Joint Diagonalization on the Oblique Manifold for Independent Component Analysis*

P.-A. Absil[†‡]   K. A. Gallivan[§]

24 January 2006

Abstract

Several blind source separation algorithms obtain a separating matrix by computing the congruence transformation that "best" diagonalizes a collection of covariance matrices. Recent methods avoid a pre-whitening phase and directly attempt to compute a non-orthogonal diagonalizing congruence. However, since the magnitude of the sources is unknown, there is a fundamental indeterminacy on the norm of the rows of the separating matrix. We show how this indeterminacy can be taken into account by restricting the separating matrix to the oblique manifold. The geometry of this manifold is developed and a trust-region-based algorithm for non-orthogonal joint diagonalization is proposed.

1 Introduction

Assume that n measured signals $x(t) = [x_1(t), \ldots, x_n(t)]^T$ are instantaneous linear mixtures of p underlying, statistically independent source signals $s(t) = [s_1(t), \ldots, s_p(t)]^T$; this can be compactly written as

$$x(t) = As(t),$$

where the matrix A is an unknown constant *mixing matrix* containing the mixture coefficients. We assume throughout that all vectors and matrices are real, and we let the superscript T denote the matrix transpose. The problem of independent component analysis (ICA) or blind source separation (BSS) is to identify the mixing matrix A or recover the source signals $s(t)$, using only the observed signals $x(t)$.

This problem is usually translated to finding a *separating matrix* (or *demixing matrix*) W such that the signals $y(t)$ given by

$$y(t) = W^T x(t)$$

are estimates of the signals $s(t)$. It is known that this problem has two basic indeterminacies: without any further information, it is impossible to recover the scaling and the order

of the source signals [Com94]. For non-Gaussian sources, these are the only indeterminacies [TLSH91]; this also holds under mild conditions for sources that are not temporally white [BAMCM97]. We will say that W is a *true* separating matrix if $W^T A$ can be expressed as the product of a permutation matrix (order indeterminacy) and a diagonal matrix (scaling indeterminacy).

In several ICA algorithms, the observed sources $x(t)$ are used to construct a set of "target matrices" $C_1, \dots, C_N$ with the following property: all the matrices $W^T C_i W$, $i = 1, \dots, N$ are diagonal if and only if W is a true separating matrix. In practice, due to the presence of noise and to the limited amount of samples of $x(t)$ available, the target set $\{C_1, \dots, C_N\}$ does not admit exact joint diagonalization (JD), and one must resort to *approximate joint diagonalization*, that is, find the matrix W that "best diagonalizes" the target set. The various JD-based ICA algorithms differ in the choice of the target matrices and in the cost function used to define the "best diagonalization". Several possibilities for the choice of the target matrices are mentioned in [Yer02], and [WSC05] lists a few possible cost functions (see also [AK05]). A frequently encountered cost function is

$$f(W) = \sum_i \|\text{off}(W^T C_i W)\|_F^2; \tag{1}$$

here $\|M\|_F^2$ denotes the square Frobenius norm of M (that is, the sum of the squares of the elements of M) and $\text{off}(M) := M - \text{ddiag}(M)$, where $\text{ddiag}(M)$ denotes the diagonal matrix whose diagonal elements are those of M. Following the notation in [JR02], we let $\text{diag}(M)$ denote the vector of diagonal elements of M.

Most joint diagonalization algorithms, such as the SOBI algorithm [BAMCM97], start with a *pre-whitening* step. First, in order to remove the scaling indeterminacy, it is assumed that $E[s(t)s^T(t)] = I$. (This is without loss of generality, since the scaling factors can be absorbed in the columns of A.) A *whitening matrix* $\tilde{W}$ is then sought such that one of the target matrices (say C_1, usually an estimation of the covariance matrix $E[x(t)x^T(t)]$) is reduced to the identity matrix; that is, $\tilde{W}^T C_1 \tilde{W} = I_p$. It follows that there exists an orthogonal matrix U such that $U^T \tilde{W}^T A = I$. In a second step, an orthogonal matrix U is sought that diagonalizes the new target set $\{\tilde{W}^T C_1 \tilde{W}, \dots, \tilde{W}^T C_1 \tilde{W}\}$. This yields a separating matrix $W = \tilde{W}U$; see, for example, [BAMCM97] for details.

