

Monte Carlo Algorithm for Least Dependent Non-Negative Mixture Decomposition

Sergey A. Astakhov,^{*,†} Harald Stögbauer,^{†,‡} Alexander Kraskov,^{†,‡} and Peter Grassberger[†]

John von Neumann Institute for Computing, Forschungszentrum Jülich, D-52425, Jülich, Germany, and Division of Biology, California Institute of Technology, Pasadena, California 91125

We propose a simulated annealing algorithm (stochastic non-negative independent component analysis, SNICA) for blind decomposition of linear mixtures of non-negative sources with non-negative coefficients. The demixing is based on a Metropolis-type Monte Carlo search for least dependent components, with the mutual information between recovered components as a cost function and their non-negativity as a hard constraint. Elementary moves are shears in two-dimensional subspaces and rotations in three-dimensional subspaces. The algorithm is geared at decomposing signals whose probability densities peak at zero, the case typical in analytical spectroscopy and multivariate curve resolution. The decomposition performance on large samples of synthetic mixtures and experimental data is much better than that of traditional blind source separation methods based on principal component analysis (MILCA, FastICA, RADICAL) and chemometrics techniques (SIMPLISMA, ALS, BTEM).

Decomposing linear mixtures (superpositions) into their components is a problem occurring in many different branches of science, such as telecommunications, seismology, image processing, and biomedical signal analysis (e.g., EEG, ECG, fMRI).¹ Blind source separation (BSS), in particular, deals with the case where neither the sources nor the mixing matrix is known, the only available data being the mixed signals.² The standard approach to BSS is independent component analysis (ICA).^{1–4} In ICA, one typically first applies a principal component transformation to “prewhiten” the data, that is, to transform the covariance matrix into a unit matrix, and then uses some statistics going beyond second order (e.g., mutual information or temporal correlations) to minimize interdependencies between the reconstructed sources by a final rotation.

A similar problem arises in analytical spectroscopy, where one wants to decompose spectra of chemical mixtures into the

contributions of individual components in order to identify the emitting (or absorbing) mixture constituents and quantify their abundances. Mixture decomposition problems of this sort are classified in analytical chemistry as quantitative analysis of “black” multicomponent systems.⁵ Compared to other applications of BSS, here we have two additional problems:

(1) Spectra of different chemically pure substances often show very large overlaps. This is in conflict with the central assumption of ICA that the pure sources (mixture components) are statistically independent. One can relax this assumption partially (see, e.g., refs 6 and 7), although one can hardly do without any assumption of this type. But, because of overlaps between spectra of similar compounds, this becomes a serious obstacle in spectral mixture decomposition.

(2) In a vast majority of spectroscopic techniques, spectra are non-negative, as are also the coefficients of the mixing matrix (concentrations). This gives rise to constraints which are not easily taken into account in typical ICA algorithms that heavily rely on linear algebra.^{8,9} This would not be problematic by itself—in other applications, signals are often non-negative, too. But it does pose a very serious problem, if, as is usually the case, there are spectral regions where some of the sources have zero (or very small) intensity.¹⁰ In this case, prewhitening will in general produce spectra violating the non-negativity constraints, and restoring them later will be far from trivial.

For these reasons, this “multivariate spectral curve resolution” (MCR) problem¹¹ has led to a variety of chemometrics methods (for recent reviews, see refs 12–16). Generally, these algorithms resemble ICA in that they perform blind recovery of spectra and concentrations from the spectra of mixtures only, using no or little a priori information about mixture composition or about pure spectra. In view of the problems mentioned above, some of these

* Corresponding author. Phone: +49 2461 612526. Fax: +49 2461 612430. E-mail: astakhov@gmail.com.

[†] John von Neumann Institute for Computing.

[‡] Division of Biology, California Institute of Technology, CA 91125, USA.

(1) Hyvärinen, A.; Karhunen, J.; Oja, E. *Independent Component Analysis*; Wiley: New York, 2001.

(2) Cichocki, A.; Amari, S. *Adaptive Blind Signal and Image Processing. Learning Algorithms and Applications*; Wiley: New York, 2002.

(3) Comon, P. *Signal Process.* **1994**, *36*, 287–314.

(4) Jutten, C.; Herault, J. *Signal Process.* **1991**, *24*, 1–10.

(5) Liang, Y.-Z.; Kvalheim, O. M.; Manne, R. *Chemom. Intell. Lab. Syst.* **1993**, *18*, 235–250.

(6) Stögbauer, H.; Kraskov, A.; Astakhov, S. A.; Grassberger, P. *Phys. Rev. E* **2004**, *70*, 066123.

(7) Astakhov, S. A.; Stögbauer, H.; Kraskov, A.; Grassberger, P. *Chemom. Intell. Lab. Syst.* Submitted (preprint at arxiv.org/abs/physics/0412029).

(8) Hesse, C. W.; James, C. J. *IEEE Signal Process. Lett.* **2005**, *12*, 792–795.

(9) Lu, W.; Rajapakse, J. C. *IEEE T. Neural Networks* **2005**, *16*, 203–212.

(10) Plumbley, M. D.; Oja, E. *IEEE T. Neural Networks* **2004**, *15*, 66–76.

(11) Lawton, W. H.; Sylvestre, E. A. *Technometrics* **1971**, *13*, 617–633.

(12) Geladi, P. *Spectrochim. Acta B* **2003**, *58*, 767–782.

(13) Hopke, P. K. *Anal. Chim. Acta* **2003**, *500*, 365–377.

(14) Jiang, T.-H.; Liang, Y.; Ozaki, Y. *Chemom. Intell. Lab. Syst.* **2004**, *71*, 1–12.

(15) de Juan, A.; Tauler, R. *Anal. Chim. Acta* **2003**, *500*, 195–210.

(16) Lavine, B.; Workman, J. J. *Anal. Chem.* **2004**, *76*, 3365–3371.

