WP04 5:10

Proceedings of the 34th
Conference on Decision & Control
New Orleans, LA - December 1995

Continuous-time algorithms for the Riemannian SVD *

Bart De Moor†

ESAT-SISTA, Dept. of Electrical Engineering, Katholieke Universiteit Leuven
Kardinaal Mercierlaan 94, B-3001 Leuven, Belgium.
tel:32/16/321715, fax:32/16/321986, email:bart.demoor@esat.kuleuven.ac.be

Abstract We define a nonlinear generalization of the singular value decomposition (SVD), which can be interpreted as a restricted SVD with Riemannian metrics in the column and row space. This so-called Riemannian SVD occurs in structured total least squares problems, for instance in the least squares approximation of a given matrix A by a rank deficient Hankel matrix B, which is an important problem in system identification and signal processing. Several algorithms to find the 'minimizing' singular triplet are suggested, both for the SVD and its nonlinear generalization. This paper reveals interesting connections between linear algebra (structured matrix problems), numerical analysis (algorithms), optimization theory, (differential) geometry and system theory (differential equations, stability, Lyapunov functions). We give some numerical examples and also point out some open problems.

1. Introduction

Since the work by Eckart-Young [12], we know how to obtain the best rank deficient least squares approximation of a given matrix $A \in \mathbf{R}^{p\times q}$ of full column rank q. This approximation follows from the SVD of A by subtracting from it the rank one matrix $u.\sigma.v^T$, where (u,σ,v) is the singular triplet corresponding to the smallest singular value σ, which satisfies

$$Av = u\sigma \ , A^T u = v\sigma \ , \ u^T u = 1 \ , \ v^T v = 1 \ . \quad (1)$$

Here $u \in \mathbf{R}^p$ and $v \in \mathbf{R}^q$ are the corresponding left resp. right singular vector. When formulated as an optimization problem, we obtain

$$\min_{B \in \mathbf{R}^{p\times q},\, y \in \mathbf{R}^q} \|A-B\|_F^2 \ \text{ s.t. } \begin{cases} B\,y = 0 \\ y^T y = 1, \end{cases} \quad (2)$$

the solution of which follows from (1) as $B = A - u.\sigma.v^T$, $y = v$. The vector y describes a linear relation between the columns of the approximating matrix B, which therefore is rank deficient as required. This problem is also known as the Total (Linear) Least Squares (TLS) problem and has a long history (see e.g. the references in [21]). Often, for applications in statistics, signal processing and system identification, the objective function in (2), $\|A-B\|_F^2 = \sum_{i=1}^p \sum_{j=1}^q (a_{ij}-b_{ij})^2$, is extended with user-defined 'elementwise' nonnegative weights w_{ij} and, in addition, the approximating rank deficient matrix B is required to have some structure (such as Hankel, Toeplitz, etc...), in which case we have *the structured total least squares problem*:

$$\min_{\substack{B \in \mathbf{R}^{p\times q} \\ \mathrm{rank}(B) = q-1 \\ B \text{ structured}}} \sum_{i=1}^{p}\sum_{j=1}^{q} (a_{ij}-b_{ij})^2.w_{ij} \ . \quad (3)$$

The results in this paper apply to so-called *affine matrix structures*, i.e. structured matrices B that can be written as an affine (linear) combination of a given set $\{B_k; k=0,\dots,N\}$ of basis matrices as $B = B_0 + B_1\beta_1 + \dots + B_N\beta_N$. Here, the coefficients $\beta_k, k=1,\dots,N$ 'parametrize' the structured matrix B. Examples of such structured matrices are (centro- and per-) symmetric matrices, (block-) Hankel, (block-) Toeplitz, (block-) circulant, Brownian, Hankel + Toeplitz, etc .... Hankel matrices in particular have important applications in systems and control theory. More specifically, when a Hankel matrix is rank deficient (of rank r, say), it contains the Markov parameters of an r-th order linear time invariant system. But in general there are many other applications where it is required to find rank deficient structured approximants [8] [9] [10]. For instance, a special case of (3) occurs when no particular structure is imposed onto B, but there are still weights w_{ij} in the objective function. An example of such a *weighted* total least squares problem is given by chosing $w_{ij} = 1/a_{ij}^2$, in which case one minimizes (a first order approximation to) the sum of relative errors squared (instead of the sum of absolute errors squared as in (2)). Another example corresponds to the choice $w_{ij} \in \{1,+\infty\}$ in which case some of the

*Research supported by the Flemish Government (Concerted Research Action MIPS: Model-based Information Processing Systems), the National Fund for Scientific Research (G.0292.95 Matrix algorithms and differential geometry for adaptive signal processing, system identification and control), the Belgian Government (Interuniversity Attraction Poles IUAP-17: Modelling and Control of Dynamical Systems, IUAP-50: Automation in Design and Production) and the European Commission (Human Capital and Mobility Network SIMONET: System Identification and Modelling Network).

