01_106

WEIGHTED TOTAL LEAST SQUARES, RANK DEFICIENCY AND LINEAR MATRIX STRUCTURES

Bart De Moor
ESAT-SISTA, Katholieke Universiteit Leuven
Kasteelpark 10, B-3001 Leuven, Belgium
bart.demoor@esat.kuleuven.ac.be web: www.esat.kuleuven.ac.be/~demoor

Abstract In this contribution we explain how the rank deficiency of a given data matrix, its linear matrix structure and its interpretation in terms of discrete-time dynamical systems, are intimately connected. This relation permits to formulate least squares dynamical system identification problems, the solution of which leads to structured total least squares (STLS). We also consider the insertion of given weights in the least squares objective function, leading to weighted TLS problems (WTLS). The Riemannian SVD, which is a 'nonlinear' generalized SVD, provides an elegant framework for these structured and/or weighted TLS problems.

Keywords: dynamic errors-in-variables, elementwise weighted total least squares, Hankel TLS, Riemannian SVD, system identification

1. Rank deficiency and linear matrix structures

There is a trivial, yet important connection between the rank deficiency of a given matrix A and the existence of linear relations between its columns. Indeed, let $A \in \mathbf{R}^{p\times q}$, with $p \geq q$. When $\text{rank}(A) = r_A < q$, there exist $q - r_A$ linearly independent linear relations between the columns of A, but in this paper we will confine ourselves to one single linear relation, in which case $r_A = q - 1$. In a certain sense, the rank deficiency of a matrix is a 'non-generic' property (a notion that can be made very precise as e.g. in [17]). This explains the fact that, in most applications, the observed data matrix $A \in \mathbf{R}^{p\times q}, p \geq q$ is of full column rank q. The central idea then of *total linear least squares* (TLLS) is to approximate the given data matrix A, by a matrix B of rank $q-1$, so that at least one linear relation exists between the columns of B:

$$\min_{B\in\mathbf{R}^{p\times q},\, v\in\mathbf{R}^{q}} \|A-B\|_F^2 \text{ subject to } \left\{ \begin{array}{rcl} B\,.\,v &=& 0\,, \\ v^T\,.\,v &=& 1\,. \end{array} \right. \tag{1}$$

A complete treatment of the TLS problem (1) and its solution via the singular value decomposition, including many references and comparisons with other fitting techniques, generic and non-generic cases, geometrical, statistical and algorithmic issues, may be found in the books [15] [16]. In this paper, we discuss so-called ***structured and/or weighted total least squares (S/W-TLS)*** problems, which are an extension of the TLS problem (1), in which the matrix A and its rank deficient approximant B are required to have a certain linear matrix structure, and/or the least squares objective function is modified by introducing weights. We call a matrix $A \in \mathbf{R}^{p\times q}$ linearly structured, when it can be written as a linear combination of given, known basis matrices $A_i, i = 1, 2, \ldots, N \in \mathbf{R}^{p\times q}$ as $A = A_1\alpha_1 + \ldots + A_N\alpha_N$, with $\alpha_i \in \mathbf{R}$. N is the number of different coefficients α_i, needed to generate the matrix A. It is assumed throughout that the matrices A_i are linearly independent. Examples of such linear structures are symmetric, centro-symmetric, per-symmetric, row-rhomboidal, negacyclic matrices, Toeplitz, Hankel and circulant Brownian matrices and their block versions. These linear structures are ubiquitous in signal processing, systems and control theory, image processing and statistics, numerical linear algebra, mechanical and electrical engineering, physics, etc... Observe that linear matrix structures are closed under addition and intersection: Examples are 'symmetric Toeplitz' matrices (intersection of symmetric and Toeplitz) or 'Hankel + Toeplitz' matrices, which are the sum of a Hankel and a Toeplitz matrix. While rank deficiency of the data matrix hints at an underlying linear model that 'explains' the data, the fact that in addition the data matrix is also linearly structured, implies that these data are generated by a dynamical system, described by linear difference or differential equations. Let us give some examples.

