Non-asymptotic state-space identification of closed-loop stochastic linear systems using instrumental variables

Szabolcs Szentpéteri [a,*], Balázs Csanád Csáji [a,b]

[a] *Institute for Computer Science and Control (SZTAKI), Eötvös Loránd Research Network (ELKH), Budapest, Hungary*
[b] *Department of Probability Theory and Statistics, Institute of Mathematics, Eötvös Loránd University (ELTE), Budapest, Hungary*

ABSTRACT

The paper suggests a generalization of the Sign-Perturbed Sums (SPS) finite sample system identification method for the identification of closed-loop observable stochastic linear systems in state-space form. The solution builds on the theory of matrix-variate regression and instrumental variable methods to construct distribution-free confidence regions for the state-space matrices. Both direct and indirect identification are studied, and the exactness as well as the strong consistency of the construction are proved. Furthermore, a new, computationally efficient ellipsoidal outer-approximation algorithm for the confidence regions is proposed. The new construction results in a semidefinite optimization problem which has an order-of-magnitude smaller number of constraints, as if one applied the ellipsoidal outer-approximation after vectorization. The effectiveness of the approach is also demonstrated empirically via a series of numerical experiments.

1. Introduction

Estimating a *mathematical model* from observations of a *dynamical system* is a fundamental problem across several fields, from system identification to signal processing and machine learning. Standard estimation techniques, such as the prediction error approach or the generalized method of moments, typically provide *point estimates* associated with *only asymptotically guaranteed* confidence sets [1].

In practical applications there is often *limited statistical knowledge* about the noises and uncertainties affecting the system and a *limited number of measurements*. Furthermore, in many situations, data can only be gathered under *feedback control*. This is especially the case, for example, for economy and biology related applications. If our problem involves strong safety or stability requirements, having guaranteed confidence regions are strongly desirable. Typical examples, where guaranteed region estimates can be essential are robust and adaptive control [2–4].

One of the first approaches to non-asymptotic system identification were [5,6]. In the past few years, due to a renewed interest in system identification methods with finite sample guarantees, non-asymptotic high probability bounds on the estimation error were investigated for linear state-space models, see for example [7–12]. The above results often use *strong statistical*

assumptions, such as joint *Gaussianity*, and [7,8,11] considered uncontrolled, while [9,12] studied open-loop LTI systems. Non-asymptotic identification of a *closed-loop* LTI system with the REDAR algorithm was investigated in [13], still assuming *Gaussian* white noise type uncertainties.

Typical examples of *distribution-free* system identification algorithms with strong non-asymptotic guarantees are the Leave-out Sign-dominant Correlation Regions (LSCR) [14–16] and the Sign-Perturbed Sums (SPS) [17–19] methods. SPS can construct *exact* confidence regions for (open-loop) *general linear systems* under mild statistical assumptions [20]. In standard SPS, the confidence set is given by its *indicator* function, which can be evaluated at any parameter. In [21] a guaranteed characterization of SPS was developed using interval analysis, while in [17] an *ellipsoidal outer approximation* was given for FIR (finite impulse response) systems. The *strong consistency* and the asymptotic shape of SPS regions were studied in [22].

Several generalizations of SPS were suggested, such as Data Perturbation (DP) methods [23] which can use other perturbations, not only sign changes; and UD-SPS that is able to detect undermodelling, in case of FIR systems [24].

The *closed-loop* applicability of SPS for general linear systems was studied in [19], where the *exact coverage* of SPS confidence regions is shown, but they are only given by their indicators. An *instrumental variables* [25] based extension of SPS for linear regression was presented in [18], making the SPS *outer ellipsoids* applicable to ARX (autoregressive with exogenous inputs) systems even under *feedback*.

All of the SPS related papers above studied *scalar* models, where the system is SISO and it is given either with transfer functions or in a linear regression form. In many control problems the system is MIMO and it is given in a *state-space* form which even includes feedback control, for which the SPS variants above are not directly applicable.

One could apply the scalar approach of [18] to a state-space model via vectorization, i.e., after reformulating it as a (potentially huge) linear regression problem [26]. This, however, needs additional assumptions on the noise vectors (i.e., that their distributions are symmetric w.r.t. each coordinate axis and even their components are independent). Furthermore, computing the resulting ellipsoidal outer approximation leads to large semidefinite programming problems.

Instead of doing so, we extend SPS to state-space models by applying *matrix-variate regression* which leads to a more compact approach with relaxed assumptions on the noises (the components of the noise vectors can be dependent and only their joint distributions should be symmetric about zero, full axial symmetry is not needed). An alternative *semidefinite program* (SDP) is suggested, as well, to compute outer ellipsoids, which has an *order-of-magnitude* smaller number of constraints than the one based on [18].

The main *contributions* of the paper are as follows:

1. A generalization of SPS is proposed to build distribution-free confidence sets for *matrix-variate regression* problems using *instrumental variables* (MIV-SPS). It is demonstrated on *closed-loop state-space* models.
2. The *exact coverage probabilities* (for any finite sample size) and the *strong consistency* of the MIV-SPS confidence regions are proved under mild assumptions.
3. A new, efficient *ellipsoidal outer approximation* algorithm is introduced for MIV-SPS, based on SDP.
4. The effectiveness of MIV-SPS is also validated experimentally and it is compared to (scalar) IV-SPS and the confidence ellipsoids of the asymptotic theory.

The paper is organized as follows. In Section 2 the problem setting and its linear and matrix-variate regression formulations as well as our main assumptions are introduced. Section 3 gives a summary of the instrumental variable (IV) method and its asymptotic theory. In Section 4 the MIV-SPS algorithm is presented, while Section 5 states the theoretical guarantees of the method. Section 6 introduces the ellipsoidal outer-approximation algorithm. The simulation experiments and comparisons are presented in Section 7. Finally, Section 8 concludes the paper.

2. Problem setting

This section introduces the closed-loop stochastic linear state-space model, presents the linear and matrix-variate regression reformulations and the main assumptions.

2.1. Stochastic linear state-space model

Consider the observable linear state-space (LSS) model

$$x_{k+1} = Ax_k + Bu_k + w_k, \qquad (1)$$

for $k = 0, 1, \ldots, n-1$, where x_k is a d_x-dimensional, u_k is a d_u-dimensional, and w_k is a d_x-dimensional real random vector; A, B are *unknown* real matrices to be estimated. The system can operate in *closed-loop* with feedback rule

$$u_k = Fx_k + Gr_k, \qquad (2)$$

where r_k is an d_r-dimensional real random vector, which can represent a reference signal, a setpoint or noise affecting the controller; while F and G are real matrices.

In this work, we will concentrate on *region estimation*, therefore, we aim at constructing a *confidence region* which contains matrices A and B with a *user-chosen* probability.

2.2. Linear regression formulation

The LSS dynamics (1) can be written as

$$
\begin{bmatrix} x_{k+1,1} \\ x_{k+1,2} \\ \vdots \\ x_{k+1,d_x} \end{bmatrix} =
\begin{bmatrix} a_1^{\mathsf{T}} \\ a_2^{\mathsf{T}} \\ \vdots \\ a_{d_x}^{\mathsf{T}} \end{bmatrix}
\begin{bmatrix} x_{k,1} \\ x_{k,2} \\ \vdots \\ x_{k,d_x} \end{bmatrix} +
\begin{bmatrix} b_1^{\mathsf{T}} \\ b_2^{\mathsf{T}} \\ \vdots \\ b_{d_u}^{\mathsf{T}} \end{bmatrix}
\begin{bmatrix} u_{k,1} \\ u_{k,2} \\ \vdots \\ u_{k,d_u} \end{bmatrix} +
\begin{bmatrix} w_{k,1} \\ w_{k,2} \\ \vdots \\ w_{k,d_x} \end{bmatrix}, \qquad (3)
$$

where $\{a_i^{\mathsf{T}}\}$ and $\{b_i^{\mathsf{T}}\}$ are the rows of A and B, respectively. Then, the system dynamics from time 1 until time n can be written in a *linear regression* form as

$$y = \varXi \, \theta^* + w, \qquad (4)$$

where $\theta^* \doteq (a_1^{\mathsf{T}}, \ldots, a_{d_x}^{\mathsf{T}}, b_1^{\mathsf{T}}, \ldots, b_{d_u}^{\mathsf{T}})^{\mathsf{T}}$, and

$$
y \doteq \begin{bmatrix} x_{1,1} \\ x_{1,2} \\ \vdots \\ x_{1,d_x} \\ x_{2,1} \\ \vdots \\ x_{n,d_x} \end{bmatrix}, \quad
\varXi \doteq \begin{bmatrix} \xi_{0,1}^{\mathsf{T}} \\ \xi_{0,2}^{\mathsf{T}} \\ \vdots \\ \xi_{0,d_x}^{\mathsf{T}} \\ \xi_{1,1}^{\mathsf{T}} \\ \vdots \\ \xi_{n-1,d_x}^{\mathsf{T}} \end{bmatrix}, \quad
w \doteq \begin{bmatrix} w_{0,1} \\ w_{0,2} \\ \vdots \\ w_{0,d_x} \\ w_{1,1} \\ \vdots \\ w_{n-1,d_x} \end{bmatrix}, \qquad (5)
$$

and the regressors $\{\xi_{k,i}\}$ are defined as follows

$$
\xi_{k,i}^{\mathsf{T}} \doteq (\overbrace{0_x^{\mathsf{T}}, \ldots, 0_x^{\mathsf{T}}, \underbrace{x_k^{\mathsf{T}}}_{i\text{th}}, 0_x^{\mathsf{T}}}^{\dim:\, d_x^2}, \overbrace{0_u^{\mathsf{T}}, \ldots, 0_u^{\mathsf{T}}, \underbrace{u_k^{\mathsf{T}}}_{i\text{th}}, 0_u^{\mathsf{T}}}^{\dim:\, d_x d_u}), \qquad (6)
$$

where 0_x and 0_u are d_x and d_u dimensional zero vectors.

The advantage of the linear regression formulation of (4) is that the previous results for scalar systems, e.g., the IV-SPS method [19], can be directly applied. Note that in this case each marginal distribution of the noise vector should satisfy the noise assumptions of the SPS, in order to allow the application of scalar SPS approaches.

2.3. Matrix-variate regression formulation

In this paper, we argue that a more natural way to handle (even closed-loop) linear state-space models is to apply a matrix-variate regression formulation. As we will see, this allows relaxed assumption on the noises and leads to computationally more efficient constructions.

