Stable Nonlinear Identification From Noisy Repeated Experiments via Convex Optimization

Mark M. Tobenkin[†,+], Ian R. Manchester[⋆], and Alexandre Megretski[†]

Abstract—This paper introduces new techniques for using convex optimization to fit input-output data to a class of stable nonlinear dynamical models. We present an algorithm that guarantees consistent estimates of models in this class when a small set of repeated experiments with suitably independent measurement noise is available. Stability of the estimated models is guaranteed without any assumptions on the input-output data. We first present a convex optimization scheme for identifying stable state-space models from empirical moments. Next, we provide a method for using repeated experiments to remove the effect of noise on these moment and model estimates. The technique is demonstrated on a simple simulated example.

Index Terms—System identification, nonlinear systems.

I. INTRODUCTION

Building nonlinear dynamical models capable of accurate long term prediction is a common goal in system identification. However, for most model structures multi-step prediction errors have a complex nonlinear dependence on the model parameters. Furthermore, assuring stability of algorithmically generated nonlinear models is a substantial challenge. In many practical situations, where data-sets are limited or under-modeling is present, widely used "one-step" prediction error minimization techniques can render models that are unstable or have poor multi-step predictions. This work presents a convex optimization method for approximating the input-output response of a nonlinear dynamical system via state-space models with stability guarantees. This paper extends recent work in [30], [3] and [18] by providing a family of consistent estimators for a class of stable nonlinear models when a small set of repeated experiments is available. We examine the problem of embedding an input-output identification task inside a state-space modeling framework. We inherit from the methods of [30], [3], [18] an unqualified guarantee of model stability and a cost function that is a convex upper bound on the "simulation error" associated with these models. However, the estimators from [30], [3], [18] are generally not consistent, and for systems that are nearly marginally stable the biasing effect of measurement noise can be quite severe. Furthermore, the complexity of these methods grows undesirably with the number of data points.

We present a modification of algorithms from [30] that mitigates these two difficulties. In particular, a technique that utilizes the problem data through empirical moments only is used. As a result, the complexity of the method generally grows linearly with data-set size. We also provide a method for asymptotically removing the effects of measurement noise on these empirical moments when a small set of repeated experiments are available, utilizing an idea which is superficially similar to instrumental variable methods [14]. We that demonstrate that this technique, a nonlinear extension of [17], recovers consistency when the data is generated by a system within a specific class of models.

A. Previous Work

The use of maximum likelihood and one-step prediction error methods is frequently motivated by the consistency and asymptotic efficiency of the resulting estimators [14]. In the face of limited data or significant under-modeling, these techniques often render models that are unstable or make poor multi-step ahead predictions [8]. Direct minimization of longer term prediction errors have appeared in several forms, including the output-error method for input-output system identification, [27], notions of "best" approximation, [20], and simulation error minimization, [4],[8]. These methods require optimization of a non-convex functional for all but the simplest model structures (e.g. finite impulse response and Volterra type models) and can suffer from local minima [27]. Appealing theoretical properties of these methods (e.g. efficiency and unbiasedness) are often predicated on finding global minima of generically hard nonlinear programming problems.

Several results are available for linear time invariant (LTI) system identification using least squares that provide stability guarantees even in the face of under-modeling (e.g. [25],[26], [31]). It is worth noting that these stability guarantees apply only as the number of available data points tends to infinity and requires an assumption that the data is generated by a (potentially under-modeled) stationary LTI process. Several modified subspace techniques have also been presented to address the issue of model stability. In [32] regularization is used to ensure model stability. In [13] and [12] a joint search over Lyapunov function and system dynamics using convex optimization was used to ensure model stability. The LTI-specific method employed by [13] and [12] is closely related to the technique by which this paper addresses stability.

Several convex relaxation techniques have recently been employed by the Set Membership (SM) identification community to address fixed order identification of LTI systems ([7], [9]). In [7] outer approximations of the set of parameters consistent with bounded noise and stability assumptions are computed. In [9] a convex relaxation approach is suggested for optimization of arbitrary polynomial objectives over the set of LTI models consistent with a given data-set and a set of stability and bounded noise assumptions. A similar approach is taken for identifying Linear Parameter Varying systems in [6]. By contrast, in this work we examine a "convex restriction" approach where inner approximations of the set of stable models are used to guarantee stability and convex upper bounds on the cost function of interest are used as a surrogate objective.

B. Outline

The paper proceeds as follows. Section II presents the notation, problem setup, and a bias elimination strategy employed in this work. Next, Section III provides a convex parameterization of stable state-space models and a convex upper bound for simulation error. This parameterization and objective are then combined with the bias elimination strategy in Section IV, wherein a system identification algorithm based on semidefinite programming is given along with asymptotic analysis of the method. Finally, a comparison of the proposed algorithm to two alternative least-squares based methods is provided in Section V.
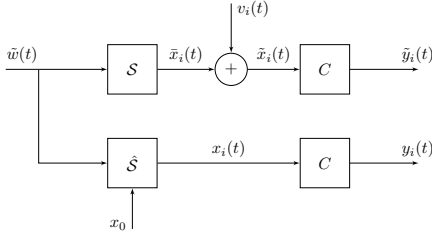
Fig. 1. The experimental setup considered in this work.

II. PRELIMINARIES

In this section we introduce basic notation, and present the problem setup to be addressed in the paper.

A. Notation

$\mathbb{C}^{k \times n}$ stands for the set of all k-by-n complex matrices, with $\mathbb{C}^n$ being a shorthand for $\mathbb{C}^{n \times 1}$. $\mathbb{R}^{k \times n}$ and $\mathbb{R}^n$ are the subsets of real matrices from $\mathbb{C}^{k \times n}$ and $\mathbb{C}^n$ respectively. $\mathbb{Z}_+^n$ is the subset of $\mathbb{R}^n$ whose elements are non-negative integers.

We use some notation from MATLAB, where A', $[A, B]$, and $[A; B]$ denote, respectively, Hermitian conjugation, horizontal concatenation, and vertical concatenation of matrices. For $R \in \mathbb{C}^{k \times n}$ we denote by $[R]_{a,b}$ the scalar element in the a-th row and b-th column of R, with the shorthand $[v]_d = v_{d,1}$ used for $v \in \mathbb{C}^n$. In addition, for $v \in \mathbb{C}^n$ and $\alpha \in \mathbb{Z}_+^n$,

$$v^\alpha := \prod_{d=1}^n [v]_d^{[\alpha]_d}$$

is the monomial function of v with *vector degree* α, and *scalar degree* $\|\alpha\|_1$, where, for $w \in \mathbb{C}^n$, $\|w\|_1 := \sum_{i=1}^n |[w]_i|$ is the ℓ_1 norm of w. For Hermitian matrices $A, B \in \mathbb{C}^{n \times n}$ (i.e. such that $A = A'$ and $B = B'$), $A \geq B$ (or $A > B$) means that $A - B$ is positive semidefinite (respectively, positive definite). For $R = R' \in \mathbb{C}^{n \times n}$ and $v \in \mathbb{C}^n$ we use the shorthand $|v|_R^2 = v'Rv$. Moreover, when $R \geq 0$, we also write $|v|_R := \sqrt{v'Rv}$. When W is a set, $\ell_T(W)$ denotes the set of all functions $w : \{0, 1, \ldots, T\} \to W$. Naturally, the elements of $\ell(W)$ are finite length *sequences* of elements from W. The notation $\xrightarrow{\text{P}}$ refers to convergence in probability.

B. Problem Setup

We define a *data set with N experiments of length T, n_w-dimensional input, and n_x-dimensional state* as a collection $(\tilde{w}, \tilde{x}_1, \ldots, \tilde{x}_N)$ of sequences $\tilde{w} \in \ell_T(\mathbb{R}^{n_w})$, $\tilde{x}_i \in \ell_T(\mathbb{R}^{n_x})$. $\mathcal{D}(n_x, n_w, N, T)$ stands for the set of all data sets of given dimensions, number of experiments, and signal length. Accordingly, $\mathcal{D}(n_x, n_w, N) = \cup_{T=0}^\infty \mathcal{D}(n_x, n_w, N, T)$ stands for the set of all data sets with unspecified signal length.

In applications, each $\tilde{x}_i(t)$ is the result of feeding the same input $\tilde{w}(t)$ into a system, $\mathcal{S}$, and measuring the sum $\tilde{x}_i(t) = \bar{x}_i(t) + v_i(t)$, where $\bar{x}_i(t)$ is the "true system response" and $v_i(t)$ is corrupting measurement noise. Additionally, in order to set a measure of quality for model predictions, we define an output signal $\tilde{y}_i(t)$ is defined by $\tilde{y}_i(t) = C\tilde{x}_i(t)$, for some fixed matrix $C \in \mathbb{R}^{n_y \times n_x}$. Informally, the identification objective will be to accurately predict the input-output behavior of this a system with $\tilde{w}(t)$ taken as input and $\tilde{y}(t)$ taken as an output (alternatively C can be seen as weighting the importance certain components of $\tilde{x}_i(t)$). This experimental setup is depicted in Figure 1.

The underlying assumption is that the collection of signals $\bar{x}_i(t)$ constitute a reasonable state, or reduced state, for a state-space model approximating the system behavior. As an example, when identifying a SISO system with input $u = u(t)$ and output $y = y(t)$, one can imagine feeding in $N \cdot D$ samples of a D-periodic input $\tilde{u}(t)$ and measuring $\tilde{y}(t) = \bar{y}(t) + v_y(t)$, where $\bar{y}(t)$ is the true system response and $v_y(t)$ is measurement noise. In this case, one could use the above setup with $n_x = n_w = n < D$ by taking

$$\tilde{w}(t) = \begin{bmatrix} \tilde{u}(t+n) \\ \vdots \\ \tilde{u}(t+1) \end{bmatrix}, \quad \tilde{x}_{i+1} = \begin{bmatrix} \tilde{y}(t+n+iD) \\ \vdots \\ \tilde{y}(t+1+iD) \end{bmatrix},$$

$$v_{i+1} = \begin{bmatrix} v_y(t+n+iD) \\ \vdots \\ v_y(t+1+iD) \end{bmatrix},$$

for $i \in \{0, \ldots, N-1\}$ and $t \in \{0, \ldots, D-n\}$. Here, the matrix $C \in \mathbb{R}^{1 \times n}$ might be $C = \begin{bmatrix} 1 & 0 & \ldots & 0 \end{bmatrix}$.