Example 1: Rank deficient Hankel matrices
Let $Z \in \mathbf{R}^{p\times q}$ be a Hankel matrix with $p \geq q$ and denote its $N = p+q-1$ different elements by $z_k, k = 0, 1, \ldots, N$. It is well known that $\text{rank}(Z) = n$ if and only if the elements z_k are generated by the output of a linear time-invariant discrete time system of the form

$$\begin{aligned} x_{k+1} &= Fx_k\,, \\ z_k &= Hx_k\,, \end{aligned} \tag{2}$$

with $F \in \mathbf{R}^{n\times n}$ the system matrix, $H \in \mathbf{R}^{1\times n}$ the output matrix, x_k the state and z_k the output at time instant k. In the case that $q = n+1$, the vector $v \in \mathbf{R}^{n+1}$ that satisfies $Zv = 0$ contains the coefficients (up to within a scalar) of the characteristic polynomial $\det(F - \lambda I_n)$ (We assume here that the characteristic and minimal polynomial of F coincide). The problem of finding the model matrices F and H and the order n is called 'the realization problem' and is, by now, a 'classical' problem in system theory. The model (2) is used in many applications, ranging from biomedical signal processing (e.g. analysis of NMR spectra), vibrational analysis of dynamical structures (e.g. characterization of flutter of wings of air planes), industrial process system identification, stochastic identification and telecommunications, etc.... We refer to [1] [4] [5] [8] for examples and details.

Example 2: Rank deficient 'double' Hankel matrices
Next we extend the model (2) with scalar inputs $w_k \in \mathbf{R}$:

$$\begin{aligned} x_{k+1} &= F x_k + G w_k \,, \\ z_k &= H x_k + J w_k \,. \end{aligned} \tag{3}$$

Here, $G \in \mathbf{R}^{n\times 1}$ is the input matrix and $J \in \mathbf{R}^{1\times 1}$ is the so-called direct-feedthrough term. Let Z be a $p \times q$ Hankel matrix with the outputs ($p \geq q$) and W be a $p \times q$ Hankel matrix with the inputs. Then, under fairly general conditions on the input sequence ('persistancy of excitation', i.e. $\operatorname{rank}(W) = q$) and the system (3) (equal number of poles and zeros, observability and controllobality) we have

$$\operatorname{rank}(\, Z \;\; W \,) = q + n \,.$$

In particular, with $q = n+1$, $\operatorname{rank}(\, Z\; W \,) = 2n+1$ and .

$$(\, Z\, W \,) \begin{pmatrix} a \\ -b \end{pmatrix} = 0 \,, \tag{4}$$

where a and b contain (up to within a scalar) the coeffients of the so-called transfer function $T(z)$ of the linear system (3), which is rational in z: $T(z) = J + H(zI_n - F)^{-1}G = \frac{b(z)}{a(z)}$. Here $a(z)$ and $b(z)$ are n-th degree polynomials in z with the coefficients of the vectors a and b. Equation (4) is nothing else then a difference equation of the form

$$z_k \alpha_n + z_{k+1}\alpha_{n-1} + \ldots + z_{k+n}\alpha_0 = w_k \beta_n + w_{k+1}\beta_{n-1} + \ldots + w_{k+n}\beta_0 \,, \tag{5}$$

for $k = 0, 1, 2, \ldots$.

Example 3: Rank deficient 'bilinear' Hankel matrices
Yet another extension of the linear model (3) is the bilinear model, in which products between inputs and states are included in the state equation by introducing an additional matrix $L \in \mathbf{R}^{n\times n}$:

$$\begin{aligned} x_{k+1} &= Fx_k + Gw_k + Lx_k w_k , \\ z_k &= Hx_k + Jw_k , \end{aligned} \tag{6}$$

The Volterra kernels of this system are of the form: 0-th order: $K_0 = HG$; 1-st order: $K_1 = (K_{11}; K_{12}) = (HFG; HLG)$; 2nd order $K_2 = (K_{21}; K_{22}; K_{23}; K_{24}) = (HF^2G; HFLG; HLFG; HL^2G)$, etc.... It can be shown (see e.g. [13]) that the following linearly structured, 'generalized' Hankel matrix

$$H = \left(\begin{array}{c|cc|cccc|c} K_{00} & K_{11} & K_{12} & K_{21} & K_{22} & K_{23} & K_{24} & \cdots \\ \hline K_{11} & K_{21} & K_{22} & K_{31} & K_{32} & K_{33} & K_{34} & \cdots \\ K_{12} & K_{23} & K_{24} & K_{35} & K_{36} & K_{37} & K_{38} & \cdots \\ \hline K_{21} & K_{31} & K_{32} & \cdots & \cdots & \cdots & \cdots & \cdots \\ K_{23} & K_{35} & K_{36} & \cdots & \cdots & \cdots & \cdots & \cdots \\ K_{22} & K_{33} & K_{34} & \cdots & \cdots & \cdots & \cdots & \cdots \\ K_{24} & K_{37} & K_{38} & \cdots & \cdots & \cdots & \cdots & \cdots \\ \hline \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \end{array}\right)$$

is rank deficient of rank n, if its elements are Volterra kernels of a bilinear system of order n of the form (6). Again, we see how the combination of a linear matrix structure and rank deficiency has a very precise dynamical system theoretic interpretation.

