

MULTI-RESOLUTION HIERARCHICAL BLIND RECOVERY OF BIOCHEMICAL MARKERS OF BRAIN CANCER IN MRSI

Shuyan Du, Paul Sajda

Department of Biomedical Engineering
Columbia University
New York, NY 10027 USA

Xiangling Mao, Dikoma Shungu

Department of Radiology
Mount Sinai School of Medicine
New York, NY 10029 USA

ABSTRACT

We present a multi-resolution hierarchical application of the constrained non-negative matrix factorization (cNMF) algorithm [1] for blindly recovering constituent source spectra in magnetic resonance spectroscopic imaging (MRSI). cNMF is an extension of non-negative matrix factorization (NMF) [2] [3] that includes a positivity constraint on amplitudes of recovered spectra. We apply cNMF hierarchically, with spectral recovery and subspace reduction constraining which observations are used in the next level of processing. The decomposition model recovers physically meaningful spectra which are tissue-specific given a processing hierarchy that proceeds coarse-to-fine. We demonstrate the decomposition procedure on ^1H long TE brain MRS data. The results show recovery of markers for normal brain tissue, high grade malignant brain tumor and removal of residual water and lipids. The coarse-to-fine hierarchy also makes the algorithm computationally efficient, thus it is potentially well-suited for use in diagnostic work-up.

1. INTRODUCTION

Magnetic resonance spectroscopic imaging (MRSI) is an imaging modality whereby high resolution magnetic resonance spectra are acquired across a volume of tissue. It has emerged in recent years as a powerful molecular imaging modality to be used as an adjunct to conventional structural MRI. MRSI allows for the non-invasive characterization and quantification of molecular markers, playing an important clinical role in the noninvasive detection, identification, treatment and monitoring of various diseases, most notably brain cancers. Clinicians are more frequently referring patients for metabolic "work up" with MRSI. The increased demand for clinical MRSI scans, and its complementarity with structural MRI, have brought about the need for effective computer-assist tools that can provide integrated biochemical and morphological features of biological tissue and disease processes critical for diagnostic or prognostic assessment.

Brain tumor ^1H MR spectra are typically characterized by the signal intensity change of several important biochemicals [4]. N-acetyl-aspartate (NAA, single resonance peak at 2.02ppm) is thought to exist primarily in viable neurons. The reduction or absence of an NAA spectral peak is attributed to a low density of viable neuronal cells in brain tumors. Choline (CHO, single resonance peak at 3.22ppm) typically is elevated in tumors compared with normal brain tissue. This is thought to be due to accelerated membrane synthesis of rapidly dividing cancer cells. Creatine (CR, singlet at 3.04 and 3.9ppm) is often reduced in tumors, but the significance of decreased CR in terms of tumor metabolism is not clear. Lactic acid (LAC, doublet at 1.33ppm) is often observed in tumor spectra, partially due to the preference for aerobic glycolysis which is often prominent in highly metabolic tumors.

In MRSI, each tissue type can be viewed as having a characteristic spectral profile or set of profiles corresponding to the chemical composition of the tissue. In tumors, for example, ^1H MRSI has shown that metabolites are heterogeneously distributed and in a given voxel multiple metabolites and tissue types may be present [5]. The observed spectra are therefore a combination of different constituent spectra. The signal measured in MRSI is the response to a coherent stimulation of the entire tissue. As a result the amplitudes of the different coherent resonators are additive. The overall gain with which a tissue type contributes to this addition is proportional to its concentration in each voxel. As a result we can explain the observed spectra \mathbf{X} as

$$\mathbf{X} = \mathbf{AS} + \mathbf{N}, \quad (1)$$

where the columns in \mathbf{A} represent the concentration, or abundance, of the constituent material and the rows in \mathbf{S} their corresponding spectra. \mathbf{N} represents additive noise. The abundance matrix \mathbf{A} has M columns (one for each material) and N rows (one for each voxel). \mathbf{X} and \mathbf{S} have L columns (one for each resonance frequency).