Since U is constrained to be orthogonal, it is a solution of an optimization problem on a manifold—the orthogonal group, or more generally the compact Stiefel manifold of matrices with orthonormal columns. This calls for the use of differential-geometric optimization techniques. There has been interest for optimization on manifolds at least since the work of Luenberger [Lue72] and Gabay [Gab82] in the 1970s and 1980s. Optimization algorithms, such as the steepest descent, Newton and conjugate gradient methods, were first proposed on general Riemannian manifolds by Smith [Smi93, Smi94] and Udrişte [Udr94]. Various optimization algorithms on several types of manifolds were further proposed and analysed in several papers; see in particular [HM94, Deh95, EAS98, Man02, MM02, AMS04, ABG04, HT04, ABG05, NA05] and references therein. Applications of differential-geometric methods to the orthogonal joint diagonalization problem have been proposed by Rahbar and Reilly [RR00], Douglas [Dou00], Joho and Mathis [JM02], Joho and Rahbar [JR02], Nikpour *et al.* [NMH02], Nishimori and Akaho [NA05].

However, the pre-whitening step that yields the orthogonality constraint has the drawback that it singles out the target matrix C_1 of which it would attain exact diagonalization at the possible cost of poor diagonalization of the other target matrices [Yer02]. Moreover,

inaccuracies in the computation of $\tilde{W}$ cannot be compensated in the sequel. Therefore, a few algorithms have been proposed that directly compute a nonorthogonal separating matrix W without resorting to a pre-whitening process; see for example [ACC00, Pha01, Yer02, AK04].

Amari *et al.* [ACC00] and Afsari and Krishnaprasad [AK04] use differential-geometric concepts for nonorthogonal JD. As was pointed out in [AK04], when the off-diagonal cost function (1) is allowed to take its argument W in the whole set $\mathbb{R}^{n \times p}$ of $n \times p$ matrices, it admits a global minimizer at the zero matrix. More generally, the cost function (1) is not scale invariant; that is, $f(WD)$ is in general different from $f(W)$ when D is a nonsingular diagonal matrix. This issue can be tackled by imposing constraints on the power of the separated signals $y(t)$. However, as argued by Amari *et al.* [ACC00], this is impractical in frequently encountered applications where the amplitude of components may change suddenly; therefore, constraints should be placed instead on the separating matrix W. Amari *et al.* [ACC00] (see also Afsari and Krishnaprasad [AK04]) propose to constrain the allowed variations of W to belong to a subspace *orthogonal* to the equivalence class $\{WD : D \text{ diagonal}\}$. The notion of orthogonality used in [ACC00] has been shown to be a nonholonomic (or nonintegrable) constraint.

Finally, we point out that most algorithms for ICA do not achieve superlinear convergence, as they are based on steepest-descent or direct-search ideas. Exceptions are the conjugate gradient on the Stiefel manifold used by Rahbar and Reilly [RR00] and the Newton method on Stiefel of Joho and Rahbar [JR02] and Nikpour *et al.* [NMH02]. Superlinear convergence is useful when a high precision is sought; this applies to situations where the noise level is low and the source signals (or the mixing matrix) want to be recovered accurately.

In this paper, we propose a superlinearly convergent algorithm for nonorthogonal joint diagonalization, based on a recently proposed trust-region method on Riemannian manifolds [ABG04, ABG05]; we dub the algorithm RTR-ICA. In comparison with the Newton method, the trust-region approach offers better global convergence properties and similar local convergence properties at a lower computational cost [ABG05]. Our approach also departs from previous work in the way constraints are imposed on the separating matrix W: we require that W be an *oblique rotation* [TL02], that is, all the columns of W have unit Euclidean norm. Instead of being nonholonomic, this constraint defines a submanifold of $\mathbb{R}^{n \times p}$ called the *oblique manifold*

$$\mathcal{OB}(n, p) = \{Y \in \mathbb{R}^{n \times p} : \mathrm{ddiag}(Y^T Y) = I_p\}. \tag{2}$$

Moreover, in contrast to the Stiefel manifold approach, a pre-whitening step is not required.