Without vectorization, state-space model (1) can be reformulated as a *matrix-variate regression* problem, i.e.,

$$Y = \varPhi \varTheta^* + W, \qquad (7)$$

where, unlike in (4), the output, Y, the true parameter, $\varTheta^*$, and the noise, W, are all matrices, as well. That is

$$
Y \doteq \begin{bmatrix} x_1^{\mathsf{T}} \\ x_2^{\mathsf{T}} \\ \vdots \\ x_n^{\mathsf{T}} \end{bmatrix}, \quad
\varPhi \doteq \begin{bmatrix} \varphi_0^{\mathsf{T}} \\ \varphi_1^{\mathsf{T}} \\ \vdots \\ \varphi_{n-1}^{\mathsf{T}} \end{bmatrix}, \quad
\varTheta^* \doteq \begin{bmatrix} A^{\mathsf{T}} \\ B^{\mathsf{T}} \end{bmatrix}, \quad
W \doteq \begin{bmatrix} w_0^{\mathsf{T}} \\ w_1^{\mathsf{T}} \\ \vdots \\ w_{n-1}^{\mathsf{T}} \end{bmatrix}, \qquad (8)
$$

where the regressors φ_k are defined as $\varphi_k \doteq (x_k^\mathrm{T}, u_k^\mathrm{T})^\mathrm{T}$ which are $d \doteq d_x + d_u$ dimensional vectors.

We will also use the notations Y_n, Φ_n, and W_n, in cases when the dependence on the sample size n is crucial.

2.4. Indirect identification

In Sections 2.2 and 2.3 the regression problems were formulated using the *direct* system identification approach, where the values $\{x_k\}$ and $\{u_k\}$ are used during estimation. The regression problems can also be formulated using the *indirect* approach [27], where the control matrices F and G are *known*, which is often the case in an adaptive control setting. By expanding the state transition equation we can reformulate (1) and (2) as follows

$$x_{k+1} = Ax_k + B(Fx_k + Gr_k) + w_k = (A + BF)x_k + BGr_k + w_k$$
$$= Cx_k + Dr_k + w_k. \qquad (9)$$

In the indirect identification scheme, using the values $\{x_k\}$ and $\{r_k\}$, the C and D matrices are estimated, from which A and B can be computed (under mild conditions on the control matrices F and G). In this case (7) should of course be modified accordingly, for example, the rows of the regressor matrix Φ_id are given by $\bar{\varphi}_k \doteq (x_k^\mathrm{T}, r_k^\mathrm{T})^\mathrm{T}$, and the true parameter Θ_id^* should contain matrices C and D.

The linear regression formulation can be obtained similarly to the approach of Section 2.2, hence omitted.

2.5. Core assumptions

Henceforth, we will directly study the matrix-variate regression problem (7), irrespectively whether it came from the direct or the indirect version of the original identification problem. Our main assumptions will be as follows:

A1. The row vectors $\{w_k\}$ of the noise matrix W are independent, and they are distributed symmetrically about zero *(but, they can have different distributions)*, i.e., for all k, random vectors w_k and $-w_k$ have the same distribution.

This assumption is very mild, even milder than the one we get from standard SPS [17–20,22], if we try to apply it to a vectorized version of (7). For example, such a vectorized approach would require that for all k and i, we have

$$(w_{k,1}, \ldots, w_{k,i-1}, w_{k,i}, w_{k,i+1}, \ldots, w_{k,d_x}) \stackrel{d}{=}$$
$$(w_{k,1}, \ldots, w_{k,i-1}, -w_{k,i}, w_{k,i+1}, \ldots, w_{k,d_x}), \qquad (10)$$

where "$\stackrel{d}{=}$" denotes equality in distribution. This is strictly stronger than $w_k \stackrel{d}{=} -w_k$, as the latter is implied by (10), but not the other way around. Moreover, A1 also allows the components of the noise vector w_k to be *correlated*.

A2. We are given a random matrix $\Psi \doteq (\psi_0, \ldots, \psi_{n-1})^\mathrm{T} \in \mathbb{R}^{n \times d}$, for which matrices Ψ and W are independent.

The rows of matrix Ψ are called the *instrumental variables*. In a typical situation, they are filtered past inputs using an auxiliary model [1]. The motivation for using instrumental variables will be discussed in Section 4.1.

A3. Matrix $\Psi^\mathrm{T}\Phi$ is full rank almost surely.

Observe that from assumption A3 it follows that matrix $\Psi^\mathrm{T}\Psi$ is also full rank (invertible) almost surely. We also introduce the notations $V_n \doteq {}^1\!/_n\, \Psi^\mathrm{T}\Phi$, and $P_n \doteq {}^1\!/_n\, \Psi^\mathrm{T}\Psi$.

Note, as well, that instrumental variable based identification methods typically also assume that the instrumental variables,

$\{\psi_k\}$, and the states, $\{x_k\}$, are correlated. In the strict sense, this assumption is not needed to construct exact confidence regions for the true parameter matrix Θ^*, but an asymptotic correlation assumption A5 is crucial to prove consistency.

3. Instrumental variable methods

Now, we briefly overview identification with *instrumental variable* (IV) methods and its asymptotic theory [1].

3.1. Instrumental variable estimate (IVE)

By defining the prediction of a particular Θ matrix by $\hat{Y} \doteq \Phi\Theta$, we compute the *prediction errors* (residuals) as $\mathcal{E}(\Theta) \doteq Y - \hat{Y} = Y - \Phi\Theta$. IVE is obtained by solving

$$\Psi^\mathrm{T}\mathcal{E}(\Theta) = \Psi^\mathrm{T}(Y - \Phi\Theta) = 0, \qquad (11)$$

then, assuming A3, the IVE can be calculated as

$$\widehat{\Theta}^\mathrm{IV} = (\Psi^\mathrm{T}\Phi)^{-1}\Psi^\mathrm{T}Y, \qquad (12)$$

which is usually a *biased* estimator of Θ^*, under A1–A3, unless we assume that Φ and W are independent.

3.2. Limiting distribution of IVE

Now, assume that Y_n is $\mathbb{R}^n$-valued, i.e., $\{x_k\}$ is a scalar process. Then, the "true" parameter is a constant vector, $\theta^* \in \mathbb{R}^{d_\theta}$, and the IVE is a random vector given by (12).

If $\{w_k\}$ are i.i.d., zero mean, and each has variance $\sigma^2 \in (0, \infty)$, then the IVE is *asymptotically Gaussian*:

$$\sqrt{n}(\hat{\theta}_n^\mathrm{IV} - \theta^*) \stackrel{d}{\longrightarrow} \mathcal{N}(0, \sigma^2 R^{-1}), \qquad (13)$$

as $n \to \infty$, with $R \doteq \lim_{n\to\infty} R_n$, where $\{R_n\}$ are

$$R_n \doteq \left[\frac{1}{n}\Psi_n^\mathrm{T}\Phi_n\right]^{-1}\left[\frac{1}{n}\Psi_n^\mathrm{T}\Psi_n\right]\left[\frac{1}{n}\Phi_n^\mathrm{T}\Psi_n\right]^{-1}, \qquad (14)$$

assuming that the limit exists and is positive definite [1], where we (explicitly) emphasized the dependences on n.

In case $\{x_k\}$ is vector-valued, we can reformulated it in a linear regression form (4). Then, the (vectorized) IVE is also asymptotically Gaussian, with appropriate modification of the instrumental variables, similarly to (6).

3.3. Asymptotic confidence regions of IVE

Assuming $\widehat{\theta}_n^\mathrm{IV}$ is (vector-valued and) asymptotically normal, that is if (13) holds, then we also have [1]

$$\frac{n}{\sigma^2}(\hat{\theta}_n^\mathrm{IV} - \theta^*)^\mathrm{T}R(\hat{\theta}_n^\mathrm{IV} - \theta^*) \stackrel{d}{\longrightarrow} \chi^2(d_\theta), \qquad (15)$$

as $n \to \infty$, where $\chi^2(d_\theta)$ is the chi-square distribution with $d_\theta = \dim(\theta^*)$ degrees of freedom. Matrix R and variance σ^2 are unknown, but they can be estimated. Matrix R_n is an empirical estimate of R and an estimate of σ^2 is

$$\hat{\sigma}_n^2 \doteq \frac{1}{n - d_\theta} \|\mathcal{E}(\hat{\theta}_n^\mathrm{IV})\|_2^2. \qquad (16)$$

Then, a p-level *confidence ellipsoid* can be defined as

$$\tilde{\Gamma}_{n,p} \doteq \left\{\theta \in \mathbb{R}^{d_\theta} : (\theta - \hat{\theta}_n^\mathrm{IV})^\mathrm{T}R_n(\theta - \hat{\theta}_n^\mathrm{IV}) \leq \frac{\mu\hat{\sigma}_n^2}{n}\right\}, \qquad (17)$$

where $p = F_{\chi^2(d_\theta)}(\mu)$ is the target confidence probability and $F_{\chi^2(d_\theta)}$ is the cumulative distribution function of the χ^2

distribution. Then, we have $\mathbb{P}(\theta^* \in \tilde{\Upsilon}_{n,p}) \approx p$. Note that these regions only have *asymptotic* guarantees.

4. Matrix-variate generalization of IV-SPS

In this section, we introduce a *matrix-variate* generalization of the SPS method using instrumental variables. Our construction can be efficiently applied to solve closed-loop state-space identification problems, formulated as (7).

The scalar variant of IV-SPS for linear regression problems was originally suggested in [18], which can be used for state-space models after suitable vectorization [26].

One of the advantages of MIV-SPS w.r.t. standard SPS methods [17,18] is that the matrix-variate regression formulation allows perturbing the components of the noise vectors simultaneously, which lead to relaxed assumptions on the noises. Moreover, the matrix-variate approach also allows a computationally more efficient ellipsoidal outer approximation construction than the original approach.

4.1. Intuitive idea of the SPS construction

Using the matrix-variate regression formulation (7), the *normal equation* (11) can be reformulated as

$$\Psi^{\mathsf{T}}\Phi(\Theta^* - \Theta) + \Psi^{\mathsf{T}}W = 0. \tag{18}$$

Following the SPS principle [17,18], we introduce $m - 1$ sign-perturbed sums, to construct the region, as follows

$$H_i(\Theta) \doteq \Psi^{\mathsf{T}}\Lambda_i(Y - \Phi\Theta) = \Psi^{\mathsf{T}}\Lambda_i\Phi(\Theta^* - \Theta) + \Psi^{\mathsf{T}}\Lambda_i W, \tag{19}$$

for $i = 1, \ldots, m - 1$, where Λ_i is a diagonal matrix, $\Lambda_i \doteq \operatorname{diag}(\alpha_{i,1}, \ldots, \alpha_{i,n})$, and $\{\alpha_{i,k}\}$ are i.i.d. Rademacher variables, i.e., $\mathbb{P}(\alpha_{i,k} = 1) = \mathbb{P}(\alpha_{i,k} = -1) = 1/2$. A reference sum is also introduced, without sign-perturbations,

$$H_0(\Theta) \doteq \Psi^{\mathsf{T}}(Y - \Phi\Theta) = \Psi^{\mathsf{T}}\Phi(\Theta^* - \Theta) + \Psi^{\mathsf{T}}W. \tag{20}$$

For the true parameter matrix, $\Theta = \Theta^*$, we have

$$H_0(\Theta^*) = \Psi^{\mathsf{T}}W, \quad \text{and} \quad H_i(\Theta^*) = \Psi^{\mathsf{T}}\Lambda_i W. \tag{21}$$

$H_0(\Theta)$ and $H_i(\Theta)$ can be compared using the Frobenius norm $\|\cdot\|_{\mathrm{F}}$. According to assumption A1, the noise vector sequence $\{w_k\}$ is independent and symmetric about zero, thus $H_0(\Theta^*)$ and $H_i(\Theta^*)$ have the same distribution, hence the probability that a particular $\|H_\ell(\Theta^*)\|_{\mathrm{F}}^2$ is ranked at a given position according to a strict total order of the values $\{\|H_j(\Theta^*)\|_{\mathrm{F}}^2\}_{j=0}^{m-1}$ is the same for all ℓ. Note that this property is not trivial, since $\{H_j(\Theta^*)\}_{j=0}^{m-1}$ are dependent.