Since we interpret \mathbf{A} as concentrations, we can assume the matrix to be non-negative. In addition, since the con-

stituent 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. The factorization of Equation 1 is therefore constrained by,

$$\mathbf{A} \geq 0 \text{ and } \mathbf{S} \geq 0. \quad (2)$$

In MRSI, spatial inhomogeneity of the magnetic field may introduce an unknown phase shift which must be separately estimated for every voxel. Inhomogeneity within a voxel and errors in the estimation of that phase as well as measurement noise may lead to violations of the positivity constraint of the observed spectra. Thus the observed spectra, \mathbf{X} , may have negative values.

cNMF is a fast algorithm [1], based on the non-negative matrix factorization (NMF) algorithm of Lee and Seung [2, 3], which enables non-negative factorization even for noisy observations which may be negative. cNMF includes a positivity constraint, forcing negative amplitude spectral values in the recovered sources and abundance distributions to a small positive value which is approximately zero.

In this paper we further develop cNMF and show how it can be applied hierarchically, automatically removing residual water and lipids, and recovering specific spectral signatures of tissue types using a coarse-to-fine pyramid decomposition strategy. Computationally, this framework is very efficient and can thus be used in near real-time, when a patient is in the MR scanner.

2. APPROACH

With a Gaussian noise assumption for MRSI data [1], the recovery of \mathbf{A} and \mathbf{S} can be formulated as maximum likelihood estimation,

$$\begin{aligned} \mathbf{A}_{ML}, \mathbf{S}_{ML} &= \underset{\mathbf{A}, \mathbf{S}}{\operatorname{argmin}} (-\log p(\mathbf{X} | \mathbf{A}, \mathbf{S})) \\ &= \underset{\mathbf{A}, \mathbf{S}}{\operatorname{argmin}} \|\mathbf{X} - \mathbf{A}\mathbf{S}\|^2 \\ \text{subject to : } & \mathbf{A} \geq 0, \mathbf{S} \geq 0. \end{aligned} \quad (3)$$

Computing the gradient updates for \mathbf{A} and \mathbf{S} and with appropriate choice of step sizes, Lee and Seung [3] showed that one could obtain multiplicative update rules for NMF,

$$\begin{aligned} \mathbf{A}_{i,m} &\leftarrow \mathbf{A}_{i,m} \frac{(\mathbf{X}\mathbf{S}^T)_{i,m}}{(\mathbf{A}\mathbf{S}\mathbf{S}^T)_{i,m}} \\ \mathbf{S}_{m,\lambda} &\leftarrow \mathbf{S}_{m,\lambda} \frac{(\mathbf{A}^T\mathbf{X})_{m,\lambda}}{(\mathbf{A}^T\mathbf{A}\mathbf{S})_{m,\lambda}} \end{aligned} \quad (4)$$

By formulating the updates as multiplicative rules in Equation 4 we can ensure non-negative \mathbf{A} and \mathbf{S} , given both

are initialized non-negative and the observations, \mathbf{X} , are non-negative. However one problem is that, due to noise, the observations can have negative values. Since all observations are used in updating \mathbf{A} and \mathbf{S} , non-negativity will not be guaranteed even if \mathbf{A} and \mathbf{S} are initialized as non-negative. This can lead to physically unrealistic solutions for the recovered spectra—i.e., spectra with negative amplitudes.