MCR methods are less ambitious than ICA and estimate only feasible ranges for spectra and concentrations (see, for instance, a recent Monte Carlo approach¹⁷ and references therein). Also, the MCR analysis is often facilitated by including some additional information such as unimodality and closure (or mass balance).¹⁵ The effect of overlap between spectra can be reduced by estimating their dependencies not from the spectra $S(f)$ directly but from their second derivatives $d^2S(f)/df^2$.^{7,18} Another approach developed in chemometrics is based on the identification of “pure variables”¹⁹ for all mixture components. A pure variable is a frequency (wavelength) at which only one of the components contributes. The pure variables, thus, approximately mark the regions where at least one of the spectral components is guaranteed to be independent from all others. This idea is central to several advanced chemometrics algorithms, for example, KSFA,²⁰ SIMPLISMA,²¹ IPCA,²² and SMAC.²³ Also, band-target entropy minimization (BTEM) has been recently proposed, which involves an explicit (made by visual inspection) choice of spectral features (target regions) to be retained in the course of constrained optimization.²⁴ While efficient and highly flexible, supervised band selection has the drawbacks of being neither fully automated nor completely blind.

To cope with the non-negativity problem, several methods have been proposed. In refs 10, 25, and 26, non-negativity is incorporated in an approximate way already in the basic decomposition algorithm. For a discussion, see ref 27. While the above are general purpose BSS methods, there exist more specialized chemometrics (MCR) algorithms designed to use non-negativity (see, e.g., refs 12–16, 24, and 28). In many of them, some variant of principal component analysis is used to first prewhiten the data, that is, to minimize the off-diagonal elements of the covariance matrix. In a class of algorithms known as MCR-ALS, prewhitening and the interdependence minimization are done as in standard ICA, but non-negativity is then enforced in a postprocessing step using the alternating least squares (ALS) technique.^{7,29,30} In general, this gives better results than the algorithms without ALS. But, since the interdependencies between the sources are not checked during the ALS step, one runs the risk of worsening them. In addition to that, applying post facto corrections to restore essential properties that were present in original data, but that were ruined by the main part of an algorithm, appears rather awkward conceptually.

Here we argue that a more straightforward strategy is not only possible but actually leads to demixing algorithms with notably improved performance. Our approach is conceptually very simple and is based on mixture decomposition into strictly non-negative least dependent components. Unlike other ICA methods, it does not resort to prewhitening. Instead, we split the linear demixing transformation into small random steps such that (i) on each step the non-negativity is preserved exactly and (ii) the dependencies between the components tend to be reduced (the mutual information (MI) is minimized nongreedily). Since greedy minimization would lead to trapping in local minima of MI, we use a Metropolis–Hastings Monte Carlo strategy.^{31,32} The algorithm (which we term SNICA, for stochastic non-negative independent component analysis) is explained in the next section. The numerical results presented here show that our new method is more efficient than previous MCR algorithms, including ICA- and ICA-ALS-type methods⁷ based on the same MI estimator.

METHOD

As in most ICA and MCR approaches, we start out with the linear mixture model

$$\mathbf{X} = \mathbf{A}\mathbf{S} \quad (1)$$

where \mathbf{X} is the matrix of mixed signals (observed spectra), \mathbf{S} are pure sources (spectra of individual mixture components), and \mathbf{A} is the matrix of their concentrations (mixing matrix). We assume that \mathbf{A} is a square $K \times K$ matrix, that is, the number of sources K is the same as the number of mixtures. \mathbf{X} and \mathbf{S} are of size $K \times N$, where the number N of frequencies in the spectra is much larger than K . The task is to estimate blindly \mathbf{S} and \mathbf{A} , given only \mathbf{X} . In our current setting, the pure spectra and concentrations are assumed to be non-negative and we seek a solution (an estimate for \mathbf{S}) in the form of statistically least dependent components $\mathbf{Y} = \mathbf{W}\mathbf{X}$, with \mathbf{W}^{-1} being an estimate for \mathbf{A} up to scaling and permutation of components (ambiguities inherent to the linear mixture model). True and estimated mixing matrices can be compared using the Amari index $P(\mathbf{W}^{-1}, \mathbf{A})$ which is a good measure of decomposition quality.^{2,6,7} The Amari index vanishes when the recovered concentrations differ from the true ones only in scaling and permutation of components, and it increases as the quality of decomposition becomes poor. Thus, small values of the Amari index are desirable. (In practice, we find that good decomposition quality roughly corresponds to Amari indices $P < 0.05$, whereas $P > 0.2$ generally characterizes unacceptably poor performance.)

We assume that each column vector \mathbf{y}_i , $i = 1, 2, \dots, N$ is a realization of the same vector-valued random variable $Y = (Y_1, Y_2, \dots, Y_K)$, and we look for solutions where its K components (i.e., the estimated individual spectra) have least mutual information. For a multivariate continuous random variable with given marginal and joint densities $\mu_j(y_j)$ and $\mu(y_1, y_2, \dots, y_K)$, the mutual information is defined as

$$I(Y_1, Y_2, \dots, Y_K) = \sum_{j=1}^K H(Y_j) - H(Y_1, Y_2, \dots, Y_K) \quad (2)$$

(17) Leger, M. N.; Wentzell, P. D. *Chemom. Intell. Lab. Syst.* **2002**, *62*, 171–188.

(18) Windig, W.; Stephenson, D. A. *Anal. Chem.* **1992**, *64*, 2735–2742.

(19) Knorr, F. J.; Futrell, J. H. *Anal. Chem.* **1979**, *51*, 1236–1241.

(20) Malinowski, E. R. *Anal. Chim. Acta* **1982**, *134*, 129–137.

(21) Windig, W.; Guilment, J. *Anal. Chem.* **1991**, *63*, 1425–1432.

(22) Bu, D. S.; Brown, C. W. *Appl. Spectrosc.* **2000**, *54*, 1214–1221.