Observe that in all these examples, the rank of the structured matrices reveals the order (= number of states) of the 'underlying' dynamical system!

2. A commuting Lemma for linear structures

Let $A \in \mathbf{R}^{p\times q}$ be a linearly structured matrix with N different elements collected in the vector $a \in \mathbf{R}^N$. For example, a $p \times q$ Hankel matrix has $N = p + q - 1$ different elements and a $p \times p$ symmetric Toeplitz matrix has $N = p$. Let $v \in \mathbf{R}^q$ be a vector. The matrix-vector product Av is bilinear in the elements of A and the elements of v, and because of this bilinearity we can reverse the order:

$$A\,v = T_v\,a\,, \tag{7}$$

where $T_v \in \mathbf{R}^{p\times N}$ is linearly structured in the elements of v. Observe that the linear structure of A and the ordering of the different elements of

A in the vector a, uniquely determine the structure of T_v. As an example, consider a $(p = 4) \times (q = 3)$ Hankel matrix A with $N = p + q - 1 = 6$ different elements. Then T_v follows from (7) as

$$\begin{pmatrix} a_0 & a_1 & a_2 \\ a_1 & a_2 & a_3 \\ a_2 & a_3 & a_4 \\ y_3 & a_4 & a_5 \end{pmatrix} \begin{pmatrix} v_0 \\ v_1 \\ v_2 \end{pmatrix} = \begin{pmatrix} v_0 & v_1 & v_2 & 0 & 0 & 0 \\ 0 & v_0 & v_1 & v_2 & 0 & 0 \\ 0 & 0 & v_0 & v_1 & v_2 & 0 \\ 0 & 0 & 0 & v_0 & v_1 & v_2 \end{pmatrix} \begin{pmatrix} a_0 \\ a_1 \\ a_2 \\ a_3 \\ a_4 \\ a_5 \end{pmatrix} .$$

Similarly we can define a linearly structured matrix $T_u \in \mathbf{R}^{N \times q}$ from $A^T u = T_u^T a$. These commuting results lead to simpler formula manipulation, in multilinear expressions of the form $\mathcal{L} = \ldots + u^T A v + \ldots$ in which they allow one to calculate expressions of the form $\frac{\partial \mathcal{L}}{\partial a} = \ldots + T_v^T u + \ldots$.. The Commuting Lemma also allows to exchange the role of $v \in \mathbf{R}^q$ and $u \in \mathbf{R}^p$ in $T_v^T u = T_u^T v$.

3. Data and noise models

In practical engineering applications, data are never 'exact', but they are always corrupted by so-called 'measurement noise'. More specifically, let $a \in \mathbf{R}^N$ be the vector of observed data. We will assume that the noise model is additive and Gaussian, i.e. $a = \hat{a} + \tilde{a}$, in which $\hat{a}$ contains the unknown exact (noisefree) data and $\tilde{a}$ the unobserved noise samples. As for the noise model, we assume that $\tilde{a} = M\tilde{\tilde{a}}$, where $M \in \mathbf{R}^{N \times N_1}$ with $N_1 \leq N$ is a given known matrix of full column rank N_1. The random vector $\tilde{\tilde{a}}$ is zero mean, normally distributed with covariance matrix equal to the identity matrix: $\tilde{\tilde{a}} \sim \mathcal{N}(0, I_{N_1})$. Hence, the noise random variable $\tilde{a}$ is zero mean normally distributed with covariance matrix $Q = MM^T$. Note that Q itself is singular whenever $N_1 < N$.

4. The Riemannian SVD

Let us now formulate structured and weighted TLS problems, the ingredients of which are a data vector $a \in \mathbf{R}^N$, a noise model $\tilde{a} = M\tilde{\tilde{a}}$ with a given specified matrix $M \in \mathbf{R}^{N \times N_1}$ with $\text{rank}(M) = N_1$, the specification of a linear matrix structure (Hankel, double Hankel, bilinear Hankel, Toeplitz, etc ...) (with corresponding linearly structured matrices T_u and T_v as defined above) and the specification of the number q of columns of the structured matrix.