In the cNMF algorithm, an initialization step is added which includes constructing a non-negative random \mathbf{A} and estimating \mathbf{S} by solving a non-negatively constrained least squares problem,

$$\underset{\mathbf{A}, \mathbf{S}}{\operatorname{argmin}} \|\mathbf{X} - \mathbf{A}\mathbf{S}\|^2 \text{ subject to } \mathbf{S} \geq 0. \quad (5)$$

Before solving this least squares problem, we must define the dimensionality of our matrices, namely choosing M , the number of recovered sources. An important element of the algorithm is the fact that the factorization in Equation 1 includes an explicit subspace reduction from an N dimensional space into a constrained M dimensional space. Such a compression or “bottleneck” has been shown to be useful in having the subspace capture statistical regularities in the data [6][7]. Except for the positivity constraints, the decomposition is completely arbitrary within that M -dimensional space. However, spectra and concentrations are non-negative and so the M -dimensional degrees of freedom within that subspace are constrained by $M(N + L)$ linear boundary constraints (Equation 2). This is the portion of the space that corresponds to realistic solutions of the factorization. We wish to further constrain the space of possible solutions by exploiting the fact resonances are non-negative, and negative spectral magnitude values are disallowed since we assume that they are due to baseline noise. We enforce this constraint by introducing a threshold on \mathbf{S} ,

$$S_{i,j} = \begin{cases} S_{i,j} & S_{i,j} > 0 \\ \epsilon & S_{i,j} \leq 0 \end{cases} \quad (6)$$

where ϵ is some very small positive value¹. We treat \mathbf{A} symmetrically, using the same positivity constraint as mentioned above, to ensure \mathbf{A} remains non-negative, given the possibility of negative values in \mathbf{X} . To summarize, the procedure for updating \mathbf{A} and \mathbf{S} is,

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.

¹Note that the spectral amplitudes cannot be set exactly to zero given the update rules for \mathbf{A} and \mathbf{S} . We therefore use $\epsilon = 2.2204 \times 10^{-16}$ (the value of floating point relative accuracy used by MATLAB 6.5)

4. Update \mathbf{S}
5. Force negative values of \mathbf{S} to be approximately zero.
6. Iterate (back to 2).

We further extend the methodology by applying cNMF hierarchically, with the spectral recovery and subspace reduction constraining which observations are used in the next hierarchical level of recovery. A flow diagram, shown in Figure 1, illustrates the basic idea. The approach enables an automatic “drilling down” of the source space, ultimately increasing the specificity of the recovered spectra based on a natural, and physically meaningful hierarchy (e.g., head, brain, tumor, etc).

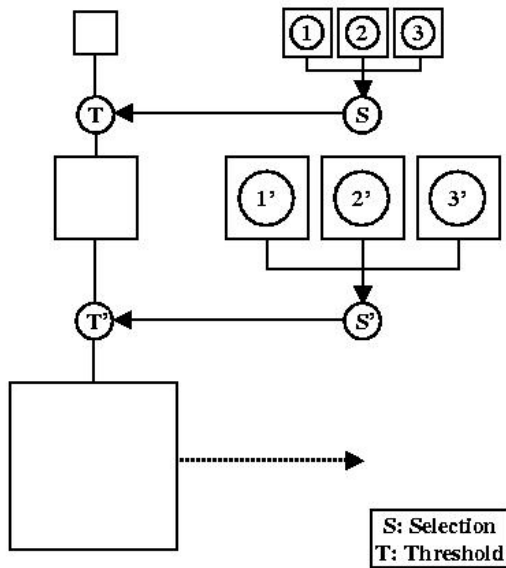


Fig. 1. Multi-resolution hierarchical blind recovery. At the first level a sub-sampled image is used as observations for recovering source spectra (e.g., 1, 2 and 3 shown in the figure). The source spectra are analyzed given prior information about the spectral signatures (i.e. spectra for brain versus muscle) and their spatial distributions (location of the head/brain in the image and its approximate shape). From the recovered spectrum which is identified as being most “brain-like” (based on shape of spatial distribution and spectral profile) a mask is constructed by thresholding the spatial distribution (abundance matrix) of that spectrum. This spatial mask is used for selecting those voxels which will be processed at the next level. At the second level, only those voxels passing through the mask are used as observations and cNMF is reapplied to recover new source spectra (e.g., 1’, 2’ and 3’). This process is iterative and continues until the highest (original) resolution. Note that the hierarchy is constructed on a pyramid, and proceeds from coarse-to-fine resolution.