When $\|\Theta^* - \Theta\|_{\mathrm{F}}$ is "sufficiently" large, the inequality $\forall i \neq 0$: $\|H_0(\Theta)\|_{\mathrm{F}}^2 > \|H_i(\Theta)\|_{\mathrm{F}}^2$, will hold, thus $\|H_0(\Theta)\|_{\mathrm{F}}^2$ will eventually be the largest of the m functions. The core idea behind the SPS method is to construct the confidence region based on the rankings of $\{\|H_j(\Theta)\|_{\mathrm{F}}^2\}_{j=0}^{m-1}$ and exclude those Θ parameter matrices for which the reference, $\|H_0(\Theta)\|_{\mathrm{F}}^2$, is among the q largest. As we will show, the so constructed confidence set has *exactly* probability $1 - q/m$ of containing the true parameter matrix. In the final formal description, the functions $\{\|H_j(\Theta)\|_{\mathrm{F}}^2\}_{j=0}^{m-1}$ will also be weighted with a suitable "shaping" matrix.

Instrumental variables (IVs) are introduced, in order to be able to handle closed-loop systems. In these cases, if Ψ was simply replaced by Φ, $\{\|H_j(\Theta^*)\|_{\mathrm{F}}^2\}_{j=0}^{m-1}$ would not have the same distribution, because the regressor matrix is *not* independent of the noise terms. The IVs counteract these dependencies, thus they ensure the validity of the confidence set construction in closed-loop setups.

Algorithm 1 MIV-SPS: Initialization (p)

1: Given the (rational) confidence probability $p \in (0, 1)$, set integers $m > q > 0$ such that $p = 1 - q/m$.
2: Calculate the outer product

$$P_n \doteq \frac{1}{n}\Psi^{\mathsf{T}}\Psi,$$

and find the principal square root $P_n^{1/2}$, such that

$$P_n^{1/2}P_n^{1/2} = P_n.$$

3: Generate $n(m - 1)$ i.i.d random signs $\{\alpha_{i,k}\}$ with

$$\mathbb{P}(\alpha_{i,k} = 1) = \mathbb{P}(\alpha_{i,k} = -1) = \frac{1}{2},$$

for $i \in \{1, \ldots, m - 1\}$, $k \in \{1, \ldots, n\}$ and construct the following matrices containing these random signs

$$\Lambda_i \doteq \begin{bmatrix} \alpha_{i,1} & & \\ & \ddots & \\ & & \alpha_{i,n} \end{bmatrix}.$$

4: Generate a uniform random permutation π of the set $\{0, \ldots, m - 1\}$, where each of the $m!$ possible permutations has the same probability $1/(m!)$ to be selected.

Algorithm 2 MIV-SPS: Indicator (Θ)

1: For the given Θ, compute the prediction errors

$$\mathcal{E}(\Theta) \doteq Y - \Phi\Theta.$$

2: Evaluate

$$S_0(\Theta) \doteq \frac{1}{n}P_n^{-\frac{1}{2}}\Psi^{\mathsf{T}}\mathcal{E}(\Theta),$$

$$S_i(\Theta) \doteq \frac{1}{n}P_n^{-\frac{1}{2}}\Psi^{\mathsf{T}}\Lambda_i\mathcal{E}(\Theta),$$

for $i \in \{1, ..., m - 1\}$.

3: Order scalars $\{\|S_i(\Theta)\|_{\mathrm{F}}^2\}$ according to $\succ_\pi$, which is the standard ">" relation with random tie-breaking: $\|S_k(\Theta)\|_{\mathrm{F}}^2 \succ_\pi \|S_j(\Theta)\|_{\mathrm{F}}^2$ if and only if $(\|S_k(\Theta)\|_{\mathrm{F}}^2 > \|S_j(\Theta)\|_{\mathrm{F}}^2) \vee (\|S_k(\Theta)\|_{\mathrm{F}}^2 = \|S_j(\Theta)\|_{\mathrm{F}}^2 \wedge \pi(k) > \pi(j))$.

4: Compute the rank $\mathcal{R}(\Theta)$ of $\|S_0(\Theta)\|_{\mathrm{F}}^2$ in the ordering:

$$\mathcal{R}(\Theta) \doteq \left[1 + \sum_{i=1}^{m-1}\mathbb{I}\left(\|S_0(\Theta)\|_{\mathrm{F}}^2 \succ_\pi \|S_i(\Theta)\|_{\mathrm{F}}^2\right) \right].$$

5: Return 1 if $\mathcal{R}(\Theta) \leq m - q$, otherwise return 0.

4.2. Formal confidence region construction

The SPS algorithm consists of two main parts, an initialization and an indicator function. In the *initialization* part, see Algorithm 1, the input is the user defined confidence probability p. The algorithm computes the main parameters and generates the random objects needed for the construction of the confidence region. In the *indicator* part, see Algorithm 2, the input of the algorithm is a particular parameter matrix Θ, and the function evaluates whether Θ is included in the confidence region.

By using the indicator function, the p-level SPS confidence region for Θ^* can be defined as follows

$$\Upsilon_n \doteq \left\{ \Theta \in \mathbb{R}^{d \times d_x} : \operatorname{Indicator}(\Theta) = 1 \right\}. \tag{22}$$

For the instrumental variable estimate $\widehat{\Theta}_n^{\mathrm{IV}}$ it holds that $S_0(\widehat{\Theta}_n^{\mathrm{IV}}) = 0$; therefore, the IV estimate is always in the SPS confidence region, assuming that it is non-empty.

5. Theoretical guarantees

5.1. Exact coverage probability

The confidence regions constructed by MIV-SPS have guaranteed coverage probabilities for the true parameter matrix, for any finite sample size. More precisely, we have:

Theorem 1. *Assuming* A1–A3, *the confidence probability of the constructed confidence region is exactly p, that is,*

$$\mathbb{P}\big(\Theta^* \in \Upsilon_n\big) = 1 - \frac{q}{m} = p. \tag{23}$$

The proof of Theorem 1, which is a generalization of the proof of *Theorem 1* in [18], can be found in Appendix B.

Observe that the coverage probability is *exact*, the confidence set is non-conservative. The assumptions on the noises are rather mild: no specific (parametric) distributions are assumed, the noise sequence can even be non-stationary (i.e., each noise vector can have a different distribution), and there are no moment or density assumptions, either.

This result is more general than the one we would get by applying the corresponding result of [18] to the vectorized version of (7), as the assumptions on the noises are milder.

5.2. Strong consistency

Similarly to standard SPS, MIV-SPS is also uniformly *strongly consistent*, under the following assumptions:

A4. There exists a positive definite matrix P such that

$$\lim_{n \to \infty} \frac{1}{n} \Psi_n^{\mathsf{T}} \Psi_n = \lim_{n \to \infty} P_n = P \quad (a.s.). \tag{24}$$

A5. There exists an invertible matrix V such that

$$\lim_{n \to \infty} \frac{1}{n} \Psi_n^{\mathsf{T}} \Phi_n = \lim_{n \to \infty} V_n = V \quad (a.s.). \tag{25}$$

A6. The following growth rate restriction holds almost surely for the rows of the regression matrices:

$$\sum_{k=1}^{\infty} \frac{\|\varphi_k\|^4}{k^2} < \infty. \tag{26}$$

A7. The following growth rate restriction holds almost surely for the rows of the instrumental variable matrices:

$$\sum_{k=1}^{\infty} \frac{\|\psi_k\|^4}{k^2} < \infty. \tag{27}$$

A8. Finally, the variance growth of the rows of the noise matrices satisfy the following condition:

$$\sum_{k=1}^{\infty} \frac{\big(\mathbb{E}[\|w_k\|^2]\big)^2}{k^2} < \infty. \tag{28}$$

Note that these assumptions are also rather mild. For example, if the regressors (or the instrumental variables) are bounded, then A6 (or, respectively, A7) is satisfied. If the noise vectors $\{w_t\}$ are square integrable and i.i.d., then A8 holds. However, these conditions allow unbounded regressors and instrumental variables as well as noise terms whose variances grow to infinity. We also note that A5 ensures the asymptotic correlation of the instrumental variables and the regressors, which is crucial for consistency.

Theorem 2. *Assuming* A1–A8, $\forall \varepsilon > 0$, *we have*

$$\mathbb{P}\left(\bigcup_{k=1}^{\infty} \bigcap_{n=k}^{\infty} \big\{ \Upsilon_n \subseteq \mathcal{B}_{\varepsilon}(\Theta^*) \big\}\right) = 1, \tag{29}$$

where $\mathcal{B}_{\varepsilon}(\Theta^*) \doteq \{ \Theta \in \mathbb{R}^{d \times d_x} : \| \Theta - \Theta^* \|_{\mathrm{F}} \leq \varepsilon \}$.

The proof can be found in Appendix C. The theorem studies the event that, given $\varepsilon > 0$, after at most finite number of observations, the confidence regions $\{\Upsilon_n\}$ will remain included in the ε ball around Θ^*, as the sample size, n, increases. This *tail event* happens almost surely.

6. Ellipsoidal outer approximation

The MIV-SPS Indicator function evaluates whether a given candidate parameter matrix Θ belongs to the confidence region, by comparing the $\{\|S_i(\Theta)\|_{\mathrm{F}}^2\}_{i=1}^{m-1}$ functions with $\|S_0(\Theta)\|_{\mathrm{F}}^2$.

In many applications, we need a compact representation of the *whole* confidence set, to make it easier working with it; for example, in robust control.

In this section we introduce a generalization of the *ellipsoidal outer approximation* (EOA) given in [17,18]. Our construction results in a semidefinite programming (convex) optimization problem with significantly less constraints than the one we would get by applying the approach of [17,18] to the vectorized system.