The rest of the paper is organized as follows. The geometry of the oblique manifold is described in Section 2. Formulas for the gradient and Hessian of the off-diagonal cost function (1) are obtained in Section 3. The workings of the RTR algorithm are briefly explained in Section 4 (we refer to [ABG05] for details). Numerical experiments are presented in Section 5.

2 Geometry of the Oblique manifold

We refer the reader to [ABG05] and references therein for the relevant notions of Riemannian geometry. Let
$$\mathcal{OB} := \{Y \in \mathbb{R}^{n \times p} : \mathrm{ddiag}(Y^T Y) = I_p\}$$
denote the set of all $n \times p$ matrices with normalized columns. With its usual vector space structure, $\mathbb{R}^{n \times p}$ is trivially a manifold. It is possible to show that $\mathcal{OB}$ is a *regular submanifold*

(or *embedded submanifold*) of $\mathbb{R}^{n \times p}$, that is, there exists a differentiable structure on $\mathcal{OB}$ that makes the inclusion mapping $i : \mathcal{OB} \to \mathbb{R}^{n \times p}$ an immersion. The theory of submanifolds states that such a differentiable structure is unique.

Proposition 2.1 *The set $\mathcal{OB}$ is a regular submanifold of $\mathbb{R}^{n \times p}$.*

Proof. Consider the function $F : \mathbb{R}^{n \times p} \to \mathbb{R}^p : Y \mapsto \mathrm{diag}(Y^T Y) - e$, where e is the vector of all ones. Notice that $\mathcal{OB} = F^{-1}(0)$ and that F is full rank on $\mathcal{OB}$. The result follows from a classical characterization of regular submanifolds; see [Boo75, Th. III.5.8]. $\qquad\square$

We consider the canonical inner product

$$\langle Z_1, Z_2 \rangle := \mathrm{trace}(Z_1^T Z_2) \tag{3}$$

in $\mathbb{R}^{n \times p}$ and view $\mathcal{OB}$ as an embedded Riemannian submanifold of $\mathbb{R}^{n \times p}$. Notice that $\mathcal{OB}$ is then isometric to the product of p unit spheres in $\mathbb{R}^n$ with its standard Riemannian structure.

The tangent space to $\mathcal{OB}$ at Y (which is defined independently of the metric) is

$$T_Y \mathcal{OB} = \{Z : \mathrm{ddiag}(Y^T Z) = 0\}$$

which means that $y_i^T z_i = 0$, $i = 1, \ldots, p$, where y_i denotes the ith column of Y. The dimension of $\mathcal{OB}$ is $\dim(\mathcal{OB}) = p(n-1)$. The normal space (which depends on the embedding in the Euclidean space $\mathbb{R}^{n \times p}$) is

$$N_Y \mathcal{OB} = \{YD : D \in \mathbb{R}^{p \times p} \text{ diagonal}\}.$$

Projections of $Z \in T_Y \mathbb{R}^{n \times p}$ into a normal and tangent component are

$$P_{N_Y}(Z) = Y\,\mathrm{ddiag}(Y^T Z) \text{ and } P_{T_Y}(Z) = Z - Y\,\mathrm{ddiag}(Y^T Z).$$

Finally, in order to apply the RTR schemes on $\mathcal{OB}$, we must define a *retraction*, which establishes a correspondence between tangent vectors and points on the manifolds. A natural choice that satisfies the required properties [ABG05] is

$$R_Y(Z) = (Y + Z)(\mathrm{ddiag}((Y + Z)^T(Y + Z)))^{-1/2}, \tag{4}$$

which simply consists of adding Z to Y and scaling the columns of the result.

3 The off-diagonal cost function on the oblique manifold

We compute the gradient and the Hessian of $f = \tilde{f}|_{\mathcal{OB}}$ where $\tilde{f}$ is the function on $\mathbb{R}^{n \times p}$ defined as

$$\begin{aligned}
\tilde{f}(Y) &= \sum_{i=1}^{N} \|Y^T C_i Y - \mathrm{ddiag}(Y^T C_i Y)\|_F^2 \\
&= \mathrm{trace}((Y^T C_i Y)^2 - \mathrm{ddiag}^2(Y^T C_i Y)) \\
&= \sum_{i=1}^{N} \mathrm{trace}(\mathrm{off}(Y^T C_i Y)Y^T C_i Y),
\end{aligned}$$

C_i symmetric. Notice that f is a function on $\mathcal{OB}$; its gradient and Hessian are thus defined in the sense of the manifold $\mathcal{OB}$ endowed with its Riemannian metric (3).

