Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Least-Squares Reconstruction of Floquet-Form Surrogates for Discrete-Time LTP Systems

Onurcan Yilmaz

onurcan.yilmaz@hacettepe.edu.tr

Hacettepe University

Ismail Uyanik

Hacettepe University

Research Article

Keywords: Linear time-periodic systems, Floquet theory, System identification, Least-squares estimation, Monodromy matrix

Posted Date: November 12th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11242964/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Least-Squares Reconstruction of Floquet-Form Surrogates for Discrete-Time LTP Systems

Onurcan Yilmaz[1*] and Ismail Uyanik[1]

[1]Department of Electrical and Electronics Engineering, Hacettepe University, Ankara, 06800, Turkiye.

*Corresponding author(s). E-mail(s): onurcan.yilmaz@hacettepe.edu.tr;
Contributing authors: ismailuyanik@hacettepe.edu.tr;

Abstract

Ill-conditioned periodic transformations can hinder dense Floquet reconstruction for discrete-time linear time-periodic systems. We construct a surrogate using a fixed real monodromy root and sequential phase-wise least squares to fit periodic input and output maps from full-state and input data. Exact input–output equivalence is not imposed. Under the stated excitation, independence, and moment assumptions, the population input regression targets the original periodic input matrices. Completion requires finite predictions with averaged deviation-normalized error below one. Tests on a structured commuting ensemble with analytical roots, dimensions **2**–**8**, and periods **3**–**501** show that the surrogate completes every trial. The dense canonical reference has lower conditional median prediction errors but completes only **1.3%** of noiseless trials at $\boldsymbol{n = 8}$, $\boldsymbol{P = 501}$. For a shared surrogate-designed gain at $\boldsymbol{n = 8}$, $\boldsymbol{P = 251}$, boundary estimates within **30%** of the numerically estimated true-system boundary are returned in **92.6%** of surrogate trials versus **2.6%** of reference trials. These results identify an accuracy–availability trade-off for the tested implementations.

Keywords: Linear time-periodic systems, Floquet theory, System identification, Least-squares estimation, Monodromy matrix

1 Introduction

Linear time-periodic (LTP) models arise in rotorcraft, power electronics, locomotion, and robotics [1–4]. When a real Floquet representation exists, it replaces the periodic state matrix by a constant latent matrix while retaining periodic input and output maps. Discrete-time existence and construction results [5, 6], Floquet–Fourier realizations [7], and periodic-systems theory [8] provide the foundations. Practical use also depends on accurate computation of the periodic maps.

Related identification approaches include periodic subspace methods [9–11], data-driven Floquet analysis [3], harmonic transfer-function methods for locomotion [4, 12], and regularized least squares for general time-varying systems [13]. Adaptive filtering, recursive least squares, and block-oriented identification address related signal-processing problems [14–16]. Periodic Schur and structure-preserving algorithms provide tools for periodic-system computations [17, 18]; nevertheless, ill-conditioned transformations can make

canonical input-map recovery sensitive to perturbations. Here we examine whether useful approximate models can be constructed without forming those transformations.

We fix a real Pth monodromy root $\bar{\boldsymbol{A}}$ and estimate periodic input and output maps from full-state and input data. Rather than identifying the complete periodic state-space model, we isolate map reconstruction once the constant latent dynamics are fixed. Sequential phase-wise least squares estimates input maps, propagates a latent trajectory, and fits output maps. The fitted maps need not define an invertible coordinate transformation: the resulting Floquet-form *surrogate* generally has a nonzero output residual, and exact input–output equivalence is not imposed.

Under open-loop i.i.d. zero-mean excitation and the independence and moment assumptions of Theorem 1, the input regression targets the original periodic input matrices rather than their canonical counterparts. We analyze transformation conditioning and compare the surrogate against a dense canonical recursion supplied with the true periodic matrices and the same analytical root. On a structured commuting ensemble, we separate completion from conditional prediction accuracy and evaluate numerical stability-boundary estimates for a shared surrogate-designed gain. The results characterize an accuracy–availability trade-off for the tested implementations. Relative to the preliminary procedure in [19], this work adds the surrogate formulation, population analysis, conditioning study, extended Monte Carlo evaluation, and closed-loop assessment.