(23) Windig, W.; Gallagher, N. B.; Shaver, J. M.; Wise, B. M. *Chemom. Intell. Lab. Syst.* **2005**, *77*, 85–96.

(24) Widjaja, E.; Li, C.; Chew, W.; Garland, M. *Anal. Chem.* **2003**, *75*, 4499–4507.

(25) Lee, D. D.; Seung, H. S. *Nature* **1999**, *401*, 788–791.

(26) Yuan, Z.; Oja, E. *Lect. Notes Comput. Sc.* **2004**, *3195*, 1–8.

(27) Cichocki, A.; Georgiev, P. *IEICE T. Fundam. Electron. EA* **2003**, *E86A*, 522–531.

(28) Duponchell, L.; Elmi-Rayaleh, W.; Ruckebusch, C.; Huvenne, J. P. J. *Chem. Inf. Comput. Sci.* **2003**, *43*, 2057–2067.

(29) Tauler, R.; Kowalski, B.; Fleming, S. *Anal. Chem.* **1993**, *65*, 2040–2047.

(30) Wang, J.-H.; Hopke, P. K.; Hanczewicz, T. M.; Zhang, S. L. *Anal. Chim. Acta* **2003**, *476*, 93–109.

where

$$H(Y_j) = - \int \mu_j \ln \mu_j dy_j \quad (3)$$

and

$$H(Y_1, Y_2, \dots, Y_K) = - \int \mu \ln \mu dy_1 dy_2 \dots dy_K \quad (4)$$

are the differential entropies.³³ Mutual information is zero if and only if the Y_j are strictly independent (i.e., their joint density factorizes, $\mu = \prod \mu_j$), and it is positive otherwise. Notice that the MI is sensitive to all types of dependencies, not only linear correlations.

To assess the dependencies between the estimated spectra (given by rows of \mathbf{Y}), we use the precise MI estimator based on nearest-neighbor statistics, denoted here as $I(\mathbf{Y})$, developed in ref 34. It is derived from the Kozachenko–Leonenko estimate for Shannon entropy³⁵

$$H(Y) = -\psi(k) + \psi(N) + \log c_D + \frac{D}{N} \sum_{i=1}^N \log \epsilon(i) \quad (5)$$

where $\psi(\cdot)$ is the digamma function, $\epsilon(i)$ is twice the distance from the point \mathbf{y}_i to its k th neighbor, D is the dimension of Y , c_D is the volume of the D -dimensional unit ball, and N is the length of sample. Mutual information could be computed by estimating the entropies 3 and 4 separately, with the same value of k , and using eq 2. But this would involve very different scales in the marginal and joint spaces, leading to different biases. Looking for nearest neighbors on the same scale in the joint (Y_1, Y_2, \dots, Y_K) and marginal Y_j spaces significantly increases the chance that the biases in $H(Y_j)$ and $H(Y_1, Y_2, \dots, Y_K)$ cancel. This reduces the bias of the MI estimate. Since eq 5 holds for any value of k , one can first look for k nearest neighbors in the joint space and then use the hyper-rectangles defined by them to calculate the number of neighbors in the marginal spaces. The estimate for MI is then

$$I(Y_1, Y_2, \dots, Y_K) = \psi(k) - (K-1)/k - \langle \psi(n_1) + \psi(n_2) + \dots + \psi(n_K) \rangle + (K-1)\psi(N) \quad (6)$$

where n_j are the numbers of nearest neighbors in the spaces Y_j , $j = 1, 2, \dots, K$; $\langle \dots \rangle$ denotes the average over all points \mathbf{y}_i , $i = 1, 2, \dots, N$. Implementation details are given in ref 34.

Statistical errors in this MI estimator decrease with k , while systematic errors increase. Empirically, we found that the decomposition performance is best for the values of k in the range 10–

20, reflecting a balance between these two types of errors. Also, spectral curve resolution by ICA was shown to be more efficient if performed in the derivative space;^{7,18} therefore, in all simulations discussed below, we estimate the MI from the second derivatives of spectra with respect to frequency.

With this numerical measure of dependence, we can formulate the mixture decomposition problem as minimization of $I(\mathbf{Y})$ under the linear demixing transformation $\mathbf{Y} = \mathbf{W}\mathbf{X}$, subject to the non-negativity constraint $y_{jk} \geq 0$. (Note that the non-negativity of estimated mixing matrix \mathbf{W}^{-1} is not enforced explicitly, but it will be nonetheless satisfied to high accuracy.)

In contrast to more traditional ICA methods that proceed by splitting the demixing matrix into a prewhitening transformation (eliminating linear correlations) and subsequent rotations,^{1,6} we approach the above constrained optimization with a Monte Carlo strategy. To this end, the demixing matrix \mathbf{W} is factorized into small steps, $\mathbf{W} = \lim_{n \rightarrow \infty} \mathbf{W}_n$ with $\mathbf{W}_n = \mathbf{M}_n \mathbf{W}_{n-1}$, where each \mathbf{M}_n is chosen randomly and is either a shear transformation \mathbf{T} in a two-dimensional (2D) subspace or a rotation \mathbf{R} in a 3D subspace (see below). The sequence of moves starts with $\mathbf{Y}_0 = \mathbf{X}$, $\mathbf{W}_0 = \mathbf{I}$ (identity matrix). A move $\mathbf{Y}_n = \mathbf{M}_n \mathbf{Y}_{n-1}$ is immediately rejected, if it results in at least one negative component $y_{n,ik} = (\mathbf{Y}_n)_{ik}$. Otherwise (i.e., if all $y_{n,ik} \geq 0$), it is accepted with probability one, if it leads to less dependent components, $\Delta I = I(\mathbf{Y}_n) - I(\mathbf{Y}_{n-1}) < 0$, while it is accepted with probability $\exp(-\Delta I/T)$, if $\Delta I > 0$. Here, T is a fictitious temperature. This is the famous Metropolis–Hastings Markov chain Monte Carlo method.^{31,32} To make the method more greedy at later stages, when \mathbf{W}_n is already close to the optimum, we use a simple annealing schedule³⁶ where we start with high values of T and decrease it successively. Since both \mathbf{R} and \mathbf{T} matrices act only on low-dimensional subspaces (and are the identity in the complements), the difference ΔI in each step can be computed from information in a 2D or 3D subspace only.⁶ The convergence to the global solution is monitored by computing the full dimensional mutual information $I(\mathbf{Y}_n)$, which guarantees that the least dependent, although not necessarily independent, components are reached.

