

HANKEL MATRIX RANK MINIMIZATION WITH APPLICATIONS TO SYSTEM IDENTIFICATION AND REALIZATION *

MARYAM FAZEL[†], TING KEI PONG[‡], DEFENG SUN[§], AND PAUL TSENG[¶]

In honor of Professor Paul Tseng, who went missing while on a kayak trip on the Jinsha river, China, on August 13, 2009, for his contributions to the theory and algorithms for large-scale optimization.

Abstract. In this paper, we introduce a flexible optimization framework for nuclear norm minimization of matrices with linear structure, including Hankel, Toeplitz and moment structures, and catalog applications from diverse fields under this framework. We discuss various first-order methods for solving the resulting optimization problem, including alternating direction methods, proximal point algorithm and gradient projection methods. We perform computational experiments to compare these methods on system identification problem and system realization problem. For the system identification problem, the gradient projection method (accelerated by Nesterov’s extrapolation techniques) and the proximal point algorithm usually outperform other first-order methods in terms of CPU time on both real and simulated data, for small and large regularization parameters respectively; while for the system realization problem, the alternating direction method, as applied to a certain primal reformulation, usually outperforms other first-order methods in terms of CPU time.

Key words. Rank minimization, nuclear norm, Hankel matrix, first-order method, system identification, system realization

1. Introduction. The matrix rank minimization problem, or minimizing the rank of a matrix subject to convex constraints, has recently attracted much renewed interest. This problem arises in many engineering and statistical modeling applications, where notions of order, dimensionality, or complexity of a model can be expressed by the rank of an appropriate matrix. Thus, choosing the “simplest” model that is consistent with observations or data often translates into finding a matrix with the smallest rank subject to convex constraints. Rank minimization is NP-hard in general, and a popular convex heuristic for it minimizes the nuclear norm of the matrix (the sum of the singular values) instead of its rank [17]. The regularized version of this problem can be written as

$$\min_X \frac{1}{2} \|\mathcal{A}(X) - b\|^2 + \mu \|X\|_*, \quad (1.1)$$

where $X \in \mathbb{R}^{m \times n}$ is the optimization variable and $\mathcal{A} : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^p$ is a linear map, $b \in \mathbb{R}^p$, and $\mu > 0$ is the tradeoff parameter between the nuclear norm and the least squares fitting error. Problem (1.1) has been widely studied and recently a variety of efficient algorithms have been developed [5, 9, 26, 29, 30, 37, 51]. A special case of this problem is the matrix completion problem [7, 8] which has applications

*This work was initiated by the first, the third and the fourth authors in May, 2009 when the third author was spending his sabbatical leave with Professor Paul Tseng at the University of Washington, Seattle. With Professor Tseng’s missing from August 2009, it seems that this work would never be finished. Fortunately, the second author, Professor Tseng’s last PhD student, has kindly agreed to join our force to get this work done.

[†]Department of Electrical Engineering, University of Washington, Seattle, Washington 98195, U.S.A. (mfazel@ee.washington.edu)

[‡] Department of Mathematics, University of Washington, Seattle, Washington 98195, U.S.A. (tkpong@uw.edu)

[§] Department of Mathematics and Risk Management Institute, National University of Singapore, 10 Lower Kent Ridge Road, Singapore. (matsundf@nus.edu.sg)

[¶] Department of Mathematics, University of Washington, Seattle, Washington 98195, U.S.A.

in collaborative filtering and machine learning. In this problem the measurements are simply a subset of the entries of the matrix. The majority of existing work on algorithms for problem (1.1) has concentrated on this special case.

In this paper, we focus on problems where we need to find a matrix X that, in addition to being low-rank, is required to have a certain *linear structure*, for example, (block-)Hankel, (block-)Toeplitz, or moment structure. Hankel (and Toeplitz) structures arise in dynamical systems problems discussed in Section 1.1, while moment structure comes up in Lasserre relaxations for minimizing polynomials [27]. We consider problem (1.1), and represent the desired structure by a linear map $X = \mathcal{H}(y)$, where y is our optimization variable. Note that if $\mathcal{H}(y)$ is a moment matrix we need to add the constraint $\mathcal{H}(y) \succeq 0$.

1.1. Motivating applications.

1.1.1. Applications in linear dynamical systems. Linear time-invariant (LTI) systems have a long and successful history in modeling dynamical phenomena in many fields, from engineering to finance. The goal of fitting an LTI model to observed data gives rise to different classes of optimization problems, depending on whether the model is parametric or black-box, given in time or frequency domain, deterministic or stochastic, as well as on the type of data, e.g., input-output or state measurements (see, e.g., [12, 19, 33]). In all these cases, picking the appropriate model order or complexity, and understanding its tradeoff with the fitting or validation errors is crucial. In the problems described in this section, the system order or complexity can be expressed as the rank of a Hankel-type matrix. We discuss some of these problems in more detail in Sections 4 and 5.

Minimal system realization with time-domain constraints. Consider the problem of designing a discrete-time, linear time-invariant (LTI) dynamical system, directly from convex specifications on the system's response in the time domain; see, e.g., [18, 31]. Such a problem arises in designing filters for signal processing and control applications. The objective is to trade-off the *order* of the linear system with how well the specifications are met. A low-order design is desired since in practice, it translates into a system that is easier and cheaper to build and analyze. Typical specifications are desired rise-time, settling-time, slew-rate, and overshoot of the filter's response to a step input signal. These specifications can be expressed as upper and lower bounds on the step response over a fixed time horizon, say n time samples. Equivalently, they can be written in terms of the impulse response, which translate into linear inequality constraints on the entries of a Hankel matrix whose rank corresponds to the system order or McMillan degree; see, e.g., [50]. Using the nuclear norm heuristic for rank, we get

$$\begin{aligned} \min_y \quad & \|\mathcal{H}(y)\|_* \\ \text{s.t.} \quad & l_i \leq \sum_{k=1}^i y_k \leq b_i, \quad i = 1, \dots, n, \end{aligned} \tag{1.2}$$

where the optimization variable is $y \in \mathbb{R}^{2n}$ with y_i corresponding to the value of the impulse response at time i , l_i and b_i denoting the bounds on the step response given by the specifications, and $\mathcal{H}(y)$ denoting an $n \times n$ Hankel matrix; see [18] for more details. Notice that this problem is not exactly in the form of (1.1); we shall discuss how algorithms proposed in this paper can be extended to tackle this model in Appendix A.

Minimal partial realization. A related problem in linear system theory is the minimal partial realization problem for multi-input, multi-output systems: given a

sequence of matrices H_k , $k = 1, \dots, n$, find a minimal state space model, described by a 3-tuple of appropriately sized matrices (A, B, C) such that $H_k = CA^{k-1}B$ [50, Chapter 6]. This problem can be viewed as a more general version of the above realization problem which handled a single input and a single output. In this problem, the order of the system (minimal size of a state-space representation) is equal to the rank of a *block-Hankel* matrix consisting of the H_k ; see [32, section II.A].

Input-output system identification (system ID). Identifying a linear dynamical system given noisy and/or partial observations of its inputs and outputs, also related to time-series analysis, is a fundamental problem studied in a variety of fields [54, 55], including signal processing, control and robotics; see, e.g., [11, 33]. We will discuss this problem and a Hankel rank formulation for it in detail in Section 4.

An interesting variation of this problem is a case where the output information is limited to a few time points, for example the case with a switched output briefly mentioned in [46, Section 1]. In this setup, our observations are the system output sampled at a fixed time T , after an input signal is applied from $t = 0$ to $t = T$. We make output measurements for several different input signals, observing $y_i(T) = \sum_{t=0}^T a_i(T-t)h(t)$, where the vector a_i is the i th input signal and $h(t)$ denotes the impulse response. Writing this compactly as $y = Ah$ where $A_{ij} = a_i(T-j)$ and $h = [h(0) \dots h(T)]^T$ is the optimization variable, and replacing rank of the Hankel matrix with its nuclear norm, we get a problem of the form (1.1). However, unlike the main system ID problem, in this setup the linear map \mathcal{A} has a nullspace so the fitting error term is not strongly convex; as discussed later, this affects the choice of algorithm for solving the problem computationally.

Stochastic realization. Another fundamental problem in linear system theory is finding a minimal stochastic ARMA (autoregressive moving average) model for a vector random process, given noisy and/or partial estimates of process covariances [12, 34]. The minimal order is the rank of a block-Hankel matrix consisting of the exact covariances. This problem is discussed in detail in Section 5.

1.1.2. Other applications.

Shape from moments estimation. Consider a polygonal region P in the complex plane with ordered vertices z_1, \dots, z_m . Complex moments of P are defined as $\tau_k := k(k-1) \int_P z^{k-2} dx dy$, $\tau_0 = \tau_1 = 0$, and can be expressed as $\tau_k = \sum_{i=1}^m a_i z_i^k$, where m is the number of vertices and a_i are complex constants. The problem of determining P given its complex moments has been studied in [16, 38, 49], and arises in many applications such as computer tomography, where X-ray is used to estimate moments of mass distribution, and geophysical inversion, where the goal is to estimate the shape of a region from external gravitational measurements. Note that the *number* of vertices is equal to the rank of the Hankel matrix consisting of the moments [16, 23]. In practice, often only noisy or partial measurements of the complex moments are available, and the challenge is to find a polygon with the minimum number of vertices that is consistent with the measurements. Formulating this problem as a rank minimization problem, we propose using the nuclear norm heuristic for rank. This leads to a problem of the form (1.1), where the optimization variable is the vector of complex moments τ .

Moment matrix rank minimization for polynomial optimization. Suppose $p(x)$, $x \in \mathbb{R}^n$ is a polynomial of degree d . Denote the corresponding moment matrix by $M(y)$, where y is the vectors of moments, i.e., y_i corresponds to the i th monomial; see [27, 28]. Moment matrices are important in Lasserre's hierarchy of relaxations for polynomial optimization, where a condition on the rank of the moment matrix

in successive relaxations in the hierarchy determines whether the relaxation is exact; see, e.g., [28, Section 5].

In the dual problem of representing a polynomial as a sum of squares of other polynomials [44], the rank of the coefficients matrix equals the minimum number of squared polynomials in the representation, thus the nuclear norm (or trace) penalty helps find simpler representations. Note that in these problems, we also have an additional positive semidefinite constraint on the desired matrix.

Further applications. Another application arises in video inpainting in computer vision, where features extracted from video frames are interpolated by finding a low-rank completion to a Hankel matrix, and help reconstruct the missing frames or occluded parts of a frame [13].

Finally, our problem formulation also gives a relaxation for the problem known as the Structured Total Least Squares problem [10] studied extensively in the controls community; see [35, 36]. Our algorithms are thus applicable to this set of problems as well. A variety of applications are discussed in [36].

1.2. Our contributions. We introduce a flexible optimization framework for nuclear norm minimization of linearly structured matrices, including Hankel, Toeplitz, and moment matrices. We identify and catalog applications from diverse fields, some of which have not been studied from this perspective before; for example, the shape from moments estimation problem. In view of the wide applicability of the model (1.1), it is important to find efficient algorithms for solving it. Along this direction, recently, Liu and Vandenberghe [31] proposed an interior-point method for solving a reformulation of problem (1.1), where they used the SDP representation for the nuclear norm and exploited the problem structure to efficiently solve the Newton system. They applied their algorithm to the system identification and system realization problems mentioned above. The cost per iteration of the algorithm grows roughly as $\mathcal{O}(pqn^2)$, where $y \in \mathbb{R}^n$ and $X = \mathcal{H}(y) \in \mathbb{R}^{p \times q}$, $q \leq p \leq n$.

In this paper, we derive various primal and dual reformulations of problem (1.1) (with $X = \mathcal{H}(y)$), and propose several first-order methods for solving the reformulations. In particular, we show that the alternating direction method and the proximal point algorithm can be suitably applied to solve reformulations of (1.1). These methods have been widely used in the literature recently for solving (1.1), when no linear structure is imposed on X ; see, e.g., [30, 56]. We discuss implementation detail of these methods in Section 3.1 and Section 3.2. For these methods, typically, each iteration involves a singular value decomposition whose cost grows as $\mathcal{O}(p^2q)$ where $X = \mathcal{H}(y) \in \mathbb{R}^{p \times q}$ and $p < q$; see, e.g., [24, Page 254]. Next, in Section 3.3, assuming that $\mathcal{A}^* \mathcal{A}$ is invertible, we show that the (accelerated) gradient projection algorithms can be efficiently applied to solve a dual reformulation of (1.1). In this approach, each iteration also involves a singular value decomposition and there is an explicit upper bound on the number of iterations required to attain a certain accuracy. This solution approach has been considered recently in [39] for solving the system identification problem.

To demonstrate the computational efficiency of our algorithms, we apply them to solve the input-output system identification problem and the stochastic system realization problem. For the system identification problem, we consider both simulated data and real data from the DaISy database [11]. Our computational results show that the accelerated gradient projection algorithm and the proximal point algorithm usually outperform other first-order methods for this application in terms of CPU time for small and large regularization parameter, respectively. We also observe that

these methods significantly outperform the interior point implementation proposed in [31]. For the system realization problem, we consider only simulated data, and our computational results show that the alternating direction method, as applied to a certain primal reformulation of (1.1), usually outperforms other first-order methods in terms of CPU time.

The rest of this paper is organized as follows. In Section 1.3, we introduce notations used in this paper. We then derive primal and dual reformulations of (1.1) in Section 2 and discuss several first-order methods for solving the reformulations in Section 3. In Section 4 and Section 5, we present computational results of our first-order methods for solving the system identification problem and system realization problem, respectively. We give some concluding remarks in Section 6. Finally, we discuss an alternative formulation for modeling the structured matrix rank minimization problem in Appendix A, and provide a convergence proof that covers all versions of proximal alternating direction methods used in this paper in Appendix B.

1.3. Notation. In this paper, \mathbb{R}^n denotes the n -dimensional Euclidean space. For a vector $x \in \mathbb{R}^n$, $\|x\|$ denotes the Euclidean norm of x . The set of all $m \times n$ matrices with real entries is denoted by $\mathbb{R}^{m \times n}$. For any $A \in \mathbb{R}^{m \times n}$, $\|A\|$ denotes the spectral norm of A , $\|A\|_F$ denotes the Fröbenius norm of A , $\|A\|_*$ denotes the nuclear norm of A , which is the sum of all singular values of A , and $\text{vec}(A)$ denotes the column vector formed by stacking columns of A one by one. For two matrices A and B in $\mathbb{R}^{m \times n}$, $A \circ B$ denotes the Hadamard (entry-wise) product of A and B . If a symmetric matrix A is positive semidefinite, we write $A \succeq 0$. Linear maps will be denoted by scripted letters. For a linear map $\mathcal{A} : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^p$, \mathcal{A}^* denotes the adjoint of \mathcal{A} , $\|\mathcal{A}\|$ denotes the spectral norm of \mathcal{A} , while $\sigma_{\max}(\mathcal{A})$ and $\sigma_{\min}(\mathcal{A})$ denote the maximum and minimum singular value of \mathcal{A} , respectively. Finally, we denote the identity matrix and identity map by I and \mathcal{I} respectively, whose dimensions should be clear from the context.