C. State Space Models

In general, a nonlinear state space model (time invariant, in discrete time) with n_w-dimensional input and n_x-dimensional state is specified by a function $a : \mathbb{R}^{n_x} \times \mathbb{R}^{n_w} \to \mathbb{R}^{n_x}$, which in turn defines the *input-output function* $G_a : \mathbb{R}^{n_x} \times \ell(\mathbb{R}^{n_w}) \to \ell(\mathbb{R}^{n_w})$ mapping *initial state* $x_0 \in \mathbb{R}^{n_x}$ and *input sequence* $w \in \ell(\mathbb{R}^{n_w})$ to the *output sequence* $x \in \ell(\mathbb{R}^{n_x})$ according to

$$x(t) = a(x(t-1), w(t)), \qquad x(0) = x_0. \qquad (1)$$

For $w : \mathbb{Z}_+ \mapsto \mathbb{R}^{n_w}$, we define $x = G_a(x_0, w)$ to be the sequence similarly defined by this recurrence. Let $x = G_a(x_0, w)$, $\hat{x} = G_a(\hat{x}_0, w)$ be two responses of system (1) to the same input $w : \mathbb{Z}_+ \mapsto \mathbb{R}^{n_w}$ and different initial conditions $x_0, \hat{x}_0$. We call system (1) ℓ^2-*incrementally stable* when $x - \hat{x}$ is square summable for all $x_0, \hat{x}_0, w$. The system (1) is *incrementally exponentially stable* if there exist constants $c > 0$ and $\rho \in (0, 1)$, independent of $x_0, \hat{x}_0$ and w, such that $|x(t) - \hat{x}(t)| \leq c\rho^t |x_0 - \hat{x}_0|$ for all $x_0, \hat{x}_0, w$ and $t \geq 0$.

This paper deals with subsets of state space models (1) which have more specific finite dimensional structure. For positive integers n_x, n_w, n_θ let Θ, Φ, Ψ be a non-empty set $\Theta \subset \mathbb{R}^{n_\theta}$ and two sequences $\Phi = \{\phi_i\}_{i=1}^{n_\theta}$, $\Psi = \{\psi_i\}_{i=1}^{n_\theta}$ of real analytical functions $\phi_i : \mathbb{R}^{n_x} \times \mathbb{R}^{n_w} \to \mathbb{R}^{n_x}$, $\psi_i : \mathbb{R}^{n_x} \to \mathbb{R}^{n_x}$. We say that the 3-tuple (Θ, Φ, Ψ) is a *stable projective parameterization with n_w inputs, n_x states, and n_θ parameters* when, for all $\theta \in \Theta$,

- the function $e_\theta : \mathbb{R}^{n_x} \to \mathbb{R}^{n_x}$ defined by $e_\theta(x) = \sum_{i=1}^{n_\theta} [\theta]_i \psi_i(x)$ is a bijection;
- the state space model (1) with $a = a_\theta$ defined by

$$a_\theta(x, w) = e_\theta^{-1}(f_\theta(x, w)) \qquad (2)$$

$$f_\theta(x, w) = \sum_{i=1}^{n_\theta} [\theta]_i \phi_i(x, w), \quad e_\theta(\xi) = \sum_{i=1}^{n_\theta} [\theta]_i \psi_i(\xi), \quad (3)$$

is ℓ^2-incrementally stable.

Once a stable projective parameterization (Θ, Φ, Ψ) is selected, a stable state space model can be defined by specifying a vector parameter $\theta \in \Theta$.

Recent discussion and applications of incremental stability and the related notions of contractive and convergent systems can be found in [11], [16], [1], [29], and [2]. This property is related to familiar "fading memory" conditions employed in other identification and system approximation papers (e.g. [5], [22], [23]), though we note

that the condition as defined above does not restrict the behavior of solutions with different input sequences.

In practice, both the particular construction of the signals $\tilde{w}(t)$ and $\tilde{x}_i(t)$ from measureable quantities, and the selection of the sequences of functions Ψ and Φ defining (e_θ, f_θ) should be guided by a model selection criteria such as cross-validation [14]. We consider both of these selections fixed for the remainder of the paper.

D. Empirical Moments

For given positive integers n_x, n_w, N let $n_z = 2n_x + n_w$ be the dimension of the vectors

$$\tilde{z}_i(t) = \begin{bmatrix} \tilde{x}_i(t) \\ \tilde{x}_i(t-1) \\ \tilde{w}(t) \end{bmatrix} \qquad (t \in \{1, \ldots, T\},\ i \in \{1, \ldots, N\}) \quad (4)$$

defined by the data set $\Xi = (\tilde{w}, \tilde{x}_1, \ldots, \tilde{x}_N) \in \mathcal{D}(n_x, n_w, N, T)$. For $\alpha \in \mathbb{Z}_+^{n_z}$ such that $\|\alpha\|_1 \leq N$ and for $z_1, \ldots, z_N \in \mathbb{R}^{n_z}$ define

$$p_\alpha(z_1, \ldots, z_N) = \prod_{i=1}^{\|\alpha\|_1} [z_i]_{\beta_\alpha(i)},$$

where

$$\beta(i) = \min\left\{ d \in \{1, \ldots, N\} \ \middle|\ \sum_{j=1}^{d} [\alpha]_j \geq i \right\}.$$

By construction, $z^\alpha = p_\alpha(z, \ldots, z)$, so that one can view p_α as a multi-linear function which generates the monomial z^α when evaluated on the multi-diagonal (note that such multi-linear functions are not uniquely defined by α). For a given data set $\Xi = (\tilde{w}, \tilde{x}_1, \ldots, \tilde{x}_N)$ and $\alpha \in \mathbb{Z}_+^{n_z}$ define the *linearized empirical moment* $\tilde{\mu}_\alpha(\Xi)$ by

$$\tilde{\mu}_\alpha = \tilde{\mu}_\alpha(\Xi) = \frac{1}{T} \sum_{t=1}^{T} p_\alpha(\tilde{z}_1(t), \ldots, \tilde{z}_N(t)). \quad (5)$$

Since it is sometimes convenient to emphasize $\tilde{\mu}_\alpha(\Xi)$ as a function of variable α with a fixed Ξ, we will also use the equivalent notation $\tilde{\mu}_\alpha(\Xi) = \hat{\mu}_\Xi(\alpha)$. According to this notation, for a given data set Ξ with N experiments, n_x states, and n_w inputs, $\hat{\mu}_\Xi$ is a real-valued function defined on the set of elements $\alpha \in \mathbb{Z}_+^{2n_x+n_w}$ such that $\|\alpha\|_1 \leq N$.

Informally speaking, linearized empirical moments represent an attempt at "de-noising" the data contained in the vectors $\tilde{z}_i(t)$, as defined by (4) in the case when $\tilde{x}_i(t) = \bar{x}(t) + v_i(t)$ for $t \in \{0, 1, \ldots, T\}$, where $\bar{x} = \bar{x}(t)$, the "true system response", does not depend on the experiment number i, and the noise variables $v_i(t)$ are suitably independent of $\bar{x}$ and of each other, to produce good estimates $\tilde{\mu}_\alpha(\Xi)$ of the standard empirical moments

$$\mu_\alpha(\bar{x}, \tilde{w}) = \frac{1}{T} \sum_{t=1}^{T} \bar{z}(t)^\alpha, \qquad \bar{z}(t) = \begin{bmatrix} \bar{x}(t) \\ \bar{x}(t-1) \\ \tilde{w}(t) \end{bmatrix}. \quad (6)$$

This approach is inspired by instrumental variable (IV) techniques, [28], with repeated experiments playing a role comparable to a specific choice of instruments. Rather than asymptotically approximate a least squares parameter estimate, as in IV methods, this work focuses on asymptotically minimizing an alternative convex loss function that depends only on empirical moments. To have a meaningful convergence of the linearized empirical moments we require both the aforementioned independence of the noise sequences, to be made more precise shortly, and that the true system responses, $\bar{x}_i(t)$, tend to one another despite their differing initial conditions.

E. Persistence of Excitation

The following notion of persistence of excitation will be used in our consistency analysis.

Definition 1: Fix two signals $w : \mathbb{Z}_+ \to \mathbb{R}^{n_w}$ and $x : \mathbb{Z}_+ \to \mathbb{R}^{n_x}$, and let $w^{(T)}$ and $x^{(T)}$ be the restriction of these signals to $\{0, \ldots, T\}$. For a given function $a : \mathbb{R}^{n_x} \times \mathbb{R}^{n_w} \to \mathbb{R}^{n_x}$, we say a pair of signals (w, x) is *persistently exciting for a* if there exists a positive measure π on $\mathbb{R}^{n_x} \times \mathbb{R}^{n_w}$ such that π is supported on an open set, and for every finite subset $\aleph = \{\alpha_j\}_{j=1}^{|\aleph|}$ of $\mathbb{Z}_+^{n_z}$

$$\liminf_T \lambda_{\min}(M_T^\aleph - M_\pi^\aleph) \geq 0,$$

where $M_T^\aleph, M_\pi^\aleph \in \mathbb{R}^{|\aleph| \times |\aleph|}$ are defined by:

$$[M_T^\aleph]_{i,j} = \mu_{\alpha_i + \alpha_j}(w^{(T)}, x^{(T)}),$$

$$[M_\pi^\aleph]_{i,j} = \int [a(x, w); x; w]^{\alpha_i + \alpha_j} d\pi(x, w).$$

Informally, this non-standard notion of persistence of excitation will be employed to establish a connection between

$$\frac{1}{T} \sum_{t=1}^{T} |e_\theta(a(x(t), w(t))) - f_\theta(x(t), w(t))|^2$$

vanishing as $T \to \infty$ and a being equivalent to $a_\theta = e_\theta^{-1} \circ f_\theta$. The use of a projective representation, i.e. a_θ being implicitly defined, renders several complications to standard consistency arguments based on strong convexity (for example, the e_θ and f_θ that define a_θ can be non-unique). The above notion of persistence will be used to circumvent these difficulties.