The following identities will be useful:

$$\mathrm{ddiag}(AD) = \mathrm{ddiag}(A)D, \ D \in \mathcal{D}$$
$$\mathrm{trace}(A) = \mathrm{trace}(A^T)$$
$$\mathrm{trace}(AB) = \mathrm{trace}(BA)$$
$$\mathrm{trace}(aA) = a\mathrm{trace}(A)$$
$$\det(AB) = \det(A)\det(B).$$

The gradient of f at Y, denoted by $\mathrm{grad}\, f(Y)$, is defined as the element of $T_Y\mathcal{OB}$ such that $\langle \mathrm{grad}\, f(Y), \xi \rangle = \mathrm{D}f(Y)[\xi]$ for all ξ in $T_Y\mathcal{OB}$, where $\mathrm{D}f(Y)[\xi]$ denotes the directional derivative of f in the direction of ξ (or the differential of f at Y applied to ξ, which is often denoted as ξf). We obtain

$$\mathrm{D}f(Y)[\xi] = \sum_{i=1}^{N} \mathrm{trace}(\mathrm{off}(Z^T C_i Y)Y^T C_i Y) + \mathrm{trace}(\mathrm{off}(Y^T C_i Z)Y^T C_i Y)$$
$$+ \mathrm{trace}(\mathrm{off}(Y^T C_i Y)Z^T C_i Y) + \mathrm{trace}(\mathrm{off}(Y^T C_i Y)Y^T C_i Z)$$
$$= \sum_{i=1}^{N} 4\mathrm{trace}(Z^T C_i Y(\mathrm{off}(Y^T C_i Y)).$$

For the gradient of $\tilde{f}$, we thus get

$$\mathrm{grad}\, \tilde{f}(Y) = \sum_{i=1}^{N} 4C_i Y \mathrm{off}(Y^T C_i Y).$$

It is easy to check that $P_{T_Y}\mathrm{grad}\, \tilde{f}(Y)$, the projection of $\mathrm{grad}\, \tilde{f}(Y)$ onto $T_Y\mathcal{OB}$, satisfies the defining properties of $\mathrm{grad}\, f(Y)$, hence

$$\mathrm{grad}\, f(Y) = P_{T_Y}\mathrm{grad}\, \tilde{f}(Y)$$
$$= \sum_{i=1}^{N} 4C_i Y \mathrm{off}(Y^T C_i Y) - 4Y\mathrm{ddiag}(Y^T C_i Y\mathrm{off}(Y^T C_i Y)).$$

Following [dC92, ABG05], we define the Hessian of f at Y as the linear operator

$$\mathrm{Hess}\, f(Y) : T_Y\mathcal{OB} \to T_Y\mathcal{OB} : Z \mapsto \nabla_Z \mathrm{grad}\, f,$$

where ∇ denotes the *Riemannian connection* (also called *Levi-Civita connection*). Since $\mathcal{OB}$ is a Riemannian submanifold of the Euclidean space $\mathbb{R}^{n \times p}$, the Riemannian connection satisfies

$$\nabla_\eta \xi = P_{T_Y}\left(\mathrm{D}\xi(Y)[\eta]\right)$$

where Y is the foot of η. It follows that

$$\mathrm{Hess}\, f(Y)[Z] = P_{T_Y}\mathrm{Dgrad}\, f(Y)[Z]$$
$$= P_{T_Y}\mathrm{Dgrad}\, \tilde{f}(Y)[Z] - Z\mathrm{ddiag}(Y^T \mathrm{grad}\, \tilde{f}(Y))$$

with

$$\mathrm{Dgrad}\,\tilde{f}(Y)[Z] = \sum_{i=1}^{N} 4C_i Z \mathrm{off}(Y^T C_i Y) + 4C_i Y \mathrm{off}(Z^T C_i Y)$$
$$+ 4C_i Y \mathrm{off}(Y^T C_i Z).$$