2. Basic problem formulations. Consider the following general Hankel matrix nuclear norm minimization problem.

$$v := \min_y f(y) := \frac{1}{2} \|\mathcal{A}(y) - b\|^2 + \mu \|\mathcal{H}(y)\|_*, \quad (2.1)$$

where $\mathcal{A} : \mathbb{R}^{m \times n(j+k-1)} \rightarrow \mathbb{R}^p$ is a linear map, $b \in \mathbb{R}^p$, $y = (y_0 \ \cdots \ y_{j+k-2})$ is an $m \times n(j+k-1)$ matrix with each y_i being an $m \times n$ matrix for $i = 0, \dots, j+k-2$, and $\mathcal{H}(y) := H_{m,n,j,k}(y)\Upsilon$ with

$$H_{m,n,j,k}(y) := \begin{pmatrix} y_0 & y_1 & \cdots & y_{k-1} \\ y_1 & y_2 & \cdots & y_k \\ \vdots & \vdots & & \vdots \\ y_{j-1} & y_j & \cdots & y_{j+k-2} \end{pmatrix} \in \mathbb{R}^{mj \times nk},$$

and $\Upsilon \in \mathbb{R}^{nk \times q}$. We assume without loss of generality that $\sigma_{\max}(\Upsilon) \leq 1$. In this section, we derive primal and dual reformulations of (2.1)

First, using the substitutions $Y = -\mathcal{H}(y)$ and $z = b - \mathcal{A}(y)$, problem (2.1) can be reformulated as

$$\begin{aligned} \min_{Y,z,y} \quad & p(Y,z) := \frac{1}{2} \|z\|^2 + \mu \|Y\|_* \\ \text{s.t.} \quad & Y + \mathcal{H}(y) = 0, \\ & z + \mathcal{A}(y) = b. \end{aligned} \quad (2.2)$$

From the reformulation (2.2), we can easily write down the Lagrange dual to (2.2) (and hence, equivalently, to (2.1)) as follows.

$$\begin{aligned}
v &= \min_{Y,z,y} \max_{\gamma,\Lambda} \left\{ \frac{1}{2} \|z\|^2 + \mu \|Y\|_* - \langle \Lambda, Y + \mathcal{H}(y) \rangle - \langle \gamma, z + \mathcal{A}(y) - b \rangle \right\} \\
&= \max_{\gamma,\Lambda} \min_{Y,z,y} \left\{ \frac{1}{2} \|z\|^2 + \mu \|Y\|_* - \langle \Lambda, Y + \mathcal{H}(y) \rangle - \langle \gamma, z + \mathcal{A}(y) - b \rangle \right\} \\
&= \max_{\gamma,\Lambda} \min_{Y,z,y} \left\{ \frac{1}{2} \|z\|^2 - \langle \gamma, z \rangle + \mu \|Y\|_* - \langle \Lambda, Y \rangle - \langle \mathcal{H}^*(\Lambda) + \mathcal{A}^*(\gamma), y \rangle + \langle b, \gamma \rangle \right\} \\
&= \max_{\gamma,\Lambda} \left\{ -\frac{1}{2} \|\gamma\|^2 + \langle \gamma, b \rangle : \mathcal{H}^*(\Lambda) + \mathcal{A}^*(\gamma) = 0, \Lambda^T \Lambda \preceq \mu^2 I \right\},
\end{aligned}$$

where the second equality holds because of strong duality [47, Corollary 28.2.2]. The dual problem can thus be rewritten as the following minimization problem.

$$\begin{aligned}
\min_{\gamma,\Lambda} \quad & d(\gamma) := \frac{1}{2} \|\gamma\|^2 - b^T \gamma \\
\text{s.t.} \quad & \mathcal{H}^*(\Lambda) + \mathcal{A}^*(\gamma) = 0, \\
& \Lambda^T \Lambda \preceq \mu^2 I.
\end{aligned} \tag{2.3}$$

Alternatively, noting the fact that the nuclear norm is just the dual norm of the spectral norm, one can derive a reduced dual problem as follows:

$$\begin{aligned}
v &= \min_y \frac{1}{2} \|\mathcal{A}(y) - b\|^2 + \mu \|\mathcal{H}(y)\|_* = \min_y \max_{\Lambda^T \Lambda \preceq \mu^2 I} \frac{1}{2} \|\mathcal{A}(y) - b\|^2 - \langle \Lambda, \mathcal{H}(y) \rangle \\
&= - \min_{\Lambda^T \Lambda \preceq \mu^2 I} d_2(\Lambda) := \sup_y \left\{ \langle \Lambda, \mathcal{H}(y) \rangle - \frac{1}{2} \|\mathcal{A}(y) - b\|^2 \right\},
\end{aligned} \tag{2.4}$$

where the third equality holds because of the compactness of the spectral norm ball [47, Corollary 37.3.2]. In the special case when $\mathcal{A}^* \mathcal{A}$ is invertible, the function d_2 has a closed form representation. Indeed, in this case, writing $\mathcal{R}(\Lambda) := (\mathcal{A}^* \mathcal{A})^{-1} \mathcal{H}^*(\Lambda)$ and $\bar{b} := (\mathcal{A}^* \mathcal{A})^{-1} \mathcal{A}^* b$, we obtain that

$$d_2(\Lambda) = \frac{1}{2} (\langle \mathcal{H}^*(\Lambda), \mathcal{R}(\Lambda) \rangle + 2 \langle \mathcal{H}^*(\Lambda), \bar{b} \rangle + \langle \mathcal{A}^* \bar{b}, \bar{b} \rangle - \|b\|^2).$$

Hence, when $\mathcal{A}^* \mathcal{A}$ is invertible, the reduced dual problem (2.4) is equivalent to

$$\begin{aligned}
\min_{\Lambda} \quad & \frac{1}{2} \langle \mathcal{H}^*(\Lambda), \mathcal{R}(\Lambda) \rangle + \langle \mathcal{H}^*(\Lambda), \bar{b} \rangle + \frac{1}{2} \langle \mathcal{A}^* \bar{b}, \bar{b} \rangle - \frac{1}{2} \|b\|^2 \\
\text{s.t.} \quad & \Lambda^T \Lambda \preceq \mu^2 I.
\end{aligned} \tag{2.5}$$

Before ending this section, we derive an upper bound on the spectral norm of \mathcal{H}^* . Notice that $\mathcal{H}^*(\Lambda) = H_{m,n,j,k}^*(\Lambda \Upsilon^T)$, where for any $W \in \mathbb{R}^{mj \times nk}$

$$\begin{aligned}
H_{m,n,j,k}^*(W) &= H_{m,n,j,k}^* \begin{pmatrix} w_{00} & w_{01} & \cdots & w_{0,k-1} \\ w_{10} & w_{11} & \cdots & w_{1,k-1} \\ \vdots & \vdots & & \vdots \\ w_{j-1,0} & w_{j-1,1} & \cdots & w_{j-1,k-1} \end{pmatrix} \\
&= (w_{00} \quad w_{01} + w_{10} \quad w_{02} + w_{11} + w_{20} \quad \cdots \quad w_{j-1,k-1}) \in \mathbb{R}^{m \times n(j+k-2)}.
\end{aligned} \tag{2.6}$$

It follows from (2.6) that

$$\begin{aligned} \|H_{m,n,j,k}^*(W)\|_F^2 &= \|w_{00}\|_F^2 + \|w_{01} + w_{10}\|_F^2 + \|w_{02} + w_{11} + w_{20}\|_F^2 + \cdots + \|w_{j-1,k-1}\|_F^2 \\ &\leq \|w_{00}\|_F^2 + 2(\|w_{01}\|_F^2 + \|w_{10}\|_F^2) + \cdots + \|w_{j-1,k-1}\|_F^2 \leq \mathbf{r}\|W\|_F^2, \end{aligned}$$

where $\mathbf{r} := \min\{j, k\}$. Combining this estimate with $\sigma_{\max}(\Upsilon) \leq 1$, we obtain that

$$\|\mathcal{H}^*(\Lambda)\|_F^2 \leq \mathbf{r}\|\Lambda\Upsilon^T\|_F^2 \leq \mathbf{r}\|\Lambda\|_F^2,$$

and thus the spectral norm of \mathcal{H}^* is less than or equal to $\sqrt{\mathbf{r}}$.

3. Algorithms. In this section, we discuss several first-order methods for solving (2.2) and (2.3) (and hence (2.1)).

3.1. Alternating direction methods. In this section, we discuss how the alternating direction method (ADM) can be applied to solve (2.2) and (2.3). To apply the ADM for solving (2.2), we first introduce the augmented Lagrangian function

$$L_\beta(Y, z, y, \gamma, \Lambda) = \frac{1}{2}\|z\|^2 + \mu\|Y\|_* - \langle \Lambda, Y + \mathcal{H}(y) \rangle - \langle \gamma, z + \mathcal{A}(y) - b \rangle + \frac{\beta}{2}\|Y + \mathcal{H}(y)\|_F^2 + \frac{\beta}{2}\|z + \mathcal{A}(y) - b\|^2$$

for each $\beta > 0$. In the classical ADM (see, e.g., [2, Section 3.4.4]), in each iteration, we minimize L_β with respect to (Y, z) and then with respect to y , followed by an update of the multiplier (γ, Λ) . While minimizing L_β with respect to (Y, z) admits an easy closed form solution, minimizing L_β with respect to y does not usually have a simple closed form solution due to the complicated quadratic terms. One way to resolve this is to add a proximal term with norm induced by a suitable positive (semi-)definite matrix to “cancel” out the complicated parts. In this approach, we update

$$y^{k+1} = \arg \min_y \left\{ L_\beta(Y^{k+1}, z^{k+1}, y, \gamma^k, \Lambda^k) + \frac{\beta}{2}\|y - y^k\|_{Q_0}^2 \right\},$$

where $\|\cdot\|_{Q_0}$ is the norm induced from the inner product $x^T Q_0 x$,

$$Q_0 := \frac{1}{\sigma} \mathcal{I} - (\mathcal{H}^* \mathcal{H} + \mathcal{A}^* \mathcal{A}) \succ 0 \text{ and } \sigma < \frac{1}{\mathbf{r} + (\sigma_{\max}(\mathcal{A}))^2}$$

so that $\sigma \|\mathcal{H}^* \mathcal{H} + \mathcal{A}^* \mathcal{A}\| \leq \sigma(\|\mathcal{H}^* \mathcal{H}\| + \|\mathcal{A}^* \mathcal{A}\|) < 1$. The convergence analysis of this approach has been considered in [25], for example, in the context of variational inequalities¹. For the sake of completeness, we discuss convergence of this approach in detail in the appendix. We now present our algorithm as follows.

Primal ADM.

Step 0. Input $(y^0, \gamma^0, \Lambda^0)$, $\beta > 0$ and $0 < \sigma < \frac{1}{\mathbf{r} + (\sigma_{\max}(\mathcal{A}))^2}$.

Step 1. Compute the SVD

$$-\mathcal{H}(y^k) + \frac{\Lambda^k}{\beta} = U \Sigma V^T,$$

where U and V have orthogonal columns, Σ is diagonal. Set

$$\begin{aligned} Y^{k+1} &= U \max \left\{ \Sigma - \frac{\mu}{\beta} I, 0 \right\} V^T, \quad z^{k+1} = \frac{1}{1 + \beta} (\gamma^k - \beta \mathcal{A}(y^k) + \beta b), \\ y^{k+1} &= y^k - \sigma \left(-\frac{1}{\beta} (\mathcal{H}^* \Lambda^k + \mathcal{A}^* \gamma^k) + \mathcal{H}^*(\mathcal{H}(y^k) + Y^{k+1}) + \mathcal{A}^*(\mathcal{A}(y^k) + z^{k+1} - b) \right), \\ \gamma^{k+1} &= \gamma^k - \beta(z^{k+1} + \mathcal{A}(y^{k+1}) - b), \quad \Lambda^{k+1} = \Lambda^k - \beta(Y^{k+1} + \mathcal{H}(y^{k+1})). \end{aligned}$$

¹The main motivation of introducing the proximal terms in [25] is to weaken the imposed convergence conditions rather than for the sake of cancellation, as was more recently explained in [57].

Step 2. If a termination criterion is not met, go to Step 1.

The ADM can also be applied to solve the dual problem (2.3). In this approach, we make use of the following augmented Lagrangian function

$$l_\beta(\gamma, \Lambda, y) = \frac{1}{2}\|\gamma\|^2 - b^T\gamma + \langle y, \mathcal{H}^*(\Lambda) + \mathcal{A}^*(\gamma) \rangle + \frac{\beta}{2}\|\mathcal{H}^*(\Lambda) + \mathcal{A}^*(\gamma)\|_F^2,$$

for each $\beta > 0$. Notice that the minimizer of l_β with respect to γ is usually not easy to find and the minimizer with respect to Λ does not always admit a simple closed form solution. Thus, as before, we add suitable proximal terms to cancel out complicated terms. More precisely, we update

$$\begin{aligned} \gamma^{k+1} &:= \arg \min_{\gamma} \left\{ l_\beta(\gamma, \Lambda^k, y^k) + \frac{\beta}{2}\|\gamma - \gamma^k\|_{Q_1}^2 \right\}, \\ \Lambda^{k+1} &:= \arg \min_{\Lambda^T \Lambda \preceq \mu^2 I} \left\{ l_\beta(\gamma^{k+1}, \Lambda, y^k) + \frac{\beta}{2}\|\Lambda - \Lambda^k\|_{Q_2}^2 \right\}, \end{aligned}$$

where

$$Q_1 := \frac{1}{\sigma_1} \mathcal{I} - \mathcal{A}\mathcal{A}^* \succ 0 \text{ and } \sigma_1 < \frac{1}{(\sigma_{\max}(\mathcal{A}))^2}, \quad Q_2 := \frac{1}{\sigma_2} \mathcal{I} - \mathcal{H}\mathcal{H}^* \succ 0 \text{ and } \sigma_2 < \frac{1}{\mathbf{r}},$$

so that $\sigma_1 \|\mathcal{A}\mathcal{A}^*\| < 1$ and $\sigma_2 \|\mathcal{H}\mathcal{H}^*\| < 1$. The algorithm is described as follows.

Dual ADM.

Step 0. Input $(y^0, \gamma^0, \Lambda^0)$, $\beta > 0$, $0 < \sigma_1 < \frac{1}{(\sigma_{\max}(\mathcal{A}))^2}$ and $0 < \sigma_2 < \frac{1}{\mathbf{r}}$.