Having specified these elements, every structured/weighted TLS problem can be phrased as

$$\min_{b\in\mathbf{R}^N,\, v\in\mathbf{R}^q,\, e\in\mathbf{R}^{N_1}} \frac{1}{2}e^Te \text{ subject to } \begin{cases} b = a - Me\,, \\ Bv = 0\,, \\ v^Tv = 1\,. \end{cases} \tag{8}$$

In these equations, $B \in \mathbf{R}^{p\times q}$ is the linearly structured matrix, generated from the elements of the vector $b \in \mathbf{R}^N$, that collects the different elements of B. This structured matrix B is required to be rank deficient, a condition that is enforced via the vector v in its null space. The vector b contains the estimates of the 'exact' data $\hat{a}$. The noise model is represented by the given matrix M. The least squares criterion (in which we introduce the factor 1/2 just for convenience) derives from the principle of maximum likelihood. It is easy to verify how the 'ordinary' TLS problem (1) is just a special instance of this general formulation. To solve the constrained optimization problem (8), we consider the Langrangean function $\mathcal{L}(m,e,v,l,\lambda,b) = \frac{1}{2}e^Te + l^TBv + \lambda(1-v^Tv) + m^T(b - a + Me)$. Setting all derivatives to zero (and repeatedly using the Commuting Lemma) results in

$$\begin{aligned} &\frac{\partial\mathcal{L}}{\partial m} = 0 \Longrightarrow b = a - Me\,, &&\frac{\partial\mathcal{L}}{\partial e} = 0 \Longrightarrow e + M^Tm = 0\,, \\ &\frac{\partial\mathcal{L}}{\partial v} = 0 \Longrightarrow B^Tl = 2v\lambda\,, &&\frac{\partial\mathcal{L}}{\partial l} = 0 \Longrightarrow Bv = 0\,, \\ &\frac{\partial\mathcal{L}}{\partial b} = 0 \Longrightarrow T_v^Tl + m = 0\,, &&\frac{\partial\mathcal{L}}{\partial \lambda} = 0 \Longrightarrow v^Tv = 1\,. \end{aligned}$$

From these equations, one can easily show that $\lambda = 0$, eliminate $m = -T_v^Tl$, eliminate $e = -M^Tm = M^TT_v^Tl$ and introduce $Q = MM^T$ to find that $b = a - QT_v^Tl$, $B^Tl = T_l^Tb = 0$ and $Bv = T_vb = 0$. Using these last two equations, we can eliminate the vector b and the matrix B, to find $T_l^Ta = (T_l^TQT_l)v$ and $T_va = (T_vQT_v^T)l$. Next we define the matrices D_{vv} and D_{ll} as $D_{vv} = (T_vQT_v^T)$ and $D_{ll} = (T_l^TQT_l)$, we normalize l as $l = u\sigma$ in which σ is such that $u^TD_{vv}u = 1$. Since now $D_{ll} = D_{uu}\sigma^2$, we then find (once again using the Commuting Lemma) that

$$\begin{aligned} Av &= D_{vv}u\sigma\,, \qquad & u^TD_{vv}u &= 1\,, \\ A^Tu &= D_{uu}v\sigma\,, \qquad & v^Tv &= 1\,. \end{aligned} \tag{9}$$

In these equations, which we have called the Riemannian SVD, the matrix A has the same affine structure as the matrix B we started from in (8). The value of σ that we need follows from the objective function as $e^Te = m^TMM^Tm = l^TT_vQT_v^Tl = u^TD_{vv}u\sigma^2 = \sigma^2$. Obviously, we need the minimal value of σ and corresponding vectors u and v that satisfy