In a multi-resolution hierarchical decomposition, each level of analysis uses different spatial resolution spectral images for cNMF source recovery. Principal component analysis (PCA) is used to align the spectra and estimate the number of source components, level-by-level. For example, when a 3-level multi-resolution hierarchy model is used with the original spectral image having a resolution of $K \times K$, the next level is constructed to be $\frac{K}{2} \times \frac{K}{2}$ by averaging four neighboring spectra and sub-sampling in each dimension by two. The third level in the hierarchy is constructed from the second level using the same procedure, resulting in a $\frac{K}{4} \times \frac{K}{4}$ spectral image.

Starting with the lowest (coarsest) hierarchical level, all $\frac{K}{4} \times \frac{K}{4}$ spectra are input into the cNMF, with the number of components having already been determined via PCA. cNMF recovery results of separated spectra are analyzed together with their corresponding spatial distributions to select only one brain-like spectrum. A mask is constructed by thresholding the spatial map corresponding to the selected brain-like spectrum. This mask, constructed from the coarsest resolution $\frac{K}{4} \times \frac{K}{4}$, is up-sampled to $\frac{K}{2} \times \frac{K}{2}$ and applied to the second level $\frac{K}{2} \times \frac{K}{2}$ image in the hierarchy. Only voxels that pass through the mask are considered to contain brain tissues, and used to estimate \mathbf{A} and \mathbf{S} at the next level. The process proceeds to the highest level/finest resolution in the hierarchy.

3. RESULTS

We show results for a single case, consisting of 32-by-32 voxel axial view ^1H long echo time (TE) MRSI human brain data with each voxel having a spectrum of 1024 points. Other cases have yielded similar results.

A three level hierarchy is used (32-by-32, 16-by-16 and 8-by-8). The number of source components is set to 3 or 4 (depending on the hierarchical level) via PCA analysis. This is consistent with 3-4 primary tissue types found in human brain (e.g., normal brain, pericranial fat and tissues, skull and pathologic brain). The recovered spectra and their corresponding spatial distributions for the 3-level multi-resolution hierarchical decomposition are shown in Figure 2. We see the multi-resolution hierarchical decomposition recovers source spectra of biochemical markers with increasing tissue specificity, level-by-level. The first level yields spatial distributions and spectral characteristics indicative of head/brain. In the second level most of residual lipids are moved (column 2, 3 and 4) and clear in the spectra are biochemical markers for brain (column 1) with peaks for CHO, CR and NAA. The third level yields separation of high grade malignant brain tumor (column 3), normal brain tissue (column 4) and more residual lipids are removed (column 1 and 2).