Step 1. Set

$$\gamma^{k+1} = \frac{\sigma_1}{\sigma_1 + \beta} \left(b + \beta \frac{\gamma^k}{\sigma_1} - \beta \left(\frac{\mathcal{A}(y^k)}{\beta} + \mathcal{A}(\mathcal{H}^*(\Lambda^k) + \mathcal{A}^*(\gamma^k)) \right) \right).$$

Compute the SVD

$$\Lambda^k - \sigma_2 \left(\frac{1}{\beta} \mathcal{H}(y^k) + \mathcal{H}(\mathcal{H}^*(\Lambda^k) + \mathcal{A}^*(\gamma^{k+1})) \right) = U \Sigma V^T,$$

where U and V have orthogonal columns, Σ is diagonal. Set

$$\begin{aligned} \Lambda^{k+1} &= U \min\{\Sigma, \mu I\} V^T, \\ y^{k+1} &= y^k + \beta(\mathcal{H}^*(\Lambda^{k+1}) + \mathcal{A}^*(\gamma^{k+1})). \end{aligned}$$

Step 2. If a termination criterion is not met, go to Step 1.

From Theorem B.1 in the appendix, for the sequence $\{(Y^k, z^k, y^k, \gamma^k, \Lambda^k)\}$ generated from Primal ADM, $\{(Y^k, z^k, y^k)\}$ converges to a solution of the problem (2.2) while $\{(\gamma^k, \Lambda^k)\}$ converges to a solution of the problem (2.3). Similarly, for the sequence $\{(y^k, \gamma^k, \Lambda^k)\}$ generated from Dual ADM, $\{y^k\}$ converges to a solution of (2.1) and $\{(\gamma^k, \Lambda^k)\}$ converges to a solution of (2.3).

3.2. Dual proximal point algorithm. In this section, we discuss how the proximal point algorithm (PPA) can be applied to solve the dual problem. We shall make use of the reduced dual problem (2.4). Fix $\lambda > 0$. For any Λ , define the Moreau-Yosida regularization of d_2 at Λ associated with λ by

$$G_\lambda(\Lambda) = \min_{\Gamma^T \Gamma \preceq \mu^2 I} d_2(\Gamma) + \frac{1}{2\lambda} \|\Gamma - \Lambda\|_F^2.$$

Using this definition and the definition of d_2 in (2.4), we obtain that

$$\begin{aligned} G_\lambda(\Lambda) &= \min_{\Gamma^T \Gamma \preceq \mu^2 I} \sup_y \left\{ \langle \Gamma, \mathcal{H}(y) \rangle - \frac{1}{2} \|\mathcal{A}(y) - b\|^2 \right\} + \frac{1}{2\lambda} \|\Gamma - \Lambda\|_F^2 \\ &= \sup_y \left\{ \min_{\Gamma^T \Gamma \preceq \mu^2 I} \left\{ \langle \Gamma, \mathcal{H}(y) \rangle + \frac{1}{2\lambda} \|\Gamma - \Lambda\|_F^2 \right\} - \frac{1}{2} \|\mathcal{A}(y) - b\|^2 \right\}, \end{aligned}$$

where the second equality holds due to the compactness of the spectral norm ball [47, Corollary 37.3.2]. Furthermore, it is not hard to show that

$$\min_{\Gamma^T \Gamma \preceq \mu^2 I} \left\{ \langle \mathcal{H}(y), \Gamma \rangle + \frac{1}{2\lambda} \|\Lambda - \Gamma\|_F^2 \right\} = \langle \mathcal{H}(y), \Lambda \rangle - \frac{\lambda}{2} \|\mathcal{H}(y)\|_F^2 + \frac{1}{2\lambda} \|\mathcal{P}_\mu(\Lambda - \lambda \mathcal{H}(y))\|_F^2,$$

where $\mathcal{P}_\mu(W)$ is the unique optimal solution to the following convex optimization problem:

$$\min_Z \|Z\|_* + \frac{1}{2\mu} \|Z - W\|_F^2.$$

Thus, $G_\lambda(\Lambda) = \sup_y \Theta_\lambda(y; \Lambda)$, where

$$\Theta_\lambda(y; \Lambda) := \langle \mathcal{H}(y), \Lambda \rangle - \frac{\lambda}{2} \|\mathcal{H}(y)\|_F^2 + \frac{1}{2\lambda} \|\mathcal{P}_\mu(\Lambda - \lambda \mathcal{H}(y))\|_F^2 - \frac{1}{2} \|\mathcal{A}(y) - b\|^2.$$

Recall from the Moreau-Yosida regularization theory that $\mathcal{P}_\mu(\cdot)$ is globally Lipschitz continuous with modulus 1 and that $\|\mathcal{P}_\mu(\cdot)\|_F^2$ is continuously differentiable with $\nabla(\|\mathcal{P}_\mu(Y)\|_F^2) = 2\mathcal{P}_\mu(Y)$. Hence, $\Theta_\lambda(\cdot; \Lambda)$ is a continuously differentiable concave function in y with

$$\nabla_y \Theta_\lambda(y; \Lambda) = \mathcal{H}^*(\Lambda - \lambda \mathcal{H}(y)) - \mathcal{H}^* \mathcal{P}_\mu(\Lambda - \lambda \mathcal{H}(y)) - \mathcal{A}^*(\mathcal{A}(y) - b).$$

In addition, we have that

$$\begin{aligned} &\|\nabla_y \Theta_\lambda(y'; \Lambda) - \nabla_y \Theta_\lambda(y; \Lambda)\|_F \\ &\leq (\lambda \|\mathcal{H}^* \mathcal{H}\| + \|\mathcal{A}^* \mathcal{A}\|) \|y' - y\|_F + \|\mathcal{H}^* \mathcal{P}_\mu(\Lambda - \lambda \mathcal{H}(y')) - \mathcal{H}^* \mathcal{P}_\mu(\Lambda - \lambda \mathcal{H}(y))\|_F \\ &\leq (\lambda \mathbf{r} + (\sigma_{\max}(\mathcal{A}))^2) \|y' - y\|_F + \lambda \mathbf{r} \|y' - y\|_F, \end{aligned}$$

which implies that $\nabla_y \Theta_\lambda(\cdot; \Lambda)$ is Lipschitz continuous with Lipschitz modulus

$$2\lambda \mathbf{r} + (\sigma_{\max}(\mathcal{A}))^2. \quad (3.1)$$

We are now ready to describe the PPA for solving the dual problem (2.3).

Dual PPA.

Step 0. Input (y^0, Λ^0) and $\lambda_0 > 0$.

Step 1. (Find an approximate maximizer $y^{k+1} \approx \arg \max \Theta_{\lambda_k}(y; \Lambda^k)$.)

Input $u^0 := y^k$. Let $s_k, \text{intol}_k > 0, 1 > t, \sigma > 0$.

While $\|\nabla_y \Theta_{\lambda_k}(u^l; \Lambda^k)\| > \text{intol}_k$ **do**

(a) Let \bar{s}_l be the largest element of $\{s_k, ts_k, t^2 s_k, \dots\}$ satisfying

$$\Theta_{\lambda_k}(u[s]; \Lambda^k) > \Theta_{\lambda_k}(u^l; \Lambda^k) + \sigma s \|\nabla_y \Theta_{\lambda_k}(u^l; \Lambda^k)\|_F^2,$$

where $u[s] = u^l + s \nabla_y \Theta_{\lambda_k}(u^l; \Lambda^k)$.

(b) Set $u^{l+1} \leftarrow u[\bar{s}_l]$, $l \leftarrow l + 1$.

End (while)

Set $y^{k+1} \leftarrow u^{l+1}$.

Step 2. Compute the SVD

$$\Lambda^k - \lambda_k \mathcal{H}(y^{k+1}) = U \Sigma V^T.$$

Set

$$\Lambda^{k+1} = U \min(\Sigma, \mu I) V^T.$$

Step 3. If a termination criterion is not met, update λ_k . Go to Step 1.

3.3. Dual gradient projection methods. In this section we assume that $\mathcal{A}^* \mathcal{A}$ is invertible. Recall that in this case, the dual of (2.1) is given by (2.5). Moreover, we have

$$\begin{aligned} \|\nabla d_2(\Lambda_1) - \nabla d_2(\Lambda_2)\|_F^2 &= \|\mathcal{H}((\mathcal{A}^* \mathcal{A})^{-1} \mathcal{H}^*(\Lambda_1 - \Lambda_2))\|_F^2 \leq \mathbf{r} \|(\mathcal{A}^* \mathcal{A})^{-1} \mathcal{H}^*(\Lambda_1 - \Lambda_2)\|_F^2 \\ &\leq \frac{\mathbf{r}}{(\sigma_{\min}(\mathcal{A}^* \mathcal{A}))^2} \|\mathcal{H}^*(\Lambda_1 - \Lambda_2)\|_F^2 \leq \left(\frac{\mathbf{r}}{\sigma_{\min}(\mathcal{A}^* \mathcal{A})} \right)^2 \|\Lambda_1 - \Lambda_2\|_F^2. \end{aligned}$$

This shows that the gradient of d_2 is Lipschitz continuous with Lipschitz constant

$$L_D := \frac{\mathbf{r}}{\sigma_{\min}(\mathcal{A}^* \mathcal{A})}.$$

Since the projection onto the feasible set of (2.5) is simple, we can apply the gradient projection (GP) methods to solve (2.5). One simple version is described as follows.

Dual GP.

Step 0. Input Λ^0 such that $\Lambda^{0T} \Lambda^0 \preceq \mu^2 I$ and $L > \frac{L_D}{2}$.

Step 1. Compute the SVD

$$\Lambda^k - \frac{1}{L} \nabla d_2(\Lambda^k) = U \Sigma V^T,$$

where U and V have orthogonal columns, Σ is diagonal. Set

$$\Lambda^{k+1} = U \min\{\Sigma, \mu I\} V^T.$$

Step 2. If a termination criterion is not met, go to Step 1.

The iterate generated by the Dual GP satisfies

$$d_2(\Lambda^k) - v = \mathcal{O}\left(\frac{L}{k}\right);$$

see, e.g., [53, Theorem 1]. Hence, for faster convergence, a smaller L is favored. Also, note that if y^* and Λ^* are solutions to (2.1) and (2.5) respectively, then we can see from (2.4) that

$$y^* = \mathcal{R}(\Lambda^*) + \bar{b},$$

where $\mathcal{R}(\Lambda) := (\mathcal{A}^* \mathcal{A})^{-1} \mathcal{H}^*(\Lambda)$ and $\bar{b} := (\mathcal{A}^* \mathcal{A})^{-1} \mathcal{A}^* b$. Hence, the sequence

$$y^k := \mathcal{R}(\Lambda^k) + \bar{b} \tag{3.2}$$

converges to a solution to (2.1), which can be used to check for termination; see Section 4.1.

The Dual GP can be accelerated using Nesterov's extrapolation techniques (see, e.g., [40–43, 53]). This method has also been used in [45, 51] for nuclear norm related problems. The method is described below.

Dual AGP.

Step 0. Input Λ^0 such that $\Lambda^{0T}\Lambda^0 \preceq \mu^2 I$ and $L \geq L_D$. Initialize $\Lambda^{-1} = \Lambda^0$ and $\theta_{-1} = \theta_0 = 1$. Go to Step 1.

Step 1. Set

$$\Psi^k = \Lambda^k + \left(\frac{\theta_k}{\theta_{k-1}} - \theta_k \right) (\Lambda^k - \Lambda^{k-1}).$$

Compute the SVD

$$\Psi^k - \frac{1}{L} \nabla d_2(\Psi^k) = U \Sigma V^T,$$

where U and V have orthogonal columns, Σ is diagonal. Update

$$\Lambda^{k+1} = U \min\{\Sigma, \mu I\} V^T, \quad \theta_{k+1} = \frac{\sqrt{\theta_k^4 + 4\theta_k^2} - \theta_k^2}{2}.$$

Step 2. If a termination criterion is not met, go to Step 1.

The sequence generated from the Dual AGP satisfies

$$d_2(\Lambda^k) - v = \mathcal{O}\left(\frac{L}{k^2}\right);$$

see, e.g., [53, Theorem 1]. Hence, for faster convergence, a smaller L is favored. To generate a suitable primal variable at each iteration, instead of setting $y^k = \mathcal{R}(\Lambda^k) + \bar{b}$ as above, we can also initialize y^0 to be the zero matrix and update the sequence in Step 1 according to

$$y^{k+1} = (1 - \theta_k)y^k + \theta_k(\mathcal{R}(\Psi^k) + \bar{b}).$$

For this choice of y^k , it can be shown, following the proof of [52, Corollary 2], that

$$0 \leq f(y^{k+1}) + d_2(\Lambda^{k+1}) = \mathcal{O}(L\theta_k^2).$$

Since $\theta_k \leq \frac{2}{k+2}$, the duality gap falls below a given tolerance $tol > 0$ after at most $\mathcal{O}\left(\sqrt{\frac{L}{tol}}\right)$ iterations. In our implementation of Dual AGP in the next section, the latter choice of y^k usually does not lead to faster convergence and thus we only report results using (3.2).

4. System identification. In this section, we consider the problem of identifying a linear dynamical system from observations of its inputs and outputs. Given a sequence of inputs $u_t \in \mathbb{R}^p$ and measured (noisy) outputs $\tilde{y}_t \in \mathbb{R}^m$, $t = 1, \dots, N$, the goal is to find a discrete-time, linear time-invariant state space model,

$$\begin{aligned} x_{t+1} &= Ax_t + Bu_t, \\ y_t &= Cx_t + Du_t, \end{aligned} \tag{4.1}$$

that satisfies $y_t \approx \tilde{y}_t$ and is low-order (i.e., corresponds to a low-dimensional state vector $x_t \in \mathbb{R}^r$). To determine the model, we need to find the A, B, C, D matrices, the initial state x_0 , and the model order r . As described in [31, Eq. (23)] (see also [32]), under reasonable assumptions, the minimal model order is equal to the rank of the matrix $H_{m,1,r+1,N+1-r}(y)U^\perp$, where $U^\perp \in \mathbb{R}^{(N+1-r) \times q}$ is a matrix whose columns

form an orthogonal basis of the null space of $H_{p,1,r+1,N+1-r}(u)$, and u is the input to the system (see [31] for further details). Upon relaxing the rank function to nuclear norm, the tradeoff between the fitting error and the model order can be captured by the optimization problem,

$$\min_y \frac{1}{2} \|y - \tilde{y}\|_F^2 + \mu \|H_{m,1,r+1,N+1-r}(y)U^\perp\|_*, \quad (4.2)$$

where $\tilde{y} \in \mathbb{R}^{m \times (N+1)}$, $N \geq 1$, $r \geq 0$, $\mu > 0$,

$$H_{m,1,r+1,N+1-r}(y) = \begin{pmatrix} y_0 & y_1 & \cdots & y_{N-r} \\ y_1 & y_2 & \cdots & y_{N-r+1} \\ \vdots & \vdots & & \vdots \\ y_r & y_{r+1} & \cdots & y_N \end{pmatrix} \in \mathbb{R}^{m(r+1) \times (N+1-r)}.$$

This problem corresponds to (2.1) with $\Upsilon = U^\perp$, $\mathcal{A}(y) = \text{vec}(y)$ and $b = \text{vec}(\tilde{y})$. Thus

$$\sigma_{\max}(\mathcal{A}) = \sigma_{\max}(\Upsilon) = 1, \quad \mathcal{H}^*(\Lambda) = H_{m,1,r+1,N+1-r}^*(\Lambda(U^\perp)^T).$$

Note that the solution set of (4.2) is nonempty since the function $y \mapsto \|y - \tilde{y}\|_F^2$ is coercive.

