Software package for mosaic-Hankel structured low-rank approximation

Konstantin Usevich and Ivan Markovsky

Abstract—This paper presents the SLRA package (`http://slra.github.io`)—C software with interface to MATLAB, Octave, and R for solving low-rank approximation problems with the following features: mosaic Hankel structured approximating matrix, weighted 2-seminorm approximation criterion, and fixed and missing elements in the approximating matrix. The package has applications in system identification, machine learning, and computer algebra. The paper gives an overview of the features of the package, including the wrapper functions for system identification (IDENT package) and approximate greatest common divisor computations. An addendum to the paper, available from `http://homepages.vub.ac.be/~imarkovs/slra-demo` includes examples that demonstrate the usage, versatility, and efficiency of the software.

Index Terms—System identification; low-rank approximation; mosaic Hankel matrix; missing data; approximate greatest common factor; software.

I. INTRODUCTION

Structured low-rank approximation is a core problem in system identification, machine learning, and computer algebra. Despite of the huge diversity of problem formulations, the following common in these application areas problem can be extracted.

> *Given* measured data, prior knowledge about the data generating system, and approximation criterion, *find* optimal in the specified sense approximation of the measured data that is consistent with the given prior knowledge.

A key observation (see [1], [2], [3]) is that in system identification, machine learning, and computer algebra the prior knowledge about the data can be expressed as a rank constraint on a structured matrix constructed from the data. Then, the common problem stated informally above becomes the following structured low-rank approximation problem.

- *The measured data* is an n_p-dimensional real vector p.
- *The prior knowledge* is

$$\text{rank}\big(\mathcal{S}(p)\big) \leq r, \quad \text{where } \mathcal{S} : \mathbb{R}^{n_p} \to \mathbb{R}^{m \times n},$$

  is the matrix structure.
- *The approximation criterion* is a weighted 2-seminorm

$$\|\Delta p\|_W := \sqrt{\Delta p^\top W \Delta p},$$

  of the approximation error vector $\Delta p := p - \widehat{p}$, where W is an $n_p \times n_p$ positive semidefinite matrix.

- *The structured low-rank approximation problem* is

$$\begin{aligned} \text{minimize} \quad & \text{over } \widehat{p} \quad \|p - \widehat{p}\|_W \\ \text{subject to} \quad & \text{rank}\big(\mathcal{S}(\bar{p})\big) \leq r. \end{aligned} \quad \text{(SLRA)}$$

Due to the rank constraint, problem (SLRA) is nonconvex. Except for a few notable exceptions (unstructured approximation in the spectral or Frobenius norm, rank-1 weight matrix W, and circulant structure) there are no polynomial time global solution methods for (SLRA). State-of-the-art methods split into convex relaxation and local optimization heuristics. In the special case of (generalized) Hankel and Toeplitz structured matrices, the class of subspace methods [4], [5], [6] provides also an effective and efficient heuristic for structured low-rank approximation.

This paper presents a local optimization-based software package for structured low-rank approximation, hosted at

<p style="text-align:center;"><code>https://github.com/slra/slra/</code></p>

The SLRA package is built upon an implementation of the following methods for structured low-rank approximation: (i) variable projection (kernel representation) [2], [7], [8], [9] and (ii) penalisation (image representation) [10]. In the variable projection approach, the problem is recast as optimization on a Grassmann manifold [11]. The package includes wrapper functions for (a) system identification (the IDENT package) described in [12] and (b) approximate greatest common divisor (AGCD) computations [13].

We give a comprehensive overview of the features implemented in the package. The previous publication [7] is outdated, as it takes into account only the methods of [7], [8], [9]). Here, we describe, in addition, the methods of [9], [10], [11], [12], [13]. Thus, the paper gathers the description of the parameters and options of the methods scattered in the literature.

The paper is organized as follows. Section II defines the problem solved by the package. Section III describes the solvers included in the package and their options. Section IV presents a MATLAB wrapper function for using the SLRA package for linear time-invariant system identification. Section V describes the wrapper functions for AGCD computations. An addendum to the paper, available from `http://homepages.vub.ac.be/~imarkovs/slra-demo`, contains numerical examples of using the software.

II. PROBLEM FORMULATION

The SLRA package solves problem (SLRA). Typical code is

```
[ph, info] = slra(p, s, r, opt)
```

K. Usevich is with Centre de Recherche en Automatique de Nancy Campus Sciences, BP 70239, 54506 Vandœuvre-lès-Nancy, France; Email: konstantin.usevich@univ-lorraine.fr and I. Markovsky is with Department ELEC, Vrije Universiteit Brussel, Pleinlaan 2, Building K, B-1050 Brussels, Belgium; Email: imarkovs@vub.ac.be.

where p is the data vector, s encodes the structure and the weighted norm, and r is the rank of the approximation.