F. Data-Matching Error

We examine the following loss function for identifying models.

Definition 2: The *T-step simulation error*, J_T^{SE}, is a function of an $a : \mathbb{R}^{n_x} \times \mathbb{R}^{n_w} \to \mathbb{R}^{n_w}$, an initial condition vector $x_0 \in \mathbb{R}^{n_x}$, and two signals $\tilde{w} \in \ell_T(\mathbb{R}^{n_w})$, and $\tilde{x} \in \ell_T(\mathbb{R}^{n_x})$, defined by

$$J_T^{SE}(a, x_0, \tilde{x}, \tilde{w}) = \frac{1}{T} \sum_{t=0}^{T-1} |C(\tilde{x}(t) - x(t))|^2, \quad (7)$$

where $x = G_a(x_0, \tilde{w})$.

G. Data Generation Mechanism

Two data generation mechanisms, defined by considering data sequences as stochastic processes, will be analyzed in this work. These mechanisms consider signals defined on an infinite horizon, i.e. $\tilde{w} : \mathbb{Z}_+ \to \mathbb{R}^{n_w}$ and $\tilde{x}_i : \mathbb{Z}_+ \to \mathbb{R}^{n_w}$, for $i \in \{1, \ldots, N\}$. We express $\tilde{x}_i$ as the sum of two signals $\bar{x}_i$ and v_i, again representing the true system response and measurement noise respectively. Let $\tilde{x}_i^{(T)}$, and $\tilde{w}^{(T)}$ be the restrictions to $\{0, \ldots, T\}$ of $\tilde{x}_i$, and $\tilde{w}$ respectively. Then we define the data set $\Xi_T \in \mathcal{D}(n_x, n_w, N, T)$ by

$$\Xi_T = (\tilde{w}^{(T)}, \tilde{x}_1^{(T)}, \ldots, \tilde{x}_N^{(T)}).$$

The following assumptions define the first data generation mechanism.

(A1) The signal $\tilde{w}(t)$ is a stochastic process for $t \in \mathbb{Z}$, which is uniformly bounded in t.

(A2) The signals $v_i(t)$ are i.i.d. zero mean bounded stochastic processes independent of one another, $\tilde{w}(t)$ and each $\bar{x}_i(t)$.

(A3) The signals $\bar{x}_i(t)$ are stochastic processes which are uniformly bounded in i and t. There exist constants $c > 0$ and $\rho \in (0, 1)$ such that.

$$|\bar{x}_i(t) - \bar{x}_j(t)| \leq c\rho^t, \qquad \forall i, j \in \{1, \ldots, D\}, t \in \mathbb{Z}_+, \quad (8)$$

almost surely;

An alternative, less general, data-generation mechanism in given by the assumptions (A1), (A2), and the following.

(A4) There exists a function $a_0 : \mathbb{R}^{n_x} \times \mathbb{R}^{n_w} \to \mathbb{R}^{n_x}$ such that (1) with $a = a_0$ defines a BIBO and incrementally exponentially stable system and $\bar{x}_i = G_{a_0}(\bar{x}_{i0}, \tilde{w})$ for some unknown $\bar{x}_{i0} \in \mathbb{R}^{n_x}$, for $i \in \{1, \ldots, N\}$. The pairs of signals $(\tilde{w}, \bar{x}_i)$ are persistently exciting with respect to a_0, as in Definition 1, with probability one.

The appendix contains practical conditions on $\tilde{w}(t)$ and a_0 that ensure (A4) holds. It is immediate that assumptions (A1) and (A4) together imply (A3).

H. Identification Objective

In this paper, we view system identification algorithms as 4-tuples $(\mathcal{A}, \Theta, \Phi, \Psi)$, where (Θ, Φ, Ψ) is a stable projective parameterization with n_w inputs, n_x states, and n_θ parameters, and $\mathcal{A}$ is a function $\mathcal{A} : \mathcal{D}(n_x, n_w, N) \to \Theta$ mapping data sets to parameter vectors from Θ.

Specifically, we are interested in generating efficient *moments-based* system identification algorithms $(\mathcal{A}, \Theta, \Phi, \Psi)$, i.e. those for which the function $\mathcal{A} : \mathcal{D}(n_x, n_w, N) \to \Theta$ has the form $\mathcal{A}(\Xi) = \mathbf{A}(\hat{\mu}_\Xi)$, which means that the resulting identified model is a function of the linearized empirical moments $\hat{\mu}_\Xi(\alpha)$ with $\alpha \in \mathbb{Z}_+^{2n_x+n_w}$ satisfying $\|\alpha\|_1 \leq N$.

The main contribution of this paper is the construction of moments-based system identification algorithms $(\mathcal{A}, \Theta, \Phi, \Psi)$ and sets $\Theta_0 \subset \Theta$ with the following properties:

(a) the set $a_{\Theta_0} = \{a_\theta : \theta \in \Theta_0\}$ of models (1) generated by Θ_0 is sufficiently broad, in the sense that every stable linear state space model $a(x, w) = Ax + Bw$ is in a_{Θ_0}, and some non-linear functions are contained in a_{Θ_0} as well;

(b) when a sequence of data sets $\{\Xi_T\}_{T=1}^\infty$ is generated by signals $(\tilde{w}, \tilde{x}_1, \ldots, \tilde{x}_N)$ satisfying assumptions (A1),(A2) and (A3), then $\theta_T = \mathcal{A}(\Xi_T)$ asymptotically (with respect to T) minimizes an upper bound for

$$\frac{1}{N} \sum_{i=1}^N J_T^{SE}(a_\theta, \tilde{x}_0, \bar{x}_i^{(T)}, \tilde{w}^{(T)}), \tag{9}$$

amongst all $\theta \in \Theta$.

(c) when a sequence of data sets $\{\Xi_T\}_{T=1}^\infty$ is generated by signals $(\tilde{w}, \tilde{x}_1, \ldots, \tilde{x}_N)$ satisfying assumptions (A1),(A2) and (A4) for some $a_0 \in a_{\Theta_0}$, then for $\theta_T = \mathcal{A}(\Xi_T)$ the convergence of $a_{\theta_T}(x, w)$ to $a_0(x, w)$ takes place uniformly on every compact subset compact subset of $\mathbb{R}^{n_x} \times \mathbb{R}^{n_w}$.

III. CONVEX PARAMETERIZATION OF MODELS

In this section we introduce the main construction of this paper: a special class of stable projective parameterizations (Θ, Φ, Ψ), in which Θ are convex sets defined by a family of linear matrix inequalities arrived at via an application of the sum-of-squares relaxation, [21]. The construction is motivated by the earlier approaches from [3], [30], and [18], and is intended to improve consistency of the associated system identification algorithms.

In the following definition, x, ξ, Δ and q are real vector variables of dimensions n_x, and w is a real vector variable of dimension n_w. In addition, $z = [\xi; x; w]$ and $v = [\xi; x; w; \Delta; q]$ are the real vector variables of dimensions $2n_x + n_w$ and $4n_x + n_w$ respectively, constructed by concatenating ξ, x, w, Δ, and q. Given an positive

integer N let

$$\mathcal{P}_N = \left\{ p(z) = \sum_{\alpha \in \mathbb{Z}_+^{n_z} \,|\, \|\alpha\|_1 \leq N} c_\alpha z^\alpha \; : \; c_\alpha \in \mathbb{R} \right\}$$

denote the set of all polynomials composed of monomials with scalar degrees no greater than N.

Given a positive integer N, a positive constant δ, and a function $\Pi : \mathbb{R}^{4n_x+n_w} \to \mathbb{R}^{n_\Pi}$, let $\Theta(N, \delta, \Pi)$ be the set of all pairs (θ, r) of vectors $\theta \in \mathbb{R}^{n_\theta}$ and $r \in \mathcal{P}_N$ for which there exist matrices $P = P' \in \mathbb{R}^{n_x \times n_x}$, $\Sigma_i = \Sigma_i' \in \mathbb{R}^{n_\Pi \times n_\Pi}$ (for $i \in \{1, 2\}$), and a positive scalar ϵ such that

$$P \geq \delta I, \; \Sigma_1 \geq 0, \; \Sigma_2 \geq 0, \tag{10}$$

$$r(z) + 2\Delta'[e_\theta(x + \Delta) - e_\theta(x)] - |\Delta|_{P+C'C}^2 \tag{11}$$
$$+ |q|_P^2 - 2q'(f_\theta(x + \Delta, w) - e_\theta(\xi)) = \Pi(v)'\Sigma_1\Pi(v),$$

$$2\Delta'[e_\theta(x + \Delta) - e_\theta(x)] - |\Delta|_{P+\epsilon I}^2 \tag{12}$$
$$+ |q|_P^2 - 2q'(f_\theta(x + \Delta, w) - f_\theta(x, w)) = \Pi(v)'\Sigma_2\Pi(v),$$

where e_θ and f_θ are defined by (4). By construction, $\Theta(N, \delta, \Pi)$ is a convex set defined by a family of linear matrix inequalities.

Remark 1: The purpose of (11) is to establish the condition

$$r(z) + |e_\theta(x + \Delta) - e_\theta(x)|_{P^{-1}}^2 \tag{13}$$
$$\geq$$
$$|f_\theta(x + \Delta, w) - e_\theta(\xi)|_{P^{-1}}^2 + |C\Delta|^2,$$

which in turn serves as a dissipation inequality used to bound simulation error when using model (1) with $a = a_\theta$. The purpose of (12) is to ensure that e_θ is a bijection and establish the condition

$$|e_\theta(x + \Delta) - e_\theta(x)|_{P^{-1}}^2 \tag{14}$$
$$\geq$$
$$|f_\theta(x + \Delta, w) - f_\theta(x, w)|_{P^{-1}}^2 + \epsilon|\Delta|^2,$$

which is a non-linear version of the Lyapunov inequality, used to prove that the model (1) with $a = a_\theta$ is ℓ^2-incrementally stable.

The following statement explains, partially, the utility of this construction.