4.1. Computational results. In this section, we compare different algorithms for solving (4.2) on random and real data. Specifically, we consider Primal and Dual ADM, Dual PPA, Dual GP, and Dual AGP. Moreover, taking into account the fact that $\mathcal{A} = \mathcal{I}$, we also consider a variant of Primal ADM (referred to as Primal ADM2) by considering the following augmented Lagrangian function

$$\ell_\beta(Y, y, \Lambda) := \frac{1}{2} \|y - \tilde{y}\|_F^2 + \mu \|Y\|_* - \langle \Lambda, Y + \mathcal{H}(y) \rangle + \frac{\beta}{2} \|Y + \mathcal{H}(y)\|_F^2,$$

and replacing the minimization with respect to y by

$$y^{k+1} := \arg \min_y \left\{ \ell_\beta(Y^{k+1}, y, \Lambda^k) + \frac{\beta}{2} \|y - y^k\|_{Q_3}^2 \right\},$$

where

$$Q_3 := \frac{1}{\sigma} \mathcal{I} - \mathcal{H}^* \mathcal{H} \succ 0 \text{ and } \sigma < \frac{1}{\mathbf{r}},$$

so that $\sigma \|\mathcal{H}^* \mathcal{H}\| < 1$; as well as a variant of Dual ADM with $Q_1 = 0$ (referred to as Dual ADM2). The convergence of these two variants follows from Theorem B.1.

We initialize all algorithms except Dual PPA at the origin, where the latter algorithm is initialized at an approximate solution obtained from Dual AGP². We terminate the algorithms by checking relative duality gap with $tol = 1e - 4$, i.e.,

$$\frac{\min\{f(y^k), f(\hat{y}^k)\} + d(-\mathcal{H}^*(\mathcal{P}(\Lambda^k)))}{\max\{1, |d(-\mathcal{H}^*(\mathcal{P}(\Lambda^k)))|\}} < 1e - 4 \text{ or } \frac{f(y^k) + d_2(\Lambda^k)}{\max\{1, |d_2(\Lambda^k)|\}} < 1e - 4,$$

²We terminate Dual AGP by checking relative duality gap with $tol = 5e - 3$, checked every 10 iterations. We also early terminate the algorithm if the change in Fröbenius norm of successive iterates is small ($< 1e - 8$ for each variable) or the maximum number of iteration hits 2000.

where $\{(y^k, \Lambda^k)\}$ are defined in each of Sections 3.1, 3.2 and 3.3, $\mathcal{P}(\Lambda^k)$ is the projection of Λ^k onto the spectral norm ball with radius μ , and $\hat{y}^k := \tilde{y} + \mathcal{H}^*(\mathcal{P}(\Lambda^k))$ ³. The termination criterion is checked every 10 iterations except for Dual PPA, where we do this every (outer) iteration, as the relative duality gap is required for updating λ_k and $intol_k$.⁴ We also early terminate the algorithms if the change in Fröbenius norm of successive iterates is small ($< 1e - 8$ for each variable) or the maximum number of iteration hits 2000. We set $\beta = \frac{\mu r}{2\sigma_{\max}(\tilde{y})}$ and $\sigma = \frac{0.95}{r+1}$ for Primal ADM. We use the same β for Primal ADM2 and set $\sigma = \frac{0.95}{r}$. Furthermore, we set $\beta = \frac{\sigma_{\max}(\tilde{y})}{16\mu r}$, $\sigma_1 = 0.95$ and $\sigma_2 = \frac{0.95}{r}$ for Dual ADM and use the same β and σ_2 for Dual ADM2. We take $L = \frac{L_D}{1.95}$ for Dual GP and $L = L_D$ for Dual AGP. All codes are written in Matlab and run on a Sunfire X4440, with 4 Opteron 8384 quad-core CPUs and 32G of RAM, equipped with CentOS 5.8 and Matlab 7.12.

4.1.1. Random data. We generate randomly a matrix $u = (u_0 \ \cdots \ u_N) \in \mathbb{R}^{p \times (N+1)}$ with standard Gaussian entries (i.e., zero mean and variance one), and let \bar{r} be the true order of the system. We then generate matrices $A \in \mathbb{R}^{\bar{r} \times \bar{r}}$, $B \in \mathbb{R}^{\bar{r} \times p}$, $C \in \mathbb{R}^{m \times \bar{r}}$ and $D \in \mathbb{R}^{m \times p}$ with i.i.d. standard Gaussian entries, and normalize them to have spectral norm 1. We also generate a vector $x_0 \in \mathbb{R}^{\bar{r}}$, again with standard Gaussian entries. The output $\bar{y} = (\bar{y}_0 \ \cdots \ \bar{y}_N) \in \mathbb{R}^{m \times (N+1)}$ is then generated using a state-space model: for each $t = 0, \dots, N$,

$$\begin{aligned} x_{t+1} &= Ax_t + Bu_t \\ \bar{y}_t &= Cx_t + Du_t. \end{aligned}$$

To model the measurement noise, we add noise to the output \bar{y} to get $\tilde{y} = \bar{y} + \sigma\epsilon$, where ϵ has Gaussian entries with mean 0 and variance 1, and $\sigma > 0$. Finally, U^\perp is a matrix whose columns form an orthonormal basis of the nullspace of $H_{p,1,2\bar{r}+2,N-2\bar{r}}(u)$. Note that in theory, we require the r used in determining the size of the Hankel matrix to be larger than the true order of the system. However in practice, we often don't know the true system order, and only have a guess or estimate for it. Therefore, when we set the size of the Hankel matrix in our problem, as a rule of thumb, we use roughly twice the estimated order; i.e., $r = 2\bar{r} + 1$.

In the tests below, we consider $p = 5$, $m = 5, 10$, $\bar{r} = 10, 20$ and $N + 1 = 2000, 4000$. We pick $\sigma = 5e - 2$, which corresponds roughly to a 5% noise. The statistics of the test problems used are reported in Table 4.1⁵. We run our algorithms for $\mu = 1e - 2, 1e - 1, 1, 10$. We observed that Primal ADM2 usually outperforms Primal ADM, and the two Dual ADMs perform similarly. Thus, we do not report the performances of Primal ADM and Dual ADM.

Our computational results are reported in Table 4.2, where **iter** stands for the number of iterations, **cpu** is the CPU time taken and **obj** is the primal objective value at termination corresponding to each algorithm. The words ‘‘Primal’’ and ‘‘Dual’’ are abbreviated as ‘‘P.’’ and ‘‘D.’’ respectively. For Dual PPA, the CPU time and number

³This is modeled after (3.2).

⁴For Dual PPA, we set $t = 0.3$, $\sigma = 1e - 4$, $s_0 = \frac{1.95}{L}$, with L given in (3.1). For each fixed k and any $l \geq 1$, we set $s_l = 1.11s_{l-1}$ if $\bar{s}_{l-1} = s_k$, and $s_l = s_{l-1}$ otherwise. The inner iterations are terminated when l hits 1000. The parameters λ_k are initialized at $\lambda_0 = 1$ and the update follows the rule that $\lambda_{k+1} = 2\lambda_k$ if $gap_k > 0.95gap_{k-1}$, where gap_k denotes the relative duality gap in the k th outer iteration. The $intol_k$ decreases from 0.04 based on the value and the decrease of gap_k , and is bounded below by $1e - 3$.

⁵Note that each of our algorithm requires an SVD of a matrix of size $m(r+1) \times q$ in each iteration.

of iterations for Dual AGP used for initialization are in parenthesis, and the CPU time not in parenthesis refers to the total runtime. The word “max” denotes the maximum number of iterations. The fastest algorithm(s) in each instance is highlighted in bold. We see that the gradient projection algorithms usually work best when μ is small, with Dual AGP usually more robust (i.e., solving all instances within 2000 iterations) than Dual GP. Furthermore, Dual PPA usually works best when μ is large.

TABLE 4.1
Parameters for the randomly generated test problems.

P	$N + 1$	\bar{r}	m	q
1	2000	10	5	1869
2	2000	10	10	1869
3	2000	20	5	1749
4	2000	20	10	1749
5	4000	10	5	3869
6	4000	10	10	3869
7	4000	20	5	3749
8	4000	20	10	3749

4.1.2. Real data from DaISy database. In this section, we consider eight benchmark problems from DaISy (Database for the Identification of Systems) [11]. A brief description of the data is given in the Table 4.3. Each data set is given in form of a y (output) and a u (input). We take the first 25% of the inputs and outputs for testing purposes. We set $\bar{r} = 20$ (hence, $r = 2\bar{r} + 1 = 41$) and generate the matrix U^\perp so that the column space forms an orthonormal basis of the nullspace of $H_{p,1,2\bar{r}+2,N-2\bar{r}}(u)$, where $u = (u_0 \ \dots \ u_N) \in \mathbb{R}^{p \times (N+1)}$. We apply the five algorithms from Table 4.2 to solve these problems for different values of μ , using the same parameters as in the previous section. The results are reported in Table 4.4, where **iter** and **cpu** are as in Table 4.2, while **nr** denotes the number of singular values of $\mathcal{H}(y^*)$ that are larger than $0.005 \sigma_{\max}(\mathcal{H}(y^*))$, and **err** denotes the fitting error measured by $\|y^* - \tilde{y}\|_F$; here, y^* is the approximate optimal solution of (4.2) obtained by the algorithms. We see that Dual AGP and Dual PPA work best. Furthermore, we see that **nr** decreases with μ , while **err** increases as μ increases.

To further illustrate the performance of the first-order methods on real data, we consider the first problem in Table 4.3 and attempt to identify the system order using solutions obtained from Dual AGP. For comparison with the work [31], we use the same number of data points for identification and validation as in [31, Table 5.1]; i.e., 200 points for identification and 600 points for validation. We draw parameters μ between $1e - 4$ and 10, set $r = \lfloor \frac{N+1}{p+m+1} \rfloor$ as in [31], solve the corresponding problem (4.2) by Dual AGP, then estimate the state-space matrices A , B , C , D and x_0 in (4.1) as done in [31, Section 5.2] and compute our identification errors (err_{id}) and validation errors (err_{v}) as in [31, Eq. (5.7)]:

$$err_{\text{id}} = \sqrt{\frac{\sum_{t=0}^{199} \|\tilde{y}_t - \hat{y}_t\|^2}{\sum_{t=0}^{199} \|\tilde{y}_t - \bar{y}_I\|^2}}, \quad err_{\text{v}} = \sqrt{\frac{\sum_{t=0}^{599} \|\tilde{y}_t - \hat{y}_t\|^2}{\sum_{t=0}^{599} \|\tilde{y}_t - \bar{y}_V\|^2}},$$

where

$$\bar{y}_I = \frac{1}{200} \sum_{t=0}^{199} \tilde{y}_t, \quad \bar{y}_V = \frac{1}{600} \sum_{t=0}^{599} \tilde{y}_t,$$

TABLE 4.2
Computational results for problems from Table 4.1.