In our method, the mutual information is used to direct the search in the global space, whereas the non-negativity constraint is important in that it restricts solutions to physically meaningful subspace. This constraint has strongest effect if large parts of the constraint boundaries $y_k = 0$, $k = 1, 2, \dots, K$, are populated by the pure sources, as then any linear transformation is likely to violate non-negativity. This happens when each pure spectrum has a high chance to be occasionally close to its zero baseline. Such sources are called well-grounded¹⁰ and are common in spectroscopy.

The idea of using 2D shear transformations in tandem with the non-negativity constraint was first proposed in the context of “positive matrix factorization”.^{37,38} The matrix $\mathbf{T}^{(i)}(\alpha)$ representing

(31) Metropolis, N.; Rosenbluth, A. W.; Rosenbluth, M. N.; Teller, A. H.; Teller, E. *J. Chem. Phys.* **1953**, *21*, 1087–1092.

(32) Hastings, W. K. *Biometrika* **1970**, *57*, 97–109.

(33) Cover, T. M.; Thomas, J. A. *Elements of Information Theory*; Wiley: New York, 1991.

(34) Kraskov, A.; Stögbauer, H.; Grassberger, P. *Phys. Rev. E* **2004**, *69*, 066138.

(35) Kozachenko, L. F.; Leonenko, N. N. *Probl. Inf. Transm.* **1987**, *23*, 95–101.

(36) Kirkpatrick, S.; Gelatt, C. D.; Vecchi, M. P. *Science* **1983**, *220*, 671–680.

(37) Paatero, P.; Tapper, U. *Environmetrics* **1994**, *5*, 111–126.

(38) Paatero, P.; Hopke, P. K.; Begum, B. A.; Biswas, S. K. *Atmos. Environ.* **2005**, *39*, 193–201.

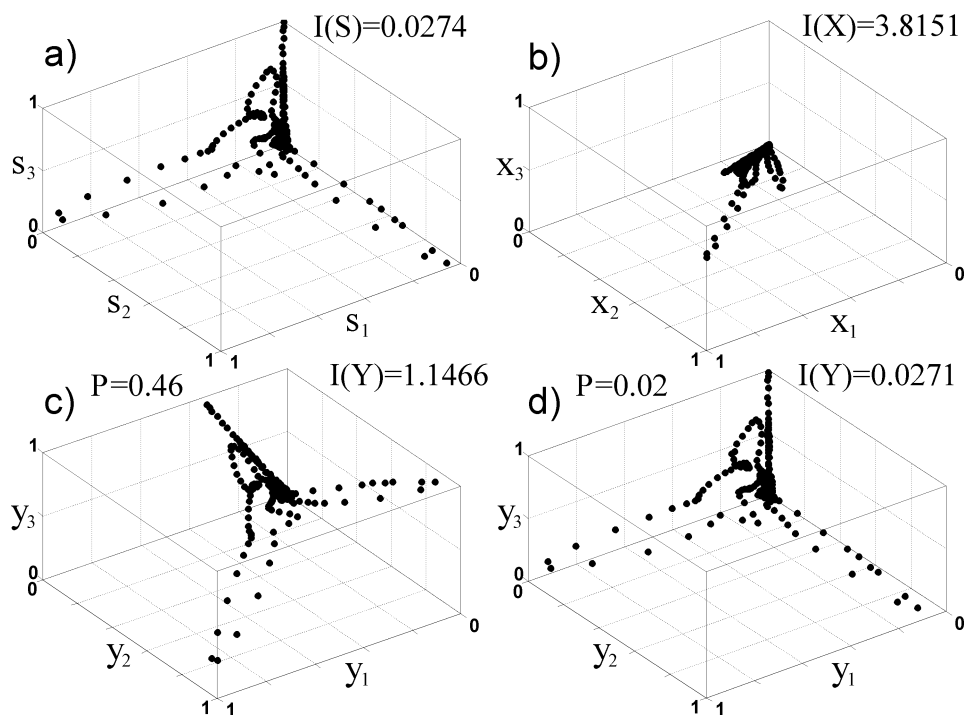


Figure 1. Demixing in three dimensions by shears and rotations. Scatterplots: pure sources (a); mixture signals (b); components (c) resolved by 2D shears (trapped in a local minimum); components (d) recovered through optimization with 2D shears and 3D rotations about the diagonal. Also shown are the values of 3D mutual information (I) and the Amari index (P) of incomplete (c) and successful (d) decompositions.

a shear $y'_i = y_i + \alpha y_j$, $y'_j = y_j$ in the (i, j) subspace is obtained from identity matrix by adding one off-diagonal element,

$$t_{lm}^{(ij)} = [\mathbf{T}^{(ij)}(\alpha)]_{lm} = \delta_{lm} + \alpha \delta_{li} \delta_{mj}, \quad i \neq j \in [1, K]; \\ l, m = 1, 2, \dots, K \quad (7)$$

While 2D shear transformations are sufficient to decompose two-component mixtures, they appear to be too restrictive to sample efficiently higher dimensional demixing matrices. In higher dimensions, the minimization through a sequence of shears can easily get trapped in local minima of the MI landscape. One such case is illustrated in Figure 1 in which a three-component mixture of exemplary infrared spectra is considered. First, decomposition was attempted using only 2D shears, and Figure 1c shows the components found by constrained minimization. This configuration corresponds to a local minimum of MI (compare to Figure 1a and notice the values of mutual information and very poor decomposition performance, $P = 0.46$). Since the components have reached the constraint boundaries and a minimum of the cost function has been found to lie on this boundary, further demixing shears can be accepted only by increasing $I(\mathbf{Y})$, which can be very slow. However, Figure 1c suggests that the configuration needs to be rotated in such a way that, afterward, the three principal “beams” of points could be brought close to the coordinate axes (and MI be decreased further) by subsequent shears. This works indeed and gives the correct pure components with high precision, as seen in Figure 1d.