†Associate Professor and Senior Research Associate with the National Fund for Scientific Research

elements of B will be the same as corresponding elements of A (namely the ones that correspond to a weight $w_{ij} = +\infty$, see the Example in Section 4). Yet other examples of weighted TLS problems are given in [8] [9].

The main result in this paper, which was derived in [8], states that the problem (3) can be solved by obtaining the singular triplet (u, σ, v) that corresponds to the smallest singular value σ in

$$\boxed{\begin{array}{ll} Av = D_v u\sigma , & u^T D_v u = 1 , \\ A^T u = D_u v\sigma , & v^T D_u v = 1 . \qquad (4) \end{array}}$$

Notice the similarity with the SVD expressions in (1). Here A is the structured data matrix that one wants to approximate by a rank deficient one. D_u and D_v are nonnegative or positive definite matrix functions of the components of the left and right singular vectors u and v. Their precise structure depends on the weights and/or the required affine structure of the rank deficient approximant B. To give just one example, let us consider the approximation of a full column rank Hankel matrix $A \in \mathbf{R}^{p\times q}, p \geq q$, $\text{rank}(A) = q$, by a rank deficient Hankel matrix B such that $\|A - B\|_F^2$ is minimized. In this case, the matrix D_v has the form $D_v = T_v W^{-1} T_v^T$ where

$$W = \text{diag}[\ 1\ 2\ 3\ \ldots\ \underbrace{q\ q\ \ldots\ q}_{(p-q+1)\text{ times}}\ \ldots 3\ 2\ 1] ,$$

and T_v is a banded Toeplitz matrix (illustrated here for the case $p = 4$, $q = 3$) of the form:

$$T_v = \begin{pmatrix} v_1 & v_2 & v_3 & 0 & 0 & 0 \\ 0 & v_1 & v_2 & v_3 & 0 & 0 \\ 0 & 0 & v_1 & v_2 & v_3 & 0 \\ 0 & 0 & 0 & v_1 & v_2 & v_3 \end{pmatrix} .$$

Similarly, $D_u = T_u W^{-1} T_u^T$. Obviously, in this example, both D_u and D_v are positive definite matrices. Observe that B has disappeared from the picture, as it does not appear in (4) (neither does it in (1)), but it can be reconstructed as

$$B = A - \text{multilinear function of } (u, \sigma, v) .$$

The modification to A no longer is a rank one matrix as is the case with the 'unstructured' TLS problem (2). Instead, the modification is a multilinear function of the 'smallest' singular triplet, the detailed formulas of which can be found in [8] [9]. We are interested in finding the *smallest* singular value in (4) because it can be shown that its square is precisely equal to the objective function:

$$\sum_{i=1}^{p} \sum_{j=1}^{q} (a_{ij} - b_{ij})^2 . w_{ij} = \sigma_{\min}^2 . \qquad (5)$$

In the special case that $D_u = I_p$ and $D_v = I_q$ we obtain the SVD expressions (1). In the case that D_u and D_v are fixed positive definite matrices that are independent of u and v, one obtains the so-called *Restricted* SVD, which is extensively studied in [6] together with some structured/weighted TLS problems for which it provides a solution. In the Restricted SVD, D_u and D_v are positive (or nonnegative) definite matrices which can be associated to a certain inner product in the column and row space of A. Here in (4), D_u and D_v are also positive (nonnegative) definite, but instead of being constant, their elements are a function of the components of u and v. It turns out that we can interprete these matrices as Riemannian *metrics*, an interpretation which might be useful when developping new (continuous-time) algorithms. For this reason, we propose to call the equations in (4), the *Riemannian SVD* [1].