6.1. Ellipsoidal outer approximation

Using the definition of the Frobenius norm, $\|S_0(\Theta)\|_{\mathrm{F}}^2$ can be reformulated as follows

$$
\begin{aligned}
\|S_0(\Theta)\|_{\mathrm{F}}^2 &= \mathrm{tr}\Big(\big[\tfrac{1}{n}\Psi^{\mathsf{T}}(Y - \Phi\Theta)\big]^{\mathsf{T}} P_n^{-1} \big[\tfrac{1}{n}\Psi^{\mathsf{T}}(Y - \Phi\Theta)\big]\Big) \\
&= \mathrm{tr}\Big(\big[\tfrac{1}{n}\Psi^{\mathsf{T}}\Phi(\Theta - \widehat{\Theta}_n^{\mathrm{IV}})\big]^{\mathsf{T}} P_n^{-1} \big[\tfrac{1}{n}\Psi^{\mathsf{T}}\Phi(\Theta - \widehat{\Theta}_n^{\mathrm{IV}})\big]\Big) \\
&= \mathrm{tr}\big((\Theta - \widehat{\Theta}_n^{\mathrm{IV}})^{\mathsf{T}} V_n^{\mathsf{T}} P_n^{-1} V_n(\Theta - \widehat{\Theta}_n^{\mathrm{IV}})\big) \\
&= \|P_n^{-1/2} V_n(\Theta - \widehat{\Theta}_n^{\mathrm{IV}})\|_{\mathrm{F}}^2.
\end{aligned} \tag{30}
$$

To obtain an ellipsoid over-bound, consider the set given by the particular values of Θ for which $\|S_i(\Theta)\|_{\mathrm{F}}^2 \geq \|S_0(\Theta)\|_{\mathrm{F}}^2$ holds true for q of the $\|S_i(\Theta)\|_{\mathrm{F}}^2$, that is

$$\Upsilon_n \subseteq \left\{ \Theta \in \mathbb{R}^{d \times d_x} : \|P_n^{-1/2} V_n(\Theta - \widehat{\Theta}_n^{\mathrm{IV}})\|_{\mathrm{F}}^2 \leq r(\Theta) \right\}, \tag{31}$$

where $r(\Theta)$ is the qth largest value of $\{\|S_i(\Theta)\|_{\mathrm{F}}^2\}_{i=1}^{m-1}$.

Our objective now is to determine an ellipsoidal over-bound with a parameter-independent radius, r, instead of the parameter-dependent $r(\Theta)$. This outer approximation is a guaranteed confidence region for any finite sample with a compact representation given by $\widehat{\Theta}_n^{\mathrm{IV}}$, V_n, P_n and r.

6.2. Convex programming formulation

Comparing $\|S_0(\Theta)\|_{\mathrm{F}}^2$ with a single $\|S_i(\Theta)\|_{\mathrm{F}}^2$, we have

$$
\begin{aligned}
&\big\{ \Theta : \|S_0(\Theta)\|_{\mathrm{F}}^2 \leq \|S_i(\Theta)\|_{\mathrm{F}}^2 \big\} \subseteq \\
&\Big\{ \Theta : \|S_0(\Theta)\|_{\mathrm{F}}^2 \leq \max_{\Theta : \|S_0(\Theta)\|_{\mathrm{F}}^2 \leq \|S_i(\Theta)\|_{\mathrm{F}}^2} \|S_i(\Theta)\|_{\mathrm{F}}^2 \Big\}
\end{aligned} \tag{32}
$$

Inequality $\|S_0(\Theta)\|_{\mathrm{F}}^2 \leq \|S_i(\Theta)\|_{\mathrm{F}}^2$ can be reformulated as

$$
\mathrm{tr}\big((\Theta - \widehat{\Theta}_n^{\mathrm{IV}})^{\mathsf{T}} V_n^{\mathsf{T}} P_n^{-1} V_n(\Theta - \widehat{\Theta}_n^{\mathrm{IV}})\big) \leq
$$

$$
\mathrm{tr}\Big(\big[\tfrac{1}{n}\Psi^{\mathsf{T}}\Lambda_i(Y - \Phi\Theta)\big]^{\mathsf{T}} P_n^{-1} \big[\tfrac{1}{n}\Psi^{\mathsf{T}}\Lambda_i(Y - \Phi\Theta)\big]\Big) =
$$

$$
\mathrm{tr}\big(\Theta^{\mathsf{T}} Q_i^{\mathsf{T}} P_n^{-1} Q_i \Theta - \Theta^{\mathsf{T}} Q_i^{\mathsf{T}} P_n^{-1} M_i - M_i^{\mathsf{T}} P_n^{-1} Q_i \Theta + M_i^{\mathsf{T}} P_n^{-1} M_i\big), \tag{33}
$$

where matrices Q_i and M_i are defined as

$$Q_i \doteq \frac{1}{n}\Psi^{\mathsf{T}}\Lambda_i\Phi, \qquad \text{and} \qquad M_i \doteq \frac{1}{n}\Psi^{\mathsf{T}}\Lambda_i Y. \tag{34}$$

Observe that

$$\max_{\Theta:\|S_0(\Theta)\|_{\mathsf{F}}^2 \le \|S_i(\Theta)\|_{\mathsf{F}}^2} \|S_i(\Theta)\|_{\mathsf{F}}^2 = \max_{\Theta:\|S_0(\Theta)\|_{\mathsf{F}}^2 \le \|S_i(\Theta)\|_{\mathsf{F}}^2} \|S_0(\Theta)\|_{\mathsf{F}}^2. \tag{35}$$

After defining matrix $Z \doteq P_n^{-1/2}V_n(\Theta - \widehat{\Theta}_n^{\mathrm{IV}})$, the quantity $\max_{\Theta:\|S_0(\Theta)\|_{\mathsf{F}}^2 \le \|S_i(\Theta)\|_{\mathsf{F}}^2} \|S_i(\Theta)\|_{\mathsf{F}}^2$ can be computed by

$$\begin{aligned}
\max \quad & \|Z\|_{\mathsf{F}}^2 \\
\text{s.t.} \quad & \mathrm{tr}(Z^{\mathsf{T}}A_i Z + Z^{\mathsf{T}}B_i + B_i^{\mathsf{T}}Z + C_i) \le 0,
\end{aligned} \tag{36}$$

where A_i, B_i and C_i are defined as follows

$$\begin{aligned}
A_i &\doteq I - P_n^{\frac{1}{2}\mathsf{T}}V_n^{-\mathsf{T}}Q_i^{\mathsf{T}}P_n^{-1}Q_i V_n^{-1}P_n^{\frac{1}{2}} \\
B_i &\doteq P_n^{\frac{1}{2}\mathsf{T}}V_n^{-\mathsf{T}}Q_i^{\mathsf{T}}P_n^{-1}(M_i - Q_i\widehat{\Theta}_n^{\mathrm{IV}}) \\
C_i &\doteq -\Theta^{\mathsf{T}}Q_i^{\mathsf{T}}P_n^{-1}Q_i\Theta + \Theta^{\mathsf{T}}Q_i^{\mathsf{T}}P_n^{-1}M_i + M_i^{\mathsf{T}}P_n^{-1}Q_i\Theta - M_i^{\mathsf{T}}P_n^{-1}M_i.
\end{aligned} \tag{37}$$

This program is *not* convex, however weak Lagrange duality holds, therefore an *upper bound* on the value of the above optimization problem can be obtained from the *dual* of the problem, for which an equivalent problem is

$$\begin{aligned}
\min \quad & \mathrm{tr}(\Gamma - \lambda C_i) \\
\text{s.t.} \quad & \lambda \ge 0 \\
& \begin{bmatrix} -I + \lambda A_i & \lambda B_i \\ \lambda B_i^{\mathsf{T}} & \Gamma \end{bmatrix} \succeq 0,
\end{aligned} \tag{38}$$

where Γ is a symmetric matrix and $\succeq 0$ denotes that the matrix is positive semidefinite. The derivation of this formulation can be found in Appendix A. The optimization problem (38) is now *convex* and can be solved using CVXPY [28], e.g., with solver CVXOPT or MOSEK.

Let γ_i^* be the solution of the program (38). We have

$$\left\{\Theta : \|S_0(\Theta)\|_{\mathsf{F}}^2 \le \|S_i(\Theta)\|_{\mathsf{F}}^2\right\} \subseteq \left\{\Theta : \|S_0(\Theta)\|_{\mathsf{F}}^2 \le \gamma_i^*\right\}, \tag{39}$$

therefore

$$\Upsilon_n \subseteq \widehat{\Upsilon}_n \doteq \left\{\Theta \in \mathbb{R}^{d \times d_x} : \|P_n^{-1/2}V_n(\Theta - \widehat{\Theta}_n^{\mathrm{IV}})\|_{\mathsf{F}}^2 \le r\right\}, \tag{40}$$

where r is the qth largest value of $\{\gamma_i^*\}_{i=1}^{m-1}$. $\widehat{\Upsilon}_n$ is an *ellipsoidal outer approximation* of Υ_n, and from (40) we have

$$\mathbb{P}\left(\Theta^* \in \widehat{\Upsilon}_n\right) \ge 1 - \frac{q}{m} = p, \tag{41}$$

for any n. The method is summarized by Algorithm 3.

An important difference between the IV-SPS [18] and the MIV-SPS ellipsoidal outer approximation computation, apart from the reduced assumptions on the noise vectors, is the size of the constraint matrix in the resulting convex programming formulations. Let $L \in \mathbb{R}^{l \times l}$ be the constraint matrix we get in the direct LSS identification problem. Then, in the IV-SPS case, l grows according to $l = d_x^2 + d_x d_u + 1$ with the dimensions of the inputs and the outputs. On the other hand, in the MIV-SPS case, it just grows according to $l = 2d_x + d_u$ with the dimensions.

7. Numerical experiments

In this section we present numerical experiments evaluating MIV-SPS for a closed-loop LSS system. The ellipsoidal outer approximations of MIV-SPS are compared to that of IV-SPS, after suitable vectorization. The effect of using different exploitation–exploration trade-offs in the feedback is also studied experimentally. Finally, the sample efficiency of the ellipsoidal over-bound of the confidence regions of MIV-SPS are investigates for different feedback rules.