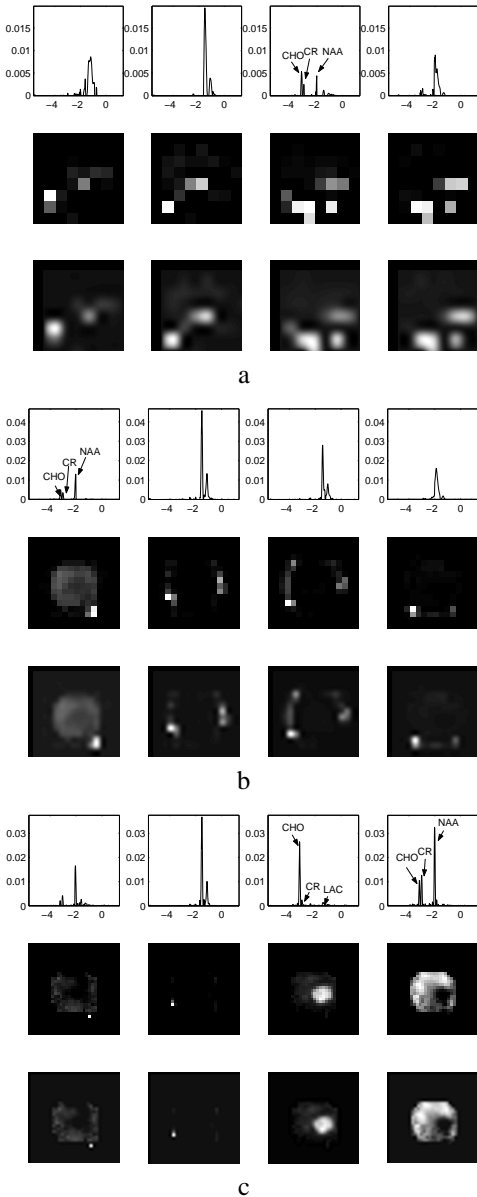


Fig. 2. (a) First level of hierarchy, resolution 8-by-8, separation results. First row is the recovered spectra, second row is the spatial distribution of each spectra and the third row is an upscaled version of the spatial distribution. (b) Second level of hierarchy separation results, resolution 16-by-16. Clear in the spectra are biochemical markers for brain (1st column) with peaks for CHO, CR and NAA. (c) Third level of hierarchy separation results, resolution 32-by-32. Note that the corresponding spatial mapping of these markers maps to the region of the brain. Column 3 spectrum is indicative of high grade malignant brain tumor: highly elevated CHO, decreased CR, almost no NAA and LAC peak appearance; Column 4 spectrum indicates normal brain tissue: low CHO, high CR, high NAA and no LAC peak. Column 1 and 2 are more residual lipids removed.

4. CONCLUSION

In this paper, we present a multi-resolution hierarchical decomposition for recovering physically meaningful source spectra in MRSI. The decomposition uses the cNMF algorithm [1] and therefore does not explicitly impose a sparsity constraint, though the algorithm recovers sparse sources quite well. The algorithm also does not assume independence or orthogonality, which avoids separation failure in the case of MRSI where source spectra are highly correlated. The multi-resolution hierarchical cNMF decomposition enables an automatic “drilling down” of source space with increasing specificity of the recovered spectra based on a natural and physically meaningful hierarchy (e.g., head, brain, tumor, etc), saving data preprocessing (e.g., PCA) and computational cost.

5. REFERENCES

- [1] P. Sajda, S. Du, T. Brown, L. Parra, and R. Stoyanova, “Recovery of constituent spectra in 3D chemical shift imaging using non-negative matrix factorization,” in *4th International Symposium on Independent Component Analysis and Blind Source Separation (ICA2003)*, 2003, pp. 71–76.
- [2] D. D. Lee and H. S. Seung, “Learning the parts of objects by non-negative matrix factorization,” *Nature*, vol. 401, pp. 788–791, 1999.
- [3] D. D. Lee and H. S. Seung, “Algorithms for non-negative matrix factorization,” in *Advances in Neural Information Processing Systems 13*. 2001, pp. 556–562, MIT Press.
- [4] F. A. Howe and K. S. Opstad, “ ^1H MR spectroscopy of brain tumours and masses,” *NMR in Biomedicine*, vol. 16, pp. 123–131, 2003.
- [5] S. Furuya, S. Naruse, M. Ide, H. Morishita, O. Kizu, S. Ueda, and T. Maeda, “Evaluation of metabolic heterogeneity in brain tumors using ^1H -chemical shift imaging method,” *NMR in Biomedicine*, vol. 10, pp. 25–30, 1997.
- [6] N. Tishby, F.C. Pereira, and W. Bialek, “The information bottleneck method,” in *Proc. of 37th Annual Allerton Conference on Communication, Control and Computing*, 1999, pp. 368–377.
- [7] N. Friedman, O. Mosenzon, N. Slonim, and N. Tishby, “Multivariate information bottleneck,” in *Proc. of 17th Conf. on Uncertainty in Artificial Intelligence (UAI)*, 2001, pp. 152–161.