In the following subsection, we give more details on the structure specification and the weighted 2-norm (contained in the parameter s). The parameter opt determines the solver being used and contains the parameters of the solvers (see Section III for more details).

The output parameter ph contains the computed approximation $\widehat{p}$. The structure info contains additional output information from the optimization, such as:

- info.fmin: the value of the cost function $\|p - \widehat{p}\|_W^2$;
- info.iter: number of iterations;
- info.time: execution time;
- info.Rh: low-rank certificate—a full row rank matrix $\widehat{R} \in \mathbb{R}^{(m-r) \times m}$, such that

$$\widehat{R}\mathcal{S}(\widehat{p}) = 0. \qquad \text{(KER)}$$

Additional information may be returned in info for some of the solvers.

A. Structure specification

The SLRA package deals with affine structures, *i.e.*, structured matrices that form an affine subset of $\mathbb{R}^{m \times n}$.

The package allows for the following matrix structures:

- *General affine structures.*

$$\mathcal{S}(p) = \Phi(S_0 + \sum_{k=1}^{n_p} p_k S_k), \qquad (\mathcal{S})$$

where $S_k \in \mathbb{R}^{m' \times n}$ are basis matrices and $\Phi \in \mathbb{R}^{m \times m'}$ is a full row rank matrix (an identity matrix by default). Moreover, S_k must be orthogonal (*i.e.*, $\langle S_k, S_j \rangle = \text{trace}(S_k S_j^\top) = 0$ for $k \neq j$) and consist only of zeros and ones. The latter condition ensures that the linear part (*i.e.*, $\sum_{k=1}^{n_p} p_k S_k$) contains in each entry either zero, or an element of the vector p, so it can be compactly represented as an integer matrix of the indices of the elements of p.

For example, the matrix structure

$$\mathcal{S}(p) = \begin{bmatrix} 4 & 5 & 6 & p_1 \\ 5 & 6 & 7 & p_2 \\ 6 & p_1 & p_2 & p_3 \end{bmatrix}.$$

can be specified as follows:
```
s.S0  = [4 5 6 0; ...
         5 6 7 0; ...
         6 0 0 0];
s.tts = [0 0 0 1; ...
         0 0 0 2; ...
         0 1 2 3];
```
An example for the matrix Φ will be given next.

- *Mosaic Hankel structures*

The mosaic Hankel [9], [14] structure

$$\mathcal{H}_{\mathbf{m},\mathbf{n}}(p) := \begin{bmatrix} \mathcal{H}_{m_1,n_1}(p^{(1,1)}) & \cdots & \mathcal{H}_{m_1,n_N}(p^{(1,N)}) \\ \vdots & & \vdots \\ \mathcal{H}_{m_q,n_1}(p^{(q,1)}) & \cdots & \mathcal{H}_{m_q,n_N}(p^{(q,N)}) \end{bmatrix},$$

is a generalization of the classical Hankel structure

$$\mathcal{H}_{m,n}(p) := \begin{bmatrix} p_1 & p_2 & p_3 & \cdots & p_n \\ p_2 & p_3 & \cdot^{\cdot^{\cdot}} & & p_{n+1} \\ p_3 & \cdot^{\cdot^{\cdot}} & & & \vdots \\ \vdots & & & & \\ p_m & p_{m+1} & \cdots & & p_{m+n-1} \end{bmatrix}.$$

A mosaic Hankel matrix $\mathcal{H}_{m,n}(p)$ is a block matrix with scalar Hankel blocks. The vectors

$$\mathbf{m} := \begin{bmatrix} m_1 & \cdots & m_q \end{bmatrix} \quad \text{and} \quad \mathbf{n} := \begin{bmatrix} n_1 & \cdots & n_N \end{bmatrix}$$

define the sizes of the scalar Hankel blocks and therefore uniquely specify the structure. The parameter vector $p \in \mathbb{R}^{n_p}$ is the concatenation of chunks $p^{(i,j)}$

$$p = (p^{(1,1)}, \ldots, p^{(q,1)}, \ldots, p^{(1,N)}, \ldots, p^{(q,N)}),$$

where $p^{(i,j)} \in \mathbb{R}^{m_i + n_j - 1}$.