Lemma 1: If Θ is the set of all $\theta \in \mathbb{R}^{n_\theta}$ such that $(\theta, r) \in \Theta(N, \delta, \Pi)$ for some r, then (Θ, Φ, Ψ) is a stable projective parameterization. Furthermore, for each $(\theta, r) \in \Theta(N, \delta, \Pi)$ and data set $\Xi = (\tilde{w}, \tilde{x}_1, \ldots, \tilde{x}_n) \in \mathcal{D}(n_x, n_w, N, T)$ the function

$$\hat{J}_r(\Xi) := \frac{1}{NT} \sum_{t=1}^T \sum_{i=1}^N r(\tilde{z}_i(t)) \tag{15}$$

satisfies $\hat{J}_r(\Xi) \geq \frac{1}{N} \sum_{i=1}^N J_T^{SE}(a_\theta, \tilde{x}_i(0), \tilde{x}_i, \tilde{w})$.

Proof. For Θ as defined above to be a valid projective parameterization requires that e_θ be a bijection for all $\theta \in \Theta$. The equality (12) holding for some $P \geq 0$ and $\Sigma_2 \geq 0$ implies that

$$2\Delta'(e_\theta(x + \Delta) - e_\theta(x)) \geq \epsilon|\Delta|^2 \tag{16}$$

holds for all $x, \Delta \in \mathbb{R}^n$. As e_θ is continuous, this condition implies e_θ is a bijection ([24], Theorem 18.15).

Next, we establish the connection between the conditions (11) and (12) and the inequalities (13) and (14). For all $a, b \in \mathbb{R}^{n_x}$ and symmetric, positive definite $P \in \mathbb{R}^{n_x \times n_x}$, the inequality

$$|a|_{P^{-1}}^2 \geq 2b'a - |b|_P^2 \tag{17}$$

holds due to the fact that $|a - Pb|^2_{P^{-1}} \geq 0$. Fixing any $(\theta, r) \in \Theta(N, \delta, \Pi)$, and applying (17) with $a = e_\theta(x + \Delta) - e_\theta(x)$ and $b = \Delta$, we see that there exists a $P \in \mathbb{R}^{n_x \times n_x}$ such that:

$$r(z) + |e_\theta(x + \Delta) - e_\theta(x)|^2_{P^{-1}} - |C\Delta|^2$$
$$+ |q|^2_P - 2q'(f_\theta(x + \Delta, w) - e_\theta(\xi)) \geq 0,$$

and

$$|e_\theta(x + \Delta) - e_\theta(x)|^2_{P^{-1}} - \epsilon|\Delta|^2$$
$$+ |q|^2_P - 2q'(f_\theta(x + \Delta, w) - f_\theta(x, w)) \geq 0,$$

hold for all x, ξ, Δ, q in $\mathbb{R}^{n_x}$ and $w \in \mathbb{R}^{n_w}$. Analytically minimizing these expressions with respect to q demonstrates these inequalities imply (13) and (14) hold for all $x, \xi, \Delta \in \mathbb{R}^{n_x}$ and $w \in \mathbb{R}^{n_w}$.

Fix $x_{01}, x_{02} \in \mathbb{R}^{n_x}$ and $w : \mathbb{Z}_+ \mapsto \mathbb{R}^{n_w}$, and let x_i be the solution $x_i = G_{a_\theta}(x_{0i}, w)$ for $i \in \{1, 2\}$. The inequality (14) with $x \equiv x_1(t)$ and $\Delta \equiv x_2(t) - x_1(t)$ implies:

$$|e_\theta(x_{02}) - e_\theta(x_{01})|^2_{P^{-1}} \geq |e_\theta(x_2(T)) - e_\theta(x_1(T))|^2_{P^{-1}}$$
$$+ \epsilon \sum_{t=0}^{T-1} |x_1(t) - x_2(t)|^2.$$

As $P \geq 0$, we conclude that (1) with $a \equiv a_\theta$ is ℓ^2-incrementally stable. Take $x_i = G_{a_\theta}(\tilde{x}_i(0), \tilde{w})$, then examining (13) with $z = \tilde{z}_i(t)$ and $\Delta = x_i(t) - \tilde{x}_i(t)$ leads to

$$\sum_{t=1}^{T} r(\tilde{z}_i(t)) \geq |e_\theta(x_i(T)) - e_\theta(\tilde{x}_i(T))|^2_{P^{-1}} + \sum_{t=0}^{T-1} |C(x_i(t) - \tilde{x}_i(t))|^2,$$

from which one can readily conclude $\hat{J}_r(\Xi) \geq J_T^{SE}(a_\theta, \Xi)$. $\qquad \square$

The following definition provides a subset of systems for which we can establish consistency results.

Definition 3: The *recoverable set* defined by $\Theta(N, \delta, \Pi)$ is the set of functions $a : \mathbb{R}^{n_x} \times \mathbb{R}^{n_w} \mapsto \mathbb{R}^{n_x}$ such that there exists a pair $(\theta, r) \in \Theta(N, \delta, \Pi)$ with $a(x, w) \equiv a_\theta(x, w)$ and

$$r([a(x, w); x; w]) = 0 \qquad (18)$$

for all $(x, w) \in \mathbb{R}^{n_x} \times \mathbb{R}^{n_w}$.

The following lemma establishes that the subset of a_Θ consisting of recoverable models can be made to include all stable, linear state-space models of appropriate dimensionality.

Lemma 2: Let Φ and Ψ be finite sequences of real analytic functions, as above, whose respective spans include all linear functions, and let $\Pi : \mathbb{R}^{4n_x + n_w} \mapsto \mathbb{R}^{n_\Pi}$ be a function such that the span of its components include all linear functions. Then for all $C \in \mathbb{R}^{n_y \times n_x}$ and $N \geq 2$ the recoverable set defined by $\Theta(N, \delta, \Pi)$ includes all functions $a : \mathbb{R}^{n_x} \times \mathbb{R}^{n_w} \mapsto \mathbb{R}^{n_x}$ given by

$$a(x, w) = Ax + Bw, \qquad (19)$$

where $A \in \mathbb{R}^{n_x \times n_x}$ is Schur (stable) and $B \in \mathbb{R}^{n_x \times n_w}$.

Proof. As A is Schur, there exists a symmetric positive definite matrix P solving the Lyapunov equation $P - A'PA = C'C + \delta I$ for any positive δ. Choose θ such that:

$$e_\theta(x) = Px, \quad f_\theta(x, w) = Pa(x, w).$$

The constraint (11) is therefore equivalent to

$$r(z) + \Delta'P\Delta + |q|^2_P$$
$$+ 2q'(PA\Delta - e(\xi) + f(x, w)) - |C\Delta|^2 = \Pi(v)'\Sigma_1\Pi(v).$$

Explicit minimization w.r.t. q shows that the polynomial on the left hand side of this equality is lower bounded by:

$$r(z) + \Delta'P\Delta - |PA\Delta - e(\xi) + f(x, w)|^2_{P^{-1}} - |C\Delta|^2.$$

This is a concave function of Δ as $\Delta'(P - A'PA - C'C)\Delta = -\delta|\Delta|^2$. Explicit minimization w.r.t Δ provides

$$r(z) - |e(\xi) - f(x, w)|^2_Q$$

as a lower bound for the original polynomial for some $Q = Q' \geq 0$. The function $r(z) = |e(\xi) - f(x, w)|^2_Q$, belongs to $\mathcal{P}_N$ and clearly satisfies (18). Furthermore, with this choice of $r(z)$ the left hand side of (11) is a non-negative quadratic polynomial so that there exists an appropriate choice of $\Sigma_1 \geq 0$ to ensure (11) holds. A similar analysis shows that an appropriate choice of $\Sigma_2 \geq 0$ also exists, thus a belongs to the recoverable set. $\qquad \square$

A simple example of a nonlinear function $a : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ belonging to such a recoverable set is given by

$$e(a(x, w)) = \frac{1}{2}x + b(w)$$

where $e(x) = \frac{3}{2}x + x^3$ and $b : \mathbb{R} \to \mathbb{R}$ is an arbitrary polynomial. That an appropriate recoverable set exists is shown in the appendix.

IV. IDENTIFICATION ALGORITHM

This section presents an algorithm for transforming data sets $\Xi = (\tilde{w}, \tilde{x}_1, \ldots, \tilde{x}_N) \in \mathcal{D}(n_x, n_w, N)$ into parameter vectors $\hat{\theta} \in \mathbb{R}^{n_\theta}$, followed by an asymptotic analysis of the algorithm. For the remainder of this section we define

$$\aleph = \{\alpha \in \mathbb{Z}_+^{n_z} \mid 2\|\alpha\|_1 \leq N\}.$$

Algorithm $\mathcal{A}(\delta, \Pi, \kappa)$:

(i) Select a constant $\delta > 0$ and a function $\Pi : \mathbb{R}^{4n_x + n_w} \to \mathbb{R}^{n_\Pi}$, as described in Section III. Additionally, select a constant $\kappa \in (0, \infty]$.

(ii) Form the matrix $\tilde{M}_\Xi \in \mathbb{R}^{|\aleph| \times |\aleph|}$ given by:

$$[\tilde{M}_\Xi]_{j_1, j_2} = \tilde{\mu}_{\alpha_{j_1} + \alpha_{j_2}}(\Xi),$$

where $\tilde{\mu}_\alpha(\cdot)$ are the linearized empirical moments defined by (5), and let $\hat{M}_\Xi$ be the projection of $\frac{1}{2}(\tilde{M}_\Xi + \tilde{M}'_\Xi)$ onto the closed convex cone of positive semidefinite matrices.

(iii) Find the $\theta \in \mathbb{R}^{n_\theta}$, $r \in \mathcal{P}_N$ and $R = R' \in \mathbb{R}^{|\aleph| \times |\aleph|}$ that minimize:

$$\text{tr}(R\hat{M}_\Xi)$$

subject to $(\theta, r) \in \Theta(N, \delta, \Pi)$,

$$r(z) = \sum_{j_1=1}^{|\aleph|} \sum_{j_2=1}^{|\aleph|} [R]_{j_1, j_2} z^{\alpha_{j_1} + \alpha_{j_2}},$$

and $\|R\|_F^2 \leq \kappa$. Take $(\hat{\theta}, \hat{r}, \hat{R})$ to be an optimizing (θ, r, R).