4 The RTR-ICA approach

We now have the necessary ingredients to apply the Riemannian trust-region (RTR) approach to the problem of minimizing the cost function (1) on the oblique manifold $\mathcal{OB}$ endowed with the Riemannian metric (3) and the retraction (4). The general idea is as follows. The RTR scheme is an iterative process that, from a current iterate W on the oblique manifold, produces a next iterate W_+ on the oblique manifold. First, a model m_W of the cost function f is constructed around W; more precisely, $m_W(Z)$ approximates $f(R_W(Z))$. When not too expensive computationally (which is the case here), it is useful to choose m_W as the *Newton model*, i.e., second-order Taylor expansion of $f \circ R_W$:

$$m_W(Z) = f(W) + \langle \mathrm{grad}\, f(W), Z \rangle + \frac{1}{2} \langle \mathrm{Hess}\, f(W)[Z], Z \rangle.$$

Next, an (approximate) minimizer of the model is sought within a region where the model is "trusted" (hence the name of the method). The size of the trust-region has to be externally specified for the first iterate; it is subsequently automatically updated: if there is a good agreement between the cost function and the model at the proposed next iterate, then the proposed iterate is accepted and the size of the trust region is possibly increased. If the agreement is poor, the size of the trust region can be reduced and the proposed iterate can even be discarded. For more details on the trust-region concept, which originates from the work of Powell in the 1970s, we refer to [ABG05] and references therein. Finally, we point out that there are several ways to approximately solve trust-region subproblems, i.e., to compute an approximate minimizer of a quadratic model within a trust region. Usually, finding a high-precision solution of each subproblem is not necessary and would constitute a waste of computational effort; on the other hand, the approximate solution has to be sufficiently precise so that strong local and global convergence properties hold. To handle this tradeoff, the use of Steihaug's truncated CG method is advocated in [ABG05], and we use it in the numerical experiments reported on in the next section. We also refer to [ABG05] for a convergence analysis of the general RTR schemes.

We now describe the resulting algorithm in detail. The trust-region subproblems take the form

$$\min_{\eta \in T_{W_k} M} m_{W_k}(\eta) = f(W_k) + \langle \mathrm{grad}\, f(W_k), \eta \rangle_{W_k} + \frac{1}{2} \langle \mathcal{H}_{W_k} \eta, \eta \rangle_{W_k} \quad \text{s.t.} \ \langle \eta, \eta \rangle_{W_k} \le \Delta_k^2 \quad (5)$$

with $\mathcal{H}_Y := \mathrm{Hess}\, f(Y)$, and we have the following *outer iteration*.

Algorithm 1 (RTR-ICA) *Parameters:* $\bar{\Delta} > 0$, $\Delta_0 \in (0, \bar{\Delta})$, *and* $\rho' \in [0, \frac{1}{4})$.
Input: initial iterate $W_0 \in \mathcal{OB}$.
Output: sequence of iterates $\{W_k\}$.
for $k = 0, 1, 2, \ldots$

Obtain η_k by (approximately) solving (5);

Evaluate $\rho_k := \frac{f(W_k) - f(R_{W_k}(\eta_k))}{m_{W_k}(0_{W_k}) - m_{W_k}(\eta_k)}$.

if $\rho_k < \frac{1}{4}$
    $\Delta_{k+1} = \frac{1}{4}\Delta_k$
else if $\rho_k > \frac{3}{4}$ *and* $\|\eta_k\| = \Delta_k$
    $\Delta_{k+1} = \min(2\Delta_k, \bar{\Delta})$
else
    $\Delta_{k+1} = \Delta_k$;
if $\rho_k > \rho'$
    $W_{k+1} = R_{W_k}\eta_k$
else
    $W_{k+1} = W_k$;
end (for).

In order to approximately solve the trust-region subproblems (5), we suggest using the truncated CG algorithm of Steihaug [Ste83] and Toint [Toi81].