The software package supports mosaic-Hankel-like structures of the form

$$\mathcal{S}(p) := \Phi\mathcal{H}_{\mathbf{m},\mathbf{n}}(p), \qquad (\Phi\mathcal{H}_{\mathbf{m},\mathbf{n}})$$

where Φ is a full row rank matrix, further extending the class of mosaic Hankel matrices to (mosaic) Hankel-like matrices.

For example, a Toeplitz matrix

$$\begin{bmatrix} p_3 & p_4 & p_5 \\ p_2 & p_3 & p_4 \\ p_1 & p_2 & p_3 \end{bmatrix}.$$

can be specified as follows:
```
s.m = 3; s.n = 3; s.Phi = [0 0 1; ...
                           0 1 0; ...
                           1 0 0];
```

A Toeplitz-plus-Hankel matrix can be defined in a similar way, as shown in [7].

B. Approximation criterion

The approximation criteria supported by the package are weighted 2-seminorms $\| \cdot \|_W$, with two types of W.

- A diagonal weight matrix $W = \text{diag}(w)$, where w is a vector with nonnegative real elements from 0 to ∞. Weight $w_i = 0$ corresponds to treating p_i as a missing value. In the other extreme $w_i = \infty$, $\widehat{p}_i$ should match the data p_i (an equality constraint in the optimization problem (SLRA)).

  In general, the weights are specified by a given vector s.w of length n_p. For mosaic-Hankel-like structure, block-wise weights can be specified in two ways:
  - if s.w is of length qN, then the values are the weights for the whole chunks $p^{(i,j)}$ in the vectorisation order;
  - if s.w is of length q, the value of the i-th element of s.w is a weight for $p^{(i,j)}$ for every j.

- Another option is to use a general symmetric positive semi-definite matrix $W \in \mathbb{R}^{n_p \times n_p}$.

III. SLRA PACKAGE: SOLVERS AND OPTIONS

This section gives an overview of methods implemented in the SLRA package. The package contains:

1) fast C++ implementation of the variable projection (VARPRO) method for mosaic Hankel matrices [9];
2) an implementation of the VARPRO method for SLRA with missing data [8]; this method is also called "experimental Matlab solver" in [7];
3) the factorization approach to SLRA based on a penalty method [10].

In Table I the supported structure and weight specification for each methods are listed.

TABLE I
STRUCTURES SUPPORTED BY SOLVERS.

	1)	2)	3)
opt.solver	'c'	'm'	'r'
General affine structure	−	+	+
Mosaic-Hankel-like ($\Phi\mathcal{H}_{\mathbf{m},\mathbf{n}}$)	+	+	+
elementwise weights $w_k \in (0, \infty]$	+	+	+
missing data $w_k = 0$	−	+	+
semi-definite weight matrix W	−	+	+

By default, the efficient VARPRO method is used, if a certain structure/weight specification is not supported, the method 2) is used. The user can select the method by using the field opt.solver.

A. Variable projection methods

In the variable projection [2] approach, the rank constraint is replaced by the equivalent constraint (KER) and problem (SLRA) is restated as an equivalent bi-level optimisation problem

$$\text{minimize} \quad \text{over } R \in \mathbb{R}^{d \times m} \quad f(R)$$
$$\text{subject to} \quad R \text{ has full row rank,} \quad (\text{SLRA}_R)$$

where $d = m - r$ and $f(R)$ is a solution of the (SLRA) problem for fixed kernel matrix R

$$f(R) := \min_{\widehat{p}} \|p - \widehat{p}\|_W^2 \quad \text{subject to} \quad R\mathcal{S}(\widehat{p}) = 0,$$

see [9], [8], [7] for more details.

The idea of the VARPRO method is that the cost function $f(R)$ and its derivatives can be found in a closed form, and the cost function can be minimized using conventional optimisation methods, see [2], [8].

1) Search space and constraints on the kernel: The cost function has an invariance property

$$f(R) = f(UR), \quad \text{for any nonsingular } U \in \mathbb{R}^{d \times d}.$$

Hence, the search space is the Grassmann manifold [11].

In addition, the package allows for linear constraints on the kernel. Two types of constraints can be used.

1) General linear constraint

$$R = \mathcal{R}'(\theta) := \text{vec}_d^{-1}(\theta\Psi), \quad \text{where } \theta \in \mathbb{R}^{n_\theta},$$

and $\Psi \in \mathbb{R}^{n_\theta \times md}$, *i.e.*, it is assumed that R belongs to a linear subspace of $\mathbb{R}^{d \times m}$.