Remark 2: Note that the algorithm is well defined if any subset of the set of vector degrees $\aleph$ is substituted in lieu of $\aleph$.

Remark 3: Examining the definition of R and $\hat{M}$, one sees that when $\tilde{x}_1 = \tilde{x}_2 = \ldots = \tilde{x}_N$, $\tilde{M}_\Xi = \hat{M}_\Xi$ and the objective function $\text{tr}(R\hat{M}_\Xi)$ is equal to $\hat{J}_{\hat{r}}(\Xi)$, the previously established upper bound on simulation error. The additional parameter κ, when finite, ensures that the optimal value of the optimization problem of step (iii) has a continuous dependence on $\hat{M}_\Xi$ (and by extension, the linearized empirical moments).

A. Asymptotic Analysis

This section analyzes the properties of algorithm $\mathcal{A}$ when data sets are generated according to one of the two data generation mechanisms described in Section II-G. By $(\tilde{w}, \tilde{x}_1, \ldots, \tilde{x}_N)$, $(v_1, \ldots, v_N)$, and $(\bar{x}_1, \ldots, \bar{x}_N)$, we mean those signals described in assumptions (A3) or (A4). Let $\tilde{w}^{(T)}$, $\tilde{x}^{(T)}$ and $\bar{x}_i^{(T)}$ are the restrictions of $\tilde{w}, \tilde{x}_i$ and

$\bar{x}_i$ to $\{0, \ldots, T\}$. We define $\bar{M}_T \in \mathbb{R}^{|\aleph| \times |\aleph|}$ to be a matrix of "noiseless" empirical moments, given by

$$[\bar{M}_T]_{j_1,j_2} = \frac{1}{N} \sum_{i=1}^{N} \mu_{\alpha_{j_1}+\alpha_{j_2}}(\bar{x}_i^{(T)}, \tilde{w}^{(T)}).$$

The following lemma demonstrates that the linearized empirical moments, under suitable assumptions, converge to these noiseless empirical moments.

Lemma 3: Let $\tilde{w} : \mathbb{Z}_+ \to \mathbb{R}^{n_w}$ and $\tilde{x}_i : \mathbb{Z}_+ \to \mathbb{R}^{n_x}$, for $i \in \{1, \ldots, N\}$, satisfy assumptions (A1)-(A3). Then

$$\hat{M}_{\Xi_T} - \bar{M}_T \overset{P}{\to} 0$$

as $T \to \infty$.

The following statement justifies the use of this algorithm under the assumptions (A1)-(A3).

Theorem 1: Let $\tilde{w} : \mathbb{Z}_+ \to \mathbb{R}^{n_w}$ and $\tilde{x}_i : \mathbb{Z}_+ \to \mathbb{R}^{n_x}$, for $i \in \{1, \ldots, N\}$, satisfy assumptions (A1)-(A3). Fix $\delta, \Pi,$ and κ as in algorithm $\mathcal{A}$, and let (θ_T, r_T, R_T) be the $(\hat{\theta}, \hat{r}, \hat{R})$ found by applying algorithm $\mathcal{A}(\delta, \Pi, \kappa)$ to the data set Ξ_T. If κ is finite, then for every $\epsilon > 0$ and $\gamma \in (0,1)$ there exists a T such that, with probability $1 - \gamma$,

$$\epsilon + \text{tr}(R\bar{M}_T) \geq \text{tr}(R_T\bar{M}_T) \geq \frac{1}{N} \sum_{i=1}^{N} J_T^{SE}(a_{\theta_T}, \bar{x}_i(0), \bar{x}_i^{(T)}, \tilde{w}^{(T)}),$$

holds for all (θ, r, R) feasible for the optimization problem in step (iii) of $\mathcal{A}(\delta, \Pi, \kappa)$.
The proof of the above theorem is in the appendix.

The following result in characterizes the consistency of $\mathcal{A}$ in terms of recoverable functions.

Theorem 2: Fix $\Pi : \mathbb{R}^{4n_x+n_w} \to \mathbb{R}^{n_\Pi}$, $\kappa \in (0, \infty)$, $\delta > 0$. Let $\{\Xi_T\}_{T=1}^{\infty}$ be a sequence of data sets defined by signals $\tilde{w} : \mathbb{Z}_+ \to \mathbb{R}^{n_w}$ and $\tilde{x}_i : \mathbb{Z}_+ \to \mathbb{R}^{n_x}$, for $i \in \{1, \ldots, N\}$, satisfying assumptions (A1), (A2), and (A4), where the function $a_0 : \mathbb{R}^{n_x} \times \mathbb{R}^{n_w} \to \mathbb{R}^{n_x}$ described in assumption (A4) is in the recoverable set defined by $\Theta(N, \delta, \Pi)$. If θ_T is the parameter vector found by applying algorithm $\mathcal{A}(\delta, \Pi, \kappa)$ to Ξ_T, then

$$\sup_{(x,w) \in K} \{|e_{\theta_T}^{-1}(f_{\theta_T}(x,w)) - a_0(x,w)|\} \overset{P}{\to} 0$$

as $T \to \infty$ for all compact sets $K \subset \mathbb{R}^{n_x} \times \mathbb{R}^{n_w}$.
The proof of this theorem and supporting lemmas are given in the appendix.

V. EXAMPLE

This section examines the performance of algorithm $\mathcal{A}$ via a simple simulated example. Two alternative identification algorithms are introduced based on least square minimization; this is followed by a comparison of these algorithms and $\mathcal{A}$ on a simulated data set.

A. Least Squares Identification Approaches

Let $\Phi = \{\phi_i\}_{i=1}^{n_\theta}$ and $\Psi = \{\psi_i\}_{i=1}^{n_\theta}$ be fixed sequences of polynomial functions, $\phi_i : \mathbb{R}^{n_x} \times \mathbb{R}^{n_w} \to \mathbb{R}^{n_x}$ and $\psi_i : \mathbb{R}^{n_x} \to \mathbb{R}^{n_x}$. Let x and Δ denote real vector variables of dimension n_x. Given a positive constant δ and function $\Lambda : \mathbb{R}^{2n_x} \to \mathbb{R}^{n_\Lambda}$, let $\Omega(\delta, \Lambda)$ be the set of $\theta \in \mathbb{R}^{n_\theta}$ such that there exists some $\Sigma \in \mathbb{R}^{n_\Lambda \times n_\Lambda}$ such that:

$$\Sigma \geq 0,$$

$$2\Delta'(e_\theta(x+\Delta) - e_\theta(x)) - \delta|\Delta|^2 = \Lambda([x;\Delta])'\Sigma\Lambda([x;\Delta]).$$

From the proof of Lemma 1 we see that $\theta \in \Omega(\delta, \Lambda)$ guarantees e_θ is a bijection.

The following two algorithms produce a parameter vector $\hat{\theta}$ from a data set $\Xi = (\tilde{w}, \tilde{x}_1, \ldots, \tilde{x}_N) \in \mathcal{D}(n_x, n_w, N, T)$.

1) Least Squares Algorithm: Take $\hat{\theta}$ to be a $\theta \in \Omega(\delta, \Lambda)$ that minimizes

$$\frac{1}{NT} \sum_{i=1}^{N} \sum_{t=1}^{T} |e_\theta(\tilde{x}_i(t)) - f_\theta(\tilde{x}_i(t-1), \tilde{w}(t-1))|^2.$$

The following algorithm adapts the least squares objective to use the linearized empirical moments $\tilde{\mu}_\alpha(\cdot)$ defined by (5) with the aim of bias elimination.

2) Modified Least Squares Algorithm:
(i) Fix a $\delta > 0$ and $\kappa \in (0, \infty]$, and let $\aleph = \{\alpha_j\}_{j=1}^{|\aleph|} \subset \mathbb{Z}_+^{n_z}$ be the smallest set of vector degrees such that for each $\theta \in \Theta$ there exists coefficients $c_\alpha \in \mathbb{R}$ satisfying $e_\theta(\xi) - f_\theta(x,w) = \sum_{\alpha \in \aleph} c_\alpha[\xi; x; w]^\alpha$.
(ii) Define $\tilde{M} \in \mathbb{R}^{|\aleph| \times |\aleph|}$ by $[\tilde{M}]_{j_1,j_2} = \tilde{\mu}_{\alpha_{j_1}+\alpha_{j_2}}(\Xi)$, and take $\hat{M}$ to be the projection of $\frac{1}{2}(\tilde{M} + \tilde{M}')$ on the cone of positive semidefinite matrices.
(iii) Find the $\theta \in \Omega(\delta, \Lambda)$, satisfying $|\theta|^2 \leq \kappa$, that minimizes

$$\text{tr}\left(\theta\theta' \sum_{i=1}^{n_x} \Gamma_i' \hat{M} \Gamma_i\right),$$

where $\Gamma_i : \mathbb{R}^{|\aleph| \times n_\theta}$ is defined so that $[e_\theta(\xi) - f_\theta(x,w)]_i = \sum_{\alpha \in \aleph} z^{\alpha_j}[\Gamma_i\theta]_j$.

B. Simulated Example

We provide a simple simulated example to compare the performance algorithm to the least squares algorithms defined above. We examine a SISO nonlinear output error data set generated in the following fashion. The input is a scalar sequence $\tilde{u} : \{0, \ldots, 800\} \mapsto \mathbb{R}$, generated as i.i.d. random variables distributed uniformly on $[-1, 1]$. We examine the response of the system:

$$\bar{x}(t+1) + \frac{1}{5}\bar{x}(t+1)^5 = \frac{1}{3}\bar{x}(t)^3 + 5\tilde{u}(t)$$

starting from $\bar{x}(0) = 0$. For $i \in \{1, \ldots, 10\}$, observed data sequences, $\tilde{x}_i(t)$, are generated according to $\tilde{x}_i(t) = \bar{x}(t) + v_i(t)$ where the $v_i(t)$ are i.i.d. zero mean Gaussian random variables with variance 0.09 and independent across trials and from the input, leading to a signal-to-noise ratio of approximately 25 dB. We take $C = 1$, Φ to contain all monomials of degree less than or equal to five and Ψ to contain all monomials affine in u and of degree less than or equal to three in x.