The remainder of this paper is organized as follows: In Section 2, we derive some continuous-time algorithms for the symmetric and non-symmetric eigenvalue problem, or more precisely, vector differential equations that converge to the smallest eigenvalue and corresponding eigenvector. We also provide some (differential) geometrical and numerical interpretations. These results are then applied in Section 3 for deriving continuous-time algorithms for the smallest singular value of a matrix, which then leads to a heuristic algorithm for the smallest singular value of a Riemannian SVD, which is the subject of Section 4. One of the goals of this paper is to point out several interesting connections between linear algebra, optimization theory, numerical analysis, differential geometry and system theory. Continuous time algorithms for solving and analyzing numerical problems have gained considerable interest the last decade or so (see e.g. [3] [4] [5] [11] [15] [14] [20]). Roughly speaking, a continuous-time method involves a system of matrix or vector differential equations. The idea that a computation can be thought as a flow that starts at a certain initial state and evolves until it reaches an equilibrium point (which then is the desired result of the computation) is a natural one when one thinks about iterative algorithms and even more, about recent developments in natural information processing systems related to artificial neural

[1] This name is slightly misleading in the sense that we do NOT want to suggest that there is a complete decomposition with $\min(p, q)$ different singular triplets, which are mutual 'independent' ('orthogonal') and which can for instance be added together to give an additive decomposition of the matrix A (the dyadic decomposition). There might be several solutions to (4) (for some examples, there is only one), but since each of these solutions goes with a different matrix D_u and D_v, it is not exaclty clear how these solutions relate to each other, let alone that they would add together in one way or another to obtain the matrix A.

networks [2]. There are several reasons why the study of continuous time algorithms is important. Continuous time methods can provide new additional insights with respect to and shed light upon existing discrete time iterative or recursive algorithms. In many cases, continuous-time algorithms provide an alternative or sometimes even better understanding of discrete-time versions (e.g. in optimization, see the continuous-time version of interior point techniques in [13, p.126], or in numerical analysis, the self-similar iso-spectral (calculating the eigenvalue decomposition) or self-equivalent (singular value decomposition) matrix differential flows (see e.g. [4] [15]). It is not our intention to claim that these continuous-time algorithms are in any sense competitive with classical algorithms from e.g. numerical analysis. Yet, there are examples in which we have only continuous-time solutions and for which the discrete-time iterative counterpart has not yet been derived.

2. Continuous-time eigenvalue algorithms

Let us now consider a continuous-time algorithm for the symmetric eigenvalue problem: Given a matrix $C = C^T \in \mathbf{R}^{p \times p}$, find at least one pair (x, λ) with $x \in \mathbf{R}^p$ and $\lambda \in \mathbf{R}$ such that

$$Cx = x\lambda \ , \ x^T x = 1 \ . \tag{6}$$

Here we will concentrate on power methods, both in continuous and discrete time. The continuous-time power methods are basically systems of vector differential equations. Some of these have been treated in the literature (see e.g. [4] [14] [20]). Others presented here are new. The *geometric* interpretation of the symmetric eigenvalue is as follows: Consider the two quadratic surfaces $x^T Cx = 1$ and $x^T x = 1$. While the second surface is the unit sphere in p dimensions, the first surface can come in many disguises. For instance, when C is positive definite, it is an ellipsoid. In three dimensions, depending on its inertia, it can be a one-sheeted or two-sheeted hyperboloid or a (hyperbolic) cilinder. In higher dimensions, there are many possibilities, the enumeration of which is not relevant right now. In each of these cases, the vectors Cx and x are the normal vectors at x to the two surfaces. Hence, when trying to solve the symmetric eigenvalue problem, we are looking for a vector x such that the normal at x to the surface $x^T Cx = 1$ is proportional to the normal at x to the unit sphere $x^T x = 1$. The constant of proportionality is precisely the eigenvalue λ. The *variational* characterization of the minimal eigenvalue of a symmetric matrix, follows from the following optimization problem:

$$\min_{x \in \mathbf{R}^p} f(x) = x^T Cx \quad \text{s.t.} \ x^T x = 1 \ . \tag{7}$$

[2] Specifically for neural nets and SVD, we refer to e.g. [1] [2] [5] [16] [17] (and the references in there).

The Lagrangian for this constrained problem is $L(x, \lambda) = x^T Cx + \lambda(1 - x^T x)$, where $\lambda \in \mathbf{R}$ is a scalar Lagrange multiplier. The necessary conditions for a stationary point follow from $\nabla_x L(x, \lambda) = 0$ and $\nabla_\lambda L(x, \lambda) = 0$, and will correspond exactly to the two equations in (6). Observe that these equations have p solutions (x, λ) while we are only interested in the minimizing one. We can also derive a so-called *gradient flow* for this optimization problem, which is a set of nonlinear differential equations, which evolves on the manifold defined by the constraints. For the symmetric eigenvalue problem (6), the set of vectors that satisfies the constraints is the unit sphere, which is known to be a manifold $M = \{ x \in \mathbf{R}^p \mid x^T x = 1 \}$. The tangent space at x is given by the vectors z that belong to $T_x M = \{ z \in \mathbf{R}^p \mid z^T x = 0 \}$. The directional derivative $D_x g(z)$ is the amount by which the objective function changes when moving in the direction z of a vector in the tangent space: $D_x g(z) = x^T Cz$. Next, we can choose a Riemannian metric represented by a smooth matrix function $W(x)$ which is positive definite for all $x \in M$. It is well known (see e.g. [14]) that, given the metric $W(x)$, the gradient ∇g can be uniquely determined from two conditions: **1. Compatibility**: $D_x g(z) = z^T Cx = z^T W(x) \nabla g$, and **2. Tangency**: $\nabla g \in T_x M \iff x^T \nabla g = 0$. The unique solution to these two equations applied to the constrained optimization problem (7) results in the (negative) gradient, given by the right hand side of the following gradient flow:

$$\dot{x} = -W(x)^{-1}(Cx - x \frac{x^T W(x)^{-1} Cx}{x^T W(x)^{-1} x}) \ . \tag{8}$$

It is easily seen that the stationary points of this system must be eigenvector-eigenvalues of the matrix C. Convergence is guaranteed because one can easily find a Lyapunov function (essentially the norm of the gradient) using the chain rule (see [14] for details). Observe that the norm $||x(t)||$ is constant for all t. This can be seen by substituting (8) into

$$\frac{1}{2} \frac{d||x||^2}{dt} = x^T \dot{x} = 0 \ .$$

Hence, if $||x(0)|| = 1$, we have $||x(t)|| = 1, \forall t > 0$. An interesting open problem is how to chose the metric $W(x)$ such that for instance the convergence speed could be optimized, or how an appropriate metric might lead to an elegant discrete-time integration algorithm. If we choose the Euclidean metric, $W(x) = I_n$, we obtain the following continuous-time power method:

$$\dot{x} = -Cx + x \frac{x^T Cx}{x^T x} \ . \tag{9}$$

This flow can also be interpreted as a special case of Brockett's [3] double bracket flow (see [15]). Another interesting interpretation is the following: Define the residual vector $r(t) = Cx - x(x^T Cx/x^T x)$. Then the differential equation (9) reads

$$\dot{x} = -r(t) \ . \tag{10}$$

Hence, the differential equation is driven by a residual error and when $r(t) = 0$, we also have $\dot{x} = 0$ so that a zero residual results in a stationary point. Moreover, using a little known fact in numerical anaylsis [18, p.69] we can give a *backward error interpretation* as follows. Define the rank one matrix $M(t) = r(t).x(t)^T$. Then we easily see that

$$(C - M(t))x(t) = x(t)\lambda(t)$$

with $\lambda(t) = (x(t))^T Cx(t)/(x(t)^T x(t))$. The interpretation is that, at any time t, the real number $\lambda(t)$ is the *exact* eigenvalue of a *modified* matrix, namely $C - M(t)$. The norm of the modification is given by $||M(t)|| = ||r(t)||$, which is the norm of the residual vector and from the convergence, we know that $||M(t)|| \rightarrow 0$ as $t \rightarrow \infty$.

The surprising fact about the nonlinear differential equation (9) is that it can be solved analytically. Its solution is

$$x(t) = e^{-Ct}x(0)/||e^{-Ct}x(0)|| \ , \tag{11}$$

which can be verified by direct substitution. If we consider the analytic solution at integer times $t = k = 0, 1, 2, \ldots$, we see that

$$x(k+1) = e^{-C}x(k)/||e^{-C}x(k)|| \ ,$$

which shows that the continuous time equation (9) interpolates the discrete time power method. Obviously, the stationary points (points where $\dot{x} = 0$) are the eigenvectors of C, but there is only one stationary point that is stable (as could be shown by linearizing (9) around all the stationary points and calculating the eigenvalues of the linearized system). The solution for the linear part in (9) (the first term of the right hand side) is $x(t) = \exp(-Ct)x(0)$. The second term is a normalization term which at each time instant projects the solution of the linear part back to the unit sphere.