the equations (9).
Observe that the matrices D_{uu} and D_{vv} are defined from the given required linear matrix structure (as 'represented' by the matrices T_u and T_v from the Commuting Lemma), and from the noise model represented by the covariance matrix $Q = MM^T$. By construction, the matrices D_{uu} and D_{vv} are symmetric, nonnegative definite matrix functions of the elements of u, resp. v. If they would be fixed matrices (i.e. independent of u and v), the Riemannian SVD would reduce to the so-called *Restricted Singular Value Decomposition* (see e.g. [3]). By analogy, we call u and v left resp. right singular vectors and σ the corresponding singular value of the Riemannian SVD. In the case of the ordinary TLS problem (1), we have $D_{uu} = I_p$ and $D_{vv} = I_q$. One can easily show that $u^T D_{vv} u = v^T D_{uu} v$ always, as from the Commuting Lemma it follows that $u^T D_{vv} u = u^T T_v Q T_v^T u = v^T T_u^T Q T_u v = v^T D_{uu} v$. The 'filtered' data (i.e. the least squares estimates of the unknown 'exact' data) can be estimated from $b = a - QT_v^T u\sigma$. The linearly structured matrix B generated from the elements of the vector b is rank deficient because $Bv = T_v b = T_v(a - QT_v^T l) = Av - D_{vv} u\sigma = 0$; The left singular vector u is the normalized vector l of Lagrange multipliers corresponding to the rank deficiency constraint. From the Riemannian SVD (9) we easily find, assuming that D_{vv} is invertible:

$$\sigma^2 = v^T A^T D_{vv}^{-1} Av \, . \tag{10}$$

This equation can be interpreted as the weighted norm of the so-called equation error. Assume that, in the noiseless case $\hat{A}\hat{v} = 0$. From the noise model specified above, we find that $A\hat{v} = (\hat{A} + \tilde{A})\hat{v} = \tilde{A}\hat{v} = T_{\hat{v}}\tilde{a} = T_{\hat{v}}M\tilde{a}$. For obvious reasons, $A\hat{v}$ is called the equation error. Its covariance matrix is given by $\mathbf{E}(A\hat{v}\hat{v}^T A) = T_{\hat{v}} Q T_{\hat{v}}^T = D_{\hat{v}\hat{v}}$. Therefore, the objective function in (10) is the weighted norm of the equation error, in which the weight is the inverse covariance matrix of the equation error. Finally, it can be shown that the vector of residuals $(b - a)$ is '$Q^\dagger$-orthogonal' to the vector with filtered data b as one can show that $b^T.Q^\dagger.(b - a) = 0$.
For more details and additional properties, we refer to [4] [5].

5. Elementwise weighted TLS

So far, we have obtained the general formulas for solving a structured and/or weighted TLS problem, leading to the Riemannian SVD (9). In this section, we derive the weighting matrices D_{uu} and D_{vv} for an elementwise weighted TLS problem. Let $A \in \mathbf{R}^{p \times q}$ be a given data matrix; Assume that the noise is additive, $A = \hat{A} + \tilde{A}$, and that the elements of $\tilde{A}$ are zero mean, Gaussian, independently distributed scalar random variables with noise variance σ_{ij}^2 for element (i,j). When approximating in a least squares sense, this given data matrix A, by a rank deficient matrix B, the principle of maximum likelihood leads to the following constrained optimization problem:

$$\min_{b_{ij}\in\mathbf{R}\,,\,v\in\mathbf{R}^q} \frac{1}{2}\sum_{i=1}^{p}\sum_{j=1}^{2}(a_{ij}-b_{ij})^2/\sigma_{ij}^2 \text{ subject to } \begin{cases} B\,.\,v=0\,, \\ v^T\,.\,v=1\,. \end{cases} \quad (11)$$

Let Σ be the $p\times q$ matrix that has σ_{ij}^2 as its (i,j)-th element. Let diag(.) be the operator that turns a vector into a diagonal matrix. By specializing the general derivation of Section 4 to this case, one obtains the Riemannian SVD for the weighted TLS problem with $D_{uu} = \text{diag}(\Sigma^T\text{diag}(u)u)$, $D_{vv} = \text{diag}(\Sigma\text{diag}(v)v)$. The rank deficient matrix B can be reconstructed from $B = A - \text{diag}(u)\Sigma\text{diag}(v)\sigma$, where σ is the smallest singular value of (9) and u and v are the corresponding left and right singular vectors. There are several 'extreme' cases that one can consider:

- Choosing $\sigma_{ij}\to 0$ implies that there is no noise on the corresponding element in A, hence that element is considered to be exact, and will not be modified, i.e. $a_{ij}=b_{ij}$; As an example, consider the matrix A

$$A = \begin{pmatrix} 1 & \boxed{2} & 3 \\ \boxed{-5} & 6 & \boxed{7} \\ 0 & 2 & -1 \\ \boxed{3} & 4 & 1 \end{pmatrix},$$

  where elements in a box can not be modified, when trying to least squares approximate A by a rank deficient matrix. The matrices D_{uu} and D_{vv} for this example are

$$D_{uu} = \text{diag}\begin{pmatrix} u_1^2+u_3^2 \\ u_2^2+u_3^2+u_4^2 \\ u_1^2+u_3^2+u_4^2 \end{pmatrix}, D_{vv} = \text{diag}\begin{pmatrix} v_1^2+v_3^2 \\ v_2^2 \\ v_1^2+v_2^2+v_3^2(=1) \\ v_2^2+v_3^2 \end{pmatrix}.$$