The identified models are computed on a subset of the available data revealing only the samples with $t \in \{0, \ldots, T_h\}$ for each $T_h = 100 \cdot 2^h$ for $h \in \{0, 1, 2, 3\}$. The parameters δ and κ were taken to be 0.01 and ∞ respectively, and the set $\aleph$ from the definition of the modified least squares algorithm was also used for algorithm $\mathcal{A}$. The sum-of-squares programs were prepared using YALMIP [15]. The choices of Π, as in algorithm $\mathcal{A}$, and Λ, as in the modified least squares algorithm, that this software makes ensure that $\Theta(N, \delta, \Pi) \subset \Omega(\delta, \Lambda)$, i.e. the modified least squares algorithm searches over a larger set of models than algorithm $\mathcal{A}$.

To validate the models we generate an additional input sequence $\bar{u} : \{0, \ldots, 800\} \mapsto \mathbb{R}$, again i.i.d. uniformly distributed on [-1,1], that is independent from the training input, and noise. We compute the response $\bar{x}_{\text{val}}(t)$ of the true system to this input from zero initial conditions and compute a normalized simulation error:

$$\sum_{t=0}^{T_h} |\bar{x}_{\text{val}}(t) - x_h(t)|^2 \Big/ \sum_{t=0}^{T_h} |\bar{x}_{\text{val}}(t)|^2$$

where $x_h(t)$ is the response of the optimized model to the same input and starting from the same initial condition. These calculations
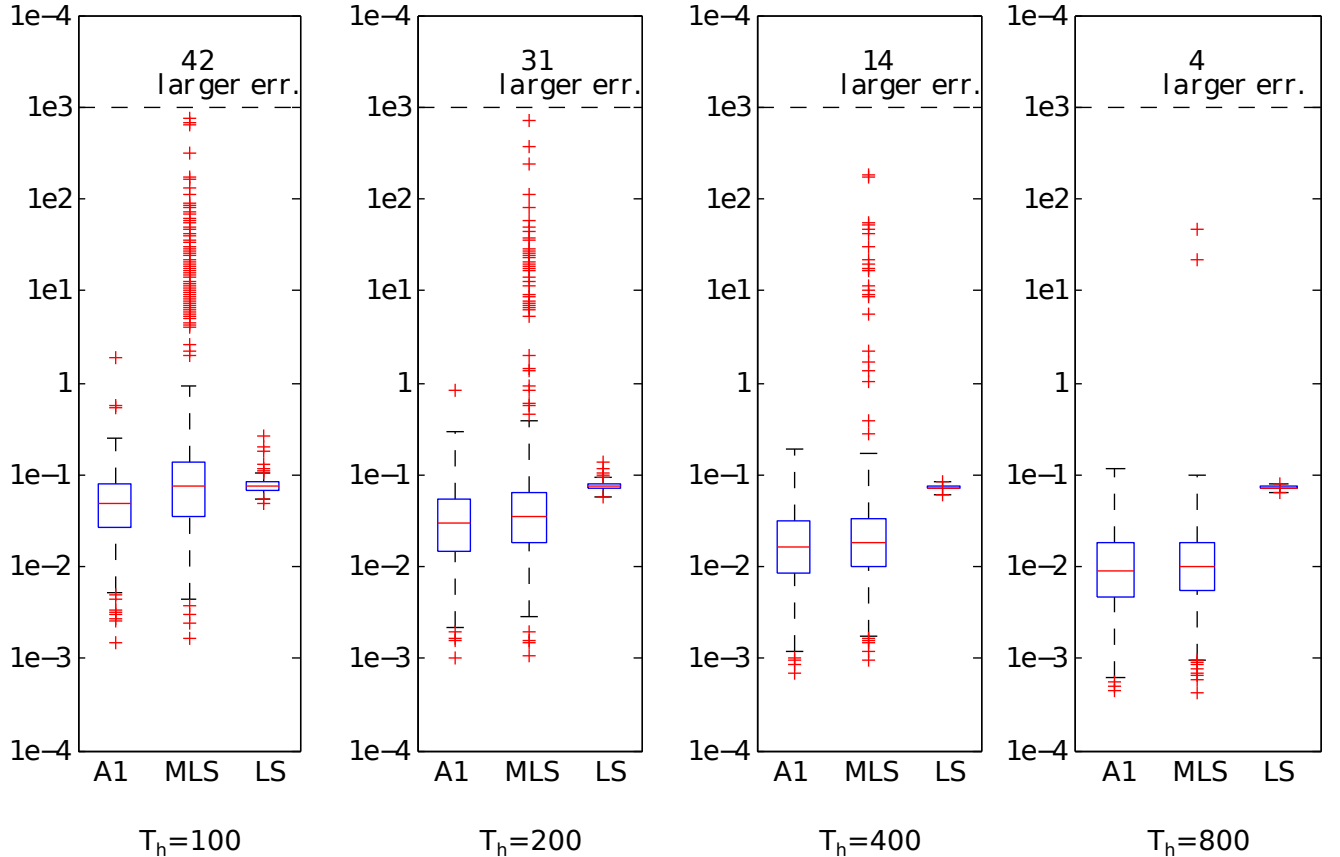
Fig. 2. Comparison of the algorithm $\mathcal{A}$ (A), the modified least squares algorithm (MLS) and the least squares algorithm (LS) for various training horizons. Plotted on a log scale is the distribution of the normalized simulation error for the validation input over 1,000 realizations. T_h indicates the number of training samples. The vertical scale of 1 indicates identical performance to the model that always outputs zero (i.e. 100 percent simulation error) and $1e-2$ indicates 1 percent simulation error. At the top of each plot is the number of MLS models having greater than 1000 percent simulation error.

were performed for 1,000 independent realizations of the problem. Figure 2 plots a comparison of the models generated by the three algorithms

As the amount of available data increases, the distribution of validation simulation errors tends toward zero for both algorithm $\mathcal{A}$ and the modified least squares approach. By contrast the result of the least squares approach without modification remains biased, though the variance of errors decreases. One sees that the modified least squares algorithm generates a large number of poorly performing models and generally under-performs algorithm $\mathcal{A}$ in terms of median as well. Note that the vertical scale in these plots is logarithmic: at $T_h = 200$ the majority of the models rendered by algorithm $\mathcal{A}$ have less that 3 percent validation simulation error.

VI. CONCLUSIONS

In this paper we have presented a convex optimization approach to nonlinear system identification from noise-corrupted data. The main contributions are a particular convex family of stable nonlinear models, and a procedure for identifying models from data based on empirical moments.

This builds upon previous work by the authors [30], [3] and [18] and offers two main advantages over the previously proposed methods: the complexity of the computation does not grow with

the length of the data, and the empirical moments can be "mixed" from different experiments to achieve a consistent estimator when the true system is in the model class. This is reminiscent of the instrumental variables approach to total least squares, although we suggest minimizing an alternative convex criterion. The advantages of the proposed method over least squares methods were illustrated via a simple simulated example.

APPENDIX
PROOFS

A. Modeling Scenario Satisfying (A1) and (A4)

The section provides an example of a class of dynamical systems and inputs that result in signals satisfying conditions (A1) and (A4).

Proposition 1: Let $\tilde{w}(t) \in \mathbb{R}^{n_w}$ be a vector valued i.i.d. stochastic process with $\tilde{w}(t)$ having an absolutely continuous distribution with a lower semicontinuous density whose support is bounded and contains 0. Let $a_0 : \mathbb{R}^{n_x} \times \mathbb{R}^{n_w}$ be a real analytic function such that (1) with $a = a_0$ is BIBO and incrementally exponentially stable and the linearization of a_0 about the unique equilibrium $(x, u) = (x_0, 0)$, with $x_0 = a_0(x_0, 0)$, is controllable. Then for any $\{\bar{x}_{i0}\}_{i=1}^N \subset \mathbb{R}^{n_x}$ independent of $\tilde{w}(t)$, the signal $\tilde{w}(t)$ and the signals $\bar{x}_i : \mathbb{Z}_+ \to$

$\mathbb{R}^{n_x}$, each defined by (1) with $w(t) \equiv \tilde{w}(t)$ and $\bar{x}_i(t) \equiv x(t)$ with $x(-1) = \bar{x}_{i0}$, satisfy (A1) and (A4).

Proof. That $\tilde{w}(t)$ satisfies (A1) is immediate. We use several results from [19]. The controllability assumptions given imply that the Markovian system generating the sequences $(\bar{x}_i(t), \tilde{w}(t))$ is weakly stochastically controllable (w.s.c) ([19], Corollary 2.2). The BIBO stability of (1) with $a = a_0$ and boundedness of $\tilde{w}(t)$ ensure $(\bar{x}_i(t), \tilde{w}(t))$ is bounded in probability. This boundedness in probability, the w.s.c. of the system, and the conditions on the distribution of $\tilde{w}(t)$ imply that the Markovian system generating $(\bar{x}_i(t), \tilde{w}(t))$ is a positive Harris recurrent Markov Chain. Thus there exists a unique stationary probability distribution π that is independent of initial condition. The w.s.c of the system and the assumptions on the distribution of $\tilde{w}(t)$ ensure that the support of π contains an open set. As the support of π is clearly bounded

$$\lim_{T \to \infty} \frac{1}{T} \sum_{t=0}^{T-1} h(\bar{x}_i(t), \tilde{w}(t)) = \int h \, d\pi, \qquad (20)$$

for all continuous $h : \mathbb{R}^{n_x} \times \mathbb{R}^{n_w} \to \mathbb{R}^{n_x}$ ([19] Proposition 2.3). From this we conclude that each pair $(\tilde{w}, \bar{x}_i)$ is persistently excited with respect to a_0, as in Definition 1. $\qquad \square$

B. A Simple Recoverable Nonlinear System

Let $C = 1$ and define $a : \mathbb{R} \times \mathbb{R} \mapsto \mathbb{R}$ by:

$$e(a(x, w)) = \frac{1}{2}x + b(w)$$

where $e(x) = \frac{3}{2}x + x^3$ and $b : \mathbb{R} \to \mathbb{R}$ is a polynomial. Let Φ and Ψ be sequences of real analytic functions as in the definition of e_θ and f_θ such that there exists a θ with $e = e_\theta$ and $f = f_\theta$. Let x, ξ, Δ, q, w be real numbers, $z = [\xi; x; w]$ and $v = [\xi; x; w; \Delta; q]$. With $r(z) = 2(e(\xi) - f(x, w))^2$ and $P = 1$, clearly (18) is satisfied and the left hand side of the equality (11) can be expressed as

$$(\sqrt{2}(e(\xi) - f(x, w)) - q/\sqrt{2})^2 + (\Delta - q/2)^2$$
$$+ \frac{1}{4}q^2 + (\sqrt{6}x\Delta + \sqrt{3/2}\Delta^2)^2 + \frac{1}{2}\Delta^4,$$

i.e. as the sum of squares of polynomials. Similarly the left hand size of (12) can be expressed as:

$$(\Delta + q/2)^2 + \frac{3}{4}q^2 + (\sqrt{6}x\Delta + \sqrt{3/2}\Delta^2)^2 + \frac{1}{2}\Delta^4.$$

These sum of squares decompositions show that there is a choice of $\Pi(v)$ such that matrices $\Sigma_1 \geq 0$ and $\Sigma_2 \geq 0$ satisfying (11) and (12) are guaranteed to exist. For this Π and $N \geq 6$, the polynomial r belongs to $\mathcal{P}_N$ so that a is in the recoverable set defined by $\Theta(N, 1, \Pi)$.