2) Matrix-product constraint

$$R = \mathcal{R}(\theta) := \Theta\Psi, \quad \text{where } \Theta \in \mathbb{R}^{d \times m''},$$

and $\Psi \in \mathbb{R}^{m'' \times m}$ is a full row rank matrix.

The matrix-product linear constraint is a special case of the general linear constraint since

$$\text{vec}^\top(\Theta\Psi) = \text{vec}^\top(\Theta)(\Psi \otimes I_d).$$

Note 1: Both linear constraints are supported by both variable projection methods (efficient C++ solver and experimental MATLAB solver). Note that in [7] it was written that only the MATLAB solver supports the general linear constraint.

Both linear constraints are handled by specifying the field opt.psi. The package determines the constraint under consideration based on the size of the matrix Ψ. By default, it is assumed that $\Psi = I_m$

B. Efficient C++ solver and its options

The efficient C++ solver (opt.solver = 'c') implements the method described in [9]. It can handle mosaic Hankel structures ($\Phi\mathcal{H}_{\mathbf{m},\mathbf{n}}$) and elementwise positive weights, with a linear complexity in the size of data evaluation complexity of cost function and its derivatives (assuming m and d fixed). General structure ($\mathcal{S}$), missing data, or general weight matrices cannot be handled. The solver can use optimization algorithms implemented in the GSL (GNU Scientific Library) library [15] and own implementation of Levenberg-Marquardt method.

1) GSL and constraints: The methods implemented in the GSL library are not designed for optimization on manifolds, hence these methods work under two assumptions:

1) matrix-product constraint;
2) and additional constraint

$$\Theta = \begin{bmatrix} X & -I_d \end{bmatrix}, \quad \text{where } X \in \mathbb{R}^{d \times (m''-d)}. \quad (X)$$

Note 2: The latter assumption parametrizes the Grassmann manifold minus a subset of measure zero. The whole Grassmann manifold can be considered by employing permutations (in Ψ), as explained in [11].

The employed methods are:

- opt.method = 'n...': derivative-free (Nelder-Mead) methods;
- opt.method = 'q...': quasi-Newton methods;
- opt.method = 'l...': Levenberg-Marquardt (see the next subsection).

The second letter of the string opt.method determines the sub-method being used. The possible sub-methods can be found in the documentation of the package, see,

cpp/OptimizationOptions.h.

2) Levenberg-Marquardt methods: The Levenberg-Marquardt methods use sum of squares representation of the cost function

$$f(R) = \|g(R)\|_2^2, \quad (1)$$

where $g : \mathbb{R}^{d \times m} \to \mathbb{R}^{n_g}$ is a vector-valued map. By default, the Cholesky factorisation is used as g (the function g_s in [9, Alg. 6]). Another option (`opt.ls_correction = 1`) is to use the function $g = \Delta p^*(R)$ (see [9, p. 435]).

The LM method from the GSL library (`opt.method = 'l...'`) is applicable only to the matrix-product constraint on the kernel. The general linear constraint can be handled by an own implementation of the 'data-driven local coordinates' modification of the Levenberg-Marquardt method, described in [11, §3.2]. This method (chosen by `opt.method = 'p...'`) automatically takes care of the linear dependencies imposed by the Grassmann manifold and of the general or matrix-product linear constraint.

Note 3: For `opt.method = 'p...'`, in the case of a matrix-product constraint, still the additional constraint (X) is used. If the constraint (X) needs to be avoided, one should specify `opt.avoid_xi = 1`.

3) Additional parameters: The available stopping criteria are:

- `opt.maxiter`: maximum number of iterations;
- `opt.epsgrad`: lower bound on $\|\nabla f\|$
- `opt.epsabs`, `opt.epsrel`: stops if

$$\|X_i - X_{i+1}\| < \text{epsabs} + \text{epsrel} \cdot \|X_i\|.$$

- `opt.maxx`: when parameterisation (X) is used, stops if $\|X\|_\infty$ exceeds `opt.maxx`.

C. Experimental (MATLAB) solver

The MATLAB solver (`opt.solver = 'm'`), based on the results of [8], uses general purpose optimisation routines such as `fmincon/fminunc`.

The constraint on R is imposed in one of the following ways:

- hard constraint (default): $RR^\top = I_d$. In that case, `fmincon` is used. Note that, in this case, the cost function is minimised on the Stiefel manifold.
- "penalty" term (if `opt.method = 'reg'`):

$$f(R) + \gamma \|RR^T - I_d\|_F^2,$$

together with the `fminunc`. In fact, as shown in [8], this penalty term is exact (enforces the hard constraint $RR^\top = I_d$).