2 Problem formulation

2.1 System model and monodromy

Consider the discrete-time LTP system

$$\boldsymbol{x}_{k+1} = \boldsymbol{A}_k \boldsymbol{x}_k + \boldsymbol{B}_k \boldsymbol{u}_k, \qquad \boldsymbol{y}_k = \boldsymbol{x}_k, \tag{1}$$

where $\boldsymbol{x}_k \in \mathbb{R}^n$, $\boldsymbol{u}_k \in \mathbb{R}^m$, and $k \geq 0$. Each $\boldsymbol{A}_k$ is nonsingular, and

$$\boldsymbol{A}_{k+P} = \boldsymbol{A}_k, \qquad \boldsymbol{B}_{k+P} = \boldsymbol{B}_k. \tag{2}$$

The phase $p = k \bmod P$ identifies one of the P matrix pairs. Full-state measurements are assumed; disturbances are introduced in Section 3. The formulation permits arbitrary m, whereas the experiments use $m = 1$.

The monodromy matrix is

$$\boldsymbol{\Phi} = \boldsymbol{A}_{P-1} \boldsymbol{A}_{P-2} \cdots \boldsymbol{A}_0. \tag{3}$$

It is nonsingular, and its eigenvalues are the characteristic multipliers. The unforced system is asymptotically stable when $\rho(\boldsymbol{\Phi}) < 1$. If $\|\boldsymbol{A}_k\|_2 \leq \gamma < 1$ at every phase, then $\|\boldsymbol{\Phi}\|_2 \leq \gamma^P$; the ensemble of Section 5.1 satisfies $\|\boldsymbol{\Phi}\|_2 \leq 0.9^P$.

2.2 Classical Floquet representation

Fix a real matrix $\bar{\boldsymbol{A}} \in \mathbb{R}^{n \times n}$ satisfying

$$\bar{\boldsymbol{A}}^P = \boldsymbol{\Phi}. \tag{4}$$

The discrete-time Floquet construction [5, 6, 8] gives nonsingular, P-periodic transformations satisfying

$$\boldsymbol{T}_{k+1} \bar{\boldsymbol{A}} = \boldsymbol{A}_k \boldsymbol{T}_k, \qquad \boldsymbol{T}_{k+P} = \boldsymbol{T}_k. \tag{5}$$

With $\boldsymbol{z}_k = \boldsymbol{T}_k^{-1} \boldsymbol{x}_k$, the canonical representation is

$$\boldsymbol{z}_{k+1} = \bar{\boldsymbol{A}} \boldsymbol{z}_k + \boldsymbol{B}_k^{\mathrm{F}} \boldsymbol{u}_k, \qquad \boldsymbol{y}_k = \boldsymbol{C}_k^{\mathrm{F}} \boldsymbol{z}_k, \tag{6}$$

where

$$\boldsymbol{B}_k^{\mathrm{F}} = \boldsymbol{T}_{k+1}^{-1} \boldsymbol{B}_k, \qquad \boldsymbol{C}_k^{\mathrm{F}} = \boldsymbol{T}_k. \tag{7}$$

A conventional recursive construction is

$$\boldsymbol{T}_{k+1} = \boldsymbol{A}_k \boldsymbol{T}_k \bar{\boldsymbol{A}}^{-1}, \qquad \boldsymbol{T}_0 = \boldsymbol{I}_n. \tag{8}$$

Its numerical sensitivity depends on the transformations and the implementation, as distinguished below; periodic Schur methods [17] reorganize the same computation through orthogonal factors.

2.3 Conditioning of the transformation

Define $\boldsymbol{\Psi}_k = \boldsymbol{A}_{k-1} \cdots \boldsymbol{A}_0$ for $k \geq 1$ and $\boldsymbol{\Psi}_0 = \boldsymbol{I}_n$. The closed form

$$\boldsymbol{T}_k = \boldsymbol{\Psi}_k \bar{\boldsymbol{A}}^{-k} \tag{9}$$

satisfies the recursion. Since $\boldsymbol{\Psi}_{k+P} = \boldsymbol{\Psi}_k \boldsymbol{\Phi}$, (4) also gives $\boldsymbol{T}_{k+P} = \boldsymbol{\Psi}_k \boldsymbol{\Phi} \bar{\boldsymbol{A}}^{-(k+P)} = \boldsymbol{T}_k$.

Throughout, $\mathrm{cond}(\boldsymbol{T}_k) = \|\boldsymbol{T}_k\|_2\|\boldsymbol{T}_k^{-1}\|_2$ denotes the 2-norm condition number. Large values make the inverse in (7) sensitive to perturbations. When $\varepsilon\,\mathrm{cond}(\boldsymbol{T}_k)$ approaches one, with unit roundoff $\varepsilon \approx 10^{-16}$ in double precision, generic normwise bounds no longer guarantee accuracy. This does not imply failure of every implementation. Section 5.1 evaluates the conditioning analytically for the structured ensemble, independently of the computed transformation recursion.