C. Proof of Lemma 3

Proof. By assumptions (A1) and (A3), the matrices $\tilde{M}_{\Xi_T}$ are uniformly bounded in T with probability one. Since each $\bar{M}_T \geq 0$ and projection onto the positive semidefinite cone is a continuous function, this boundedness implies it is sufficient to show that

$$\tilde{\mu}_\alpha(\Xi_T) - \mu_\alpha(\bar{x}_i^{(T)}, \tilde{w}^{(T)}) \xrightarrow{\text{P}} 0,$$

as $T \to \infty$, for each $i \in \{1, \dots, N\}$ and $\alpha \in \mathbb{Z}_+^{n_z}$ with $\|\alpha\|_1 \leq N$. Let $\bar{z}_i(t) = [\bar{x}_i(t); \bar{x}_i(t-1); \tilde{w}(t-1)]$ and $\zeta_i(t) = \tilde{z}_i(t) - \bar{z}_i(t)$. For all $\alpha \in \mathbb{Z}_+^{n_z}$ with $\|\alpha\|_1 \leq N$ we have:

$$\prod_{i=1}^{\|\alpha\|_1} [\tilde{z}_i(t)]_{\beta_\alpha(i)} = \sum_\kappa \nu_\kappa(t), \qquad (21)$$

where the sum runs over all $\kappa : \{1, \dots, \|\alpha\|_1\} \to \{0, 1\}$ and $\nu_\kappa(t)$ is defined by:

$$\nu_\kappa(t) := \prod_{i=1}^{\|\alpha\|_1} [\bar{z}_i(t)]_{\beta_\alpha(i)}^{(1-\kappa(i))} [\zeta_i(t)]_{\beta_\alpha(i)}^{\kappa(i)}. \qquad (22)$$

We see:

$$\frac{1}{T} \sum_{t=1}^{T} \nu_0(t) = \frac{1}{T} \sum_{t=1}^{T} \prod_{i=1}^{\|\alpha\|_1} [\bar{z}_i(t)]_{\beta_\alpha(i)}. \qquad (23)$$

Condition (8) of (A3) implies that $\lim_{t \to \infty} |\bar{z}_i(t) - \bar{z}_j(t)| = 0$ which, combined with the boundedness of $\bar{z}_i(t)$, allows us to conclude:

$$\lim_{T \to \infty} \left| \frac{1}{T} \sum_{t=1}^{T} \nu_0(t) - \mu_\alpha(\bar{x}_i^{(T)}, \tilde{w}^{(T)}) \right| = 0. \qquad (24)$$

Each $\nu_\kappa(t)$, for $\kappa \neq 0$, is obtained as a multi-linear function of the noise sequences $v_i(t)$, whose coefficients depend only on $\{\bar{z}_i(t)\}_{i=1}^N$. As the $v_i(t)$ are zero mean, bounded, and independent from one another and each $\bar{z}_i(t)$, we see $\nu_\kappa(t)$ is a zero mean, bounded random variable. In addition, condition (8) of (A3) and the boundedness of $\bar{x}_i(t)$ imply the correlation between $\nu_\kappa(t)$ and $\nu_\kappa(\tau)$ decays exponentially as $|t - \tau| \to \infty$, hence $\frac{1}{T} \sum_{t=1}^{T} \nu_\kappa(t) \xrightarrow{\text{P}} 0$. $\square$

D. Proof of Theorem 1

Proof. Optimality of (θ_T, r_T, R_T) implies that:

$$\text{tr}(R(\hat{M}_{\Xi_T} - \bar{M}_T)) + \text{tr}(R_T(\bar{M}_T - \hat{M}_{\Xi_T})) + \text{tr}(R\bar{M}_T) \geq \text{tr}(R_T\bar{M}_T) \qquad (25)$$

for all feasible (θ, r, R). From this, and the fact that $\kappa \geq \|R\|_F^2$ for all feasible R, we see:

$$2\sqrt{\kappa}\|\bar{M}_T - \hat{M}_T\|_F + \text{tr}(\bar{M}_T R) \geq \text{tr}(\bar{M}_T R_T). \qquad (26)$$

Fixing $\gamma \in (0, 1)$, Lemma 3 guarantees there exists a T sufficiently large such that $\|\bar{M}_T - \hat{M}_T\|_F \leq \frac{\epsilon}{2\sqrt{\kappa}}$ with probability $1 - \gamma$. Taking this T, the result follows by examining the definition of $r(\cdot)$, $\bar{M}_T$ and Lemma 1. $\qquad \square$

E. Proof of Theorem 2 and Supporting Lemmas

We begin by proving several supporting Lemmas.

Lemma 4: For any feasible point (θ, r, R) of the optimization of step (iii) in algorithm $\mathcal{A}$, the function $a_\theta = e_\theta^{-1} \circ f_\theta$ is real analytic. Furthermore, we have:

$$r(z) \geq \frac{1}{\delta}|e_\theta(\xi) - f_\theta(x, w)|^2 \geq \frac{\delta}{4}|\xi - a_\theta(x, w)|^2, \qquad (27)$$

where $z = [\xi; x; w]$, with $\xi, x \in \mathbb{R}^{n_x}$ and $w \in \mathbb{R}^{n_w}$.

Proof. Examining (12) with $x = x_1$ and $\Delta = x_1 - x_2$ we see that:

$$2\Delta'(e_\theta(x + \Delta) - e_\theta(x)) \geq |\Delta|_P^2 \qquad (28)$$

for all $x, \Delta \in \mathbb{R}^{n_x}$. Letting $E(x) = \frac{\partial e_\theta}{\partial x}$, the above implies $E(x) + E(x)' \geq P \geq \delta I$ for all $x \in \mathbb{R}^{n_x}$ as $e_\theta(\cdot)$ is continuously differentiable. Thus $\det(E(x)) \neq 0$ for all $x \in \mathbb{R}^{n_x}$, and as $e_\theta(\xi) - f_\theta(x, w)$ is real analytic the Inverse Function Theorem for holomorphic maps ([10] Chapter I, Theorem 7.6) ensures $e_\theta^{-1}(f_\theta(x, w))$ is real analytic.

The condition (11) for $x_1 = x_2 = x$ gives

$$r(z) \geq |e_\theta(\xi) - f_\theta(x, w)|_{P^{-1}}^2 \geq \frac{1}{\delta}|e_\theta(\xi) - f_\theta(x, w)|^2, \qquad (29)$$

as $P \geq \delta I$. Furthermore,

$$|e_\theta(\xi) - f_\theta(x, w)|^2 = |e_\theta(\xi) - e_\theta(a_\theta(x, w))|^2$$
$$= |\xi - a_\theta(x, w)|_{\bar{E}'\bar{E}}^2,$$

where $\bar{E} = \int_0^1 E(a_\theta(x,w) + \tau(\xi - a_\theta(x,w)))d\tau$. We see that

$$\bar{E}'\bar{E} \geq \frac{\delta}{2}(\bar{E} + \bar{E}') - \frac{\delta^2}{4}I \geq \frac{\delta^2}{4}I,$$

where the first inequality follows from $(\bar{E} - \frac{\delta}{2}I)'(\bar{E} - \frac{\delta}{2}I) \geq 0$ and the second inequality from $\bar{E}$ being a convex combination of point evaluations of $E(x)$. From this we can conclude (27) $\quad\square$

Now we present the proof of Theorem 2.

Proof. Fix any compact set $K \subset \mathbb{R}^{n_x} \times \mathbb{R}^{n_w}$. By assumption there exists a (θ_0, r_0, R_0), feasible for the optimization in step (iii) of algorithm $\mathcal{A}(\delta, \Pi, \kappa)$ such that $a_0 = a_{\theta_0}$ and $r_0(a_0(x,w), x, w) \equiv 0$. This implies $\text{tr}(R_0 \bar{M}_T) = 0$, so that, by Lemma 3, $\text{tr}(R_0 \hat{M}_{\Xi_T}) \xrightarrow{\text{P}} 0$.

As (θ_0, r_0, R_0) is in the feasible set of the optimization in step (iii) of $\mathcal{A}(\delta, \Pi, \kappa)$ and (θ_T, r_T, R_T) is optimal,

$$\text{tr}(R_0 \hat{M}_{\Xi_T}) \geq \text{tr}(R_T \hat{M}_{\Xi_T}), \tag{30}$$

so that $\text{tr}(R_T \hat{M}_{\Xi_T}) \xrightarrow{\text{P}} 0$ as well. Moreover, as $\|R_T\|_F^2 \leq \kappa$,

$$\sqrt{\kappa}\|\hat{M}_{\Xi_T} - \bar{M}_T\|_F + \text{tr}(R_T \hat{M}_{\Xi_T}) \geq \text{tr}(R_T \bar{M}_T), \tag{31}$$

by Cauchy-Scharwz. Lemma 3 now implies $\text{tr}(R_T \bar{M}_T) \xrightarrow{\text{P}} 0$.