Algorithm 2 (tCG) *Set $\eta^0 = 0$, $r_0 = \operatorname{grad} f(W_k)$, $\delta_0 = -r_0$;*

for $j = 0, 1, 2, \dots$ *until a stopping criterion is satisfied, perform the iteration:*
    if $\langle \delta_j, \mathcal{H}_{W_k}\delta_j \rangle_{W_k} \leq 0$
        Compute τ such that $\eta = \eta^j + \tau\delta_j$ minimizes $m_{W_k}(\eta)$ in (5)
          and satisfies $\|\eta\| = \Delta$;
        return η;
    Set $\alpha_j = \langle r_j, r_j \rangle_{W_k} / \langle \delta_j, \mathcal{H}_{W_k}\delta_j \rangle_{W_k}$;
    Set $\eta^{j+1} = \eta^j + \alpha_j\delta_j$;
    if $\|\eta^{j+1}\| \geq \Delta$
        Compute $\tau \geq 0$ such that $\eta = \eta^j + \tau\delta_j$ satisfies $\|\eta\| = \Delta$;
        return η;
    Set $r_{j+1} = r_j + \alpha_j\mathcal{H}_{W_k}\delta_j$;
    Set $\beta_{j+1} = \langle r_{j+1}, r_{j+1} \rangle_{W_k} / \langle r_j, r_j \rangle_{W_k}$;
    Set $\delta_{j+1} = -r_{j+1} + \beta_{j+1}\delta_j$;
end (for).

With a view towards achieving superlinear convergence, we stop as soon as an iteration j is reached for which

$$\|r_j\| \leq \|r_0\| \min(\|r_0\|^\theta, \kappa). \tag{6}$$

Concerning the computation of τ, it can be shown that when $\langle \delta_j, \mathcal{H}_{W_k}\delta_j \rangle \leq 0$, $\arg\min_{\tau \in \mathbb{R}} m_k(\eta^j + \tau\delta_j)$ is equal to the positive root of $\|\eta^j + \tau\delta_j\| = \Delta$, which is explicitly given by

$$\frac{-\langle \eta^j, \delta_j \rangle_x + \sqrt{\langle \eta^j, \delta_j \rangle_x^2 - (\Delta^2 - \langle \eta^j, \eta^j \rangle_x)\langle \delta_j, \delta_j \rangle_x}}{\langle \delta_j, \delta_j \rangle_x}.$$

5 Simulation results

The following simulation follows closely the one in [RR00]; we simply increased the number of data points and modified some of the signals to improve the spatial uncorrelatedness of the sources. We use four synthetic signals with 10^6 sample points, whose 100 first points

are shown on Figure 1. The sources are mixed using a four-by-four randomly generated mixing matrix A (Figure 2). Three target matrices C_1, C_2, C_3 are chosen as lagged sample covariance matrices of $x(t)$. The iteration is initialized with W equal to the identity matrix. The performance of separation displayed in Figure 4 is measured using the formula

$$ P_{\text{index}} = 20 \log_{10} \left(\frac{1}{n} \left(\sum_{i=1}^{n} \left(\sum_{j=1}^{n} \frac{|q_{ij}|}{\max_\ell |q_{i\ell}|} - 1 \right) \right) \right), $$

where q_{ij} is the (i,j)th element of $Q := W^T A$. The source signals are recovered precisely, modulo the scaling and ordering indeterminacies (Figure 3).

To illustrate the benefit of taking the model m_W as the Newton model, we applied to the same problem the linearly convergent method obtained by defining m_W as the first-order Taylor expansion; the algorithm then took more than 6000 (inner) iterations to only reach a performance index of -100.


Figure 1: The four source signals.

With a view towards a comparison with the simulation results in [RR00, JR02], notice that one iteration of nonlinear CG involves a line-search process with possibly several evaluations of the cost function, and that an iteration of Newton's method involves solving the Newton equation. In contrast, in the truncated CG process, the major work consists in one application of the Hessian operator per inner iteration.

6 Discussion and future work

We have proposed a Riemannian trust-region method on the oblique manifold to perform joint diagonalization with superlinear convergence. Numerical experiments have shown the efficacy of the method for high accuracy blind source separation with few sources and mixtures. The observed behaviour, however, became less satisfactory as the number of sources and mixtures grew. In future work, we will study the properties of the off-diagonal cost function (1) and other cost functions, and investigate the numerical stability of the truncated CG inner iteration.

Figure 2: The four mixed signals.

Acknowledgement

Special thanks to Bijan Afsari for helpful discussions about ICA and joint diagonalization.

References

[ABG04]   P.-A. Absil, C. G. Baker, and K. A. Gallivan, *Trust-region methods on Riemannian manifolds with applications in numerical linear algebra*, Proceedings of the 16th International Symposium on Mathematical Theory of Networks and Systems (MTNS2004), Leuven, Belgium, 5–9 July 2004, 2004.