P	μ	D. GP	D. AGP	D. PPA	P. ADM2	D. ADM2
		iter/cpu/obj	iter/cpu/obj	iter/cpu/obj	iter/cpu/obj	iter/cpu/obj
1	0.01	6/0.7/3.76e+0	9/1.1/3.76e+0	1(9)/1.4(1.1)/3.76e+0	10/1.3/3.76e+0	10/1.1/3.76e+0
	0.10	10/1.1/2.87e+1	10/1.1/2.87e+1	1(10)/1.5(1.2)/2.87e+1	30/3.6/2.87e+1	20/2.2/2.87e+1
	1.00	150/15.0/1.70e+2	70/7.4/1.70e+2	10(20)/5.6(2.2)/1.70e+2	50/5.9/1.70e+2	50/5.2/1.70e+2
	10.00	1310/130.7/1.01e+3	240/24.9/1.01e+3	33(50)/21.2(5.4)/1.01e+3	200/23.7/1.01e+3	220/22.8/1.01e+3
2	0.01	7/1.9/5.19e+0	10/2.7/5.19e+0	1(10)/4.2(2.7)/5.19e+0	10/3.0/5.19e+0	10/2.8/5.19e+0
	0.10	20/5.1/3.29e+1	20/5.3/3.29e+1	1(10)/4.1(2.7)/3.29e+1	50/14.5/3.29e+1	40/10.5/3.29e+1
	1.00	210/50.4/1.05e+2	90/22.9/1.05e+2	18(20)/19.5(5.2)/1.05e+2	70/20.0/1.05e+2	70/17.7/1.05e+2
	10.00	900/221.5/4.26e+2	420/105.6/4.26e+2	17(70)/36.6(17.7)/4.26e+2	250/71.5/4.26e+2	280/71.2/4.26e+2
3	0.01	9/2.0/4.57e+0	10/2.4/4.57e+0	1(10)/3.0(2.4)/4.57e+0	10/2.6/4.57e+0	10/2.4/4.57e+0
	0.10	380/81.2/2.06e+1	70/15.6/2.06e+1	22(20)/19.5(4.6)/2.06e+1	390/98.1/2.06e+1	550/122.3/2.06e+1
	1.00	730/155.5/8.60e+1	180/39.5/8.60e+1	17(40)/27.6(8.9)/8.60e+1	140/34.7/8.60e+1	130/28.6/8.60e+1
	10.00	max/423.1/3.70e+2	420/91.9/3.70e+2	16(140)/48.0(31.0)/3.70e+2	650/161.4/3.70e+2	580/126.7/3.70e+2
4	0.01	9/4.8/8.69e+0	10/5.6/8.69e+0	1(10)/7.1(5.6)/8.69e+0	10/6.1/8.69e+0	10/5.6/8.69e+0
	0.10	460/240.5/3.65e+1	80/43.7/3.65e+1	22(20)/50.0(11.0)/3.65e+1	470/289.4/3.65e+1	660/359.3/3.65e+1
	1.00	400/207.9/1.33e+2	170/90.6/1.33e+2	13(30)/60.8(16.3)/1.33e+2	130/79.3/1.33e+2	110/59.3/1.33e+2
	10.00	max/1010.8/6.19e+2	470/250.7/6.19e+2	20(100)/162.7(53.7)/6.19e+2	550/333.4/6.19e+2	490/259.6/6.19e+2
5	0.01	6/1.7/4.67e+0	8/2.3/4.67e+0	1(8)/3.8(2.3)/4.67e+0	10/3.1/4.67e+0	10/2.9/4.67e+0
	0.10	10/2.8/3.73e+1	10/2.9/3.73e+1	1(10)/3.6(2.9)/3.73e+1	10/3.1/3.73e+1	10/2.9/3.73e+1
	1.00	150/36.3/1.75e+2	60/15.4/1.75e+2	10(20)/13.3(5.4)/1.75e+2	60/17.2/1.75e+2	50/13.1/1.75e+2
	10.00	1070/258.0/9.39e+2	270/69.0/9.39e+2	31(50)/80.0(13.0)/9.39e+2	180/52.2/9.39e+2	190/49.5/9.39e+2
6	0.01	6/3.4/7.57e+0	9/5.4/7.57e+0	1(9)/8.4(5.2)/7.57e+0	10/6.3/7.57e+0	10/5.9/7.57e+0
	0.10	10/5.7/5.56e+1	10/6.1/5.56e+1	1(10)/7.5(6.0)/5.56e+1	20/12.5/5.56e+1	10/5.9/5.56e+1
	1.00	200/105.0/1.63e+2	80/45.0/1.63e+2	10(20)/33.4(11.7)/1.63e+2	80/49.9/1.63e+2	60/33.4/1.63e+2
	10.00	470/244.9/7.75e+2	250/139.7/7.75e+2	13(50)/59.5(27.8)/7.75e+2	200/124.3/7.75e+2	230/126.9/7.75e+2
7	0.01	7/3.7/6.91e+0	10/5.6/6.91e+0	1(10)/8.2(5.6)/6.91e+0	10/5.9/6.91e+0	10/5.5/6.91e+0
	0.10	170/82.7/3.66e+1	50/25.8/3.66e+1	10(10)/38.1(5.5)/3.66e+1	250/144.4/3.66e+1	350/178.2/3.66e+1
	1.00	360/173.0/1.32e+2	130/66.3/1.32e+2	12(30)/47.5(15.7)/1.32e+2	100/57.6/1.32e+2	90/46.4/1.32e+2
	10.00	980/476.0/6.09e+2	410/209.8/6.09e+2	17(90)/96.9(46.5)/6.09e+2	520/300.0/6.09e+2	460/233.4/6.09e+2
8	0.01	7/9.2/1.29e+1	10/13.9/1.29e+1	1(10)/17.3(13.9)/1.29e+1	10/14.5/1.29e+1	10/13.9/1.29e+1
	0.10	230/289.1/6.34e+1	60/77.6/6.34e+1	13(10)/127.0(13.7)/6.34e+1	340/500.1/6.34e+1	470/623.4/6.34e+1
	1.00	320/402.1/1.77e+2	140/180.8/1.77e+2	14(30)/139.0(39.5)/1.77e+2	100/146.3/1.77e+2	90/116.9/1.77e+2
	10.00	850/1065.9/8.53e+2	430/553.9/8.53e+2	15(80)/223.8(104.3)/8.53e+2	460/685.2/8.53e+2	410/535.0/8.53e+2

TABLE 4.3
Description of test problems, taken from DaISy [11].

P	Description	$N + 1$	m	q
1	[96-007] CD Player arm	513	2	388
2	[98-002] Continuous Stirring Tank Reactor	1876	2	1793
3	[96-006] Hair Dryer	251	1	168
4	[97-002] Steam Heat Exchanger	1001	1	918
5	[96-011] Heat Flow density	421	1	296
6	[97-003] Industrial Winding Process	626	2	375
7	[96-002] Glass Furnace	312	6	145
8	[96-016] Industrial Dryer	217	3	50

and \hat{y}_t is the output generated from (4.1) using the estimated A, B, C, D and x_0 .

We identify the corresponding parameter μ^* at which a “branching” occurs; see Figure 4.1, which plots identification and validation errors against the nuclear norm of $\mathcal{H}(y)$. The estimated system order is then set to be the rank of $\mathcal{H}(y^*)$, where y^* is the approximate solution of (4.2) with $\mu = \mu^*$. Table 4.5 shows some identification errors

TABLE 4.4
Computational results for problems from Table 4.3.


P	μ	D. GP	D. AGP	D. PPA	P. ADM2	D. ADM2
		iter/cpu/nr/err	iter/cpu/nr/err	iter/cpu/nr/err	iter/cpu/nr/err	iter/cpu/nr/err
1	0.01	60/1.1/35/3.2e-1	30/0.6/35/3.2e-1	9(10)/0.6(0.2)/35/3.1e-1	160/3.3/35/3.2e-1	220/4.1/35/3.2e-1
	0.10	1470/26.2/23/1.6e+0	180/3.3/23/1.6e+0	73(30)/3.9(0.6)/23/1.6e+0	470/9.8/23/1.6e+0	680/12.6/23/1.6e+0
	1.00	max/35.2/10/2.8e+0	1200/21.7/5/2.8e+0	45(180)/11.3(3.3)/5/2.8e+0	420/8.7/5/2.8e+0	530/9.7/5/2.8e+0
	10.00	max/35.1/26/3.0e+0	max/35.9/14/3.2e+0	55(max)/47.6(35.8)/3/3.2e+0	max/41.1/3/3.2e+0	max/36.5/4/3.2e+0
2	0.01	10/0.8/6/2.7e-1	10/0.8/6/2.7e-1	1(10)/1.0(0.8)/6/2.7e-1	max/171.3/6/2.5e-1	max/149.1/6/2.9e-1
	0.10	30/2.2/6/2.0e+0	10/0.8/6/2.0e+0	1(10)/1.1(0.8)/6/2.0e+0	1060/91.1/6/2.0e+0	1250/93.3/6/2.0e+0
	1.00	50/3.6/3/1.4e+1	20/1.6/3/1.4e+1	1(10)/2.9(0.8)/3/1.4e+1	770/68.5/3/1.4e+1	970/71.9/3/1.4e+1
	10.00	130/9.7/1/6.5e+1	50/3.9/1/6.5e+1	6(10)/15.9(0.8)/1/6.5e+1	540/46.6/1/6.5e+1	710/53.0/1/6.5e+1
3	0.01	20/0.1/19/3.0e-1	20/0.1/19/3.0e-1	1(10)/0.1(0.1)/19/3.0e-1	180/1.1/19/3.0e-1	230/1.2/19/3.0e-1
	0.10	240/1.2/6/9.9e-1	80/0.4/6/9.9e-1	32(10)/0.5(0.1)/6/9.8e-1	410/2.4/6/9.9e-1	560/2.9/6/9.9e-1
	1.00	550/2.7/3/2.6e+0	160/0.8/3/2.6e+0	17(40)/0.6(0.2)/3/2.6e+0	110/0.6/3/2.6e+0	130/0.7/3/2.6e+0
	10.00	390/1.9/1/1.2e+1	390/2.0/1/1.2e+1	25(130)/2.0(0.7)/1/1.2e+1	500/2.9/1/1.2e+1	530/2.7/1/1.2e+1
4	0.01	7/0.1/2/4.0e-1	10/0.2/2/4.0e-1	1(10)/0.3(0.2)/2/4.0e-1	80/1.6/2/4.0e-1	110/2.0/2/4.0e-1
	0.10	30/0.5/2/3.3e+0	10/0.2/2/3.4e+0	1(10)/0.3(0.2)/2/3.4e+0	280/5.7/2/3.3e+0	330/6.0/2/3.3e+0
	1.00	100/1.7/2/1.1e+1	20/0.4/2/1.1e+1	2(10)/1.0(0.2)/2/1.1e+1	190/3.9/2/1.1e+1	240/4.3/2/1.1e+1
	10.00	160/2.7/1/6.0e+1	80/1.4/1/6.0e+1	7(10)/4.8(0.2)/1/6.0e+1	110/2.2/1/6.0e+1	150/2.7/1/6.0e+1
5	0.01	9/0.1/42/3.8e-1	10/0.1/42/3.8e-1	1(10)/0.1(0.1)/42/3.8e-1	10/0.1/42/3.8e-1	10/0.1/42/3.8e-1
	0.10	410/3.0/10/2.0e+0	80/0.6/10/2.0e+0	22(20)/0.6(0.2)/11/2.0e+0	260/2.2/10/2.0e+0	370/2.8/10/2.0e+0
	1.00	850/6.1/3/2.8e+0	360/2.6/3/2.8e+0	39(60)/3.6(0.4)/3/2.8e+0	140/1.2/3/2.8e+0	130/1.0/3/2.8e+0
	10.00	700/5.0/1/7.1e+0	700/5.1/1/7.1e+0	27(210)/2.4(1.5)/1/7.1e+0	860/7.2/1/7.1e+0	810/6.0/1/7.1e+0
6	0.01	10/0.2/79/4.9e-1	10/0.2/79/4.9e-1	1(10)/0.3(0.2)/79/4.9e-1	20/0.4/79/5.0e-1	20/0.4/79/5.0e-1
	0.10	190/3.6/49/3.3e+0	60/1.2/49/3.3e+0	17(10)/1.5(0.2)/49/3.3e+0	120/2.6/49/3.3e+0	160/3.2/49/3.3e+0
	1.00	max/37.5/5/5.8e+0	320/6.1/5/5.8e+0	40(70)/4.1(1.3)/5/5.8e+0	230/5.0/5/5.8e+0	210/4.1/5/5.8e+0
	10.00	990/18.6/2/1.4e+1	830/15.7/2/1.4e+1	26(210)/7.8(4.0)/2/1.4e+1	1540/33.9/2/1.4e+1	1290/25.2/2/1.4e+1
7	0.01	10/0.2/100/5.9e-1	10/0.2/100/5.9e-1	1(10)/0.3(0.2)/100/5.9e-1	20/0.5/100/5.9e-1	20/0.5/100/5.9e-1
	0.10	340/7.9/37/3.4e+0	70/1.7/37/3.4e+0	17(20)/1.8(0.5)/37/3.3e+0	230/6.2/37/3.4e+0	320/7.7/37/3.4e+0
	1.00	1640/36.6/8/8.5e+0	250/5.7/8/8.5e+0	32(40)/5.7(0.9)/8/8.5e+0	190/5.0/8/8.6e+0	180/4.2/8/8.5e+0
	10.00	max/44.5/3/2.8e+1	530/12.1/2/2.8e+1	27(130)/6.2(3.0)/2/2.8e+1	890/23.5/2/2.8e+1	810/18.7/2/2.8e+1
8	0.01	10/0.1/38/2.7e-1	10/0.1/38/2.7e-1	1(10)/0.1(0.1)/38/2.7e-1	110/0.7/38/2.7e-1	140/0.8/38/2.7e-1
	0.10	250/1.3/23/1.8e+0	70/0.4/23/1.8e+0	15(10)/0.3(0.1)/23/1.8e+0	490/3.0/23/1.8e+0	670/3.7/23/1.8e+0
	1.00	max/10.5/12/6.9e+0	380/2.0/11/6.9e+0	66(50)/1.5(0.3)/12/6.9e+0	770/4.6/12/6.9e+0	1150/6.3/12/6.9e+0
	10.00	max/10.2/2/1.6e+1	1460/7.6/2/1.6e+1	46(260)/13.7(1.4)/2/1.6e+1	990/5.9/2/1.6e+1	1310/7.0/2/1.6e+1

(err_{id}) and validation errors (err_v), as well as the corresponding rank and nuclear norm of $\mathcal{H}(y^*)$. From this and Figure 4.1, we conclude that the estimated system order in this case is 3, which agrees with the result obtained in [31, Table 5.2].

TABLE 4.5
Identification errors and validation errors for CD player arm data.

μ	rank($\mathcal{H}(y^*)$)	err_{id}	err_v	$\ \mathcal{H}(y^*)\ _*$
1.233	6	0.141	0.230	0.185
1.385	5	0.157	0.230	0.127
1.556	3	0.178	0.187	0.104
1.748	3	0.178	0.187	0.097
1.963	3	0.177	0.186	0.090

Another commonly used heuristic for identifying system order is the subspace algorithm, which is implemented in the Matlab function `n4sid`. A detailed comparison between the nuclear norm heuristics described above and the Matlab function `n4sid` has been presented in [31, Table 5.2]. Since our first-order method usually obtains a

FIG. 4.1. Plot of err_{id} and err_v against $\|\mathcal{H}(y)\|_*$.

solution to (4.2) with lower accuracy than the interior point method used in [31], and the Matlab function `n4sid` has been updated since the publication of [31], it is worth repeating the comparison for the problems in Table 4.3. In the test below, we use the same number of data points for identification and validation as in [31, Table 5.1]. We report the system orders (`ord`), identification errors (err_{id}) and validation errors (err_v) obtained via the nuclear norm heuristics (`nn`), the `n4sid` with default order (`n4sidd`), and the `n4sid` using the order obtained from the nuclear norm heuristics (`n4sidn`) in Table 4.6. We observe that the orders and errors obtained from the nuclear norm heuristics are usually slightly different from those in [31, Table 5.2]. This can be attributed to the solution accuracy in solving (4.2), and the threshold for determining the rank of $\mathcal{H}(y^*)$. On the other hand, we still observe that the nuclear norm heuristics usually yields lower identification and validation errors than the two versions of `n4sid`.⁶

Comparison with interior point solver from [31]. Finally, we apply the interior point method in [31] to solve the first problem in Table 4.3 (i.e., with $N = 512$). We use the code (written in Python and calling `cvxopt 1.1.3`⁷) available on the authors'

⁶We observe that the errors for problems 2 and 4 obtained from `n4sid` are exceptionally large. In view of this, we also tried to vary the input order for `n4sid` from 1 to 12 for these two problems and observed the corresponding identification and validation errors. For problem 2, the smallest $(err_{id}, err_v) = (0.19, 0.39)$, which occurs when order is 10; while for problem 4, the smallest $(err_{id}, err_v) = (0.14, 0.22)$, which occurs when order is 8.

⁷The package `cvxopt` is installed from source using Python 2.6.5.