So far we have only considered symmetric matrices. However most of the results still hold true *mutatis mutandis* if C is a non-symmetric matrix. For instance, the analytic solution to (9) is still given by (11) even if C is nonsymmetric. For the convergence proof, we can find a Lyapunov function using the left and right eigenvectors of the non-symmetric matrix C. Let $X \in \mathbf{R}^{p\times p}$ be the matrix of right eigenvectors of C while Y is the matrix of left eigenvectors, normalized such that

$$AX = X\Lambda \ , \qquad Y^T X = I_p \ ,$$
$$A^T Y = Y\Lambda \ , \qquad XY^T = I_p \ .$$

For simplicity we assume that all the eigenvalues of C are real (although this is not really a restriction). The vector $y_{\min} \in \mathbf{R}^p$ is the left eigenvector corresponding to the smallest eigenvalue. It can be shown that the scalar function $L(x) = (x^T y_{\min})^2 \ / \ (x^T YY^T x)$ is a Lyapunov function for the differential equations (9) because $\dot{L} > 0, \forall t$. Note that $x = X(Y^T x)$, which says that the vector $Y^T x$ contains the components of x with respect to the basis generated by the column vectors in X (which are the right eigenvectors). The denominator is just the norm squared of this vector of components. The numerator is the component of x along the last eigenvector (last column of X), which corresponds to the smallest eigenvalue. Since $\dot{L} > 0$, this component grows larger and larger relative to all other ones, which proves convergence to the 'smallest' eigenvector.

3. Continuous-time singular value flows

We can exploit the insights obtained so far to construct a flow for the smallest singular triplet of a matrix $A \in \mathbf{R}^{p\times q}$, which 'solves' the total least squares problem (2). Consider the 'asymmetric' continuous-time power method:

$$\begin{pmatrix} \dot{u} \\ \dot{v} \end{pmatrix} = -\begin{pmatrix} \alpha I_p & -A \\ A^T & -\alpha I_q \end{pmatrix}\begin{pmatrix} u \\ v \end{pmatrix} + \begin{pmatrix} u \\ v \end{pmatrix}\sigma$$

with $\sigma = \alpha(u^T u - v^T v) \ / \ (u^T u + v^T v)$. Here α is a user-defined real scalar. When it is chosen such that $\alpha > \sigma_{\min}(A)$, u and v will converge to the left resp. right singular vector corresponding to the smallest singular value of A. This can be understood by realizing that the $2q$ (we assume that $p \geq q$) eigenvalues of the matrix $\begin{pmatrix} \alpha I_p & -A \\ A^T & -\alpha I_q \end{pmatrix}$ are $\lambda_i = \pm\sqrt{\alpha^2 - \sigma_i^2(A)}$ and that there are $p-q$ eigenvalues equal to α. Hence, if $\alpha > \sigma_i$, the corresponding eigenvalues λ_i are real, else they are pure imaginary. So if $\alpha > \sigma_{\min}$, the smallest eigenvalue is $\lambda = -\sqrt{\alpha^2 - \sigma_{\min}^2}$ to which the flow will converge. This is illustrated in Figure 1.

The problem with this asymmetric continuous power method is that we have to know a scalar α that is an upper bound to the smallest singular value. Here we propose a new continuous-time algorithm for which we have *strong indications* that it always converges to the 'smallest' singular value, but for which there is no formal proof of convergence yet.

Let $A \in \mathbf{R}^{p\times q}$ $(p \geq q)$ and $\phi \in \mathbf{R}$ be a given (user-defined) scalar satisfying $0 < \phi < 1$. Consider the following system of differential equations, which we call *Christiaan's flow* [3]:

$$\begin{pmatrix} \dot{u} \\ \dot{v} \end{pmatrix} = -\begin{pmatrix} \sigma I_p & -A \\ \phi A^T & -\sigma\phi I_q \end{pmatrix}\begin{pmatrix} u \\ v \end{pmatrix} \tag{12}$$

[3]... after one of our PhD students Christiaan Moons who one day just tried it out and to his surpise found out that it seems to converge always.

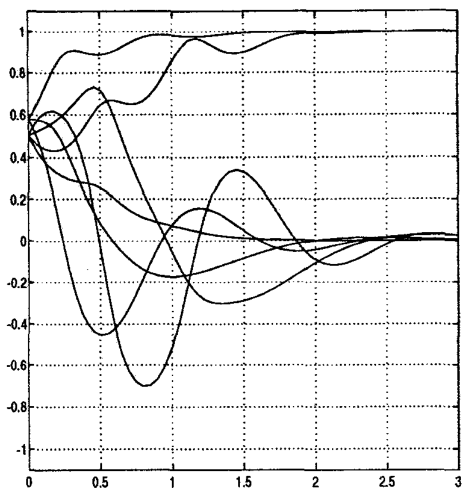

Figure 1: Convergence as a function of time of the components of $u(t)$ and $v(t)$ of the asymmetric continuous-time power method of Section 3, for a 4×3 'diagonal' matrix A with diagonal elements $(5, 3, 1)$. The initial vectors $u(0)$ and $v(0)$ are random. Shown is the behavior for $\alpha = 2$, which is larger than the smallest singular value and therefore there is convergence. Here as well as in the other figures, the simulation was done with Matlab's numerical integration routine 'ode45'.