- By taking $\sigma_{ij}\to\infty$, we take the noise variance on element (i,j) in A to be extremely large; In this way, one can tackle *missing observations*, i.e. b_{ij} will be filled in 'automatically', without taking into account the corresponding element a_{ij} in A (Said in other words, because $\sigma_{ij}\to\infty$, the corresponding term does not contribute anything in the weighted objective function (11), leaving the choice for b_{ij} independent of a_{ij} (but constrained to make the matrix B rank deficient)).

- By taking $\sigma_{ij} = a_{ij}$, we get an objective function that minimizes the sum-of-relative-errors-squared, instead of 'absolute' errors.
- One can also consider the case in which every noise element $\tilde{a}_{ij}$ is correlated with every other noise element $\tilde{a}_{kl}$, as expressed in a correlation matrix $Q \in \mathbf{R}^{pq \times pq}$. For details we refer to [4].
- When Σ is a rank one matrix, i.e. $\sigma_{ij} = \xi_i \cdot \eta_j$, for positive scalars $\xi_i, \eta_j \in \mathbf{R}_0^+$, the solution to the weighted TLS problem (11) can be derived from the SVD of the diagonally scaled matrix with elements $a_{ij}/(\xi_i \cdot \eta_j)$.

6. Double Hankel structured TLS

The 'errors-in-variables' formulation of the linear dynamic system identification problem, for models of the form (3) is the following: Given noise corrupted scalar observations (u_k, y_k) of the input, resp. output, of a single-input single-output system. A least squares estimate w_k of u_k and z_k of y_k will be determined as follows. Let w, z, u, y be vectors with the scalar w_k, z_k, y_k, u_k. The structured TLS problem is then

$$\min_{w,z,a,b} \|y - z\|_2^2 + \|u - w\|_2^2 \text{ subject to } \begin{cases} (\, Z \;\, W \,) \begin{pmatrix} a \\ -b \end{pmatrix} = 0 \,, \\ a^T a + b^T b = 1 \,. \end{cases} \qquad (12)$$

Here, the vectors a and b contain the coefficients of the difference equation that relates the 'filtered' inputs w_k to the 'filtered' outputs z_k as in (5). From a similar derivation as the general one in Section 4, one can show that one has to find the minimal singular triplet of the Riemannian SVD

$$(\, Y \;\, U \,) \begin{pmatrix} a \\ -b \end{pmatrix} = (D_{aa} + D_{bb})\, u\, \sigma \,, \qquad a^T a + b^T b = 1 \,,$$

$$\begin{pmatrix} Y^T \\ U^T \end{pmatrix} u = \begin{pmatrix} D_{uu} & 0 \\ 0 & D_{uu} \end{pmatrix} \begin{pmatrix} a \\ -b \end{pmatrix} \sigma \,, \qquad u^T (D_{aa} + D_{bb}) u = 1 \,.$$

Here $D_{aa} = T_a T_a^T$, $D_{bb} = T_b T_b^T$ and $D_{uu} = T_u T_u^T$, in which the matrices T_a, T_b and T_u are banded Toeplitz. More details, such as expressions for the filtered input and output sequences, the proof that the residuals and filtered data are orthogonal, and the fact that the residuals themselves can be described by a linear difference equation, can be found in [6] [7]. As an extension to these results, one might also include weights in the objective function (12). In this way, one can treat missing observations, time-varying observation noise variances, etc.... A general framework for doing so, including the incorporation of 'unobserved latent' inputs, has been formulated in [12].