The rotations needed to improve the efficiency are around the space diagonals in randomly chosen 3D subspaces. Let us denote by $\mathbf{R}^{(ijk)}(\alpha)$ the matrix that performs a rotation by the angle α

around the diagonal in the positive octant of the subspace spanned by the components (i, j, k) . It has the elements

$$r_{ii}^{(ijk)} = r_{jj}^{(ijk)} = r_{kk}^{(ijk)} = (1 + 2 \cos \alpha) / 3 \\ r_{ij}^{(ijk)} = r_{jk}^{(ijk)} = r_{ki}^{(ijk)} = \frac{1}{3}(1 - \cos \alpha) + \frac{\sin \alpha}{\sqrt{3}} \\ r_{ik}^{(ijk)} = r_{ji}^{(ijk)} = r_{kj}^{(ijk)} = \frac{1}{3}(1 - \cos \alpha) - \frac{\sin \alpha}{\sqrt{3}} \quad (8)$$

in the subspace spanned by (i, j, k) , while it is the unit matrix in the complement, $r_{lm}^{(ijk)} = \delta_{lm}$ for $l, m \notin \{i, j, k\}$. In the results shown below, random rotations and shears are used in ratio 1:1.

Of course, more sophisticated combinations of affine transformations^{39,40} can be designed to sample the demixing matrix. This might improve the performance of the algorithm further, especially in higher dimensions. But it is not clear whether the added complexity would be worthwhile.

As seen from eqs 7 and 8, the proposed moves \mathbf{T} and \mathbf{R} are parametrized by the “angle” α . For each Monte Carlo step, its value is chosen randomly from some suitable range $\alpha \in [-h, h]$. To realize an adaptive step-size control (where h is adjusted according to the progress of optimization), we use the following update rule

$$h_{n+1} = \begin{cases} f_1 h_n & \text{if the move } M_n \text{ is accepted} \\ f_2 h_n & \text{if it is rejected} \end{cases} \quad (9)$$

with the factors $f_1 > 1, f_2 < 1$ to be determined empirically. In our simulations, we took $f_1 = 1.06, f_2 = 0.98$ with initial $h_0 = 0.2-0.3$,

although the decomposition performance was found to be practically insensitive to precise values of h_0 (see next section).

For the temperature annealing, we used rather simple schemes with only two or three cooling steps, with the temperature in the last step chosen close to zero. This is because we found that the detailed annealing schedule had no significant effect on the final results. It means, however, that the CPU times needed in the following simulations might not be optimal. The initial value of temperature T_1 was chosen of the order of the mutual information of pure mixture components. Choosing T_1 too high leads to unfocused search, while choosing it too small increases the risk to get trapped in false minima. Both can result in poor performance, whereas our empirical choice of T_1 ensured efficient decomposition in most cases.

At each temperature, the sequence of demixing moves is terminated when the mutual information has reached its global minimum, according to some stopping criterion. In our code, we kept track of the minimum of $I(\mathbf{Y}_n)$ during the entire run,

$$I(\mathbf{Y}_n)_{\min} = \min_{n' < n} I(\mathbf{Y}_{n'}) \quad (10)$$

Minimization is stopped when $I(\mathbf{Y}_n)_{\min}$ does not decrease any more during the last M Monte Carlo steps, and the demixing matrix \mathbf{W}_{\min} that corresponds to $I(\mathbf{Y}_n)_{\min}$ is given as output. The chosen values of M were some compromise between the speed and the quality of decomposition.

With all the above together, the stochastic non-negative demixing procedure with two cooling steps is then as follows:

1. initialize $n = 0$, $\mathbf{Y}_n = \mathbf{X}$, $\mathbf{W}_n = \mathbf{I}$; choose $h_n = h_0$, $T = T_1$, $M = M_1$
2. pick random $i, j, k \in [1, K]$, $\alpha \in [-h_n, h_n]$
3. if n is even, compute shearing matrix $\mathbf{M}_n = \mathbf{T}^{(ij)}(\alpha)$ (eq 7); if else, compute rotation matrix $\mathbf{M}_n = \mathbf{R}^{(ijk)}(\alpha)$ (eq 8)
4. propose the move $\mathbf{Z} = \mathbf{M}_n \mathbf{Y}_n$, $\mathbf{V} = \mathbf{M}_n \mathbf{W}_n$
5. if $\exists z_{im} < 0$, reject the move, keep $\mathbf{Y}_{n+1} = \mathbf{Y}_n$, $\mathbf{W}_{n+1} = \mathbf{W}_n$ and go to (9)
6. estimate $\Delta I = I(\mathbf{Z}) - I(\mathbf{Y}_n)$; pick random $p \in [0, 1]$
7. if $(\Delta I < 0)$ or $(\Delta I > 0 \text{ and } e^{-\Delta I/T} > p)$, accept the move $\mathbf{Y}_{n+1} = \mathbf{Z}$, $\mathbf{W}_{n+1} = \mathbf{V}$; if else, reject the move, and keep $\mathbf{Y}_{n+1} = \mathbf{Y}_n$, $\mathbf{W}_{n+1} = \mathbf{W}_n$
8. if $I(\mathbf{Y}_{n+1})$ reaches a new minimum, store it together with $\mathbf{Y}_{\min} = \mathbf{Y}_{n+1}$, $\mathbf{W}_{\min} = \mathbf{W}_{n+1}$, and $n_{\min} = n + 1$
9. adjust the step size h_{n+1} according to eq 9
10. $n = n + 1$; if $n < n_{\min} + M$, go to (2); or else, if $T \neq T_2$, choose $T = T_2$ and $M = M_2$, set $n = n_{\min}$, and go to (2); or else, output \mathbf{Y}_{\min} and \mathbf{W}_{\min} and stop.