[ABG05]   _____, *Trust-region methods on Riemannian manifolds*, http://www.inma.ucl.ac.be/~absil/, submitted, 2005.

[ACC00]   Shun-ichi Amari, Tian-Ping Chen, and Andrzej Cichocki, *Nonholonomic orthogonal learning algorithms for blind source separation*, Neural Computation **12** (2000), 1463–1484.

[AG06]   P.-A. Absil and K. A. Gallivan, *Joint diagonalization on the oblique manifold for independent component analysis*, Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal processing (ICASSP), to appear, 2006.

[AK04]   Bijan Afsari and P. S. Krishnaprasad, *Some gradient based joint diagonalization methods for ICA*, Proceedings of the 5th International Conference on Independent Component Analysis and Blind Source Separation (Springer LCNS Series, ed.), 2004.

[AK05]   _____, *A novel non-orthogonal joint diagonalization cost function for ica*, Tech. Report TR 2005-106, Institute for Systems Research, University of Maryland, http://www.isr.umd.edu, 2005.

[AMS04]   P.-A. Absil, R. Mahony, and R. Sepulchre, *Riemannian geometry of Grassmann manifolds with a view on algorithmic computation*, Acta Appl. Math. **80** (2004), no. 2, 199–220.

[BAMCM97]   Adel Belouchrani, Karim Abed-Meraim, Jean-François Cardoso, and Eric Moulines, *A blind source separation technique using secton-order statistics*, IEEE Trans. Signal Process. **45** (1997), no. 2, 434–444.

[Boo75]   W. M. Boothby, *An introduction to differentiable manifolds and Riemannian geometry*, Academic Press, 1975.

Figure 3: The four recovered signals.

[Com94]   Pierre Comon, *Independent component analysis, A new concept?*, Signal Processing **36** (1994), 287–314.

[dC92]   M. P. do Carmo, *Riemannian geometry*, Mathematics: Theory & Applications, Birkhäuser Boston Inc., Boston, MA, 1992, Translated from the second Portuguese edition by Francis Flaherty.

[Deh95]   J. Dehaene, *Continuous-time matrix algorithms, systolic algorithms and adaptive neural networks*, Ph.D. thesis, Katholieke Universiteit Leuven, Faculteit Toegepaste Wetenschappen, Departement elektrotechniek-ESAT, Kard. Mercierlaan 94, 3001 Leuven, Belgium, October 1995, accessible from ftp://ftp.esat.kuleuven.ac.be/pub/SISTA/dehaene/phd/.

[Dou00]   Scott C. Douglas, *Self-stabilized gradient algorithms for blind source separation with orthogonality constraints*, IEEE Trans. Neural Networks **11** (2000), no. 6, 1490–1497.

[EAS98]   A. Edelman, T. A. Arias, and S. T. Smith, *The geometry of algorithms with orthogonality constraints*, SIAM J. Matrix Anal. Appl. **20** (1998), no. 2, 303–353.

[Gab82]   D. Gabay, *Minimizing a differentiable function over a differential manifold*, Journal of Optimization Theory and Applications **37** (1982), no. 2, 177–219.

[HM94]   U. Helmke and J. B. Moore, *Optimization and dynamical systems*, Springer, 1994.

[HT04]   Knut Hüper and Jochen Trumpf, *Newton-like methods for numerical optimization on manifolds*, Proc. 38th IEEE Asilomar Conference on Signals, Systems, and Computers, Pacific Grove, CA, November 7-10, 2004, 2004.

[JM02]   Marcel Joho and Heinz Mathis, *Joint diagonalization of correlation matrices by using gradient methods with application to blind signal separation*, Proceedings of IEEE Sensor Array and Multichannel Signal Processing Workshop SAM, 2002, pp. 273–277.

[JR02]   Marcel Joho and Kamran Rahbar, *Joint diagonalization of correlation matrices by using Newton methods with applications to blind signal separation*, Proceedings of IEEE Sensor Array and Multichannel Signal Processing Workshop SAM, 2002, pp. 403–407.

[Lue72]   David G. Luenberger, *The gradient projection method along geodesics*, Management Sci. **18** (1972), 620–631.