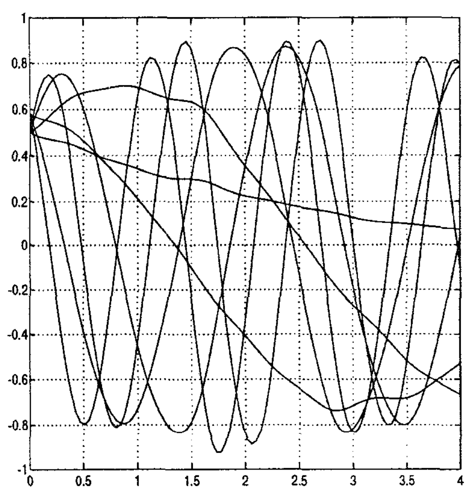

Figure 2: This picture shows the dynamic behavior of the asymmetric continuous-time power method of Section 3, applied to the same matrix A as in Figure 1, with the same initial random vectors, but for $\alpha = 0.5$. Because this is smaller than the smallest singular value of A, there is no convergence but instead, there is a nonlinear oscillation.

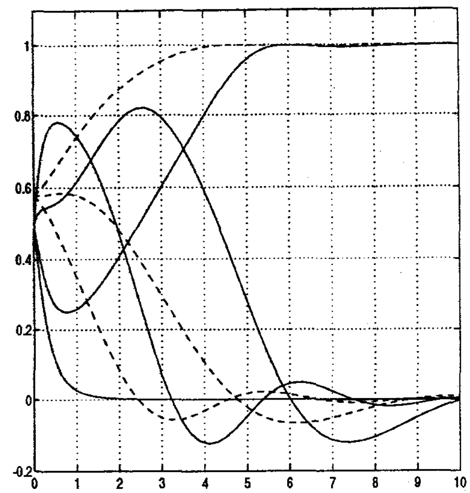

Figure 3: Example of the convergence behavior of Christiaan's flow (12) for the same matrix A and initial vectors $u(0)$ and $v(0)$ as in the first Figure. The full lines are the components of u, the dashed ones those of v. Both vectors converge asymptotically to the left and right singular vectors of A corresponding to the smallest singular value.

with $\sigma = (u^T A v)/(u^T u)$. Numerical simulations show that the triplet $(u(t), \sigma(t), v(t))$ converges to the smallest singular triplet of A. The convergence behavior can be influenced by ϕ (for instance, when $\phi \rightarrow 1$, there are many oscillations (bad for numerical integration) but convergence is quite fast in time; when $\phi \rightarrow 0$, there are no oscillations but convergence is slow in time). There are many similarities with the continuous-time algorithms discussed so far. When we rewrite these equations as $\dot{u} = Av - u\sigma$ and $\dot{v} = -\phi(A^T u - v\sigma)$, we easily see from (1) that both equations are 'driven' by the (scaled) residual error (compare to the interpretation of equation (10)). Another intriguing connection is seen by rewriting (12) as

$$\begin{pmatrix} \dot{u} \\ \dot{v} \end{pmatrix} = W.[-\begin{pmatrix} 0 & A \\ A^T & 0 \end{pmatrix}\begin{pmatrix} u \\ v \end{pmatrix} + \begin{pmatrix} u \\ v \end{pmatrix}\sigma]$$

with $W = \begin{pmatrix} -I & 0 \\ 0 & \phi I \end{pmatrix}$, which compares very well to (8), except that the metric $W(x)$ here is constant and indefinite! As for a formal convergence proof [4], there are the following facts: It is readily verified that the singular triplets of A are the stationary points of the flow. When we linearize the system around these stationary points, it is readily verified that they are all unstable, except for the stationary point corresponding to the 'smallest' singular triplet of A.

[4]We offer a chique dinner in the most exquis restaurant of Leuven if somebody solves our problem.