RESULTS AND DISCUSSION

This section reports the results of extensive tests of the new algorithm performed on a large sample of spectral resolution problems. Here we stick to the strategy of statistical validation of decomposition performance developed earlier.^{6,7} We stress that performing such tests on a statistically representative set of precisely known mixtures is virtually the only way to examine real efficiency of a demixing technique.

To compose a test set of randomized mixtures we first collected a pool of 99 experimental infrared absorption spectra in the range 550–3830 cm^{-1} (822 data points each) selected from the NIST database.⁴¹ This pool of chemical species was designed to contain organic compounds that have common structural groups (halogen-, alkyl-, and nitro-substituted benzene derivatives, phenols, alkanes; alcohols, thiols, amines, esters) as well as a number of “outliers”, i.e. structures dissimilar from all the rest. This was done to be able to test how well the algorithm can treat both—mixtures of highly dependent sources⁷ and cases in which very deep minima of MI are to be located. After that, four sets with increasing dimensionality $K = 3, 4, 7, \text{ or } 10$ (6000 mixtures each) were constructed by randomly choosing K normalized pure spectra from the pool and applying random square mixing matrices of proper dimension K . Concentrations a_{ij} were generated uniformly randomly in the interval $[0, 1]$. Thus, our assessment of performance is based on 24 000 decompositions in total.

We examine the performance of the new decomposition technique (SNICA) in comparison with the recently developed MILCA algorithm, which uses the same MI estimator.⁶ In our previous study,⁷ MILCA was shown to be very efficient when compared to other (MCR) decomposition algorithms on typical spectral mixture problems (both synthetic and real). MILCA can be easily combined with derivative preprocessing (to deal with overlapping spectra) and ALS postprocessing (to improve non-negativity of estimates). Decompositions by MILCA were performed in the second-derivative space,⁷ and the postprocessing ALS iterations were stopped after 1000 steps (we verified that longer runs did not result in noticeable improvements).

Monte Carlo demixing was done with a two-step annealing scheme. The initial temperatures and stopping criteria were $T_1 = 0.02, 0.05, 0.5, 1$ and $M_1 = 1000, 1000, 2500, 8000$ (for $K = 3, 4, 7, \text{ and } 10$, respectively). For the subsequent “cooling” stage, these were set to $T_2 = 10^{-7}$ in all four cases and $M_2 = 500, 500, 1500, 3500$.

The Amari index $P(\mathbf{W}^{-1}, \mathbf{A})$ was computed for each decomposition. Figure 2 presents the distributions of the Amari indices with increasing dimensionality K of the mixture problem. Judging from the most probable performance (peaks of the P distributions) and the minimal achieved values of P , SNICA clearly outperformed the other two algorithms. Also, the Monte Carlo demixing seems to gain advantage as the number of mixture components K increases. In the current implementation (with a very simple annealing scheme), SNICA is of order 10 times slower than MILCA and MILCA+ALS for $K = 3$, the algorithms have comparable speeds at $K = 7$, and SNICA is on average two times faster for $K = 10$ (with our stopping criteria).

One of the mixture sets ($K = 3$) was also processed by two representative ICA algorithms: the widely used FastICA⁴² (for references on its chemometrics applications see ref 7) and recently developed RADICAL⁴³ (both algorithms rely on decorrelation of data by PCA). (Second-derivative spectra were used. FastICA ran with its nonlinearity parameter set to “skew”. RADICAL worked

(41) NIST Mass Spec Data Center. S. E. Stein, director. Infrared Spectra. In *NIST Chemistry WebBook*; Linstrom, P. J., Mallard, W. G., Eds.; NIST Standard Reference Database 69; National Institute of Standards and Technology: Gaithersburg, MD, 20899, March 2003 (<http://webbook.nist.gov>).

(42) Hyvärinen, A.; Oja, E. *Neural Comput.* **1997**, *9*, 1483–1492.

(43) Learned-Miller, E. G.; Fisher, J. W. *J. Machine Learn. Res.* **2003**, *4*, 1271–1295.

(39) Chen, B.; Kaufman, A. *Graph. Models* **2000**, *62*, 308–322.

(40) Toffoli, T.; Quick, J. *Graph. Model. Im. Proc.* **1997**, *59*, 89–95.




Figure 2. Statistics of performance of MI-based algorithms (MILCA, MILCA+ALS, SNICA) at decomposing 3-, 4-, 7-, and 10-component mixtures.

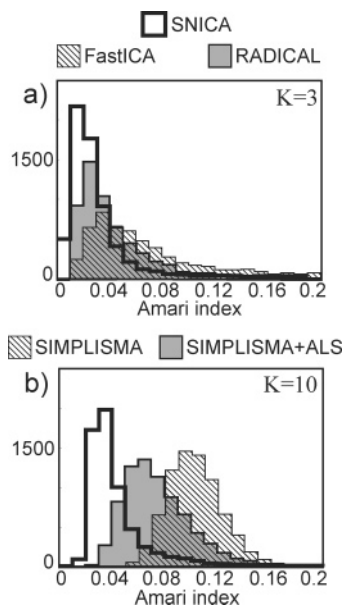


Figure 3. Performance of SNICA in comparison to traditional ICA (a) and MCR (b) algorithms (3- and 10-component mixtures, respectively).

without augmentation of data.) RADICAL is similar to MILCA in many respects, notably in speed and performance,⁶ but uses a nearest-neighbor estimator for entropies $H(Y)$,⁴⁴ not for MI. We observe that the performance of SNICA is superior to these two ICA techniques (Figure 3a). In particular, a significant fraction ($\sim 10\%$) of decompositions by FastICA have $P > 0.2$ (not shown), indicating relatively low performance.