Note that:

- If the linear constraints on the kernel are present, they are substituted in $RR^\top = I_d$.
- Standard optimisation options (such as `opt.maxiter`, `opt.tolx`, etc.) can be passed in the structure `opt`.

D. Factorization (alternating least squares) methods

The method `opt.solver = 'r'` is an implementation of the method in [10]. (SLRA) is reformulated as

$$\min_{P \in \mathbb{R}^{m \times r}, L \in \mathbb{R}^{r \times n}} \|p - \mathcal{S}^\dagger \text{vec}(PL)\|_W^2 \\ + \lambda_k \|PL - P_{\mathcal{S}}(PL)\|_F^2, \quad (2)$$

where $P_{\mathcal{S}}$ is the orthogonal projection on the space of structured matrices, $\mathcal{S}^\dagger$ is the orthogonal projection expressed

in terms of a parameter vector in $\mathbb{R}^{n_p}$, and λ_k is the penalty parameter controlling how close the approximation is to the set of structured matrices (for more details, see [10, p. 1186]).

The method proceeds by alternately minimising the cost function for fixed P and L. The hyper parameter λ_k depends on the iteration and is gradually increasing. The additional input parameters for this method are:

- `opt.P_init`: initial approximation for P (default: based on the SVD);
- `opt.lamba_init`: initial value of λ_k;
- `opt.lambda_max`: maximal value of λ_k;
- `opt.max_inner_iter`: maximal number of inner iterations (the iterations when λ_k does not change);

Additional output parameters include:

- `info.P`, `info.L`: computed P and L matrices;
- `info.lambda`: last value of lambda;
- `info.outer_iterations`: number of of outer iterations (changes of λ_k);
- `info.inner_iterations`: number of inner iterations for each outer iteration;

IV. WRAPPER FUNCTIONS FOR SYSTEM IDENTIFICATION

This section describes the functions of the IDENT package [12] for linear time-invariant system identification. System identification problems can be solved via mosaic Hankel structured low-rank approximation, but a wrapper function is needed in order to convert the system identification problem to the input representation for the SLRA solver.

A. The identification problem

We use the behavioral language [16], where a dynamical system is viewed as a collection of trajectories. Let $\mathcal{L}_{\mathtt{m},\ell}^q$ be the model class of linear-time invariant systems with q manifest variables (inputs and outputs) of bounded complexity (at most $\mathtt{m}$ inputs and lag at most ℓ). Given a system $\mathcal{B} \in \mathcal{L}_{\mathtt{m},\ell}^q$, time series $w_{\mathrm{d}} = \{w_{\mathrm{d}}^1, \ldots, w_{\mathrm{d}}^N\}$, and weights $v = \{v^1, \ldots, v^N\}$, the lack of fit (misfit) between the model $\mathcal{B}$ and the data w_{d} is defined as

$$M(w_{\mathrm{d}}, \mathcal{B}) := \min_{\widehat{w}^1, \ldots, \widehat{w}^N \in \mathcal{B}} \sqrt{\sum_{k=1}^N \|w_{\mathrm{d}}^k - \widehat{w}^k\|_{v^k}^2},$$

Note that $\widehat{w}^i$ is the projection of w_{d}^i on $\mathcal{B}$.

The identification problem solved by the IDENT package is: given a (set of) time series w_{d}, complexity specification $(\mathtt{m}, \ell)$, and weights v, find the model that minimizes the misfit criterion over all models with bounded complexity:

$$\widehat{\mathcal{B}} := \arg \min_{\mathcal{B} \in \mathcal{L}_{\mathtt{m},\ell}^q} M(w_{\mathrm{d}}, \mathcal{B}). \quad \text{(SYSID)}$$

As shown in [1, Section 3.1] (SYSID) yields the maximum likelihood estimator in the errors-in-variables setting [17]

$$w_{\mathrm{d}}^k = \bar{w}^k + \widetilde{w}^k,$$

where the true data $\bar{w}^k$ is a trajectory of a true model $\bar{\mathcal{B}} \in \mathcal{L}_{\mathtt{m},\ell}^q$ and the measurement noise $\widetilde{w}^k$ is zero mean normally distributed measurement noise with variance $\sigma^2(v^k)^{-1}$, *i.e.*,

up to a scaling factor, the weights v are the inverse of the noise covariance. Under additional mild assumptions, the estimator $\widehat{\mathcal{B}}^*$ is consistent and the parameters have asymptotically normal joint distribution [18], [19].