4. Algorithms for the Riemannian SVD

In this section, we present some undeveloped and premature ideas about continuous time algorithms for the Riemannian SVD (4) and hence for the structured total least squares problem (3). The ultimate challenge would consist of deriving a *gradient flow* for the Riemannian SVD, as we have done for the eigenvalue problem in Section 2. As this seems more involved than one would think at first sight, we will only present here the heuristic idea of Christiaan's flow of the previous section, this time applied to the Riemannian SVD. It should be noted that we have derived a heuristic iterative discrete-time algorithm for the Riemannian SVD in [8], the basic inspiration of which is the classical power method. While this algorithm works well in most cases, there is no formal proof of convergence nor is there any guarantee that it will convergence to (an at least local) minimum. Of course, all of the insights presented here ideally would lead to a discrete-time algorithm with guaranteed convergence to a local minimum. It seems that the following generalization of Christiaan's flow (12) works very well to find the minimal singular triplet of the Riemannian SVD (4).

$$\begin{pmatrix} \dot{u} \\ \dot{v} \end{pmatrix} = -\begin{pmatrix} \sigma D_v & -A \\ \phi A^T & -\sigma\phi D_u \end{pmatrix}\begin{pmatrix} u \\ v \end{pmatrix} \quad (13)$$

with $\sigma = (u^T Av)/(u^T u)$, and with $\phi \in \mathbf{R}$ a user-defined number satisfying $0 < \phi < 1$. The only difference between (12) and (13) is the introduction of the positive definite metric matrices D_u and D_v. It can be readily verified that in a stationary point, the necessary conditions (4) are satisfied. From numerical experiments, we have noticed that the system of differential equations converges to a local minimum of the objective function. As an example, let $A \in \mathbf{R}^{6\times 5}$ be a given matrix (we took a random matrix), that will be approximated in Frobenius norm by a rank deficient matrix B, by not modifying all of its elements but only those elements that are 'flagged' by a '1' in the matrix

$$V = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 & 0 \\ 1 & 1 & 1 & 0 & 1 \\ 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 \end{pmatrix} ;$$

Let W be be the elementwise inverse of V (with the convention that $1/0 = +\infty$). It contains the weights w_{ij} of (3) and the element $+\infty$ means that we impose an infinite weight on the modification of the corresponding element in A (which implies that it will not be modified and that the corresponding element in B will be equal to that in A). It can be shown [9]

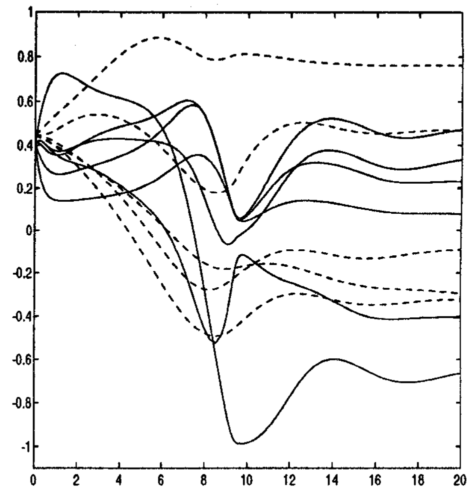

Figure 4: Vectors $u(t)$ and $v(t)$ as a function of time for Christiaan's flow of Section 4, which solves the problem of least squares approximation of a given matrix A by a rank deficient one, while not all of its elements can be modified as specified by the elements in the (0,1)-matrix V. The vector differential equation converges to the same solution as the one provided by the discrete-time algorithm of [9], which on its turn is inspired by the discrete-time power method.

that the metric matrices D_u and D_v of (4) for this problem are diagonal matrices given by

$$\begin{aligned} D_v &= \operatorname{diag}(V(v_1^2\ v_2^2\ v_3^2\ v_4^2\ v_5^2)^T)\ , \\ D_u &= \operatorname{diag}(V^T(u_1^2\ u_2^2\ u_3^2\ u_4^2\ u_5^2\ u_6^2)^T)\ . \end{aligned}$$

Here v_i and u_i denote the i-th component of v, resp. u and u_i^2 and v_i^2 are their squares. As an initial vector for Christiaan's flow (13) we took $[u(0)^T\ v(0)^T] = [e_6^T/\sqrt{6}\ \ e_5^T/\sqrt{5}]$ where $e_k \in \mathbf{R}^k$ is a vector with all ones as its components. The resulting convergent behavior as a function of time is shown in Figure 4.

5. Conclusions and future research

In this paper, we have discussed several continuous-time algorithms to find a minimizing solution to the total (2) or the structured total least squares problem (3). There are many applications for structured total least squares problems in statistics, system theory and signal processing, which makes it an important problem (see e.g. [7] [8] [9] [10] [11] [19]). Of course, since (3) is a nonlinear constrained optimization problem, one could apply 'classical' algorithms, such as Gauss-Newton, to find a local minimum of the objective function. However, the ultimate challenge here is to find a continuous-time or discrete-time (iterative) algorithm, that exploits the inherent structure of the structured total least squares problem (3), as it is for instance revealed in the Riemannian SVD (4). One possible idea to pursue is the following: If D_v is nonsingular, it is straightforward to show that the objective function (5) is equal to $\sigma^2 = v^T A^T D_v^{-1} A v$. The gradient with respect to v would be the first term in a gradient flow, the second term being the projection on the manifold generated by the constraints in (4), which are quadratic in the components of u and v. This idea however remains to be explored in future research.