As each r_T is a polynomial and the matrices R_T are uniformly bounded,

As each $(\bar{x}_i, \tilde{w})$ is persistently exciting with respect to a_0 (Definition 1), there exists of a positive measure π on the space $\mathbb{R}^{n_x} \times \mathbb{R}^{n_w}$, supported on an open set, such that for all $\epsilon > 0$ there exists a $T_0 \in \mathbb{Z}_+$ with

$$\epsilon\|R_T\|_F + \text{tr}(R_T \bar{M}_T) \geq \int r_T([a_0(x,w); x; w])d\pi(x,w),$$

for all $T \geq T_0$. As $\|R_T\|_F \leq \sqrt{\kappa}$ for all T, this sequence of integrals converges to zero. Lemma 4 now implies

$$\int r_T([a_0(x,w); x; w])d\pi(x,w)$$
$$\geq$$
$$\frac{1}{\delta}\int |e_{\theta_T}(a_0(x,w)) - f_{\theta_T}(x,w)|^2 d\pi(x,w).$$

From that same lemma we see that the map $(x,w) \mapsto e_{\theta_T}(a_0(x,w)) - f_{\theta_T}(x,w)$ is real analytic.

Let L be the subspace of $\mathbb{R}^{n_\theta}$ defined by

$$L = \{\theta \mid 0 = e_\theta(a_0(x,w)) - f_\theta(x,w), \quad \forall (x,w) \in \mathbb{R}^{n_x} \times \mathbb{R}^{n_w}\}.$$

Let V be the quotient space $V = \mathbb{R}^{n_\theta}/L$, i.e. the set of equivalence classes where $\theta \cong 0$ if $e_\theta(a_0(x,w)) - f_\theta(x,w) \equiv 0$. The following functions are norms on V:

$$\|\theta\| = \sqrt{\int |e_\theta(a_0(x,w)) - f_\theta(x,w)|^2 d\pi(x,w)},$$
$$|\theta|_{\infty,U} = \sup_{(x,w) \in U} \{|e_\theta(a_0(x,w)) - f_\theta(x,w)|\},$$

where U is any bounded open set. Both functions are clearly bounded, homogeneous, sub-additive and positive for all θ. $\|\theta\| = |\theta|_{\infty,U} = 0$ when $\theta \cong 0$ so that the functions are well-defined functions on V. As U is open and π is supported on an open set we see that $\theta \not\cong 0$ implies both $\|\theta\| > 0$ and $|\theta|_{\infty,U} > 0$, by the Identity Theorem for analytic maps [10]. Now, fix U to be a bounded open set with $U \supset K$. As all norms are equivalent for finite dimensional spaces, we know there exists a constant $c > 0$ such that $c\|\theta\| \geq |\theta|_{\infty,U}$.

That $\theta_0 \cong 0$ implies $\|\theta_T\| \xrightarrow{\text{P}} 0$, so that $|\theta_T|_{\infty,U} \xrightarrow{\text{P}} 0$. Lemma 4 then yields

$$\frac{2}{\sqrt{\delta}}|\theta_T|_{\infty,U} \geq \sup_{(x,w) \in K}\{|e_{\theta_T}^{-1}(f_{\theta_T}(x,w)) - a_0(x,w)|\},$$

which establishes the claim.

$\quad\square$

References

[1] D. Angeli. A Lyapunov approach to incremental stability properties. *Automatic Control, IEEE Transactions on*, 47(3):410 –421, mar 2002.

[2] B. Besselink, N. van de Wouw, and H. Nijmeijer. Model reduction for a class of convergent nonlinear systems. *Automatic Control, IEEE Transactions on*, 57(4):1071 –1076, april 2012.

[3] B.N. Bond, Z. Mahmood, Yan Li, R. Sredojevic, A. Megretski, V. Stojanovi, Y. Avniel, and L. Daniel. Compact modeling of nonlinear analog circuits using system identification via semidefinite programming and incremental stability certification. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, 29(8):1149 –1162, aug. 2010.

[4] M. Bonin, V. Seghezza, and L. Piroddi. NARX model selection based on simulation error minimisation and LASSO. *Control Theory Applications, IET*, 4(7):1157 –1168, july 2010.

[5] S. Boyd and L. O. Chua. Fading memory and the problem of approximating nonlinear operators with volterra series. *Circuits and Systems, IEEE Transactions on*, 32(11):1150–1171, nov. 1985.

[6] V. Cerone, D. Piga, and D. Regruto. Set-membership error-in-variables identification through convex relaxation techniques. *Automatic Control, IEEE Transactions on*, 57(2):517 –522, feb. 2012.

[7] Vito Cerone, Dario Piga, and Diego Regruto. Enforcing stability constraints in set-membership identification of linear dynamic systems. *Automatica*, 47(11):2488 – 2494, 2011.

[8] M. Farina and L. Piroddi. Simulation error minimization identification based on multi-stage prediction. *International Journal of Adaptive Control and Signal Processing*, 25(5):389–406, 2011.

[9] C. Feng, C.M. Lagoa, and M. Sznaier. Identifying stable fixed order systems from time and frequency response data. In *American Control Conference (ACC), 2011*, pages 1253 –1259, 29 2011-july 1 2011.

[10] Klaus Fritzsche and Hans Grauert. *From Holomorphic Functions to Complex Manifolds*. Springer, Englewood Cliffs, 2002.

[11] V. Fromion and G. Scorletti. The behaviour of incrementally stable discrete time systems. In *American Control Conference, 1999. Proceedings of the 1999*, volume 6, pages 4563 –4567 vol.6, 1999.

[12] J.B. Hoagg, S.L. Lacy, R.S. Erwin, and D.S. Bernstein. First-order-hold sampling of positive real systems and subspace identification of positive real models. In *American Control Conference, 2004. Proceedings of the 2004*, volume 1, pages 861 –866 vol.1, 30 2004-july 2 2004.

[13] S.L. Lacy and D.S. Bernstein. Subspace identification with guaranteed stability using constrained optimization. In *American Control Conference, 2002. Proceedings of the 2002*, volume 4, pages 3307 – 3312 vol.4, 2002.

[14] Lennart Ljung. *System Identification: Theory for the User*. Prentice Hall, Englewood Cliffs, New Jersey, USA, 3 edition, 1999.

[15] J. Löfberg. Pre- and post-processing sum-of-squares programs in practice. *IEEE Transactions on Automatic Control*, 54(5):1007–1011, 2009.

[16] Winfried Lohmiller and Jean-Jacques E. Slotine. On contraction analysis for non-linear systems. *Automatica*, 34(6):683 – 696, 1998.

[17] I.R. Manchester, M.M. Tobenkin, and A.. Megretski. Stable nonlinear system identification: Convexity, model class, and consistency. In *System Identification, 16th IFAC Symposium on*, volume 16, jul. 2012.

[18] A. Megretski. Convex optimization in robust identification of nonlinear feedback. In *Decision and Control, 2008. CDC 2008. 47th IEEE Conference on*, pages 1370 –1374, dec. 2008.

[19] S. Meyn and P. Caines. Asymptotic behavior of stochastic systems possessing markovian realizations. *SIAM Journal on Control and Optimization*, 29(3):535–561, 1991.

[20] Johan Paduart, Lieve Lauwers, Jan Swevers, Kris Smolders, Johan Schoukens, and Rik Pintelon. Identification of nonlinear systems using polynomial nonlinear state space models. *Automatica*, 46(4):647 – 656, 2010.

[21] Pablo A. Parrilo. Semidefinite programming relaxations for semialgebraic problems. *Mathematical Programming*, 96:293–320, 2003. 10.1007/s10107-003-0387-5.

[22] J.R. Partington and P.M. Makila. Modeling of linear fading memory systems. *Automatic Control, IEEE Transactions on*, 41(6):899 –903, jun 1996.

[23] G. Pillonetto, Minh Ha Quang, and A. Chiuso. A new kernel-based approach for nonlinearsystem identification. *Automatic Control, IEEE Transactions on*, 56(12):2825 –2840, dec. 2011.

[24] Alexander S. Poznyak. *Advanced Mathematical Tools for Automatic Control Engineers Volume 1: Deterministic Techniques*. Elsevier, Amsterdam, The Netherlands, 2008.

[25] P.A. Regalia. An unbiased equation error identifier and reduced-order approximations. *Signal Processing, IEEE Transactions on*, 42(6):1397 –1412, jun 1994.

[26] P.A. Regalia and P. Stoica. Stability of multivariable least-squares models. *Signal Processing Letters, IEEE*, 2(10):195 –196, oct. 1995.

[27] T. Söderström and P. Stoica. Some properties of the output error method. *Automatica*, 18(1):93 – 99, 1982.

[28] Torsten Söderström and Petre Stoica. Instrumental variable methods for system identification. *Circuits, Systems, and Signal Processing*, 21:1–9, 2002. 10.1007/BF01211647.

[29] Eduardo Sontag. Contractive systems with inputs. In Jan Willems, Shinji Hara, Yoshito Ohta, and Hisaya Fujioka, editors, *Perspectives in Mathematical System Theory, Control, and Signal Processing*, volume 398 of *Lecture Notes in Control and Information Sciences*, pages 217–228. Springer Berlin / Heidelberg, 2010.

[30] M.M. Tobenkin, I.R. Manchester, J. Wang, A. Megretski, and R. Tedrake. Convex optimization in identification of stable non-linear state space models. In *Decision and Control (CDC), 2010 49th IEEE Conference on*, pages 7232 –7237, dec. 2010.

[31] J.K. Tugnait and C. Tontiruttananon. Identification of linear systems via spectral analysis given time-domain data: consistency, reduced-order approximation, and performance analysis. *Automatic Control, IEEE Transactions on*, 43(10):1354 –1373, oct 1998.

[32] T. Van Gestel, J.A.K. Suykens, P. Van Dooren, and B. De Moor. Identification of stable models in subspace identification by using regularization. *Automatic Control, IEEE Transactions on*, 46(9):1416 –1420, sep 2001.