B. Functions of the package

The function `ident` solves the approximate identification problem (SYSID), and `misfit` computes the misfit $M(w_d, \mathcal{B})$. They implement the following mappings:

$$\text{ident: } (w_d, \mathtt{m}, \ell) \mapsto \widehat{\mathcal{B}}, \text{ where } \widehat{\mathcal{B}} \text{ (specified by}$$

either input/state/output or kernel representation) is a locally optimal solution of (SYSID)

```
[sysh, info, wh, xini] = ...
                  ident(w, m, ell, opt)
```

$\text{misfit: } (w_d, \mathcal{B}) \mapsto (M, \widehat{w})$, where M is the misfit between $\mathcal{B}$ (specified by either input/state/output or kernel representation) and w_d, and $\widehat{w}$ is the optimal approximation of w_d within $\mathcal{B}$ (the smoothed trajectory)

```
[M, wh, xini] = misfit(w, sysh, opt)
```

Main parameters:

- `w` is the given set of time series w_d — a real MATLAB array of dimension $T \times q \times N$, where T is the number of samples, q is the number of variables, and N is the number of time series. In case of multiple experiments of different duration, `w` should be specified as a cell array with N cells, each one of which is a $T_i \times q$ matrix containing the ith time series, *i.e.*,

$$\mathtt{w}(\mathtt{t}, :, \mathtt{k}) = \left(w^k(t)\right)^\top \quad \text{or} \quad \mathtt{w\{k\}}(\mathtt{t},:) = \left(w^k(t)\right)^\top.$$

- (`m`, `ell`) is the complexity specification (input dimension `m` and lag ℓ).
- `sysh` is an input/state/output representation of the identified or validated system $\widehat{\mathcal{B}}$, given by an `ss` object, or a parameter $\widehat{R}$ of a kernel representation(KER) if `opt.ss` is set to zero.
- `info` is a structure, containing output information from the structured low-rank approximation solver: `info.M` is the misfit $M(w_d, \widehat{\mathcal{B}})$, `info.time` is the execution time, and `info.iter` is the number of iterations.
- `M` is the misfit $M(w_d, \widehat{\mathcal{B}})$.
- `wh` is the optimal approximating set of time series $\widehat{w}$.
- `xini` is the initial condition for `wh`.

C. Optional parameters

`opt` is an optional argument specifying exact variables, exact initial conditions, and options for the optimization solver, used by the `ident` function. The options are passed to the functions `ident` and `misfit` as fields of a structure.

- `'exct'` (default value `[]`) — q-dimensional vector or N-dimensional cell array with q-dimensional vector elements, specifying the indices of the exact variables.
- `'wini'` — specifies exact initial conditions. If `wini` = 0, exact zero initial conditions are specified, *i.e.*, $\mathrm{col}(0, \widehat{w}^k) \in \widehat{\mathcal{B}}$, More generally, `wini` $= w_{\mathrm{ini}}$ is an ℓ samples long trajectory (specified by an $\ell \times q \times N$

array or a N-dimensional cell array of $\ell \times q$ matrices), defining initial conditions for the time series $\widehat{w}$, *i.e.*, $\mathrm{col}(w_{\mathrm{ini}}^k, \widehat{w}^k) \in \widehat{\mathcal{B}}$.

- `'sys0'` — initial approximation: an input/state/output representation of a system, given as an `ss` object, with `m` inputs, `p` $:= q - $ `m` outputs, and order `n` $:= \ell$`p`; or specified by a minimal kernel parameter, *i.e.*, a `p` $\times$ $(\ell+1)q$ matrix Default value is computed by the `slra` function, using unstructured low-rank approximation.
- Arguments allowing the user to specify different optimization algorithms (`'solver'` and `'method'`), control the displayed information (`'disp'`), and change the convergence criteria, see Section III.
- `'ss'` set to zero disables the conversion of $\widehat{R}$ to the state-space representation $(\widehat{A}, \widehat{B}, \widehat{C}, \widehat{D})$.
- `'n'` — order of the identified model. For multiple output system (`p` > 1), the order of the identified model $\widehat{\mathcal{B}} \in \mathcal{L}_{\mathtt{m}, \ell}$ is a multiple of the number of outputs `n` $= $ `p`ℓ. A model of order `p`$(\ell-1) < $ `n` $< $ `p`ℓ can be specified by the optional parameter `opt.n`. In this case, the IDENT package does balanced model order reduction of the identified model $\widehat{\mathcal{B}}$ to obtained a model of order `n`.