2.4 Monodromy root

Every real root of the nonsingular $\boldsymbol{\Phi}$ supports (9). Different roots can yield different surrogate dynamics, so the branch convention must be fixed. General-purpose algorithms for fractional matrix powers are available [20, 21]; here the root is instead constructed analytically in the modal basis of the ensemble (Section 5.1), so that root computation is not a source of variation between the two reconstructions, and the commutator residual

$$\max_{0\le p<P} \frac{\|\bar{\boldsymbol{A}}\boldsymbol{A}_p - \boldsymbol{A}_p\bar{\boldsymbol{A}}\|_F}{\|\boldsymbol{A}_p\|_F} \tag{10}$$

is reported as a consistency check of that construction. A real root exists for every real nonsingular $\boldsymbol{\Phi}$ when P is odd; for even P, existence depends on the Jordan structure at negative real eigenvalues. We use odd periods. Alternatively, a $2P$-periodic description has monodromy $\boldsymbol{\Phi}^2$ and root equation $\bar{\boldsymbol{A}}^{2P} = \boldsymbol{\Phi}^2$.

2.5 Surrogate objective

Superscripts F and S denote canonical and surrogate quantities, respectively. For fixed $\bar{\boldsymbol{A}}$, the surrogate has P-periodic input and output maps and obeys

$$\boldsymbol{z}_{k+1} = \bar{\boldsymbol{A}}\boldsymbol{z}_k + \boldsymbol{B}_k^{\mathrm{S}}\boldsymbol{u}_k, \qquad \hat{\boldsymbol{y}}_k = \boldsymbol{C}_k^{\mathrm{S}}\boldsymbol{z}_k. \tag{11}$$

The output residual $\boldsymbol{r}_k = \boldsymbol{y}_k - \hat{\boldsymbol{y}}_k$ is generally nonzero. Given measured outputs $\boldsymbol{y}_0, \dots, \boldsymbol{y}_{N-1}$ and inputs $\boldsymbol{u}_0, \dots, \boldsymbol{u}_{N-2}$, Section 4 constructs the maps by sequential phase-wise least squares without forming $\{\boldsymbol{T}_k\}$. The output fit approximates the measured trajectory for the applied input and prescribed latent initialization. The matrices need not equal their canonical counterparts, and exact input-output equivalence is not imposed. A canonical representation gives zero residual for noiseless data and consistent initialization.

3 Algebraic analysis of the regression

For noiseless data, the canonical residual $\boldsymbol{T}_{k+1}^{-1}\boldsymbol{y}_{k+1} - \bar{\boldsymbol{A}}\boldsymbol{T}_k^{-1}\boldsymbol{y}_k$ equals $\boldsymbol{B}_k^{\mathrm{F}}\boldsymbol{u}_k$ by (5). This oracle relation requires the transformations; the implemented regression instead uses measured outputs directly.

Consider the noisy model

$$\boldsymbol{x}_{k+1} = \boldsymbol{A}_k\boldsymbol{x}_k + \boldsymbol{B}_k\boldsymbol{u}_k + \boldsymbol{w}_k, \quad \boldsymbol{y}_k = \boldsymbol{x}_k + \boldsymbol{v}_k, \tag{12}$$

where $\boldsymbol{w}_k$ and $\boldsymbol{v}_k$ are zero-mean process and measurement disturbances. With $\bar{\boldsymbol{A}}$ fixed, define

$$\boldsymbol{e}_k := \boldsymbol{y}_{k+1} - \bar{\boldsymbol{A}}\boldsymbol{y}_k. \tag{13}$$

Substitution gives

$$\boldsymbol{e}_k = (\boldsymbol{A}_k - \bar{\boldsymbol{A}})\boldsymbol{x}_k + \boldsymbol{B}_k\boldsymbol{u}_k + \boldsymbol{\eta}_k, \tag{14}$$

with

$$\boldsymbol{\eta}_k := \boldsymbol{w}_k + \boldsymbol{v}_{k+1} - \bar{\boldsymbol{A}}\boldsymbol{v}_k. \tag{15}$$

For each phase $p \in \{0, \dots, P-1\}$, use the transition indices

$$\mathcal{I}_p := \{\, 0 \le k \le N-2 : k \bmod P = p \,\}. \tag{16}$$

Let $N_p = |\mathcal{I}_p|$ and collect the corresponding columns as

$$\boldsymbol{E}_p := [\,\boldsymbol{e}_k\,]_{k\in\mathcal{I}_p}, \qquad \boldsymbol{U}_p := [\,\boldsymbol{u}_k\,]_{k\in\mathcal{I}_p}. \tag{17}$$

Step 1 estimates

$$\hat{\boldsymbol{B}}_p^{\mathrm{S}} = \boldsymbol{E}_p\boldsymbol{U}_p^{\dagger}, \tag{18}$$

where $\dagger$ denotes the Moore–Penrose pseudoinverse. The term $(\boldsymbol{A}_p - \bar{\boldsymbol{A}})\boldsymbol{x}_k$ is omitted from this regression. Its correlation with the input, rather than regressor-noise bias addressed by adaptive compensation schemes [22, 23], determines its effect on the population target.