TABLE 4.6
Comparing nuclear norm heuristics and $n4sid$ on problems from Table 4.3.

	nn	$n4sid_d$	$n4sid_n$
\mathbf{P}	ord/err _{id} /err _v	ord/err _{id} /err _v	ord/err _{id} /err _v
1	3/0.18/0.19	2/0.21/0.23	3/0.16/0.18
2	3/0.19/0.20	3/0.21/0.58	3/2.20/5.77
3	4/0.068/0.11	2/0.13/0.18	4/0.080/0.13
4	4/0.12/0.24	2/3.36/5.38	4/11.68/11.98
5	6/0.14/0.13	1/0.18/0.17	6/0.14/0.13
6	3/0.18/0.18	2/0.52/0.58	3/0.18/0.17
7	4/0.19/0.39	2/0.41/0.53	4/0.60/0.64
8	6/0.28/0.28	1/0.47/0.29	6/0.31/0.27

webpage, compiled with Python 2.6.5. While the interior point method provides a solution of higher accuracy (the relative duality gap is usually around $1e - 6$, rather than the $1e - 4$ as required in our first order methods), it is much slower. We observe that the number of iterations and CPU time for the interior point method are not sensitive to the change of μ , and are usually around 9 iterations and 1000 seconds respectively. In particular, when $\mu \leq 1$, this method is significantly slower than Dual AGP, which takes between 1 and 40 seconds for this range of μ . Thus, the first order methods appear more suitable for larger scale identification problems.

5. Stochastic system realization. Another fundamental problem in linear system theory is finding a minimal stochastic ARMA (autoregressive moving average) model for a vector random process, given noisy and/or partial estimates of process covariances. Covariance estimates are often obtained from sample averages of a sequence of noisy observations of the process, hence including both measurement noise and error due to the finite sample size. Here we focus on a form of this problem, described in [32, section II.B]; see also [34]. Consider a state-space model of an ARMA process $y_t \in \mathbb{R}^n$,

$$\begin{aligned}x_{t+1} &= Ax_t + Be_t, \\y_t &= Cx_t + e_t,\end{aligned}$$

where $x_t \in \mathbb{R}^r$, and e_t is white noise with covariance matrix Q . The process covariances $h_i = \mathbf{E}(y_t y_{t+i}^T)$ satisfy

$$h_0 = CPC^T + Q, \quad h_t = CA^{t-1}D, \quad t \geq 1,$$

where $D = APC^T + BQ$, and $P = \mathbf{E}(x_t x_t^T)$ satisfies the Lyapunov equation $P = APA^T + BQB^T$. In a stochastic realization problem, we are given noisy estimates of h_i , denoted by \tilde{h}_i , $i = 1, \dots, T - 1$, and the goal is to find the minimal model order r as well as the model parameters A, B, C, Q . It is known that the minimal order is equal to the rank of the block-Hankel matrix $H_{n,n,j,k}$ consisting of the exact process covariances [34] (as in the system ID problem, we need the Hankel matrix to be large enough, that is, j and k should be larger than the rank).

A general form of this problem, allowing for both noise and missing covariance information, can be stated as follows,

$$\min_y \frac{1}{2} \|w \circ (y - \tilde{h})\|_F^2 + \mu \|H_{m,n,j,k}(y)\|_*, \quad (5.1)$$

where \circ denotes the Hadamard (or entry-wise) product, and $w = (w_0 \ \cdots \ w_{j+k-2})$ is an $m \times n(j+k-1)$ matrix, with each w_i being an $m \times n$ zero matrix or a matrix of all ones (denoting the case where some covariance blocks are unknown or missing), or a matrix with 0, 1 entries (denoting the case where some covariance entries are missing). The problem corresponds to (2.1) with $\Upsilon = I$, $\mathcal{A}(y) = \text{vec}(w \circ y)$ and $b = \text{vec}(w \circ \tilde{h})$. Thus,

$$\sigma_{\max}(\mathcal{A}) = \sigma_{\max}(\Upsilon) = 1, \quad \mathcal{H}^*(\Lambda) = H_{m,n,j,k}^*(\Lambda).$$

Notice that the solution set of (5.1) is nonempty since the function $y \mapsto \|\mathcal{H}(y)\|_*$ is coercive.

5.1. Computational results. In this section, we compare different algorithms for solving (5.1). Specifically, we consider Primal ADM, Dual ADM and Dual PPA. Moreover, since the action of $(\mathcal{H}^*\mathcal{H} + \mathcal{A}^*\mathcal{A})^{-1}$ and $(\mathcal{I} + \mathcal{A}^*\mathcal{A})^{-1}$ on vectors can be easily computed, we also consider a variant of Primal ADM (referred to as Primal ADM3) obtained from the following augmented Lagrangian function

$$\frac{1}{2}\|w \circ (y - \tilde{h})\|_F^2 + \mu\|Y\|_* - \langle \Lambda, Y + \mathcal{H}(y) \rangle + \frac{\beta}{2}\|Y + \mathcal{H}(y)\|_F^2;$$

as well as a variant of Dual ADM with $Q_1 = 0$ (referred to as Dual ADM3). The convergence of these two variants follows from [2, Proposition 4.2] and Theorem B.1, respectively. Furthermore, notice that the quadratic term in (5.1) is *not* strictly convex in y , the dual problem will then have additional linear constraints that makes gradient projection expensive computationally. Thus, we do not consider gradient projection algorithms for solving (5.1).

We initialize all algorithms except Dual PPA at the origin, where the latter algorithm is initialized at an approximate solution obtained from Primal ADM3⁸, and terminate the algorithms by checking

$$\max \left\{ \frac{f(y^k) + d(-w \circ \mathcal{H}^*(\mathcal{P}(\Lambda^k)))}{\max\{1, |d(-w \circ \mathcal{H}^*(\mathcal{P}(\Lambda^k)))|\}}, \frac{\|\mathcal{H}^*(\mathcal{P}(\Lambda^k)) - w \circ \mathcal{H}^*(\mathcal{P}(\Lambda^k))\|_F}{\max\{1, \|\mathcal{H}^*(\mathcal{P}(\Lambda^k))\|_F\}} \right\} < 1e - 4,$$

with $\{(y^k, \Lambda^k)\}$ defined as in Sections 3.1 and 3.2, and \mathcal{P} is the projection onto the spectral norm ball with radius μ . For the ADMs, we set $\beta = \frac{\mu r}{2\sigma_{\max}(\tilde{h})}$ and $\sigma = \frac{0.95}{r+1}$ for Primal ADM and use the same β for Primal ADM3. We set $\beta = \frac{\sigma_{\max}(\tilde{h})}{8\mu r}$, $\sigma_1 = 0.95$ and $\sigma_2 = \frac{0.95}{r}$ for Dual ADM and use the same β and σ_2 for Dual ADM3. For Dual PPA, we use the same parameters as in the previous section.

Following [32, Section II(B)], we generate matrices $A \in \mathbb{R}^{r \times r}$, $B \in \mathbb{R}^{r \times n}$ and $C \in \mathbb{R}^{n \times r}$ with i.i.d. standard Gaussian entries (i.e., with mean 0 and variance 1). These matrices are then normalized to have spectral norm 1. We also randomly generate an initial state $x_0 \sim N(0, I)$ and noise vectors $e_t \sim N(0, I)$ for $t = 0, \dots, T-1$, again with i.i.d. standard Gaussian entries. We then generate an output \bar{y}_t , $t = 0, \dots, T-1$, according to the state space model:

$$\begin{aligned} x_{t+1} &= Ax_t + Be_t, \\ \bar{y}_t &= Cx_t + e_t. \end{aligned}$$

⁸We terminate Primal ADM3 by checking relative duality gap and relative dual feasibility with $tol = 5e - 3$, checked every 10 iterations. We also early terminate the algorithm if the change in Fröbenius norm of successive iterates is small ($< 1e - 8$ for each variable) or the maximum number of iteration hits 2000.

To model measurement noise, we further add noise to \bar{y} and get $\tilde{y} = \bar{y} + \sigma\epsilon$, where ϵ has i.i.d. Gaussian entries with mean 0 and variance 1. We then set, for each $i = 0, \dots, k-1$,

$$\tilde{h}_i = \frac{1}{T} \sum_{t=0}^{T-1-i} \tilde{y}_{t+i} \tilde{y}_t^T,$$

and \tilde{h}_i is zero for $i \geq k$. Finally, set $w = (w_0 \ \dots \ w_{j+k-2})$ such that $w_0 = \dots = w_{k-1}$ equals the matrix of all ones, and zero otherwise. In the tests below, we consider $T = 1000$, $n = 20, 30$ and $k = 100, 200$. We use $r = 10$ and hence set $j = 21$. We further pick $\sigma = 5e - 2$. The statistics of the test problems used are reported in Table 5.1; recall that $q = nk$ in this case. We run our algorithms for a range of values of μ , namely $\mu = 1e - 2, 1e - 1, 1, 10$, in our simulations below to study the performance of the algorithms for different values of μ . The computational results are reported in Table 5.2. We see that Primal ADM3 works best in all instances ⁹.

TABLE 5.1
Parameters for the randomly generated test problems.

P	k	n
1	100	20
2	100	30
4	200	20
5	200	30

TABLE 5.2
Computational results for problems from Table 5.1.

P	μ	P. ADM	P. ADM3	D. ADM3	D. PPA
		iter/cpu/obj	iter/cpu/obj	iter/cpu/obj	iter/cpu/obj
1	0.01	170/98.2/5.38e+0	30/16.8/5.38e+0	110/54.5/5.38e+0	135(10)/455.9(5.4)/5.38e+0
	0.10	120/67.3/2.73e+1	110/61.6/2.73e+1	150/74.3/2.73e+1	53(10)/74.3(5.4)/2.73e+1
	1.00	220/123.8/4.18e+1	30/16.5/4.18e+1	160/78.7/4.18e+1	32(10)/56.3(5.4)/4.18e+1
	10.00	210/119.1/4.38e+1	7/4.0/4.38e+1	220/108.6/4.38e+1	1(7)/4.9(3.7)/4.38e+1
2	0.01	220/350.4/9.12e+0	30/47.1/9.12e+0	160/223.9/9.12e+0	128(10)/1252.5(14.9)/9.12e+0
	0.10	80/125.8/4.99e+1	70/110.8/4.99e+1	90/124.0/4.99e+1	136(10)/570.5(14.9)/4.99e+1
	1.00	200/315.5/7.04e+1	20/30.8/7.04e+1	140/193.9/7.04e+1	9(10)/63.2(15.0)/7.04e+1
	10.00	210/330.1/7.16e+1	7/10.6/7.16e+1	190/263.7/7.16e+1	1(7)/776.9(10.2)/7.16e+1
3	0.01	190/227.3/7.16e+0	30/35.1/7.16e+0	140/147.2/7.16e+0	130(10)/376.6(11.4)/7.16e+0
	0.10	70/83.0/4.02e+1	40/48.1/4.02e+1	50/54.5/4.02e+1	47(10)/304.1(11.7)/4.02e+1
	1.00	180/214.5/5.29e+1	20/23.2/5.29e+1	130/136.6/5.29e+1	18(10)/76.3(11.3)/5.29e+1
	10.00	180/213.8/5.33e+1	7/8.4/5.33e+1	190/206.2/5.33e+1	1(7)/10.5(8.0)/5.33e+1
4	0.01	220/766.2/1.35e+1	30/101.6/1.35e+1	150/466.9/1.35e+1	136(10)/2765.8(32.1)/1.35e+1
	0.10	100/341.9/8.74e+1	40/137.5/8.74e+1	70/212.7/8.74e+1	28(10)/962.5(33.1)/8.74e+1
	1.00	140/482.8/1.23e+2	20/67.2/1.23e+2	110/336.9/1.23e+2	18(10)/159.4(33.0)/1.23e+2
	10.00	180/616.4/1.25e+2	8/26.6/1.25e+2	190/574.5/1.25e+2	1(8)/31.9(24.9)/1.25e+2

We also adapted Cadzow's method [4], an approach based on alternating projections, that has been used in engineering applications for similar problems. We start by specifying an estimate \tilde{r} of the order and $\tilde{\epsilon} > 0$ of the noise level. In each iteration,

⁹The two versions of Dual ADM perform similarly. Hence, we only report test results for the variant with $Q_1 = 0$

we first project onto the set of matrices with a Hankel structure to obtain H^k , then project onto the ball centered at \tilde{h} with radius $\tilde{\epsilon}$ to obtain Y^k , and finally project onto the (nonconvex) set of matrices with rank less than \tilde{r} to obtain R^k . The algorithm is terminated when

$$\frac{\max\{\|H^k - Y^k\|_F, \|R^k - Y^k\|_F\}}{\max\{1, \|H^k\|_F\}} < 1e - 3.$$

In the examples we tested, the performance and convergence of this algorithm is very sensitive to the values of the two parameters, the noise level and the order estimate. Indeed, it is known to be suboptimal as discussed in [10, Section V.A]. On the other hand, the convex optimization approach presented here only depends on one parameter (μ), and the performance does not change significantly as μ varies.

Finally, we mention that another common approach to solving the system realization problem is via subspace methods [34]. For a comparison between the nuclear norm approach and the subspace method applied to this problem, we refer the readers to [32, Section 2B], where problem (5.1) was solved via an interior-point solver.

6. Concluding remarks. As a systematic approach to capturing the tradeoff between a model's order and complexity with fitting (and validation) errors, a tradeoff that commonly arises in diverse modeling applications, we studied the optimization problem of minimizing the nuclear norm of matrices with linear structure, including Hankel, Toeplitz, and moment structures. We then focused on first-order methods for solving the resulting optimization problem. In our computational experiments, the gradient projection method (accelerated by Nesterov's extrapolation techniques) and the dual proximal point algorithm usually outperform other first-order methods in terms of CPU time for large and small regularization parameters, respectively, for the system identification problem, while for the system realization problem, the alternating direction method, as applied to a certain primal reformulation, usually outperforms other first-order methods in terms of CPU time. In our tests, we also observe that these methods outperform the interior point implementation proposed in [31] for system identification problems.

Finally, we remark that most of the algorithms proposed in this work can be easily adapted to solve (1.1) with a general linear structure $X = \mathcal{L}(y)$, as long as $\mathcal{L}(y)$ can be computed efficiently and the norm of \mathcal{L} can be estimated. We chose to focus our discussion on the Hankel structure to relate more closely to our motivating applications discussed in the introduction, which mainly concern Hankel structure.

An interesting direction for future work on this problem is whether there are conditions under which the nuclear norm heuristic can be theoretically guaranteed to find the minimum rank solution. In order to make analogies with the existing low-rank matrix recovery framework, we can consider the following problem: how many generic, random linear measurements of a rank- r , $n \times n$ Hankel matrix suffice for correct recovery of the Hankel matrix? If we ignore the Hankel structure, existing results on recovery from random Gaussian measurements require $\mathcal{O}(nr)$ measurements [6]; however, it is expected that the true number would be much lower due to the Hankel structure.

Another future research direction involves applying our algorithms to a broader range of applications identified in the introduction. Some of these problems have further structure that the algorithms can exploit.

Appendix A. An alternative formulation. In this appendix, we describe an alternative formulation for modeling the structured matrix rank minimization.