V. APPROXIMATE GREATEST COMMON DIVISOR COMPUTATIONS

The problem of finding the AGCD of several polynomials appears, for example, in blind deblurring and distance to controllability problems. As shown in [13], the AGCD problem can be reformulated as a mosaic-Hankel structured low-rank approximation.

Let $\mathcal{P}_n$ denote the space of (real or complex) polynomials of degree n. The AGCD problem solved by the package is the following. Given N polynomials $p_1 \in \mathcal{P}_{n_1}, \ldots, p_N \in \mathcal{P}_{n_N}$ of and degree d, find the approximating polynomials, which have a common divisor of degree at least d:

$$\text{minimize over } \widehat{p}_1 \in \mathcal{P}_{n_1}, \ldots, \widehat{p}_N \in \mathcal{P}_{n_N} \sum_{k=1}^{N} \|p_k - \widehat{p}_k\|_{v_k}^2$$

$$\text{subject to} \quad \deg \gcd(\widehat{p}_1, \ldots, \widehat{p}_N) \geq d,$$

where $\| \cdot \|_v$ is a weighted norm on the coefficients of the polynomial. Note that $\deg \gcd(\widehat{p}_1, \ldots, \widehat{p}_N) \geq d$ if there exists a polynomial of $\widehat{h}$ degree d (common divisor) such that

$$\widehat{p}_1 = \widehat{g}_1 \widehat{h}, \quad \ldots, \quad \widehat{p}_N = \widehat{g}_N \widehat{h}.$$

The solvers are contained in a sub-directory `agcd` of the SLRA package. The common interface for calling these functions is

```
[ph, info] = gcd_xxx(p, v, d, opt)
```

where the input arguments are:

- `p` is a cell array containing vectors of coefficients of polynomials p_k; the coefficients are given from lowest to highest degree of the monomials;
- `v` is a cell array of the weight vectors of corresponding length;

- d is the degree of the GCD;
- `opt` is the structure containing options, which are passed to the `slra` solver; also the starting point for local optimization can be given by either:
  - `opt.hini`: the initial value for $\widehat{h}$,
  - `opt.gini`: the cell array of initial values for $\widehat{g}_k$.

and the output arguments are:
- `ph` is a cell array of approximating polynomials;
- `info` as for the `slra` solver.

The implemented optimization methods are:
- `gcd_nls`: uses variable projection representation with respect to $\widehat{h}$, handles real polynomials;
- `gcd_nls_complex`: uses variable projection representation with respect to $\widehat{h}$, handles complex polynomials;
- `gcd_cofe`: uses variable projection representation with respect to $\widehat{g}_k$, handles real polynomials;
- `gcd_syl`: uses Sylvester subresultant low-rank approximation, handles real polynomials.

The default initial approximation is computed from a kernel of a Sylvester subresultant matrix. There is alternative way to compute the initial approximation using the matrix pencil, which is invoked by the function `h_ini_mp`. For more details on these methods, see [13].

VI. ACKNOWLEDGMENTS

REFERENCES

[1] I. Markovsky, "Structured low-rank approximation and its applications," *Automatica*, vol. 44, no. 4, pp. 891–909, 2008.

[2] ——, "Recent progress on variable projection methods for structured low-rank approximation," *Signal Processing*, vol. 96PB, pp. 406–419, 2014.

[3] ——, *Low-Rank Approximation: Algorithms, Implementation, Applications*, 2nd ed. Springer, 2019.

[4] P. Van Overschee and B. De Moor, *Subspace identification for linear systems: Theory, implementation, applications*. Boston: Kluwer, 1996.

[5] M. Verhaegen and V. Verdult, *Filtering and system identification: a least squares approach*. Cambridge university press, 2007.

[6] I. Markovsky, J. C. Willems, S. Van Huffel, and B. De Moor, *Exact and Approximate Modeling of Linear Systems: A Behavioral Approach*. SIAM, 2006.

[7] I. Markovsky and K. Usevich, "Software for weighted structured low-rank approximation," *J. Comput. Appl. Math.*, vol. 256, pp. 278–292, 2014.

[8] ——, "Structured low-rank approximation with missing data," *SIAM J. Matrix Anal. Appl.*, vol. 34, no. 2, pp. 814–830, 2013.

[9] K. Usevich and I. Markovsky, "Variable projection for affinely structured low-rank approximation in weighted 2-norms," *J. Comput. Appl. Math.*, vol. 272, pp. 430–448, 2014.

[10] M. Ishteva, K. Usevich, and I. Markovsky, "Factorization approach to structured low-rank approximation with applications," *SIAM J. Matrix Anal. Appl.*, vol. 35, no. 3, pp. 1180–1204, 2014.