Theorem 1 (population target of Step 1) *Fix $\bar{\boldsymbol{A}}$. Assume that $\{\boldsymbol{u}_k\}$ is i.i.d., zero mean, with covariance $\boldsymbol{\Sigma}_u := E[\boldsymbol{u}_k\boldsymbol{u}_k^\top] \succ 0$, and that the entire input*

sequence is independent of the joint collection of $\boldsymbol{x}_0$, $\{\boldsymbol{w}_\ell\}$ *and* $\{\boldsymbol{v}_\ell\}$. *Assume also that the state is generated causally from past inputs and disturbances, and that* $\boldsymbol{x}_k$ *and* $\boldsymbol{\eta}_k$ *have finite second moments. At any fixed index* $k \in \mathcal{I}_p$, *let* $E_p[\cdot]$ *denote expectation over trajectories. The population cost*

$$J_p(\boldsymbol{M}) := E_p\left[\|\boldsymbol{e}_k - \boldsymbol{M}\boldsymbol{u}_k\|_2^2\right] \tag{19}$$

has the unique minimizer $\boldsymbol{M}_p^\star = \boldsymbol{B}_p$ *over* $\boldsymbol{M} \in \mathbb{R}^{n\times m}$.

Proof Causality and joint independence make both $\boldsymbol{x}_k$ and $\boldsymbol{\eta}_k$ independent of the zero-mean current input. Hence $E_p[\boldsymbol{x}_k\boldsymbol{u}_k^\top] = E_p[\boldsymbol{\eta}_k\boldsymbol{u}_k^\top] = \boldsymbol{0}$. Using (14),

$$E_p[\boldsymbol{e}_k\boldsymbol{u}_k^\top] = (\boldsymbol{A}_p - \bar{\boldsymbol{A}})E_p[\boldsymbol{x}_k\boldsymbol{u}_k^\top] + \boldsymbol{B}_p\boldsymbol{\Sigma}_u + E_p[\boldsymbol{\eta}_k\boldsymbol{u}_k^\top] = \boldsymbol{B}_p\boldsymbol{\Sigma}_u.$$

The normal equations therefore give $\boldsymbol{M}_p^\star = E_p[\boldsymbol{e}_k\boldsymbol{u}_k^\top]\boldsymbol{\Sigma}_u^{-1} = \boldsymbol{B}_p$; positive definiteness of $\boldsymbol{\Sigma}_u$ ensures uniqueness. □

Thus Step 1 targets the original input matrices, not the canonical $\boldsymbol{B}_p^{\mathrm{F}}$. This is a population characterization, not a finite-sample unbiasedness or consistency result. Zero-mean i.i.d. Gaussian inputs and independent equiprobable binary (Rademacher) inputs with nonsingular covariance satisfy the excitation assumptions. Temporal correlation or feedback can violate the required state–input orthogonality [24]; independent white excitation is sufficient, not universally necessary. Disturbances affect finite-sample accuracy without changing the stated population target, and are not premultiplied by inverse Floquet transformations.

4 Floquet reconstruction methods

Both methods use the same monodromy root $\bar{\boldsymbol{A}}$, constructed analytically as described in Section 2.4.

4.1 Direct canonical reference

The reference uses the true periodic matrices $\{\boldsymbol{A}_k, \boldsymbol{B}_k\}$, evaluates (8), and obtains the canonical maps from (7). This is an oracle-information reference, not an accuracy bound: the surrogate receives only the shared root and measured data. Canonical equivalence holds in exact arithmetic but does not guarantee finite-precision accuracy.

4.2 Proposed least-squares method

The three-step procedure of [19] avoids constructing $\{\boldsymbol{T}_k\}$.

Step 1: input maps. Assemble $\boldsymbol{E}_p$ and $\boldsymbol{U}_p$ using (17), and estimate $\hat{\boldsymbol{B}}_p^{\mathrm{S}}$ from (18). Theorem 1 characterizes its population target.