[Man02]   J. H. Manton, *Optimization algorithms exploiting unitary constraints*, IEEE Trans. Signal Process. **50** (2002), no. 3, 635–650.

[MM02]   R. Mahony and J. H. Manton, *The geometry of the Newton method on non-compact Lie groups*, J. Global Optim. **23** (2002), no. 3, 309–327.
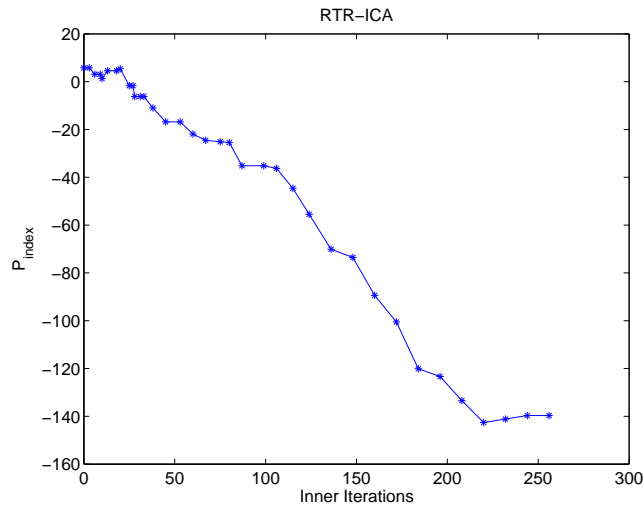
Figure 4: Convergence of RTR-ICA.

[NA05]   Yasunori Nishimori and Shotaro Akaho, *Learning algorithms utilizing quasi-geodesic flows on the stiefel manifold*, Neurocomputing **67** (2005), 106–135.

[NMH02]   Maziar Nikpour, Jonathan H. Manton, and Gen Hori, *Algorithms on the Stiefel manifold for joint diagonalization*, Proc. ICASSP, 2002, pp. II–1481–1484.

[Pha01]   Dinh Tuan Pham, *Joint approximate diagonalization of positive definite Hermitian matrices*, SIAM J. Matrix Anal. Appl. **22** (2001), no. 4, 1136–1152.

[RR00]   Kamran Rahbar and James P. Reilly, *Geometric optimization methods for blind source separation of signals*, International Conference on Independent Component Analysis ICA2000,Helsinki, Finland, June 2000.

[Smi93]   S. T. Smith, *Geometric optimization methods for adaptive filtering*, Ph.D. thesis, Division of Applied Sciences, Harvard University, Cambridge, Massachusetts, 1993.

[Smi94]   Steven T. Smith, *Optimization techniques on Riemannian manifolds*, Hamiltonian and gradient flows, algorithms and control, Fields Inst. Commun., vol. 3, Amer. Math. Soc., Providence, RI, 1994, pp. 113–136.

[Ste83]   T. Steihaug, *The conjugate gradient method and trust regions in large scale optimization*, SIAM J. Numer. Anal. **20** (1983), 626–637.

[TL02]   N. T. Trendafilov and R. A. Lippert, *The multimode Procrustes problem*, Linear Algebra Appl. **349** (2002), 245–264.

[TLSH91]   L. Tong, R. W. Liu, V. C. Soon, and Y. F. Huang, *Indeterminacy and identifiability in blind identification*, IEEE Trans. Circuits Syst. **38** (1991), 499–509.

[Toi81]   Ph. L. Toint, *Towards an efficient sparsity exploiting Newton method for minimization*, Sparse Matrices and Their Uses (I. S. Duff, ed.), Academic Press, London, 1981, pp. 57–88.

[Udr94]   C. Udrişte, *Convex functions and optimization methods on Riemannian manifolds*, Kluwer Academic Publishers, 1994.

[WSC05]   Wenwu Wang, Saeid Sanei, and Jonathon A. Chambers, *Penalty function-based joint diagonalization approach for convolutive blind separation of nonstationary sources*, IEEE Trans. Signal Process. **53** (2005), no. 5, 1654–1669.

[Yer02]   Arie Yeredor, *Non-orthogonal joint diagonalization in the least-squares sense with application in blind source separation*, IEEE Trans. Signal Process. **50** (2002), no. 7, 1545–1553.