References

[1] Baldi P., Hornik K. *Neural networks and principal component analysis: Learning from examples without local minima.* Neural Networks, Vol.2, pp.53-58, 1989.

[2] Bourlard H., Kamp Y. *Auto-association by multilayer perceptrons and singular value decomposition.* Biol. Cybern., 59, pp.291-294, 1988.

[3] Brockett R.W. *Dynamical systems that sort lists and solve linear programming problems.* Linear Algebra and its Applications, 146, pp.79-91, 1991.

[4] M.T. Chu, *On the continuous realization of iterative processes*, SIAM Review, 30, 3, pp.375-387, September 1988.

[5] Dehaene J. PhD Thesis, Department of Electrical Engineering, Katholieke Universiteit Leuven, Belgium, October 1995. (subjects: Differential geometry, continuous-time algorithms, analog implementation).

[6] De Moor B., Golub G.H. *The restricted singular value decomposition: properties and applications.* Siam Journal on Matrix Analysis and Applications, Vol.12, no.3, July 1991, pp.401-425.

[7] De Moor B., Van Overschee P., Schelfhout G. *H_2-model reduction for SISO systems.* Proc. of the 12th World Congress International Federation of Automatic Control, Sydney, Australia, July 18-23 1993, Vol. II pp.227-230. [5]

[8] De Moor B. *Structured total least squares and L_2 approximation problems.* Special issue of Linear Algebra and its Applications, on Numerical Linear Algebra Methods in Control, Signals and Systems (eds: Van Dooren, Ammar, Nichols, Mehrmann), Volume 188-189, July 1993, pp.163-207.

[5] Reprinted in the Automatic Control World Congress 1993 5-Volume Set, Volume 1: Theory.

[9] De Moor B. *Total least squares for affinely structured matrices and the noisy realization problem* . IEEE Transactions on Signal Processing, Vol.42, no.11, November 1994.

[10] De Moor B., Roorda B. *L_2-optimal linear system identification: Dynamic total least squares for SISO systems.* Proc. of 33rd IEEE CDC, pp.2874-2880, Florida, December 14-16 1994.

[11] De Moor B. *The Riemannian singular value decomposition,* Proc. 3rd Int. Workshop on SVD and Signal Processing (Eds. M. Moonen, De Moor B.), 1995, pp. 61-78, Elsevier Science Publisher.

[12] Eckart C., Young G. *The approximation of one matrix by another of lower rank.* Psychometrika, 1, pp.211-218, 1936.

[13] Fiacco A.V., McCormick G.P. *Nonlinear programming; Sequential unconstrained minimization techniques.* SIAM Classics in Applied Mathematics, 1990.

[14] Helmke U., Moore J.B. *Optimization and dynamical systems.* CCES, Springer Verlag, London, 1994.

[15] Helmke U. *Isospectral matrix flows for numerical analysis.* Proc. 3rd Int. Workshop on SVD and Signal Processing (eds. Moonen M., De Moor B.), pp. 33-60, 1995, Elsevier Science Publisher.

[16] Kung S.Y., Diamantaras K.I., Taur J.S. *Neural networks for extracting pure, constrained, oriented principal components.* Proc. of the 2nd International Workshop on SVD and Signal Processing, June 1990, pp. 57-82, Elsevier Science Publisher.

[17] Oja E. *A simplified neuron model as a principal component analyzer.* J. Math. Biology, 15, pp.267-273, 1982.

[18] Parlett B. *The symmetric eigenvalue problem.* Prentice Hall, Englewood Cliffs, NJ, 1980.

[19] Schelfhout G. *Model reduction for control design.* PhD Thesis, Department of Electrical Engineering, Katholieke Universiteit Leuven, Belgium, November 1995.

[20] S.T. Smith, *Geometric Optimization Methods for Adaptive Filtering,* PhD Thesis, Harvard University, May 1993 Cambridge, Massachusetts, USA.

[21] Van Huffel S., Vandewalle J. *The Total Least Squares Problem: Computational Aspects and Analysis.* Frontiers in Applied Mathematics 9, SIAM, Philadelphia, 300 pp., 1991.