Step 2: latent propagation. For $0 \le k \le N-2$, propagate

$$\hat{\boldsymbol{z}}_0 = \boldsymbol{y}_0, \qquad \hat{\boldsymbol{z}}_{k+1} = \bar{\boldsymbol{A}}\hat{\boldsymbol{z}}_k + \hat{\boldsymbol{B}}^{\mathrm{S}}_{k \bmod P}\boldsymbol{u}_k. \tag{20}$$

This measured-state initialization is the convention used throughout the experiments.

Step 3: output maps. Use all output indices $\mathcal{J}_p := \{\, 0 \le k \le N-1 : k \bmod P = p \,\}$, and collect $\boldsymbol{Y}_p = [\,\boldsymbol{y}_k\,]_{k\in\mathcal{J}_p}$ and $\boldsymbol{Z}_p = [\,\hat{\boldsymbol{z}}_k\,]_{k\in\mathcal{J}_p}$. Then

$$\hat{\boldsymbol{C}}_p^{\mathrm{S}} = \boldsymbol{Y}_p\boldsymbol{Z}_p^\dagger. \tag{21}$$

These maps minimize the phase-wise sum of squared output errors for the fixed latent record, but need not be invertible or equal $\boldsymbol{C}_p^{\mathrm{F}}$. The fitted residual is generally nonzero.

4.3 Rank conditions and cost

Unique unregularized fits require each $\boldsymbol{U}_p$ to have row rank m and each $\boldsymbol{Z}_p$ to have row rank n for Steps 1 and 3, respectively. The pseudoinverse formulas still define minimizers when these conditions fail. A long record and Gaussian excitation alone do not guarantee full latent rank.

For a fixed root in the single-input, overdetermined setting considered here, the dense reference requires $O(Pn^3)$ operations, whereas the surrogate costs $O(Nn^2)$, including residual evaluation, latent propagation, and phase-wise regressions. In the experiments, $N-1 = PN_p$ counts transitions; Step 3 uses all N output samples. These operation counts do not establish a measured speed advantage.

5 Monte Carlo framework

5.1 Structured random ensemble

We use the structured LTP ensemble of [19]. For each realization, QR factorization of a Gaussian random matrix provides an orthogonal $\boldsymbol{F}$, and $\boldsymbol{A}_k = \boldsymbol{F}\boldsymbol{\Lambda}_k\boldsymbol{F}^\top$. The fixed block structure

of $\mathbf{\Lambda}_k$ comprises real scalar modes and planar rotation–scaling modes. Their moduli vary as

$$r_{i,k} = 0.6 + a_i \sin(2\pi k/P + \varphi_i),$$

where $a_i \sim \mathcal{U}[0, 0.3]$ and $\varphi_i \sim \mathcal{U}[0, 2\pi)$ are drawn once per realization. Rotation angles also vary sinusoidally with phase. Thus $r_{i,k} \in [0.3, 0.9]$, $\|\boldsymbol{A}_k\|_2 \leq 0.9$, and $\rho(\mathbf{\Phi}) < 1$. The single-input maps are

$$\boldsymbol{B}_k = 0.2 \sin \boldsymbol{\phi} + 1.3 \sin(2\pi k/P - \boldsymbol{\phi}),$$

with entrywise sine and $\boldsymbol{\phi} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I}_n)$ drawn once per realization. Their phase average is $0.2 \sin \boldsymbol{\phi}$, which is generically nonzero.

Both methods receive the same real root, constructed analytically in the modal basis from the geometric-mean moduli $\rho_i = (\prod_{p=0}^{P-1} r_{i,p})^{1/P}$ and accumulated modal angles. It commutes with every $\boldsymbol{A}_k$ by construction; the residual in (10) is of order 10^{-16}. This comparison is conditional on the analytically supplied root, rather than on numerical extraction of a root from the dense monodromy.

For this commuting ensemble, (9) has the form $\boldsymbol{T}_k = \boldsymbol{F}\,\mathrm{blockdiag}_i(d_{i,k}\boldsymbol{Q}_{i,k})\boldsymbol{F}^\top$, where each $\boldsymbol{Q}_{i,k}$ is orthogonal within its modal block and

$$d_{i,k} = \rho_i^{-k} \prod_{j=0}^{k-1} r_{i,j}, \qquad \mathrm{cond}(\boldsymbol{T}_k) = \frac{\max_i d_{i,k}}{\min_i d_{i,k}}.$$

For each realization, we calculate $\max_{0 \leq k < P} \log_{10} \mathrm{cond}(\boldsymbol{T}_k)$ analytically and report its median at each grid point, pooling the two noise conditions. This structured family permits conditioning to be assessed independently of the dense recursion; it does not represent arbitrary noncommuting LTP systems.