Algorithm 3 MIV-SPS: Ellipsoidal Outer Approximation

1: Compute the instrumental variable estimate,
   $\widehat{\Theta}_n^{\mathrm{IV}} \doteq (\Psi^{\mathsf{T}}\Phi)^{-1}\Psi^{\mathsf{T}}Y$.
2: For $i \in \{1, \ldots, m-1\}$ solve the optimization problem (38) and let γ_i^* be the optimal value of the program.
3: Let r be the qth largest value of $\{\gamma_i^*\}_{i=1}^{m-1}$.
4: The ellipsoidal outer approximation of the SPS confidence region is given by
   $$\widehat{\Upsilon}_n \doteq \left\{\Theta \in \mathbb{R}^{d \times d_x} : \|P_n^{-1/2}V_n(\Theta - \widehat{\Theta}_n^{\mathrm{IV}})\|_{\mathsf{F}}^2 \le r\right\}.$$

7.1. Experimental setup

We consider the following closed-loop LSS system

$$\begin{aligned}
x_{k+1} &= Ax_k + Bu_k + w_k, \\
u_k &= \varepsilon Kx_k + (1-\varepsilon)r_k,
\end{aligned} \tag{42}$$

where A is a stable matrix (i.e., its eigenvalues are inside the unit circle), the elements of the square matrix B are sampled from the uniform distribution $\mathcal{U}(1, 10)$ and K is the optimal controller for $\varepsilon = 1$, given the quadratic cost

$$J(K) = \sum_{k=0}^{n-1} q\|x_k\|^2 + v\|u_k\|^2, \tag{43}$$

i.e., an LQR problem [2]. In the cost function the values of the scalar cost weights are $q = v = 1$, unless stated otherwise. The scalar $\varepsilon \in (0, 1)$ is the *exploitation* parameter. Sequences $\{w_k\}$ and $\{r_k\}$ contain independent random vectors, where the references are standard normal, while the noise distributions vary with the experiments. We consider a finite data sample of size n that contains input-output-reference tuples, $\{\langle u_k, x_k, r_k \rangle\}_{k=0}^{n-1}$.

The instrumental variables $\{\psi_k\}$ are generated from the data. We present the direct and the indirect cases side-by-side. First, least squares estimates of the parameter matrices are computed. The system parameters A and B are estimated in the direct case, while in the indirect case C and D are estimated. Let $\widehat{\Theta}^{\mathrm{LS}}$ and $\widehat{\Theta}_{\mathrm{id}}^{\mathrm{LS}}$ denote the LS estimate in the direct and the indirect case, respectively:

$$\widehat{\Theta}^{\mathrm{LS}} \doteq (\Phi^{\mathsf{T}}\Phi)^{-1}\Phi^{\mathsf{T}}Y, \qquad \widehat{\Theta}_{\mathrm{id}}^{\mathrm{LS}} \doteq (\Phi_{\mathrm{id}}^{\mathsf{T}}\Phi_{\mathrm{id}})^{-1}\Phi_{\mathrm{id}}^{\mathsf{T}}Y, \tag{44}$$

where

$$\widehat{\Theta}^{\mathrm{LS}} \doteq \begin{bmatrix} \hat{A}^{\mathsf{T}} \\ \hat{B}^{\mathsf{T}} \end{bmatrix}, \qquad \widehat{\Theta}_{\mathrm{id}}^{\mathrm{LS}} \doteq \begin{bmatrix} \hat{C}^{\mathsf{T}} \\ \hat{D}^{\mathsf{T}} \end{bmatrix}. \tag{45}$$

Using the estimates $\hat{A}$ and $\hat{B}$, or $\hat{C}$ and $\hat{D}$ in the indirect case, noiseless state sequences are generated as follows

$$\bar{x}_{k+1} \doteq \hat{A}\bar{x}_k + \hat{B}r_k, \qquad \bar{x}_{k+1}^{\mathrm{id}} \doteq \hat{C}\bar{x}_k + \hat{D}r_k. \tag{46}$$

Using the noiseless state sequences, the instrumental variables can be defined in the direct and indirect cases as

$$\psi_k \doteq (\bar{x}_k^{\mathsf{T}}, r_k^{\mathsf{T}})^{\mathsf{T}}, \qquad \psi_k^{\mathrm{id}} \doteq ((\bar{x}_k^{\mathrm{id}})^{\mathsf{T}}, r_k^{\mathsf{T}})^{\mathsf{T}} \tag{47}$$

In the strict sense, the IVs defined in (47) are not completely independent of the noise, because the least squares estimates $\widehat{\Theta}^{\mathrm{LS}}$ and $\widehat{\Theta}_{\mathrm{id}}^{\mathrm{LS}}$ depend on the noise. However, in both estimates the noises are "averaged out", hence their effect is minimal. If the LS estimates were constructed from a data sample that is independent of the sample used by the SPS algorithm, then the obtained regions would be guaranteed to have exact coverage. The difference between these two constructions considering the experimental results is negligible, therefore we only used one data sample.

Table 1
Flops and CPU time needed to compute the ellipsoidal outer approximation of MIV-SPS relative to IV-SPS, $n = 500$, $s = 500$.

Dim.	Params.	Rel. flops	Rel. time
1	2	0.839	1.03
2	8	0.427	0.923
3	18	0.146	0.528
4	32	0.056	0.195

Table 2
Comparison of empirical probabilities of covering the true parameter matrix for standard normal noises, $n = 500$, $s = 500$.

Dim.	Params.	$\hat{p}_{AS}$	$\hat{p}_{IN}$	$\hat{p}_{IV}$	$\hat{p}_{MIV}$
1	2	0.892	0.894	0.926	0.926
2	8	0.892	0.894	0.954	0.946
3	18	0.89	0.888	0.982	0.978
4	32	0.902	0.89	0.996	0.996

Table 3
Comparison of empirical probabilities of covering the true parameter for mixture of two Gaussian noises, $n = 500$, $s = 500$.

Dim.	Params.	$\hat{p}_{AS}$	$\hat{p}_{IN}$	$\hat{p}_{IV}$	$\hat{p}_{MIV}$
1	2	0.9	0.9	0.932	0.932
2	8	0.876	0.9	0.95	0.958
3	18	0.852	0.908	0.984	0.98
4	32	0.847	0.898	0.982	0.982

Table 4
Comparison of empirical coverage probabilities for bimodal mixture of two non-stationary Laplacian noises, $n = 500$, $s = 500$.

Dim.	Params.	$\hat{p}_{AS}$	$\hat{p}_{IN}$	$\hat{p}_{IV}$	$\hat{p}_{MIV}$
1	2	0.888	0.896	0.928	0.928
2	8	0.872	0.906	0.942	0.96
3	18	0.846	0.908	0.968	0.976
4	32	0.834	0.898	0.976	0.982

In the experiments the conservatism of the ellipsoidal outer approximations of the confidence regions is investigated using Monte Carlo simulations. In these simulations we set the confidence probability to $p = 0.9$ and computed an empirical probability $\hat{p} = s_{in}/s$, where s_{in} is the number of simulations, where the true parameter matrix Θ^* is included in the ellipsoid, and s is the number of simulations.

7.2. Comparison of MIV-SPS and IV-SPS

As discussed in Section 6.2, a crucial difference between the MIV-SPS and IV-SPS ellipsoidal outer approximation (EOA) computations is the size of the constraint matrices in the SDP formulations. As a consequence, solving the convex program in the MIV-SPS case should be faster than in the IV-SPS case. Table 1 presents the number of flops and the required CPU time for solving the convex programs in the MIV-SPS EOA *relative* to those of the IV-SPS EOA, for different dimensions. The obtained results confirm that EOAs can be computed more efficiently by the MIV-SPS approach, especially in higher dimensions.

The experiments discussed in this section, i.e., Tables 1, 2, 3 and 4, are all based on $s = 500$ Monte Carlo simulations, for each dimension, with sample size $n = 500$ using an open-loop system ($\varepsilon = 0$) and direct identification.

Although, strong duality is not proven yet for the reformulation of (38), only weak duality is exploited, the experimental evidences are indicative of the phenomenon that it holds. The empirical coverage probabilities (i.e., the ratio of cases when the regions cover the true parameter) in the MIV-SPS EOA case $\hat{p}_{MIV}$ are very close to those of IV-SPS EOA $\hat{p}_{IV}$, for which strong duality is known.

We have compared the empirical coverage probabilities $\hat{p}$ of the asymptotic ellipsoids $\hat{p}_{AS}$, see Section 3, the empirical probabilities $\hat{p}_{IN}$ for the MIV-SPS Indicator algorithm, presented in Section 4 (the ratio of cases when the Indicator algorithms return 1 for the true parameter), and the empirical probabilities $\hat{p}_{IV}$, $\hat{p}_{MIV}$ for three noise scenarios. The target confidence probability was always 0.9.

The empirical coverage probabilities in case the process noise $\{w_k\}$ is a sequence of i.i.d. (multivariate) standard normal vectors is given in Table 2. It can be observed that the asymptotic and the MIV-SPS Indicator confidence regions are very precise and the SPS confidence ellipsoidal outer approximations become more conservative as the dimension increases, but the empirical coverage probabilities of MIV-SPS and IV-SPS stay close to each other.

In the case, where the process noise $\{w_k\}$ is a sequence of i.i.d. Gaussian random vectors, the asymptotic confidence region gives

accurate estimation. SPS builds confidence regions under milder conditions, even when the noise is non-stationary. Simulation experiments were performed, where $\{w_k\}$ is a sequence of i.i.d. bimodal mixture of two Gaussians $w_k = \frac{1}{2}Z_1 + \frac{1}{2}Z_2, Z_1 \sim \mathcal{N}(1, \sigma_w I), Z_2 \sim \mathcal{N}(-1, \sigma_w I)$, and where $\{w_k\}$ is a sequence of time dependent mixture of two Laplacians $w_k = \frac{1}{2}Z_3 + \frac{1}{2}Z_4, Z_3 \sim \mathcal{L}(\frac{5(k+1)}{n}, (\frac{(k+1)}{n}+\sigma_w)I), Z_4 \sim \mathcal{L}(\frac{-5(k+1)}{n}, (\frac{(k+1)}{n}+\sigma_w)I)$, with $\sigma_w = 1$.

The results of these experiments are presented in Tables 3 and 4. In both experiments the empirical coverage probabilities of the asymptotic confidence regions are less than $p = 0.9$ by a significant margin. In the bimodal Gaussian case, this deviation is due to the insufficient sample size, while in the bimodal Laplacian noise case, it is also due to the heavier-tail and the non-stationarity of the noise.

On the other hand, the MIV-SPS Indicator algorithm constructs proper confidence regions in both cases, and the ellipsoidal outer approximations give similar coverage probabilities in the IV-SPS and MIV-SPS cases. The ellipsoidal outer approximations introduce some conservatism, but compared to the confidence regions based on the asymptotic theory, they provide better results when a non-asymptotic lower bound on the probability is required.

7.3. Exploration vs exploitation

The signal $\{r_k\}$ can be seen as an exploration noise affecting the controller instead of a reference signal. In this case the parameter ε controls the exploitation–exploration trade-off. The system operates in an open-loop setting when $\varepsilon = 0$ and without exploration in a closed-loop when $\varepsilon = 1$. In our experimental setting, given in (42), the choice of exploitation rate affects the size of the EOA of the confidence region in a way that with more exploitation the outer approximation of the confidence region gets more conservative. This phenomenon is due to the construction of our instrumental variables detailed in Section 7.1. The instrumental variables are created by generating a noiseless state sequence using the exploration noise $\{r_k\}$. With less exploration, the correlation between the instrumental variable sequence and the state sequence is decreasing, thus the IV estimate gets less accurate, therefore the EOA gets more conservative. In a real-world settings, instrumental variables can often be created from signals that are observable, but not controllable. In such cases no conservatism would be introduced with more exploitation.