Instead of minimizing a least square fitting error, we constrain the difference $\mathcal{A}(y) - b$ in a set modeling uncertainty. The problem is then formulated as follows:

$$\begin{aligned} \min \quad & \|\mathcal{H}(y)\|_* \\ \text{s.t.} \quad & \mathcal{A}(y) - b \in \Omega, \end{aligned} \quad (\text{A.1})$$

where Ω is the closed convex set modeling uncertainty. For instance, problem (1.2) can be modeled as (A.1) with $\Omega = [l_1, b_1] \times [l_2, b_2] \times \cdots \times [l_n, b_n]$.

As in Section 2, we can rewrite (A.1) as

$$\begin{aligned} v_1 := \min \quad & \|Y\|_* \\ \text{s.t.} \quad & Y + \mathcal{H}(y) = 0, \\ & z + \mathcal{A}(y) = b, \\ & z \in \Omega. \end{aligned} \quad (\text{A.2})$$

It is then not hard to show that the dual to (A.2) is equivalent to solving

$$\begin{aligned} -v_1 = \min \quad & s_\Omega(\gamma) - b^T \gamma \\ \text{s.t.} \quad & \mathcal{H}^*(\Lambda) + \mathcal{A}^*(\gamma) = 0, \\ & \Lambda^T \Lambda \preceq I, \end{aligned} \quad (\text{A.3})$$

where $s_\Omega(\gamma) := \max_{y \in \Omega} y^T \gamma$ is the support function of the set Ω .

Unlike (2.3), the objective function of (A.3) is in general not differentiable, hence gradient projection methods are not easily applicable. However, the ADMs and the dual PPA can be suitably applied to solve (A.2) and (A.3). The efficiency in solving the subproblems depends on the specific form of Ω .

Appendix B. Convergence of the proximal alternating direction method.

In this appendix, we provide a convergence proof of a proximal alternating direction method, which covers all versions of proximal ADMs used in this paper. Consider the convex optimization problem with the following separable structure

$$\begin{aligned} \min \quad & f(x) + g(y) \\ \text{s.t.} \quad & Ax + By = c, \end{aligned} \quad (\text{B.1})$$

where $f : \mathcal{X} \rightarrow (-\infty, +\infty]$ and $g : \mathcal{Y} \rightarrow (-\infty, +\infty]$ are closed proper convex functions, $A : \mathcal{X} \rightarrow \mathcal{Z}$ and $B : \mathcal{Y} \rightarrow \mathcal{Z}$ are linear operators, \mathcal{X}, \mathcal{Y} and \mathcal{Z} are real finite dimensional Euclidean spaces with inner product $\langle \cdot, \cdot \rangle$ and its induced norm $\|\cdot\|$. The proximal alternating direction method (proximal ADM) for solving (B.1) takes the following form:

Proximal ADM

Step 0. Input $(x^0, y^0, z^0) \in \mathcal{X} \times \mathcal{Y} \times \mathcal{Z}$.

Step 1. Set

$$\begin{cases} x^{k+1} = \arg \min_{x \in \mathcal{X}} f(x) - \langle z^k, Ax \rangle + \frac{\lambda}{2} \|Ax + By^k - c\|^2 + \frac{1}{2} \|x - x^k\|_S^2, \\ y^{k+1} = \arg \min_{y \in \mathcal{Y}} g(y) - \langle z^k, By \rangle + \frac{\lambda}{2} \|Ax^{k+1} + By - c\|^2 + \frac{1}{2} \|y - y^k\|_T^2, \\ z^{k+1} = z^k - \lambda(Ax^{k+1} + By^{k+1} - c), \end{cases} \quad (\text{B.2})$$

where $\lambda > 0$, S and T are self-adjoint positive semidefinite operators on \mathcal{X} and \mathcal{Y} respectively.

Step 2. If a termination criterion is not met, go to Step 1.

When $S = 0$ and $T = 0$, the proximal ADM (B.2) reduces to the classical ADM introduced by Glowinski and Marroco [22] and Gabay and Mericier [21]. It was shown by Eckstein and Bertsekas [15] that the ADM, as a special case of Douglas-Rachford splitting [20], is actually an application of the proximal point algorithm on the dual problem by means of a specially-constructed splitting operator. Based on the same argument by further applying a change of variable to the operators, Eckstein [14] presented the first proximal ADM as in (B.2) with $S = \mu_1 I$ and $T = \mu_2 I$ for positive constants $\mu_1 > 0$ and $\mu_2 > 0$. Later, He et al. [25] further extended the idea of Eckstein [14] to monotone variational inequalities to allow λ , S , and T to be replaced by different parameters λ_k , S_k , and T_k in each iteration. The convergence results provided in [14] and [25] for the proximal ADM both need S and T to be positive definite, which limits the applications of the method. However, we notice that by slightly revising the proof provided by He et al. in [25], one may readily prove the following theorem. For clarity and completeness, we include a proof here.

For technical reasons, we consider the following constraint qualification:

CQ There exists $(x_0, y_0) \in \text{ri}(\text{dom } f \times \text{dom } g) \cap P$, where P is the constraint set in (B.1).

Under **CQ**, it follows from [47, Corollary 28.2.2] and [47, Corollary 28.3.1] that $(\bar{x}, \bar{y}) \in \mathcal{X} \times \mathcal{Y}$ is an optimal solution to problem (B.1) if and only if there exists a Lagrange multiplier $\bar{z} \in \mathcal{Z}$ such that

$$A^* \bar{z} \in \partial f(\bar{x}), \quad B^* \bar{z} \in \partial g(\bar{y}), \quad A\bar{x} + B\bar{y} - c = 0, \quad (\text{B.3})$$

where ∂f and ∂g are the subdifferential mappings of f and g respectively. Moreover, any $\bar{z} \in \mathcal{Z}$ satisfying (B.3) is an optimal solution to the dual problem of (B.1).

THEOREM B.1. *Assume that the solution set of (B.1) is nonempty and **CQ** holds. Let $\{(x^k, y^k, z^k)\}$ be generated from the proximal ADM. Then $\{(x^k, y^k)\}$ converges to an optimal solution to (B.1) and $\{z^k\}$ converges to an optimal solution to the dual problem of (B.1) if one of the following conditions holds:*

- (a) f and g are strongly convex;
- (b) f is strongly convex and $B^*B + T$ is positive definite;
- (c) g is strongly convex and $A^*A + S$ is positive definite;
- (d) S is positive definite and B is injective;
- (e) T is positive definite and A is injective;
- (f) S and T are positive definite.

Proof. We first show that the sequence generated by proximal ADM is bounded. Notice that the iteration scheme (B.2) of the proximal ADM can be rewritten as: for $k = 0, 1, 2, \dots$,

$$\begin{cases} 0 \in \partial f(x^{k+1}) - A^*[z^k - \lambda(Ax^{k+1} + By^k - c)] + S(x^{k+1} - x^k), \\ 0 \in \partial g(y^{k+1}) - B^*[z^k - \lambda(Ax^{k+1} + By^{k+1} - c)] + T(y^{k+1} - y^k), \\ z^{k+1} = z^k - \lambda(Ax^{k+1} + By^{k+1} - c). \end{cases} \quad (\text{B.4})$$

On the other hand, since the subdifferential mappings of the closed proper convex functions are maximal monotone [48, Theorem 12.17], there exist some constants $\sigma_1, \sigma_2 \geq 0$ such that

$$\begin{aligned} \langle u - \bar{u}, x^{k+1} - \bar{x} \rangle &\geq \sigma_1 \|x^{k+1} - \bar{x}\|^2, \quad \forall u \in \partial f(x^{k+1}), \bar{u} \in \partial f(\bar{x}), \\ \langle v - \bar{v}, y^{k+1} - \bar{y} \rangle &\geq \sigma_2 \|y^{k+1} - \bar{y}\|^2, \quad \forall v \in \partial g(y^{k+1}), \bar{v} \in \partial g(\bar{y}), \end{aligned} \quad (\text{B.5})$$

where $\sigma_1 > 0$ if f is strongly convex, $\sigma_2 > 0$ if g is strongly convex, and $(\bar{x}, \bar{y}, \bar{z})$ satisfies (B.3). Combining (B.5) with (B.3) and (B.4), we have

$$\begin{aligned} \langle A^*[z^k - \lambda(Ax^{k+1} + By^k - c)] - S(x^{k+1} - x^k) - A^*\bar{z}, x^{k+1} - \bar{x} \rangle &\geq \sigma_1 \|x^{k+1} - \bar{x}\|^2, \\ \langle B^*[z^k - \lambda(Ax^{k+1} + By^{k+1} - c)] - T(y^{k+1} - y^k) - B^*\bar{z}, y^{k+1} - \bar{y} \rangle &\geq \sigma_2 \|y^{k+1} - \bar{y}\|^2. \end{aligned}$$

Let $w_e^k := w^k - \bar{w}$ for notational simplicity, where w represents x, y, z respectively. Then the two inequalities above can be rewritten as follows:

$$\begin{aligned} \langle z_e^k - \lambda(Ax_e^{k+1} + By_e^k), Ax_e^{k+1} \rangle - \langle S(x^{k+1} - x^k), x_e^{k+1} \rangle &\geq \sigma_1 \|x_e^{k+1}\|^2, \\ \langle z_e^k - \lambda(Ax_e^{k+1} + By_e^{k+1}), By_e^{k+1} \rangle - \langle T(y^{k+1} - y^k), y_e^{k+1} \rangle &\geq \sigma_2 \|y_e^{k+1}\|^2. \end{aligned}$$

Adding up these two inequalities we obtain that

$$\begin{aligned} \langle z_e^k, Ax_e^{k+1} + By_e^{k+1} \rangle - \lambda \langle Ax_e^{k+1} + By_e^k, Ax_e^{k+1} \rangle - \lambda \langle Ax_e^{k+1} + By_e^{k+1}, By_e^{k+1} \rangle \\ - \langle S(x^{k+1} - x^k), x_e^{k+1} \rangle - \langle T(y^{k+1} - y^k), y_e^{k+1} \rangle \geq \sigma_1 \|x_e^{k+1}\|^2 + \sigma_2 \|y_e^{k+1}\|^2. \end{aligned}$$

Using the relation $z^{k+1} = z^k - \lambda(Ax_e^{k+1} + By_e^{k+1})$ and the elementary relations $\langle u, v \rangle = \frac{1}{2}(\|u\|^2 + \|v\|^2 - \|u - v\|^2) = \frac{1}{2}(\|u + v\|^2 - \|u\|^2 - \|v\|^2)$, we obtain further that

$$\begin{aligned} \|x_e^{k+1}\|_S^2 + \|y_e^{k+1}\|_T^2 + \lambda \|By_e^{k+1}\|^2 + \frac{1}{\lambda} \|z_e^{k+1}\|^2 &\leq \|x_e^k\|_S^2 + \|y_e^k\|_T^2 + \lambda \|By_e^k\|^2 + \frac{1}{\lambda} \|z_e^k\|^2 \\ &- (2\sigma_1 \|x_e^{k+1}\|^2 + 2\sigma_2 \|y_e^{k+1}\|^2 + \|x^{k+1} - x^k\|_S^2 + \|y^{k+1} - y^k\|_T^2 + \lambda \|Ax_e^{k+1} + By_e^k\|^2). \end{aligned} \quad (\text{B.6})$$

From (B.6), we see immediately that the sequence $\{\|x_e^k\|_S^2 + \|y_e^k\|_T^2 + \lambda \|By_e^k\|^2 + \frac{1}{\lambda} \|z_e^k\|^2\}$ is monotonically nonincreasing (and thus bounded), and

$$\lim_{k \rightarrow \infty} 2\sigma_1 \|x_e^{k+1}\|^2 + 2\sigma_2 \|y_e^{k+1}\|^2 + \|x^{k+1} - x^k\|_S^2 + \|y^{k+1} - y^k\|_T^2 + \lambda \|Ax_e^{k+1} + By_e^k\|^2 = 0. \quad (\text{B.7})$$

As an immediate consequence of these two facts, we see that the sequences

$$\begin{aligned} \{\|x_e^{k+1}\|_S^2\}, \{\|y_e^{k+1}\|_T^2\}, \{\|By_e^{k+1}\|^2\}, \{\|z_e^{k+1}\|^2\}, \\ \{\sigma_1 \|x_e^{k+1}\|^2\}, \{\sigma_2 \|y_e^{k+1}\|^2\}, \{\|Ax_e^{k+1} + By_e^k\|^2\} \end{aligned}$$

are bounded. Thus, the sequence $\{\|By_e^{k+1}\|^2 + \|y_e^{k+1}\|_T^2\}$ is bounded. Moreover, since

$$\|Ax_e^{k+1}\| \leq \|Ax_e^{k+1} + By_e^k\| + \|By_e^k\|, \quad (\text{B.8})$$

it follows that the sequence $\{\|Ax_e^{k+1}\|^2 + \|x_e^{k+1}\|_S^2\}$ is also bounded. From the boundedness of the above sequences, it is now routine to deduce that the sequence $\{(x^k, y^k, z^k)\}$ is bounded under the assumption of the theorem.

Since the sequence $\{(x^k, y^k, z^k)\}$ is bounded, there exists a subsequence $\{(x^{k_i}, y^{k_i}, z^{k_i})\}$ that converges to a limit point, say $(x^\infty, y^\infty, z^\infty)$. We next show that (x^∞, y^∞) is an optimal solution to problem (B.1) and z^∞ is a corresponding Lagrange multiplier.

To this end, we first note from (B.7) that

$$\begin{aligned} \lim_{k \rightarrow \infty} \sigma_1 \|x_e^{k+1}\| = 0, \quad \lim_{k \rightarrow \infty} \sigma_2 \|y_e^{k+1}\| = 0, \quad \lim_{k \rightarrow \infty} \|Ax_e^{k+1} + By_e^k\| = 0, \\ \lim_{k \rightarrow \infty} \|S(x^{k+1} - x^k)\| = 0, \quad \lim_{k \rightarrow \infty} \|T(y^{k+1} - y^k)\| = 0. \end{aligned} \quad (\text{B.9})$$

Thus, if either S or T is positive definite, we have

$$\lim_{k \rightarrow \infty} \|x^{k+1} - x^k\| = 0 \quad \text{or} \quad \lim_{k \rightarrow \infty} \|y^{k+1} - y^k\| = 0. \quad (\text{B.10})$$

Similarly, if either f or g is strongly convex, then (B.10) also holds. Thus, (B.10) holds under the assumption of the theorem. Since

$$\begin{aligned} \|Ax_e^{k+1} + By_e^{k+1}\| &\leq \|Ax_e^{k+1} + By_e^k\| + \|B(y^{k+1} - y^k)\|, \\ \|Ax_e^{k+1} + By_e^{k+1}\| &\leq \|Ax_e^{k+2} + By_e^{k+1}\| + \|A(x^{k+2} - x^{k+1})\|, \end{aligned}$$

we obtain from (B.9), (B.10) and the definition of z^{k+1} in (B.4) that

$$\lim_{k \rightarrow \infty} \|z^{k+1} - z^k\| = \lim_{k \rightarrow \infty} \lambda \|Ax_e^{k+1} + By_e^{k+1}\| = 0. \quad (\text{B.11})$$

Taking limits on both sides of (B.4) along the subsequence $\{(x^{k_i}, y^{k_i}, z^{k_i})\}$, using (B.9), (B.11) and invoking the closedness of the graphs of ∂f and ∂g [3, Page 80], we obtain that

$$A^*z^\infty \in \partial f(x^\infty), \quad B^*z^\infty \in \partial g(y^\infty), \quad Ax^\infty + By^\infty - c = 0,$$

i.e., $(x^\infty, y^\infty, z^\infty)$ satisfies (B.3). From the discussion prior to the theorem, we conclude that (x^∞, y^∞) is an optimal solution to problem (B.1) and z^∞ is a corresponding Lagrange multiplier.