5.2 Trajectory generation and evaluation

Each realization has independent training and test trajectories, both starting from $\boldsymbol{x}_0 = \mathbf{0}$ and driven by i.i.d. Gaussian inputs. Each record contains $N = 100P + 1$ output samples. The $100P$ transitions give $N_p = 100$ samples per phase in Step 1; Step 3 uses 101 output samples at phase zero and 100 at every other phase. These fits are overdetermined, and prediction accuracy is assessed only on the independent test record.

Training uses the noisy model (12), with i.i.d. Gaussian disturbances $\boldsymbol{w}_k, \boldsymbol{v}_k \sim \mathcal{N}(\mathbf{0}, \sigma^2 \boldsymbol{I}_n)$, independent of each other and of the input, at $\sigma^2 \in \{0, 0.1\}$. Test trajectories remain noiseless to isolate deterministic reconstruction error. For a fixed system and test input, the reference is unaffected by training noise. Independent system realizations and test inputs are drawn for the two noise conditions, so differences between their reference statistics reflect Monte Carlo variation.

The grid comprises $n \in \{2, \ldots, 8\}$ and

$$P \in \{3, 13, 51, 101, 145, 175, 201, 251, 351, 501\},$$

with 1000 realizations per configuration and noise level, giving 70 configurations at each noise level.

5.3 Outcome categories and metrics

The analytical root is supplied to both methods, so completion is assessed at the reconstruction stage. For each test-output component, we divide the RMS prediction error by its standard deviation about its test-record mean, using divisor N in both quantities. The componentwise ratios are averaged over components with standard deviation greater than 10^{-10}, yielding $\mathrm{NRMSE}_{\mathrm{dev}}$. The range-normalized values in the conference version [19] are not directly comparable.

A method *completes* a trial when its prediction is finite and $\mathrm{NRMSE}_{\mathrm{dev}} < 1$. The componentwise mean predictor has score one, so completion requires more than a finite numerical output. We report completion rates over all attempted trials and conditional median errors over completed trials, keeping availability separate from prediction accuracy on the successful subset.

6 Results

6.1 Completion

The median phase-maximum logarithmic condition statistic defined in Section 5.1 grows approximately linearly with P for $n \geq 3$ ($R^2 > 0.9998$). At $n = 8$, $P = 501$, it reaches 40.37, corresponding to a condition number of approximately $10^{40.37}$. At $n = 2$, 47.6% of realizations contain a single modal block and have $\mathrm{cond}(\boldsymbol{T}_k) =$

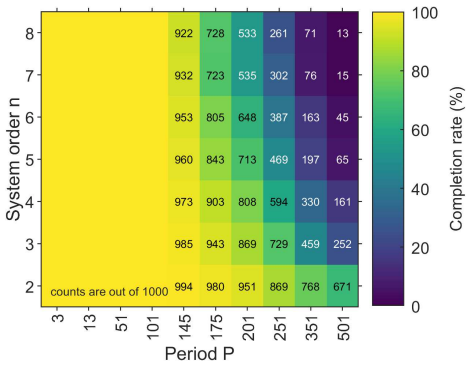


Fig. 1 Noiseless completion rate of the dense canonical reference under $\mathrm{NRMSE}_{\mathrm{dev}} < 1$. Annotations show completed trials out of 1000. The surrogate completes every trial and is omitted.

1 throughout the period. These outcomes characterize the selected ensemble, not a universal numerical failure threshold.

The surrogate completes every trial at both noise levels. The dense reference completes all trials through $P = 101$, but completion decreases at longer periods (Fig. 1). At $n = 8$, $P = 501$, it completes 13/1000 noiseless trials (1.3%) and 12/1000 trials at $\sigma^2 = 0.1$ (1.2%). By contrast, its noiseless completion rate at $n = 2$, $P = 501$ is 67.1%. Reference completion falls below 50% in sixteen of the seventy noiseless configurations.

6.2 Conditional accuracy

At $\sigma^2 = 0.1$, conditional median test errors range from 7.1×10^{-16} to 2.3×10^{-2} for the reference and 9.8×10^{-2} to 3.0×10^{-1} for the surrogate (Fig. 2). The reference has lower conditional medians in every configuration at both noise levels. This does not imply superiority on every successful trial: among paired trials where both methods complete, it is more accurate in 66–100% of cases per noiseless configuration. Small reference medians must also be interpreted alongside the limited completion counts.

Surrogate median errors vary by at most about a factor of three across periods, with $N_p = 100$ transitions per phase. At $\sigma^2 = 0.1$, errors are generally lower at larger state dimensions; the median at $n = 2$, $P = 501$ reaches 0.30. Training noise shifts the lower end of the surrogate median-error range from 1.3×10^{-2} to 9.8×10^{-2}.