7. Rank 2 reduction

What if we want to reduce the rank of a linearly structured matrix A by 2, instead of by 1 ? Obviously, we now have to extend the constrained optimization problem (8) with some additional constraints:

$$\min_{b\in\mathbf{R}^N,\, v_1,v_2\in\mathbf{R}^q,\, e\in\mathbf{R}^{N_1}} \frac{1}{2}e^Te \text{ subject to } \begin{cases} b = a - Me\,, \\ B\,.\,(v_1\; v_2) = 0\,, \\ v_1^T\,.\,v_1 = 1\,, \\ v_2^T\,.\,v_2 = 1\,, \\ v_1^T\,.\,v_2 = 0\,. \end{cases}$$

We leave it to the reader to derive that the solution follows from the singular triplet corresponding to the smallest singular value of the 'double sized' Riemannian SVD

$$\begin{pmatrix} A & 0 \\ 0 & A \end{pmatrix}\begin{pmatrix} v_1 \\ v_2 \end{pmatrix} = \begin{pmatrix} D_{v_1v_1} & D_{v_1v_2} \\ D^T_{v_1v_2} & D_{v_2v_2} \end{pmatrix}\begin{pmatrix} u_1 \\ u_2 \end{pmatrix}\sigma\,,$$

$$\begin{pmatrix} A^T & 0 \\ 0 & A^T \end{pmatrix}\begin{pmatrix} u_1 \\ u_2 \end{pmatrix} = \begin{pmatrix} D_{u_1u_1} & D_{u_1u_2} \\ D^T_{u_1u_2} & D_{u_2u_2} \end{pmatrix}\begin{pmatrix} v_1 \\ v_2 \end{pmatrix}\sigma\,,$$

with obvious definitions for the matrix $D_{v_iv_j}$ and $D_{u_iu_j}$ (e.g. $D_{v_1v_2} = T_{v_1}QT^T_{v_2}$), and with the normalization constraints

$$(\,u_1^T\; u_2^T\,)\begin{pmatrix} D_{v_1v_1} & D_{v_1v_2} \\ D^T_{v_1v_2} & D_{v_2v_2} \end{pmatrix}\begin{pmatrix} u_1 \\ u_2 \end{pmatrix} = 1\,,\; \begin{pmatrix} v_1^T \\ v_2^T \end{pmatrix}(\,v_1\; v_2\,) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}\,.$$

8. Conclusions and some research challenges

In this contribution, we have presented a short survey of recent results on the formulation of S/W-TLS problems via the so-called Riemannian SVD, which is a 'nonlinear' generalized' SVD. A lot of work remains to be done in constructing efficient algorithms to find the 'smallest' singular triplet (u,σ,v) of (9), from which the optimal solution to the S/W-TLS can be calculated. These algorithms also will have to cope with the fact that the optimization problem (8) is non-convex, hence will typically have several local minima. In [4] [5] we have described an algorithm, the basic inspiration of which is the classical power method:

1. Initialize $u^{[0]}, v^{[0]}$;
2. At iteration step k :
   2.1. Calculate $D_{u^{[k]}u^{[k]}}$, $D_{v^{[k]}v^{[k]}}$;
   2.2. Solve sets of linear equations for $u^{[k+1]}$, $v^{[k+1]}$;
   2.3. Normalize $u^{[k+1]}$, $v^{[k+1]}$ properly;
3. Test for convergence .

When convergent, the convergence rate is linear, just like the power method. For sure, a lot of progress could be made here. For instance, the power method can be interpreted as an alternating weighted least squares optimization method, with 'variable metrics' D_{uu} and D_{vv} in each iteration step. Maybe this interpretation helps in analysing algorithms and their convergence behavior. IQML-like algorithms have been analysed in [11]. One could also try to gain insight by analysing the so-called gradient flows that can be derived from the optimization formulation (8) (see e.g. [2]), in which ideas from differential geometry could play a role (the matrices D_{uu} and D_{vv} define positive definite metrics in the tangent spaces of certain manifolds). For large problems, one can also exploit the structure of the matrices D_{uu} and D_{vv} and of the data matrix A (e.g. by using displacement rank notions), in order to speeden up the calculations. Ideas in this direction and in this context have been proposed in [14]. Also algorithms to solve the rank-2 or higher reduction still need to be developed.

As for the statistical properties of the S/W-TLS estimates, we have to distinguish between the so-called 'incidental parameters' (the elements of the structured least squares rank deficient matrix B) and the 'model parameters', the elements of the vector v. We conjecture that, under the assumptions put forward in Section 3, the error $\tilde{v}$ on the estimate v, for fixed q and $p \to \infty$, asymptotically behaves like $\tilde{v} \sim \mathcal{N}(0, (A^T D_{vv}^{-1} A - D_{uu}\sigma^2)^\dagger)$. Some inspiration may be collected from [10]. As $p \to \infty$, it could also be interesting to investigate recursive algorithms to update the solutions of (9).