Finally, we show that $(x^\infty, y^\infty, z^\infty)$ is actually the unique limit of $\{(x^k, y^k, z^k)\}$ and hence complete the proof. To this end, recall that $(x^\infty, y^\infty, z^\infty)$ satisfies (B.3). Hence, we could replace $(\bar{x}, \bar{y}, \bar{z})$ with $(x^\infty, y^\infty, z^\infty)$ in the above arguments, starting from (B.5). As a consequence, the subsequence $\{\|x_e^{k_i}\|_S^2 + \|y_e^{k_i}\|_T^2 + \lambda \|By_e^{k_i}\|^2 + \frac{1}{\lambda} \|z_e^{k_i}\|^2\}$ now converges to 0 as $i \rightarrow \infty$. Since this sequence is also non-increasing, we conclude that the whole sequence converges to zero, i.e.,

$$\lim_{k \rightarrow \infty} \|x_e^k\|_S^2 + \|y_e^k\|_T^2 + \lambda \|By_e^k\|^2 + \frac{1}{\lambda} \|z_e^k\|^2 = 0. \quad (\text{B.12})$$

From this, we see immediately that $\lim_{k \rightarrow \infty} z^k = z^\infty$. Moreover, if g is strongly convex or $B^*B + T$ is positive definite, then we see from (B.12) and (B.9) that

$$\lim_{k \rightarrow \infty} \sigma_2 \|y_e^{k+1}\| = 0 \quad \text{and} \quad \lim_{k \rightarrow \infty} \|y_e^k\|_T^2 + \|By_e^k\|^2 = 0,$$

and hence $\lim_{k \rightarrow \infty} y^k = y^\infty$. On the other hand, if f is strongly convex or $A^*A + S$ is positive definite, then we see from (B.12), (B.8) and (B.9) that

$$\lim_{k \rightarrow \infty} \sigma_1 \|x_e^{k+1}\| = 0 \quad \text{and} \quad \lim_{k \rightarrow \infty} \|x_e^k\|_S^2 + \|Ax_e^k\|^2 = 0,$$

and hence $\lim_{k \rightarrow \infty} x^k = x^\infty$. Therefore, we have shown that the whole sequence $\{(x^k, y^k, z^k)\}$ converges to $(x^\infty, y^\infty, z^\infty)$ under the assumption of the theorem. \square

Remark B.1. Without the two terms $\sigma_1 \|x_e^{k+1}\|^2$ and $\sigma_2 \|y_e^{k+1}\|^2$, the inequality (B.6) is actually a special case of the one used in Theorem 1 in [25].

Remark B.2. The conditions (a)–(f) for the convergence of $\{(x^k, y^k, z^k)\}$ in Theorem B.1 can be replaced by assuming that the following two slightly stronger conditions hold:

- (i) f is strongly convex, or A is injective, or S is positive definite;

(ii) g is strongly convex, or B is injective, or T is positive definite; except for the case that both A and B are injective.

Remark B.3. Theorem B.1 provides general conditions for the convergence of the proximal ADM. It includes some recent results in the literature as special cases. For example, it has been considered by Attouch and Soueycatt [1] for the case $S = \frac{1}{\lambda}I$ and $T = \frac{1}{\lambda}I$. Moreover, Zhang et al. considered the case when $B = I$ and S is chosen to be positive definite in [57], and established a slight weaker convergence result, that is, all the cluster points of $\{(x^k, y^k)\}$ and $\{z^k\}$ are optimal solutions to the primal and dual problems, respectively. Furthermore, Yang and Zhang [56] applied the proximal ADM to solve l_1 -norm minimization problems in compressive sensing, allowing the step length to be chosen under some conditions rather than only 1. It is notable that when the step length is chosen to be 1, the iteration scheme in [56] simply reduces to (B.2) with $A = I$ and a symmetric and positive definite matrix T . The proof of the convergence in [56] is very similar to the one provided in this appendix but particularly for the l_1 -norm minimization problems.

Acknowledgement. The first and second authors would like to thank Zhang Liu and Lieven Vandenberghé for many useful discussions. The third author would like to thank his PhD student Mr Weimin Miao at the National University of Singapore for discussions on the convergence of alternating direction methods. We would also like to thank the anonymous referees for their many suggestions that help improve the manuscript.

REFERENCES

- [1] H. Attouch and M. Soueycatt. Augmented Lagrangian and proximal alternating direction methods of multipliers in Hilbert spaces. Applications to games, PDE's and control. *Pac. J. Optim.* 5, pp. 17–37 (2009).
- [2] D. Bertsekas and J. N. Tsitsiklis. *Parallel and Distributed Computation: Numerical Methods*. Prentice Hall (1989).
- [3] J. M. Borwein and A. S. Lewis. *Convex Analysis and Nonlinear Optimization*. Springer, 2nd edition (2006).
- [4] J. A. Cadzow. Signal enhancement – a composite property mapping algorithm. *IEEE T. Acoust. Speech* 36, pp. 49–62 (1988).
- [5] J.-F. Cai, E. J. Candés, and Z. Shen. A singular value thresholding algorithm for matrix completion. *SIAM J. Optim.* 20, pp. 1956–1982 (2010).
- [6] E. J. Candés and Y. Plan. Tight oracle bounds for low-rank matrix recovery from a minimal number of random measurements. *IEEE Trans. Inf. Theory* 57, pp. 2342–2359 (2011).
- [7] E. J. Candés and Y. Plan. Matrix completion with noise. *P. IEEE* 98, pp. 925–936 (2010).
- [8] E. J. Candés and B. Recht. Exact matrix completion via convex optimization. *Found. Comput. Math.* 9, pp. 717–772 (2009).
- [9] S. Ma, D. Goldfarb and L. Chen. Fixed point and Bregman iterative methods for matrix rank minimization. *Math. Program.* 128, pp. 321–353 (2011).
- [10] B. De Moor. Total least squares for affinely structured matrices and the noisy realization problem. *IEEE Trans. Signal Process.* 42, pp. 3104–3113 (1994).
- [11] B. De Moor, P. De Gersem, B. De Schutter, and W. Favoreel. DAISY: A database for the identification of systems. *Journal A* 38, pp. 4–5 (1997).
- [12] B. De Schutter. Minimal state-space realization in linear system theory: an overview. *J. Comput. Appl. Math.* 121, pp. 331–354 (2000).
- [13] T. Ding, M. Sznajder, and O. I. Camps. A rank minimization approach to video inpainting. *Proc. IEEE Conf. Comput. Vision*, Rio de Janeiro, Brazil, pp. 1–8, Oct (2007).
- [14] J. Eckstein. Some saddle-function splitting methods for convex programming. *Optim. Method Softw.* 4, pp. 75–83 (1994).
- [15] J. Eckstein and D. P. Bertsekas. On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators. *Math. Prog.* 55, pp. 293–318 (1992).
- [16] M. Elad and P. Milanfar. Shape from moments – an estimation theory perspective. *IEEE Trans. Signal Process.* 52, pp. 1814–1829 (2004).

- [17] M. Fazel. *Matrix Rank Minimization with Applications*. Ph. D. Thesis, Stanford University (2002).
- [18] M. Fazel, H. Hindi and S. Boyd. Log-det heuristic for matrix rank minimization with applications to Hankel and Euclidean distance matrices. *Proc. American Control Conf.*, Denver, Colorado, June (2003).
- [19] M. Fazel, H. Hindi and S. Boyd. A Rank Minimization Heuristic with Application to Minimum Order System Approximation. *Proc. American Control Conf.*, Arlington, Virginia, pp. 4734–4739, June (2001).
- [20] D. Gabay. Applications of the method of multipliers to variational inequalities. In M. Fortin and R. Glowinski, eds., *Augmented Lagrangion Methods: Applications to the Solution of Boundary Problems*. North-Holland, Amsterdam, 1983.
- [21] D. Gabay and B. Mercier. A dual algorithm for the solution of nonlinear variational problems via finite element approximations. *Comput. Math. Appl.* 2, pp. 17-40 (1976).
- [22] R. Glowinski and A. Marroco. Sur l'approximation, par elements finis d'ordre un, et la resolution, par penalisation-dualite, d'une classe de problemes de Dirichlet non lineares. *Revue Francaise d'Automatique, Informatique et Recherche Op'erationelle*. 9 (R-2), pp. 41-76 (1975).
- [23] G. H. Golub, P. Milanfar and J. Varah. A stable numerical method for inverting shape from moments. *SIAM J. Sci. Comput.* 21, pp. 1222–1243 (1999).
- [24] G. H. Golub and C. F. Van Loan. *Matrix Computation*. Johns Hopkins University Press, Baltimore, 3rd edition (1996).
- [25] B. He, L. Liao, D. Han and H. Yang. A new inexact alternating directions method for monotone variational inequalities. *Math. Program.* 92, pp. 103–118 (2002).
- [26] R. H. Keshavan, A. Montanari and S. Oh. Matrix completion from a few entries. *IEEE Trans. Inf. Theory* 56, pp. 2980–2998 (2010).
- [27] J. B. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM J. Optim.* 11, pp. 796–817 (2001).
- [28] M. Laurent. Sums of squares, moment matrices and optimization over polynomials. In *Emerging Applications of Algebraic Geometry*, Vol. 149 of IMA Volumes in Mathematics and its Applications, M. Putinar and S. Sullivant (eds.), Springer, pp. 157–270 (2009).
- [29] K. Lee and Y. Bresler. ADMiRA: atomic decomposition for minimum rank approximation. *IEEE Trans. Inf. Theory* 56, pp. 4402–4416 (2010).
- [30] Y. Liu, D.F. Sun, and K.-C. Toh. An implementable proximal point algorithmic framework for nuclear norm minimization. *Math. Program.* 133, pp. 399–436 (2012).
- [31] Z. Liu and L. Vandenberghe. Interior-point method for nuclear norm approximation with application to system identification. *SIAM. J. Matrix Anal. A.* 31, pp. 1235–1256 (2009).
- [32] Z. Liu and L. Vandenberghe. Semidefinite programming methods for system realization and identification. *Proc. 48th IEEE Conference on Decision and Control*, pp. 4676–4681 (2009).
- [33] L. Ljung. *System Identification: Theory for the User*. Prentice Hall (1999).
- [34] J. Mari, P. Stoica and T. McKelvey. Vector ARMA estimation: a reliable subspace approach. *IEEE Trans. Signal Process.* 48, pp. 2092–2104 (2000).
- [35] I. Markovsky, J. C. Willems, S. Van Huffel and B. De Morr. *Exact and Approximate Modeling of Linear Systems: A Behavioral Approach*. SIAM (2006).
- [36] I. Markovsky, J. C. Willems, S. Van Huffel, B. De Morr and R. Pintelon. Application of structured total least squares for system identification and model reduction. *IEEE T. Automat. Contr.* 50, pp. 1490–1500 (2005).
- [37] R. Meka, P. Jain and I. S. Dhillon. Guaranteed rank minimization via singular value projection. *Proc. Neural Information Process. Systems Conf.*, pp. 937-945, December (2010).
- [38] P. Milanfar, G. Verghese, W. Karl and A. Willsky. Reconstructing polygons from moments with connections to array processing. *IEEE Trans. Signal Process.* 43, pp. 432–443 (1995).
- [39] K. Mohan and M. Fazel. Reweighted nuclear norm minimization with application to system identification. *Proc. American Control Conference* (2010).
- [40] Y. Nesterov. A method for solving a convex programming problem with convergence rate $O(1/k^2)$. *Soviet Math. Dokl.* 27(2), pp. 372–376 (1983).
- [41] Y. Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Kluwer Academic Publishers (2003).
- [42] Y. Nesterov. Excessive gap technique in nonsmooth convex minimization. *SIAM J. Optim.* 16, pp. 235–249 (2005).
- [43] Y. Nesterov. Smooth minimization of non-smooth functions. *Math. Program.* 103, pp. 127–152 (2005).
- [44] P. A. Parrilo. Semidefinite programming relaxations for semialgebraic problems. *Math. Program.* 96, pp. 293–320 (2003).

- [45] T. K. Pong, P. Tseng, S. Ji and J. Ye. Trace norm regularization: reformulations, algorithms, and multi-task learning. *SIAM J. Optim.* 20, pp. 3465–3489 (2010).
- [46] B. Recht, M. Fazel and P. Parrilo. Guaranteed minimum rank solutions of matrix equation via nuclear norm minimization. *SIAM Review* 52, pp. 471–501 (2010).
- [47] R. T. Rockafellar. *Convex Analysis*. Princeton University Press, Princeton (1970).
- [48] R. T. Rockafellar. and R. J-B. Wets. *Variational Analysis*. Springer (1998).
- [49] M. Schuermans, P. Lemmerling, L. De Lathauwer and S. Van Huffel. The use of total least squares data fitting in the shape from moments problem. *Signal Process.* 86, pp. 1109–1115 (2006).
- [50] E. D. Sontag. *Mathematical Control Theory: Deterministic Finite Dimensional Systems*. Second Edition, Springer (1998).
- [51] K.-C. Toh and S. Yun. An accelerated proximal gradient algorithm for nuclear norm regularized least squares problems. *Pac. J. Optim.* 6, pp. 615–640 (2010).
- [52] P. Tseng. On accelerated proximal gradient methods for convex-concave optimization. Submitted to *SIAM J. Optim.* (2008).
- [53] P. Tseng. Approximation accuracy, gradient methods, and error bound for structured convex optimization. *Math. Program.* 125, pp. 263–295 (2010).
- [54] J. C. Willems. From time series to linear system – part I. finite dimensional linear time invariant systems. *Automatica* 22, pp. 561–580 (1986).
- [55] J. C. Willems. From time series to linear system – part II. exact modelling. *Automatica* 22, pp. 675–694 (1986).
- [56] J. Yang and Y. Zhang. Alternating direction algorithms for ℓ_1 -problems in compressive sensing. *SIAM J. Sci. Comput.* 33, pp. 250–278 (2011).
- [57] X. Zhang, M. Burger and S. Osher. A unified primal-dual algorithm framework based on Bregman iteration. *J. Sci. Comput.* 46, pp. 20–46 (2011).