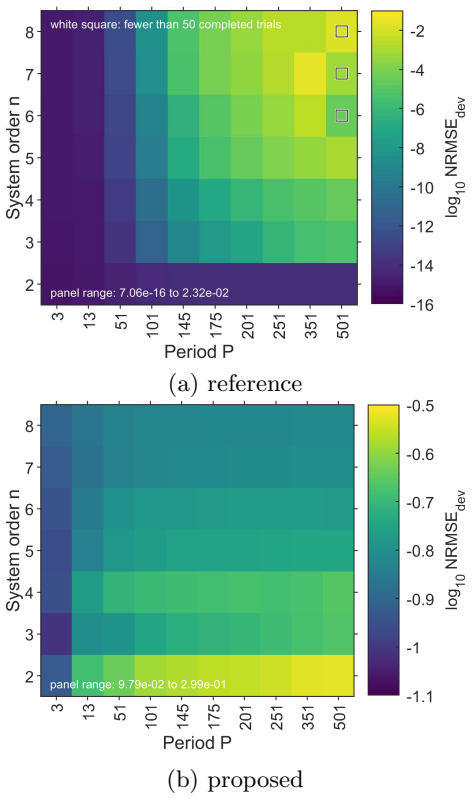


Fig. 2 Conditional median test $\mathrm{NRMSE}_{\mathrm{dev}}$ over completed trials at $\sigma^2 = 0.1$. The separate $\log_{10} \mathrm{NRMSE}_{\mathrm{dev}}$ color scales span $[-16, -1]$ in (a) and $[-1.1, -0.5]$ in (b). White squares in (a) mark fewer than fifty completed trials.

Finite-sample error is not the only possible contribution. For an input applied $j \geq 1$ samples earlier, the surrogate propagates it through $\bar{\boldsymbol{A}}^{j-1}$ between its input and output maps, rather than the true product of $j - 1$ phase-dependent state matrices. A fitted output map need not match these coefficients at every lag, so structural approximation error can persist with more training data.

6.3 Closed-loop stability transfer

To complement open-loop prediction, we compare the models' estimates of the stability boundary of one shared feedback gain. Related LTP controller design is considered in [25]. After open-loop identification, $\boldsymbol{K}$ is obtained by unit-weight Riccati iteration on $(\bar{\boldsymbol{A}}, P^{-1} \sum_{p=0}^{P-1} \hat{\boldsymbol{B}}_p^{\mathrm{S}})$. This averaged-model heuristic uses neither the true system nor the canonical reference. Applying $\boldsymbol{u}_k = -\alpha \boldsymbol{K} \boldsymbol{y}_k$

through each output map gives

$$\begin{aligned}\boldsymbol{F}_k^{\mathrm{true}}(\alpha) &= \boldsymbol{A}_k - \alpha \boldsymbol{B}_k \boldsymbol{K},\\ \boldsymbol{F}_k^{\mathrm{S}}(\alpha) &= \bar{\boldsymbol{A}} - \alpha \hat{\boldsymbol{B}}_k^{\mathrm{S}} \boldsymbol{K} \hat{\boldsymbol{C}}_k^{\mathrm{S}},\\ \boldsymbol{F}_k^{\mathrm{F}}(\alpha) &= \bar{\boldsymbol{A}} - \alpha \boldsymbol{B}_k^{\mathrm{F}} \boldsymbol{K} \boldsymbol{C}_k^{\mathrm{F}}.\end{aligned}$$

For each loop, let $\boldsymbol{\Phi}_{\mathrm{cl}}(\alpha) = \boldsymbol{F}_{P-1}(\alpha)\cdots\boldsymbol{F}_0(\alpha)$. When the loop is stable at $\alpha = 0$, its first stability boundary is

$$\alpha^\star := \inf\{\alpha > 0 : \rho(\boldsymbol{\Phi}_{\mathrm{cl}}(\alpha)) \geq 1\}. \tag{22}$$

The reported search uses a 71-point logarithmic grid up to $\alpha = 200$ and 45 bisection steps to refine the first detected stable-to-unstable bracket. These numerical estimates can miss earlier unstable intervals; all reported margins below refer to this search, not certified first boundaries. Products are evaluated using stepwise rescaling and logarithmic accumulation to limit underflow, without eliminating other conditioning effects. We use $n \in \{4, 8\}$, $P \in \{51, 251\}$, and 500 realizations per configuration. Errors compare the numerical model and true-system boundaries for the same gain.