[11] K. Usevich and I. Markovsky, "Optimization on a Grassmann manifold with application to system identification," *Automatica*, vol. 50, pp. 1656–1662, 2014.

[12] I. Markovsky, "A software package for system identification in the behavioral setting," *Control Eng. Practice*, vol. 21, no. 10, pp. 1422–1436, 2013.

[13] K. Usevich and I. Markovsky, "Variable projection methods for approximate (greatest) common divisor computations," *Theoretical Computer Science*, vol. 681, pp. 176–198, 2017.

[14] G. Heinig, "Generalized inverses of Hankel and Toeplitz mosaic matrices," *Linear Algebra Appl.*, vol. 216, no. 0, pp. 43–59, Feb. 1995.

[15] "GSL — GNU Scientific Library," 2017. [Online]. Available: www.gnu.org/software/gsl/

[16] J. Polderman and J. C. Willems, *Introduction to Mathematical Systems Theory*. New York: Springer-Verlag, 1998.

[17] T. Söderström, "Errors-in-variables methods in system identification," *Automatica*, vol. 43, pp. 939–958, 2007.

[18] R. Pintelon and J. Schoukens, *System Identification: A Frequency Domain Approach*. Piscataway, NJ: IEEE Press, 2001.

[19] A. Kukush, I. Markovsky, and S. Van Huffel, "Consistency of the structured total least squares estimator in a multivariate errors-in-variables model," *J. Statist. Plann. Inference*, vol. 133, no. 2, pp. 315–358, 2005.

[20] N. Guglielmi and I. Markovsky, "An ODE based method for computing the distance of co-prime polynomials to common divisibility," *SIAM Journal on Numerical Analysis*, vol. 55, pp. 1456–1482, 2017.

[21] A. Fazzi, N. Guglielmi, and I. Markovsky, "An ODE based method for computing the approximate greatest common divisor of polynomials," *Numerical Algorithms*, vol. 81, pp. 719–740, 2018.

[22] I. Markovsky, "A missing data approach to data-driven filtering and control," *IEEE Trans. Automat. Contr.*, vol. 62, pp. 1972–1978, 2017.

[23] D. Knuth, *Literate programming*. Cambridge University Press, 1992.

[24] N. Ramsey, "Literate programming simplified," *IEEE Software*, vol. 11, pp. 97–105, 1994.

[25] C. Dominik, *The org mode 7 reference manual*. Network theory ltd, 2010. [Online]. Available: http://orgmode.org/

[26] J. Buckheit and D. Donoho, *Wavelets and statistics*. Springer-Verlag, 1995, ch. Wavelab and reproducible research.

[27] P. Vandewalle, J. Kovacevic, and M. Vetterli, "Reproducible research in signal processing - what, why, and how," *IEEE Signal Proc. Magazine*, vol. 26, no. 3, pp. 37–47, May 2009.

[28] B. De Moor, P. D. Gersem, B. D. Schutter, and W. Favoreel, "DAISY: A database for identification of systems," *Journal A*, vol. 38, no. 3, pp. 4–5, 1997.

[29] L. Ljung, *System identification toolbox: User's guide*, The MathWorks, 2013.

[30] C. Taylor, D. Pedregal, P. Young, and W. Tych, "Environmental time series analysis and forecasting with the CAPTAIN toolbox," *Env. Modelling&Software*, vol. 22, pp. 797–814, 2007.

[31] J.-P. Marmorat and M. Olivi, "RARL2 software (realizations and rational approximation in L_2 norm)," 2012.

[32] S. Van Huffel, V. Sima, A. Varga, S. Hammarling, and F. Delebecque, "High-performance numerical software for control," *IEEE Control Systems Magazine*, vol. 24, pp. 60–76, 2004.

[33] G. Golub and V. Pereyra, "Separable nonlinear least squares: the variable projection method and its applications," *Institute of Physics, Inverse Problems*, vol. 19, pp. 1–26, 2003.

[34] J. A. Nelder and R. Mead, "A simplex method for function minimization," *Computer J.*, vol. 7, pp. 308–313, 1965.

[35] P.-A. Absil, R. Mahony, R. Sepulchre, and P. V. Dooren, "A Grassmann–Rayleigh quotient iteration for computing invariant subspaces," *SIAM Review*, vol. 44, no. 1, pp. 57–73, 2002.

[36] P.-A. Absil, R. Mahony, and R. Sepulchre, *Optimization Algorithms on Matrix Manifolds*. Princeton, NJ: Princeton University Press, 2008.