Finally, while in this paper, we have presented the 'scalar' versions of the identification problems (i.e. for systems with scalar inputs and/or outputs), the derivation and analysis of the S/W-TLS problems corresponding to systems with vector inputs and outputs, remains to be done. There is a close relation with so-called global total least squares problems in the 'behavioral' framework (see e.g. [9]).

References

[1] Bultheel A., De Moor B. *Rational approximations in linear systems and control.* Journal of Computational and Applied Mathematics, Special Issue on Numerical Analysis in the 20-th Century, Vol. I: Approximation Theory, Vol. 121, September 2000, pp.355-378.

[2] Cambre J., De Moor B. *Solving the structured total least squares problem using the Riemannian SVD.* Mathematical Theory of Networks and Systems, A. Beghi, L. Finesso, G. Picci (Editors), MTNS 1998, Padua, Italy, pp.715-718.

[3] 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.

[4] De Moor B. *Structured total least squares and L_2 approximation problems,* Lin. Alg. Appl., Spec. issue on Numerical Linear Algebra in Systems & Control, (eds: Van Dooren, Ammar, Nichols, Mehrmann), 1993, vol. 188-189, pp. 163-207.

[5] De Moor B. *Total least squares for affinely structured matrices and the noisy realization problem,* IEEE Transactions in Signal Processing, 1994, vol. 42, no. 11, pp. 3104-3113.

[6] 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.

[7] De Moor B. *Linear system identification, structured total least squares and the Riemannian SVD.* in *Recent Advances in Total Least Squares Techniques and Errors-in-Variables modelling,* ed. S. Van Huffel, SIAM, 1997, pp.225-239.

[8] De Moor B. *Structured total least squares for Hankel matrices.* in 'Communications, Computation, Control and Signal Processing', eds: A. Paulraj, V. Roychowdhury, C. Schaper, Kluwer Academic Publishers, 1997, pp.243-258,

[9] Heij C., Scherrer W. *Consistency of system identification by global total least squares.* Automatica, 35, 1999, pp.993-1008.

[10] Karlsson E., Söderström T., Stoica P. *The Cramer-Rao lower bound for noisy input-output systems.* Signal Processing, 80, 2000, pp.2421-2447.

[11] Lemmerling P., Vanhamme L., Van Huffel S., De Moor B. *IQML-like algorithms for solving structured total least squares problems.* Signal Processing, 81, 2001, pp.1935-1945.

[12] Lemmerling P., De Moor B. *Misfit versus latency.* Dept. EE, Katholieke Universiteit Leuven, SISTA Tech.Rep. 1999-76, Accepted for publication in Automatica 2001.

[13] Isidori A., Ruberti A. *Realization theory of bilinear systems*, in 'Geometric Methods in System Theory', edited by D.Q.Mayne and R.W. Brockett, Nato Advanced Study Institute Series, 1973, pp.83-130.

[14] Mastronardi N. *Fast Algorithms for Hankel/Toeplitz structured total least norm problems and applications in parameter estimation.* PhD Thesis, Department of Electrical Engineering, Katholieke Universiteit Leuven, May 2001.

[15] Van Huffel S., Vandewalle J. *The Total Least Squares Problem: Computational Aspects and Analysis.* SIAM Philadelphia, 1991.

[16] Van Huffel S. (Ed.). *Recent Advances in Total Least Squares Techniques and Errors-in-Variables modelling*, SIAM, 1997, pp.225-239.

[17] Wonham M. *Linear Multivariable Control. A Geometric Approach.* New York Springer, 1974.

Acknowledgments

Dr.Bart De Moor is a full professor at the Katholieke Universiteit Leuven. This work was supported by grants and funding from the Research Council of the K.U.Leuven (PhD/postdoc grants, IDO, GOA Mefisto-666), the Foundation for Scientific Research Flanders (FWO)(PhD and postdoc grants and projects (G.0115.01, G.0197.02, G.0407.02), research communities (ICCoS, ANMMM)), the Flemish regional government (Bilateral Coll., IWT-project Soft4s, Eureka projects Synopsis, Impact, FLiTE, STWW-project Genprom, GBOU-project McKnow, IWT PhD grants), the Belgian Federal Government (DWTC, IUAP IV-02, Sustainable Development MD/01/024), the European Commission (TMR Alapedes, Ernsi) and industry supported direct contract research (Electrabel, ELIA, ISMC, Data4s).