A similar comparison was made with a state-of-the-art curve resolution algorithm SIMPLISMA²¹ with and without ALS corrections (Figure 3b). (We used the SIMPLISMA code given in ref 18, choosing the “offset” parameter randomly from the interval $[0.01, 0.1]$, which is about 1–10% of the standard deviation of the mixed signals. This is consistent with the choice made in ref 24.) We see that ALS significantly improves SIMPLISMA decomposi-




Figure 4. Five hundred SNICA decompositions ($K = 3$) represented by the gray points in coordinates $(I(S), P)$. For 10 selected decompositions, the one standard deviation error bars of estimated $I(Y)$ and Amari indices P were computed from 250 to 500 runs with randomized method parameters in each case (see text).

tions, but there is still a remarkable gap in performance between SIMPLISMA+ALS and SNICA. Figure 3b shows the results for a high dimensional test set ($K = 10$). We should mention, however, that the advantage of SNICA over SIMPLISMA+ALS is less pronounced for smaller values of K , but it is still significant.

Interestingly, decomposing mixtures of weakly and highly dependent sources by SNICA turns out to be almost equally efficient (Figure 4, where the mutual information scale spans the whole range of $I(S)$ for triples of spectra in our test set). This is in contrast to traditional ICA methods whose inaccuracies correlate with the dependencies between pure sources, making decompositions of overlapping spectra particularly difficult.⁷

To examine the sensitivity of the proposed technique to its parameters, we gathered the statistics of 10 decompositions with varied annealing temperature T_1 , initial Monte Carlo step size h_0 and stopping criteria M_1, M_2 . Each decomposition was attempted with these parameters taken uniformly randomly from their feasible ranges for $K = 3$: $T_1 \in [0.01, 0.03]$, $h_0 \in [0.1, 0.8]$, $M_1 \in [500, 1500]$, $M_2 \in [0, 1000]$ (the “zero” temperature at the second annealing step was kept fixed at $T_2 = 10^{-7}$). As reflected by the error bars on Figure 4, the performance of SNICA as measured by the Amari index is rather robust with respect to the parameter choices, suggesting that Figures 2 and 3 are also robust. The deviations from the mean final values of the cost function $I(Y)$ are also small for all decompositions indicating that in all cases the global minima were reached. We note that these deviations may be simply due to statistical errors in the MI estimation.³⁴

A separate set of high-statistics simulations with SNICA (for $K = 3, 4$) was performed, in which the MI estimator of ref 34 was replaced by that of ref 45 leaving the rest of the SNICA code, all parameters, and data unchanged. (We used the implementation by P. Tichavsky available at http://www.utia.cas.cz/user_data/scientific/SI_dept/Tichavsky.html.) These estimators were found to be comparable in speed and efficiency when used in ICA on spectral mixtures. However, while the statistics of the Amari index (denoted P_K and P_D) were very similar for three-component mixtures, it seems that the estimator of ref 34 performs better in higher dimensions. Notably, in 70% of the cases, the decompositions of four-component mixtures using this estimator were better ($P_K < P_D$), with the average values $\langle P_K \rangle = 0.022$ and $\langle P_D \rangle = 0.035$.

Now we turn to an experimental mixture problem to test SNICA in a realistic application of spectral mixture decomposition.

(44) Vasicek, O. *J. R. Stat. Soc. B Met.* **1976**, *38*, 54–59.

(45) Darbellay, G. A.; Vajda, I. *IEEE T. Inf. Theory* **1999**, *45*, 1315–1321.

Table 1. Decomposition of Six-Component Experimental Mixtures by MCR and ICA Methods^a

	SIMPLISMA (+ ALS)	BTEM	MILCA (+ ALS)	SNICA
toluene	0.971(0.973)	0.954	0.987(0.994)	0.973
<i>n</i> -hexane	0.994(0.995)	0.992	0.990(0.991)	0.981
acetone	0.866(0.899)	0.886	0.933(0.943)	0.972
aldehyde	0.943(0.953)	0.899	0.901(0.902)	0.928
33DMB	0.576(0.963)	0.983	0.964(0.948)	0.953
DCM	0.969(0.967)	0.904	0.909(0.966)	0.964
average	0.887(0.958)	0.936	0.947(0.957)	0.962

^a The entries are normalized inner products between recovered and pure spectra of mixture constituents and their average values. Numerical data for SIMPLISMA and BTEM are from an independent study (ref 24, Table 3, which gives comparison with some other MCR methods—IPCA,²² OPA-ALS⁴⁶).

To facilitate comparison, the example we chose was the one of ref 24 where the performances of several well-established curve resolution methods were compared. The experimental data intended for blind separation consisted of FT-IR spectra (950–3200 cm⁻¹ at 5626 wavelengths) of 14 mixtures of 6 solvents (toluene, *n*-hexane, acetone, 3-phenylpropionaldehyde (aldehyde), 3,3-dimethylbut-1-ene (33DMB), dichloromethane (DCM)). In addition, the pure spectra of mixture constituents were measured in separate experiments. Thus, this mixture problem offered a test case complicated by experimental features typical to analytical practice, such as nonlinearities, instrumental noise, and the presence of spectrometer background and impurities (residual water and CO₂ in this case²⁴).