The choice of the controller K also influences the conservatism of the EOA. We have investigated the empirical coverage probabilities of MIV-SPS EOAs for different q and v instances using the direct and indirect identification approaches. The results are illustrated in Fig. 1 for the direct case and in Fig. 2 for the indirect case. The results are based on $s = 500$ Monte Carlo simulations for each ε on a 2-dimensional LSS system with sample size $n = 500$. Standard normal process noise is applied.
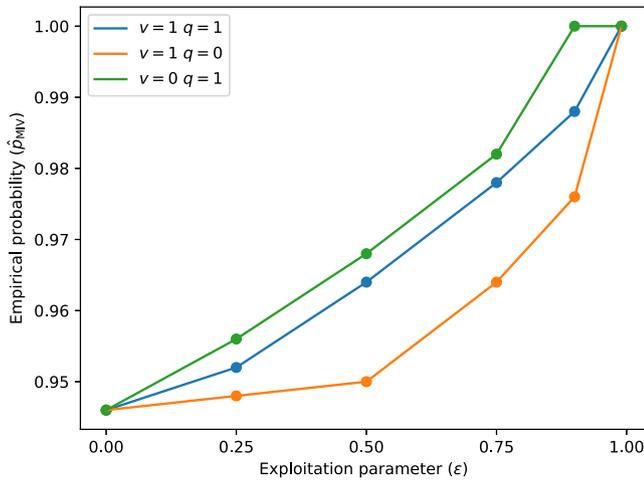
Fig. 1. Empirical coverage probability of MIV-SPS EOA as a function of the exploitation parameter for different feedback setups, using *direct* identification, for a 2-dimensional LSS, $n = 500$, $s = 500$.
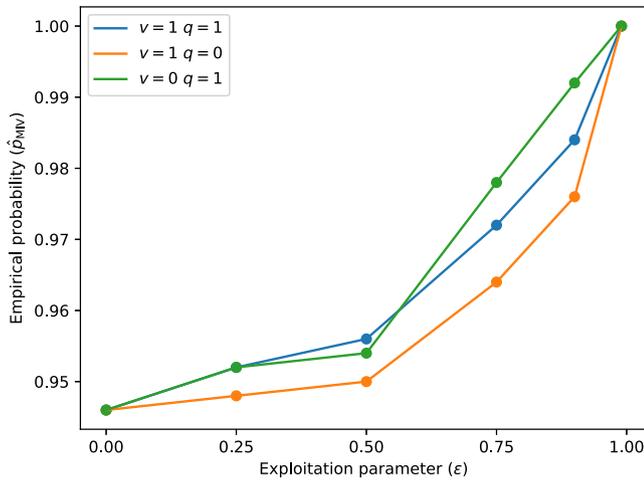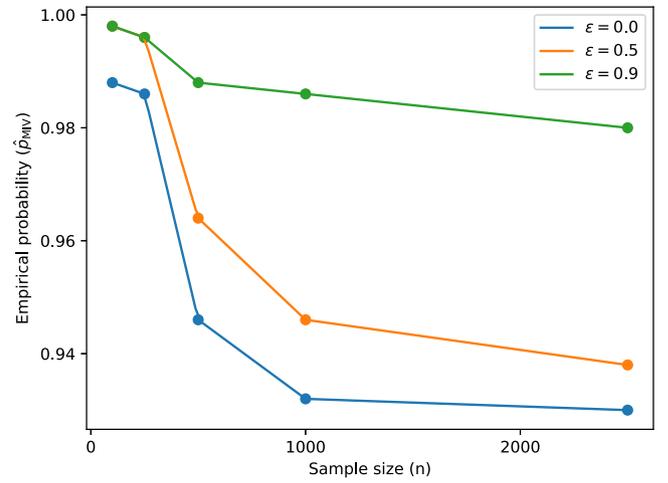

Fig. 2. Empirical coverage probability of MIV-SPS EOA as a function of the exploitation parameter for different feedback setups, using *indirect* identification, for a 2-dimensional LSS, $n = 500$, $s = 500$.

7.4. Sample complexity

In the last experiments, the sample complexity of the MIV-SPS EOA method is investigated. As the sample size n increases, the EOA becomes less conservative, hence it provides more accurate confidence regions. Simulation results are presented for different exploitation parameter ε values. 500 Monte Carlo simulations were run for each sample size on a 2-dimensional LSS system, using standard normal noise and with direct identification. The results are illustrated in Fig. 3. It can be observed that with more exploitation (as ε increases), the empirical coverage probability of EOA converges slower to the target confidence, due to the inaccuracy of the IV estimate.

8. Conclusions

The paper introduced a *matrix variate regression* and *instrumental variables* based generalization of the *Sign-Perturbed Sums* (SPS) identification method, called MIV-SPS. It was illustrated on *closed-loop linear state-space* (LSS) models. The MIV-SPS method constructs *non-asymptotic*, *distribution-free* confidence regions around the instrumental variable estimate under mild statistical

Fig. 3. Empirical coverage probability of MIV-SPS EOA as a function of the sample size n in an open-loop ($\varepsilon = 0$) setting, using direct identification, for a 2-dimensional LSS, $s = 500$.

assumptions on the process noise. The *exact confidence* and the *strong consistency* of MIV-SPS regions were proven.

A computationally efficient *ellipsoidal outer approximation* (EOA) method was suggested, as well, which is a compact over-bound of the MIV-SPS confidence region. Simulation experiments demonstrated that, for LSS models, our proposed method computes the EOA more efficiently than the scalar IV-SPS, developed earlier, while keeping a similar size. Experiments comparing MIV-SPS and the confidence ellipsoids based on the asymptotic theory were also concluded, showing the superiority of MIV-SPS, e.g., for *heavy-tailed* and *non-stationary* noise sequences.

Extending MIV-SPS to be able to handle *partially observable* systems and combining the ideas with *robust control* approaches are subject of future research.

CRediT authorship contribution statement

Szabolcs Szentpéteri: Conceptualization, Methodology, Software, Investigation, Visualization, Validation, Formal analysis, Writing – original draft, Writing – review & editing. **Balázs Csanád Csáji:** Conceptualization, Methodology, Validation, Formal analysis, Writing – original draft, Writing – review & editing, Supervision, Project administration, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Only computer generated (synthetic) data was used in the experiments.

Acknowledgments

Appendix A. Semidefinite programming formulation of the outer approximation

The optimization problem

$$\max \quad \|Z\|_F^2$$
$$\text{s.t.} \quad \text{tr}(Z^TA_iZ + Z^TB_i + B_i^TZ + C_i) \le 0, \tag{A.1}$$

has the equivalent form

$$\min \quad -\text{tr}(Z^TZ)$$
$$\text{s.t.} \quad \text{tr}(Z^TA_iZ + Z^TB_i + B_i^TZ + C_i) \le 0. \tag{A.2}$$

The Lagrangian of the problem (A.2) can be written as

$$L(Z, \lambda) \doteq -\text{tr}(Z^TZ) + \lambda \text{tr}(Z^TA_iZ + Z^TB_i + B_i^TZ + C_i)$$
$$= \text{tr}(Z^T(-I + \lambda A_i)Z + \lambda Z^TB_i + \lambda B_i^TZ + \lambda C_i). \tag{A.3}$$

The Lagrange dual function can be formulated as

$$g(\lambda) \doteq \inf_Z L(Z, \lambda). \tag{A.4}$$

If $L(Z, \lambda)$ takes its minimum with respect to Z, then

$$\nabla_Z L(Z, \lambda) = (-I + \lambda A_i)Z + \lambda B_i = 0. \tag{A.5}$$

From the above equation it can be seen that $L(Z, \lambda)$ takes its minimum value, where $Z = -(-I + \lambda A_i)^{-1}\lambda B_i$, thus the Lagrange dual function can be written as

$$g(\lambda) = \text{tr}(-\lambda^2 B_i^T(-I + \lambda A_i)^{-1}B_i + \lambda C_i), \tag{A.6}$$

therefore, the dual problem can be formulated as

$$\max \quad \text{tr}(-\lambda^2 B_i^T(-I + \lambda A_i)^{-1}B_i + \lambda C_i)$$
$$\text{s.t.} \quad \lambda \ge 0. \tag{A.7}$$

The optimization problem (A.7) can be reformulated as

$$\min \quad \text{tr}(\Gamma - \lambda C_i)$$
$$\text{s.t.} \quad \lambda \ge 0 \tag{A.8}$$
$$\Gamma \succeq \lambda^2 B_i^T(-I + \lambda A_i)^{-1}B_i,$$

where Γ is a symmetric matrix and relation "$\succeq$" denotes the Loewner partial ordering of semidefinite matrices.

Finally, the constraint $\Gamma \succeq \lambda^2 B_i^T(-I + \lambda A_i)^{-1}B_i$ can be reformulated by resorting to the Schur complement,

$$\begin{bmatrix} -I + \lambda A_i & \lambda B_i \\ \lambda B_i^T & \Gamma \end{bmatrix} \succeq 0, \tag{A.9}$$

which leads to the convex optimization problem (38).

Appendix B. Proof of Theorem 1 (Exact Confidence)

Definition 1. Let $Z_1, \ldots, Z_k$ be a finite collection of random variables and $\succ$ a strict total order. If for all permutations $i_1, \ldots, i_k$ of indices $1, \ldots, k$ we have

$$\mathbb{P}(Z_{i_k} \succ Z_{i_{k-1}} \succ \cdots \succ Z_{i_1}) = \frac{1}{k!}, \tag{B.1}$$

then we call $\{Z_i\}$ uniformly ordered w.r.t. order $\succ$.

Lemma 1. Let $\alpha, \beta_1, \ldots, \beta_k$ be i.i.d. random signs (random variables taking ± 1 with probability ½ each), then the variables $\alpha, \alpha \cdot \beta_1, \ldots, \alpha \cdot \beta_k$ are i.i.d. random signs.

Lemma 2. Let X and Y be two independent, $\mathbb{R}^d$-valued and $\mathbb{R}^k$-valued random vectors, respectively. Let us consider a (measurable) function $g : \mathbb{R}^d \times \mathbb{R}^k \to \mathbb{R}$ and a (measurable) set $A \subseteq \mathbb{R}$. If we have $\mathbb{P}(g(x, Y) \in A) = p$, for all (constant) $x \in \mathbb{R}^d$, then we also have $\mathbb{P}(g(X, Y) \in A) = p$.

Lemma 3. Let $Z_1, \ldots, Z_k$ be real-valued, i.i.d. random variables. Then, they are uniformly ordered w.r.t. $\succ_\pi$.

