: Choose dimensions of $\mathbf{A}$ and $\mathbf{S}$ (i.e., M) and initialize with non-negative values (e.g., random $\mathbf{A}$ and constrained least-squares for $\mathbf{S}$).
2	Update $\mathbf{A}$.
3	Force negative values of $\mathbf{A}$ to be approximately zero.
4	Update $\mathbf{S}$.
5	Force negative values of $\mathbf{S}$ to be approximately zero.
6	Iterate (back to 2) until convergence.

Table 2: Procedure for cNMF with data selection

1	Set $\hat{\mathbf{X}} = \mathbf{X}$ (use all voxels).
2	Apply cNMF to the voxels $\hat{\mathbf{X}}$.
3	Analyze peak resonance bands of the spectra and prior information of the spatial distribution to select a target source m.
4	Construct a spatial mask $\mathbf{t}$: thresholding the concentration matrix to select voxels which have a significant concentration of the target spectrum for $A_{i,m}$, $i = 1, \ldots, N$, $t_i = \begin{cases} 1 & if \quad A_{i,m} > threshold \\ 0 & if \quad A_{i,m} \leq threshold \end{cases}$, and apply morphology analysis to improve the spatial mask $\mathbf{t}$ by, removing small objects, smoothing edges, filling "holes" and etc.
5	voxels pass through the mask: $[rzr, rzc] = find(\mathbf{X} * repmat(\mathbf{t}, 1, L)); \quad \hat{\mathbf{X}} = \mathbf{X}(rzr(1:\sum \mathbf{t}),:)$
6	Continue to 2 if additional specificity is required, else stop.

Fig. 2(a) shows the spectral patterns recovered without data selection (Table 1), and the corresponding spatial distribution images. Three spectral patterns for residual pericranial lipids and one for brain have been recovered; however, no markers for tumor are recovered at this stage. From the spatial distribution of the recovered spectral patterns, it can be seen that the lipids are concentrated near the skull, while the spectral pattern for the brain is internal

to the lipids. Due to the large magnitude and breadth of the lipid peaks, the subtle differences between tumor and normal brain spectral patterns are masked. Fig. 2(b) and Fig. 3(b) show scatter plots of the recovered spectral patterns for

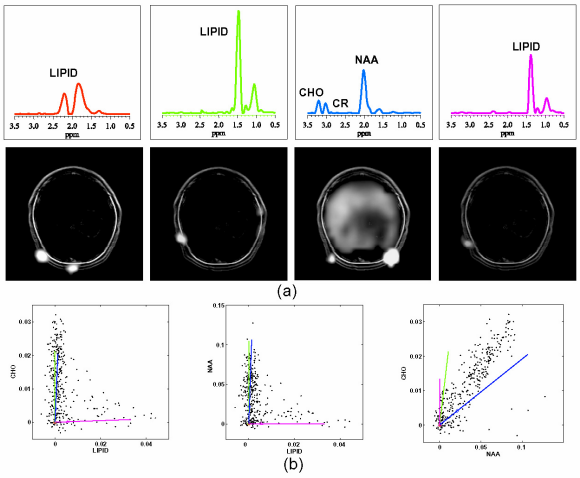

Fig. 2: (a) Basic cNMF results using all voxels for ^{1}H MRSI data of human brain with a tumor. (b) Scatter plots for selected resonances: Lipids, NAA, and tCho.

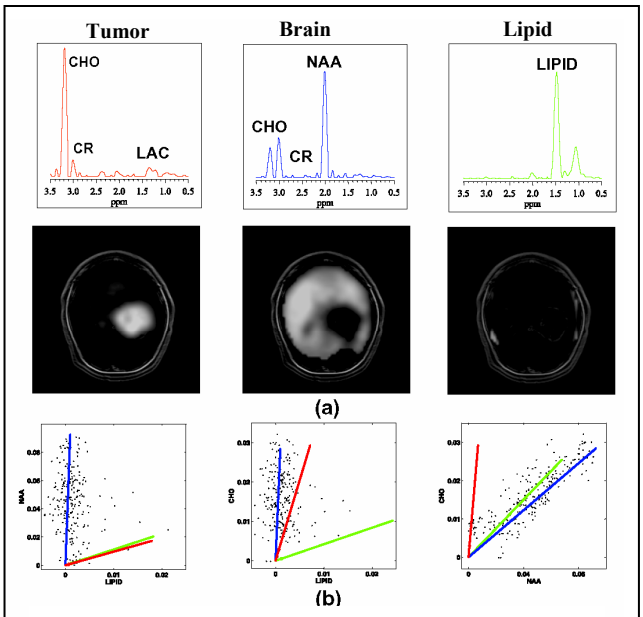

Fig. 3: (a) cNMF with data selection uncovers three spectral patterns for tumor, normal brain, and pericranial lipids. (b) Scatter plots of recovered spectra enveloping the selected data (see text for details.)

lipids, tCho, and NAA. For each case, the recovered spectral patterns form a hyper-cone in the positive hyper-quadrant that envelopes the observed data — consistent with the cNMF objective function. In Fig. 2(b), none of the cone edges corresponds to a tumor spectral pattern. Rather, prior to data selection, the edge of the cone is defined by the lipid pattern (Fig. 2(b) right: green and magenta cone edges).