First, we performed decomposition by a PCA-based ICA algorithm (MILCA, with and without ALS postprocessing) reducing the dimensionality of the data from 14 to 6 on the prewhitening step. The estimated components were then compared to the pure sources, computing the normalized inner products between them.^{7,24} The results are given in Table 1, which also contains the respective values obtained by two MCR methods—SIMPLISMA²¹ and BTEM.²⁴ SNICA was applied to this data set with a three-step annealing schedule taking $T_1 = 1.5$, $T_2 = 0.1$, $T_3 = 10^{-7}$, and $M = 6000$ for all three steps. The Monte Carlo decomposition was done in the full 14 dimensional space (using second-order Savitzky–Golay smoothing differentiation⁴⁷ with a fifth-order polynomial and a 51 point window); thus 14 components were resolved. Then, we selected only six components that had the largest average contributions to the mixed signals (this is similar to retaining only the components with largest eigenvalues in the PCA dimension reduction). These six most intense components were compared to the pure sources (Table 1). On this data set, the estimates produced by SNICA are more accurate than those obtained by SIMPLISMA and BTEM without ALS and compare favorably with the results obtained by MILCA and ALS versions of these algorithms. Notice that in order to produce accurate estimates, SIMPLISMA needs ALS postprocessing corrections^{24,28} (see also Figure 3b). With these corrections, its performance seems to match closely with SNICA, although higher statistics simulations (Figure 3b) suggest that this might not be

(46) Sánchez, F. C.; Toft, J.; van den Bogaert, B.; Massart, D. L. *Anal. Chem.* **1996**, *68*, 79–85.

(47) Savitzky, A.; Golay, M. J. E. *Anal. Chem.* **1964**, *36*, 1627–1639.

typical. Also, SNICA resolved all six components equally well, whereas the other methods produced two to three relatively poorly estimated spectra (notably, acetone and aldehyde appeared to be the most problematic due to their highly overlapping spectra).

CONCLUSIONS

In this paper, we have proposed a new stochastic technique, SNICA, for blind recovery of non-negative well-grounded sources from their linear mixtures. Avoiding the PCA decorrelation as being counterproductive in certain cases, our algorithm is based on a Metropolis Monte Carlo constrained minimization of mutual information between recovered sources. The demixing transformation is obtained as a sequence of random shears and rotations that resolves the least dependent strictly non-negative components.

The SNICA approach can be related to some other PCA-free ICA methods (e.g., Infomax⁴⁸ or TDSEP⁴⁹) in that it does not resort to the somewhat arbitrary decorrelation constraint^{50,51} introduced by PCA. Instead, it is replaced in SNICA by the non-negativity constraint which reflects a natural property of spectral signals in the present application. PCA simplifies and speeds up the decomposition considerably, but should be taken with caution. If the sources are strongly correlated, then PCA produces components that, despite being uncorrelated, are completely unphysical. If an ICA method (consisting of a rotation of prewhitened components) is used without ALS postprocessing, it has no chance to undo the errors made by PCA prewhitening. For a specific example, see ref 7. ALS postprocessing improves the situation, but only partially.

Here we have tested SNICA in spectral curve resolution on a representative set of mixture problems and experimental data. Our results indicate that the algorithm is superior in performance to its PCA-based counterpart (MILCA, with the same mutual information estimator) including the ALS-assisted version. We also found our algorithm to outperform SIMPLISMA+ALS, which represents a good reference point among the state-of-the-art curve resolution chemometrics techniques. The method presented in this paper performs blind decomposition, but the use of additional a priori information about pure components (such as, for example, sparseness of spectral signals) might possibly improve the performance further. In blind separation of multivariate time sequences, the performance of MILCA was greatly enhanced if mutual information at different time lags was also included in the minimization.⁶ In principle, the same can be done for SNICA, but we have not yet made any tests to see how practical this is.

The assumptions of non-negativity and statistical independence of source signals are met in a variety of practical problems in analytical spectroscopy and beyond. The range of applications of blind non-negative spectral mixture decomposition spans almost all quantitative analytical approaches that rely on spectroscopic measurements and subsequent mixture analysis, including, for example, hyperspectral remote imaging,^{52,53} microscopy of nano-

(48) Bell, A. J.; Sejnowski, T. J. *Neural Comput.* **1995**, *7*, 1129–1159.

(49) Ziehe, A.; Laskov, P.; Nolte, G.; Müller, K. R. *J. Machine Learn. Res.* **2004**, *5*, 777–800.

(50) Cardoso, J.-F. *J. Machine Learn. Res.* **2003**, *4*, 1177–1203.

(51) Naanaa, W.; Nuzillard, J. M. *Signal Process.* **2005**, *85*, 1711–1722.

(52) Nascimento, J. M. P.; Dias, J. M. B. *IEEE T. Geosci. Remote Sens.* **2005**, *43*, 175–187.

(53) Yuen, P. W. T.; Blagg, A.; Bishop, G. *Proc. SPIE* **2005**, *5988*, 107–118.

structures⁵⁴ and biomedical samples,^{7,55} and separation of independent light sources in multifrequency astrophysical observational data.⁵⁶ SNICA can be applied to all of them.

The source codes of SNICA, MILCA, and the MI estimator are freely available online at <http://www.fz-juelich.de/nic/cs/software>.

(54) Lidke, K. A.; Rieger, B.; Jovin, T. M.; Heintzmann, R. *Opt. Express* **2005**, *13*, 7052–7062.

(55) Krafft, C.; Sobottka, S. B.; Schackert G.; Salzer, R. *Analyst* **2005**, *130*, 1070–1077.

(56) Maino, D.; Farusi, A.; Baccigalupi, C.; Perrotta, F.; Banday, A. J.; Bedini, L.; Burigana, C.; De Zotti, G.; Górski, K. M.; Salerno, E. *Mon. Not. R. Astron. Soc.* **2002**, *334*, 53–68.

ACKNOWLEDGMENT

We thank Prof. M. Garland for providing us with the data set of ref 24. S.A.A. is grateful to Dr. Y. Alaverdyan and Prof. S. P. Mushtakova for discussions. We also thank the anonymous reviewers for their useful suggestions.

Received for review September 23, 2005. Accepted December 8, 2005.

AC051707C