The proofs of Lemmas 1–3 can be found in [17].

Lemma 4. Let w be a random vector with $w \overset{d}{=} -w$, that is w is symmetrically distributed about zero. Then, there exist a random sign s and a random vector v with the properties: s is independent of v and $w = s \cdot v$.

Proof of Lemma 4. Let s be a random sign (that is: a Rademacher random variable) which is independent of w, and let $v \doteq s \cdot w$. Then, it holds for every event A, that

$$\mathbb{P}(s = 1 \wedge v \in A) = \mathbb{P}(s = 1 \wedge s \cdot w \in A)$$
$$= \mathbb{P}(s = 1 \wedge w \in A)$$
$$= \mathbb{P}(s = 1) \cdot \mathbb{P}(w \in A) \tag{B.2}$$
$$= \mathbb{P}(s = 1) \cdot \mathbb{P}(s \cdot w \in A)$$
$$= \mathbb{P}(s = 1) \cdot \mathbb{P}(v \in A).$$

$$\mathbb{P}(s = -1 \wedge v \in A) = \mathbb{P}(s = -1 \wedge s \cdot w \in A)$$
$$= \mathbb{P}(s = -1 \wedge -w \in A)$$
$$= \mathbb{P}(s = 1) \cdot \mathbb{P}(-w \in A) \tag{B.3}$$
$$= \mathbb{P}(s = 1) \cdot \mathbb{P}(s \cdot w \in A)$$
$$= \mathbb{P}(s = 1) \cdot \mathbb{P}(v \in A).$$

Therefore, s and v are independent and $w = s \cdot v$. □

Proof of Theorem 1. By definition, the true parameter, Θ^*, is included in the confidence region if $\mathcal{R}(\Theta^*) \le m - q$. In this case $\|S_0(\Theta^*)\|_F^2$ is ranked $1, \ldots, m-q$ according to the ascending order (w.r.t. $\succ_\pi$) of variables $\{\|S_i(\Theta^*)\|_F^2\}$.

We will prove that $\{\|S_i(\Theta^*)\|_F^2\}$ are uniformly ordered, thus the probability that $\|S_0(\Theta^*)\|_F^2$ takes any position in the ordering equals $1/m$, therefore the probability that its rank is at most $m - q$ is $1 - q/m$. Throughout the proof we will work with a *fixed realization* of the instrumental variables, the general case follows from Lemma 2.

For $\Theta = \Theta^*$, the $S_i(\cdot)$ functions can be formulated as

$$S_i(\Theta^*) = \frac{1}{n}P_n^{-\frac{1}{2}}\Psi^T\Lambda_iW, \tag{B.4}$$

$\forall i \in \{0, \ldots, m - 1\}$, where $\Lambda_0 = I \in \mathbb{R}^{n \times n}$.

From (B.4) it can be seen that functions $\{S_i(\cdot)\}$ depend on the perturbed noise matrices, $\{\Lambda_iW\}$, via the same function for all i, which is denoted by $S(\Lambda_iW) \doteq S_i(\Theta^*)$.

By using the result of Lemma 4, the noise vector sequence $\{w_k\}$ can be decomposed as $w_k = s_k\tilde{w}_k$, where s_k is a Rademacher random variable independent of $\tilde{w}_k$. Let $\gamma_{i,k} \doteq \alpha_{i,k}s_k$, where $\alpha_{0,k} = 1$, $\forall k$, $\forall i \in 0, \ldots, m - 1$. Lemma 1 can be used to show that $\gamma_{i,k}$ are i.i.d. random signs, since $\{\alpha_{i,k}\}$ are random signs and $\tilde{s}_k$ is a Rademacher random variable. As a consequence of Lemma 4, the sequence $\{\gamma_{i,k}\}$ are independent of $\{\tilde{w}_k\}$.

By fixing a realization of $\{\tilde{w}_k\}$, called $\{v_k\}$, real-valued variables $\{Z_i\}$ can be defined as

$$Z_i \doteq \|S(\Gamma_iM)\|_F^2, \tag{B.5}$$

where

$$\Gamma_i \doteq \begin{bmatrix} \gamma_{i,1} & & \\ & \ddots & \\ & & \gamma_{i,n} \end{bmatrix}, \quad M \doteq \begin{bmatrix} v_0^T \\ v_1^T \\ \vdots \\ v_{n-1}^T \end{bmatrix}. \tag{B.6}$$

Random variables $\{Z_i\}$ are i.i.d., and by Lemma 3, it follows that they are uniformly ordered w.r.t. $\succ_\pi$.

In order to relax the assumption of a fixed realization of $\{\tilde{w}_k\}$, Lemma 2 can be applied. By doing so, it follows that the $\{\|S_i(\Theta^*)\|_F^2\}$ variables are uniformly ordered (unconditionally), which completes the proof of the theorem. $\quad\square$

Appendix C. Proof of Theorem 2 (Strong Consistency)

Proof of Theorem 2. The proof consists of two parts. In the first part we are going to prove that for any $\Theta' \neq \Theta^*$:

$$\|S_0(\Theta')\|_F^2 \xrightarrow{a.s.} \|P^{-\frac{1}{2}}V\tilde{\Theta}\|_F^2 > 0, \tag{C.1}$$

where matrices P and V are defined in Assumptions A4 and A5, and $\tilde{\Theta} \doteq (\Theta^* - \Theta')$. On the other hand, for $i \neq 0$, we have

$$\|S_i(\Theta')\|_F^2 \xrightarrow{a.s.} 0, \tag{C.2}$$

as $n \to \infty$. Consequently, as n grows, the rank $\mathcal{R}(\Theta')$ of $\|S_0(\Theta')\|_F^2$ will be eventually equal to m, therefore Θ' will be (a.s.) excluded from the confidence region, as $n \to \infty$.

The results will be derived for a fixed realization of the instrumental variables, just like in the proof of Theorem 1. Since the matrix of IVs Ψ_n and the noise W are independent A2, the obtained results hold true (almost surely) on the whole probability space, ensured by Lemma 2. In the second part, we will prove that the confidence region converge to Θ^* uniformly, not just pointwise.

Function $S_0(\Theta')$ can be formulated as

$$S_0(\Theta') = \frac{1}{n}P_n^{-\frac{1}{2}}\Psi_n^T\mathcal{E}_n(\Theta') = \frac{1}{n}P_n^{-\frac{1}{2}}\Psi_n^T(Y_n - \Phi_n\Theta')$$
$$= \frac{1}{n}P_n^{-\frac{1}{2}}\Psi_n^T\Phi_n\tilde{\Theta} + \frac{1}{n}P_n^{-\frac{1}{2}}\Psi_n^TW_n. \tag{C.3}$$

The two terms can be examined separately. By observing that $(\cdot)^{\frac{1}{2}}$ is a continuous matrix function and applying A4 and A5 the convergence of the first part follows, thus

$$\frac{1}{n}P_n^{-\frac{1}{2}}\Psi_n^T\Phi_n\tilde{\Theta} = P_n^{-\frac{1}{2}}V_n\tilde{\Theta} \xrightarrow{a.s.} P^{-\frac{1}{2}}V\tilde{\Theta}, \tag{C.4}$$

as $n \to \infty$. The convergence of the second term is proved from the element-wise application of Kolmogorov's strong law of large numbers (SLLN) for independent variables [29]. Note that $\{P_n^{-\frac{1}{2}}\} \xrightarrow{a.s.} \{P^{-\frac{1}{2}}\}$ as $n \to \infty$, thus it is enough to prove that $\frac{1}{n}\Psi_n^TW_n \xrightarrow{a.s.} 0$. By applying the Cauchy–Schwarz inequality, A7, and A8, we have

$$\sum_{k=0}^{\infty} \frac{\mathbb{E}[(\psi_{k,j}w_{k,l})^2]}{k^2} \leq \sum_{k=0}^{\infty} \frac{\|\psi_k\|^2}{k}\frac{\mathbb{E}[\|w_k\|^2]}{k} \leq$$
$$\sqrt{\sum_{k=0}^{\infty}\frac{\|\psi_k\|^4}{k^2}}\sqrt{\sum_{k=0}^{\infty}\frac{\mathbb{E}[\|w_k\|^2]^2}{k^2}} < \infty. \tag{C.5}$$

Therefore the condition of the SLLN holds, thus

$$\frac{1}{n}P_n^{-\frac{1}{2}}\Psi_n^TW_n \xrightarrow{a.s.} 0, \quad \text{as} \quad n \to \infty. \tag{C.6}$$

From the above results we obtain

$$\|S_0(\Theta')\|_F^2 \xrightarrow{a.s.} \|P^{-\frac{1}{2}}V\tilde{\Theta}\|_F^2 > 0, \tag{C.7}$$

since V is full rank, $P^{-\frac{1}{2}}$ is positive definite and $\Theta' \neq \Theta^*$. The limit of $S_i(\Theta')$ can be derived similarly,

$$S_i(\Theta') = \frac{1}{n}P_n^{-\frac{1}{2}}\Psi_n^T\Lambda_i(Y_n - \Phi_n\Theta')$$
$$= \frac{1}{n}P_n^{-\frac{1}{2}}\Psi_n^T\Lambda_i\Phi_n(\tilde{\Theta}) + \frac{1}{n}P_n^{-\frac{1}{2}}\Psi_n^T\Lambda_iW_n, \tag{C.8}$$

We will examine the asymptotic behavior of the two terms separately again. The convergence of $\frac{1}{n}P_n^{-\frac{1}{2}}\Psi_n^T\Lambda_iW_n$ follows from (C.5),

since the variance of $\alpha_{i,k}\psi_{k,j}w_{k,l}$ equals the variance of $\psi_{k,j}w_{k,l}$, therefore

$$\frac{1}{n}P_n^{-\frac{1}{2}}\Psi_n^T\Lambda_iW_n \xrightarrow{a.s.} 0, \quad \text{as} \quad n \to \infty. \tag{C.9}$$

In the first term, it holds that, $\{P_n^{-\frac{1}{2}}\} \xrightarrow{a.s.} \{P^{-\frac{1}{2}}\}$ and $\tilde{\Theta}$ is constant, thus it is enough to prove that $\frac{1}{n}\Psi_n^T\Lambda_i\Phi_n$ converges almost surely to 0 element-wise. First, we fix a realization of every random variable except the random signs (Λ_i). For this realization the assumptions A6 and A7 hold. Then, $\{\alpha_{i,k}\psi_{k,j}\varphi_{k,l}\}$ becomes a sequence of conditionally independent random variables with conditional covariances $(\psi_{k,j}\varphi_{k,l})^2$. Using A6 and A7, we get

$$\sum_{k=0}^{\infty}\frac{(\psi_{k,j}\varphi_{k,l})^2}{k^2} \leq \sum_{k=0}^{\infty}\frac{\|\psi_k\|^2}{k}\frac{\|\varphi_k\|^2}{k} \leq$$
$$\sqrt{\sum_{k=0}^{\infty}\frac{\|\psi_k\|^4}{k^2}}\sqrt{\sum_{k=0}^{\infty}\frac{\|\varphi_k\|^4}{k^2}} < \infty, \tag{C.10}$$

therefore, by applying Kolmogorov's SLLN element-wise,

$$\frac{1}{n}P_n^{-\frac{1}{2}}\Psi_n^T\Lambda_i\Phi_n\tilde{\Theta} \xrightarrow{a.s.} 0, \tag{C.11}$$

as $n \to \infty$ holds for (almost) any realization, therefore holds true unconditionally. From the previous derivation

$$\|S_i(\Theta')\|_F^2 \xrightarrow{a.s.} 0, \tag{C.12}$$

as $n \to \infty$, for each $i \in \{1, \ldots, m-1\}$.