Next, data selection (Table 2) is applied to select "brain-only" voxels, which are used in a second level of cNMF analysis to reduce the strong lipid component. Note that after lipid removal, there is a clear change in the hyper-cone (Fig. 3(b) right: red cone edge), which is now defined by a cone edge corresponding to the tumor spectral pattern. As a result, cNMF has achieved, with higher sensitivity and specificity, recovery of the spectral patterns (Fig. 3(a), three top panels, respectively) and associated spatial concentration distributions (Fig. 3(a), three bottom panels, respectively) for the tumor (very high tCho and little NAA), normal brain (high NAA and normal tCho), and residual pericranial lipids. Particularly noteworthy, however, is the manner in which cNMF has automatically "segmented" [7] the tumor mass from all adjacent normal brain tissue to permit unobstructed visualization of the location and spatial extent of the tumor in the corresponding spatial distribution image. Clearly, presenting such single spectral patterns and the associated spatial distribution images, rather than thousands of spectra, would allow clinicians to just *visually* "read" and interpret ^{1}H MRSI data in a manner analogous to the routine reading of standard clinical MRI data.

V. Conclusion

Starting with nearly 1000 individual spectra in the original ^{1}H MRSI data matrix, cNMF has automatically reduced this complex data to just three essential spectral patterns (Fig. 3a) that clearly provide the spatial distribution and concentrations of the tumor markers, which can facilitate rapid visual interpretation and aid in diagnosis and clinical decision-making.

References

[1] D.F. Nelson, J.V McDonald, L.W. Lapham, R. Qazi, P. Rubin (1993). Central Nervous System Tumors. *In*: Rubin P, ed. Clinical Oncology, 7th Edition. Philadelphia: W.B.Saunders Company, 617-644.

[2] S.J. Nelson, D.B. Vigneron, W.P. Dillon (1999). "Assessment of Brain Tumors Using Volume MRI and MRSI." *NMR Biomed* **12**:123-38.

[3] A. Lin,S. Bluml, A.N. Mamelak (1999). "Efficacy of Proton Magnetic Resonance Spectroscopy in Clinical Decision Making for Patients with Suspected Malignant Brain Tumors." *J Neurooncol* **45**: 69-81.

[4] C. Balmaceda, D. Critchell, X. Mao, K. Cheung, S. Pannullo, R.L. DeLaPaz, D.C. Shungu (2006). "Multisection ^{1}H magnetic resonance spectroscopic imaging assessment of glioma response to chemotherapy." *J Neurooncol* **76**: 185-191.

[5] A. Lin, B.D. Ross, K. Harris, W. Wong (2005). "Efficacy of Proton Magnetic Resonance Spectroscopy in Neurological Diagnosis and Neurotherapeutic Decision Making." *NeuroRx* **2**: 197-214.

[6] P. Sajda, S. Du, T.R. Brown, R. Stoyanova, D.C. Shungu, X. Mao, L.C. Parra (2004). "Non-negative Matrix Factorization for Rapid Recovery of Constituent Spectra in Magnetic Resonance Chemical Shift Imaging of the Brain." *IEEE Trans Med Imag*, **23**: 1453-1465.

[7] S. Du, X. Mao, P. Sajda, D.C. Shungu (2007). "Automated Tissue Segmentation and Blind Recovery of ^{1}H MRS Imaging Spectral Patterns of Normal and Diseased Human Brain." *NMR Biomed* (Epub, ahead of print, March 8, 2007).

[8] S. Du , P. Sajda , R. Stoyanova , T.R. Brown (2005). "Recovery of Metabolomic Spectral Sources Using Non-negative Matrix Factorization." *Proc. 27th Annual Int. Conf. EMBS* **2**: 1095-1098.

[9] D.D. Lee, H.S. Seung (1999). "Learning the Parts of Objects by Non-Negative Matrix Factorization." *Nature* **401**: 788-791.

[10] D. D. Lee, H. S. Seung (1996). "Unsupervised Learning by Convex and Conic Coding." *NIPS* **9:** 515-521.