Previous results showed that for each i, the function $\|S_i(\Theta')\|_F^2$ converges with probability 1. As a consequence, for each realization $\omega \in \Omega$ (from an event with probability one, where $(\Omega, \mathcal{F}, \mathbb{P})$ is the underlying probability space), and for each $\delta > 0$, there exists a (realization dependent) $N(\omega) > 0$ such that for $n \geq N$ (from now on, $i \neq 0$),

$$\|P_n^{-\frac{1}{2}}V_n - P^{-\frac{1}{2}}V\|_F \leq \delta, \quad \|\frac{1}{n}P_n^{-\frac{1}{2}}\Psi_n^TW_n\|_F \leq \delta, \tag{C.13}$$

$$\|\frac{1}{n}P_n^{-\frac{1}{2}}\Psi_n^T\Lambda_i\Phi_n\|_F \leq \delta, \quad \|\frac{1}{n}P_n^{-\frac{1}{2}}\Psi_n^T\Lambda_iW_n\|_F \leq \delta. \tag{C.14}$$

Then, for all $n \geq N$, we have

$$\|S_0(\Theta')\|_F = \|P_n^{-\frac{1}{2}}V_n\tilde{\Theta} + \frac{1}{n}P_n^{-\frac{1}{2}}\Psi_n^TW_n\|_F =$$
$$\|(P_n^{-\frac{1}{2}}V_n - P^{-\frac{1}{2}}V)\tilde{\Theta} + P^{-\frac{1}{2}}V\tilde{\Theta} + \frac{1}{n}P_n^{-\frac{1}{2}}\Psi_n^TW_n\|_F =$$
$$\| - P^{-\frac{1}{2}}V\tilde{\Theta} - (P_n^{-\frac{1}{2}}V_n - P^{-\frac{1}{2}}V)\tilde{\Theta} - \frac{1}{n}P_n^{-\frac{1}{2}}\Psi_n^TW_n\|_F \geq$$
$$\| - P^{-\frac{1}{2}}V\tilde{\Theta}\|_F - \|(P_n^{-\frac{1}{2}}V_n - P^{-\frac{1}{2}}V)\tilde{\Theta}\|_F - \|\frac{1}{n}P_n^{-\frac{1}{2}}\Psi_n^TW_n\|_F \geq$$
$$\|P^{-\frac{1}{2}}V\tilde{\Theta}\|_F - \delta\|\tilde{\Theta}\|_F - \delta = \|U_\sigma\Sigma V_\sigma^T\tilde{\Theta}\|_F - \delta\|\tilde{\Theta}\|_F - \delta \geq$$
$$\sigma_{\min}(P^{-\frac{1}{2}}V)\|\tilde{\Theta}\|_F - \delta\|\tilde{\Theta}\|_F - \delta, \tag{C.15}$$

where $U_\sigma\Sigma V_\sigma$ is the SVD decomposition of $P^{-\frac{1}{2}}V$ and $\sigma_{\min}(\cdot)$ denotes the smallest singular value. We also have

$$\|S_i(\Theta')\|_F = \|\frac{1}{n}P_n^{-\frac{1}{2}}\Psi_n^T\Lambda_i\Phi_n\tilde{\Theta} + \frac{1}{n}P_n^{-\frac{1}{2}}\Psi_n^T\Lambda_iW_n\|_F$$
$$\leq \|\frac{1}{n}P_n^{-\frac{1}{2}}\Psi_n^T\Lambda_i\Phi_n\|_F\|\tilde{\Theta}\|_F + \|\frac{1}{n}P_n^{-\frac{1}{2}}\Psi_n^T\Lambda_iW_n\|_F$$
$$\leq \delta\|\tilde{\Theta}\|_F + \delta. \tag{C.16}$$

Hence, we have $\|S_i(\Theta')\|_F < \|S_0(\Theta')\|_F$, $\forall \Theta'$ that satisfy

$$\delta\|\tilde{\Theta}\|_F + \delta < \sigma_{\min}(P^{-\frac{1}{2}}V)\|\tilde{\Theta}\|_F - \delta\|\tilde{\Theta}\|_F - \delta, \tag{C.17}$$

which can be reformulated as

$$\kappa_0(\delta) \doteq \frac{2\delta}{\sigma_{\min}(P^{-\frac{1}{2}}V) - 2\delta} < \|\tilde{\Theta}\|_F, \tag{C.18}$$

therefore, those Θ' for which $\kappa_0(\delta) < \|\Theta^* - \Theta'\|_F$ are not in the confidence region Υ_n, for $n \geq N$. For any $\varepsilon > 0$, by setting $\delta = (\varepsilon \sigma_{\min}(P^{-\frac{1}{2}}V))/(2+2\varepsilon)$, we have $\Upsilon_n \subseteq \mathcal{B}_\varepsilon(\Theta^*)$, therefore, the claim of the theorem follows. $\square$

References

[1] L. Ljung, System Identification: Theory for the User, second ed., Prentice Hall, Upper Saddle River, 1999.

[2] D.P. Bertsekas, Dynamic Programming and Optimal Control, Vol. 2, third ed., Athena Scientific, 2007.

[3] P.R. Kumar, P. Varaiya, Stochastic Systems: Estimation, Identification, and Adaptive Control, SIAM, 2015.

[4] I.S. Khalil, J. Doyle, K. Glover, Robust and Optimal Control, Prentice Hall, New Jersey, 1996.

[5] E. Weyer, R. Williamson, I. Mareels, Finite sample properties of linear model identification, IEEE Trans. Automat. Control 44 (7) (1999) 1370–1383.

[6] M.C. Campi, E. Weyer, Finite sample properties of system identification methods, IEEE Trans. Automat. Control 47 (8) (2002) 1329–1334.

[7] M. Faradonbeh, A. Tewari, G. Michailidis, Finite time identification in unstable linear systems, Automatica (2018) 342–353.

[8] T. Sarkar, A. Rakhlin, Near optimal finite time identification of arbitrary linear dynamical systems, in: International Conference on Machine Learning, PMLR, 2019, pp. 5610–5618.

[9] Y. Jedra, A. Proutiere, Sample complexity lower bounds for linear system identification, in: 2019 IEEE 58th Conference on Decision and Control, CDC, IEEE, 2019, pp. 2676–2681.

[10] S. Oymak, N. Ozay, Non-asymptotic identification of LTI systems from a single trajectory, in: 2019 American Control Conference, ACC, IEEE, 2019, pp. 5655–5661.

[11] A. Tsiamis, G.J. Pappas, Finite sample analysis of stochastic system identification, in: 2019 IEEE 58th Conference on Decision and Control, CDC, IEEE, 2019, pp. 3648–3654.

[12] Y. Zheng, N. Li, Non-asymptotic identification of linear dynamical systems using multiple trajectories, IEEE Control Syst. Lett. 5 (5) (2020) 1693–1698.

[13] B. Lee, A. Lamperski, Non-asymptotic closed-loop system identification using autoregressive processes and hankel model reduction, in: 2020 59th IEEE Conference on Decision and Control, CDC, IEEE, 2020, pp. 3419–3424.

[14] M.C. Campi, E. Weyer, Guaranteed non-asymptotic confidence regions in system identification, Automatica (2005) 1751–1764.

[15] M.C. Campi, S. Ko, E. Weyer, Non-asymptotic confidence regions for model parameters in the presence of unmodelled dynamics, Automatica 45 (2009) 2175–2186.

[16] M.C. Campi, E. Weyer, Non-asymptotic confidence sets for the parameters of linear transfer functions, IEEE Trans. Autom. Control 55 (2010) 2708–2720.

[17] B.C. Csáji, M.C. Campi, E. Weyer, Sign-Perturbed Sums: A new system identification approach for constructing exact non-Asymptotic Confidence Regions in linear regression models, IEEE Trans. Signal Process. 63 (1) (2015) 169–181.

[18] V. Volpe, B.C. Csáji, A. Carè, E. Weyer, M.C. Campi, Sign-Perturbed Sums (SPS) with instrumental variables for the identification of ARX systems, in: 54th IEEE Conference on Decision and Control (CDC), Osaka, Japan, 2015, pp. 2115–2120.

[19] B.C. Csáji, E. Weyer, Closed-loop applicability of the Sign-Perturbed Sums method, in: 54th IEEE Conference on Decision and Control, Osaka, Japan, 2015, pp. 1441–1446.

[20] B.C. Csáji, M.C. Campi, E. Weyer, Sign-Perturbed Sums (SPS): A method for constructing exact finite-Sample Confidence Regions for general linear systems, in: 51st IEEE Conference on Decision and Control, Maui, Hawaii, 2012, pp. 7321–7326.

[21] M. Kieffer, E. Walter, Guaranteed characterization of exact non-asymptotic confidence regions as defined by LSCR and SPS, Automatica 50 (2) (2014) 507–512.

[22] E. Weyer, M.C. Campi, B.C. Csáji, Asymptotic properties of SPS confidence regions, Automatica 82 (2017) 287–294.

[23] S. Kolumbán, I. Vajk, J. Schoukens, Perturbed datasets methods for hypothesis testing and structure of corresponding confidence sets, Automatica 51 (2015) 326–331.

[24] A. Carè, M.C. Campi, B.C. Csáji, E. Weyer, Facing undermodelling in Sign-Perturbed-Sums system identification, Systems Control Lett. 153 (2021).

[25] T. Söderström, P. Stoica, System Identification, Prentice Hall International, Hertfordshire, UK, 1989.

[26] B. Giacomo, A. Carè, G. Pillonetto, Finite-sample guarantees for state-space system identification under full state measurements, in: 61st IEEE Conference on Decision and Control (CDC), Cancún, Mexico, 2022, pp. 2789–2794.

[27] U. Forssell, L. Ljung, Closed-loop identification revisited, Automatica 35 (7) (1999) 1215–1241.

[28] S. Diamond, S. Boyd, CVXPY: A python-embedded modeling language for convex optimization, J. Mach. Learn. Res. 17 (83) (2016) 1–5.

[29] A.N. Shiryaev, Probability, second ed., Springer, 1995.