Canonical transformations satisfy

$$\bar{\boldsymbol{A}} - \alpha \boldsymbol{B}_k^{\mathrm{F}} \boldsymbol{K} \boldsymbol{C}_k^{\mathrm{F}} = \boldsymbol{T}_{k+1}^{-1}(\boldsymbol{A}_k - \alpha \boldsymbol{B}_k \boldsymbol{K})\boldsymbol{T}_k. \tag{23}$$

Telescoping over one period makes the canonical and true closed-loop monodromies similar in exact arithmetic. The surrogate has no such guarantee, and finite-precision canonical reconstruction need not preserve this equivalence.

Table 1 and Fig. 3 summarize the results. At $P = 251$, the reference returns no margin in 268/500 realizations for $n = 4$ and 453/500 for $n = 8$, owing to failed or numerically degenerate transformations. Other tracked outcomes –no detected crossing within the search range, failed eigenvalue evaluation, and initial instability– were not observed. At $n = 8$, $P = 251$, a margin within 30% of the true-system estimate is returned in 2.6% of reference trials versus 92.6% of surrogate trials.

The surrogate's reported median margin-error magnitudes are approximately 14–20%. It overpredicts the true-system boundary estimate in 9.0% of the 2000 realizations, with a largest

Table 1 Boundary estimates for the shared surrogate-designed gain. Percentages use all 500 attempted realizations per configuration; "within 30%" denotes absolute relative error at most 0.30. Prop. is the surrogate and Ref. the dense canonical reference.

	$n = 4$		$n = 8$	
Period P	51	251	51	251
Within 30%, Prop.	95.6%	97.8%	95.8%	92.6%
Within 30%, Ref.	100%	25.6%	100%	2.6%
No margin, Ref.	0	268	0	453

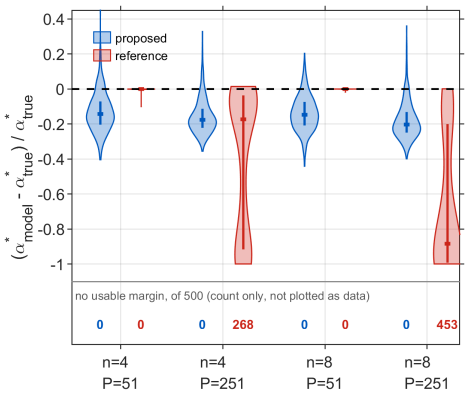


Fig. 3 Signed relative errors $(\alpha^\star_{\mathrm{model}} - \alpha^\star_{\mathrm{true}})/\alpha^\star_{\mathrm{true}}$ of numerical boundary estimates from 500 attempted realizations per configuration. Positive values over-predict the true-system estimate. Only finite returned estimates enter the violins; unavailable estimates are counted below.

over-prediction of 47%; 4.6% exceed 30% absolute relative error. No over-prediction is reported among finite reference estimates, but unavailable margins are not conservative predictions. These results concern the selected gain, ensemble, and numerical search, not a guarantee of closed-loop stability.

7 Discussion and conclusion

For a fixed real monodromy root, sequential least squares constructs a Floquet-form surrogate without forming periodic coordinate transformations. Under the assumptions of Theorem 1, Step 1 targets the original periodic input matrices, while Step 3 fits output maps for the constructed latent trajectory. These maps are not constrained to define an invertible coordinate transformation, so exact input-output equivalence is not guaranteed.

Finite-sample effects and structural approximation error can both contribute to the prediction residuals.

The experiments demonstrate an accuracy–availability trade-off. The tested dense canonical reference has lower conditional median prediction errors but loses availability in long-period configurations, whereas the surrogate meets the completion criterion throughout the tested grid. For the shared surrogate-designed gain, its numerical stability-boundary estimates remain available, but occasional over-prediction prevents interpreting them as conservative guarantees. Open-loop prediction accuracy therefore does not replace validation of the resulting closed-loop system.

The analysis assumes full-state measurements and open-loop independent white excitation. Evaluation uses odd periods, analytically supplied roots, and a structured commuting ensemble. It neither tests root identification from noisy data nor establishes superiority over structure-preserving canonical algorithms. The finite-grid search also does not certify the first stability boundary. Extensions to noncommuting systems and estimated roots would assess broader applicability. Within the present scope, sequential least squares provides an approximate alternative when the tested dense reconstruction fails, rather than a replacement for exact Floquet coordinates.

Statements and Declarations

Funding. This work was supported by the Scientific and Technological Research Council of Türkiye (TÜBİTAK) under Grant No. 124E232 awarded to I.U. and by its 2211-A National PhD Scholarship Programme awarded to O.Y.

Author contributions. O.Y.: methodology, software, validation, and writing original draft. I.U.: conceptualization, supervision, funding acquisition, and writing review and editing. Both authors read and approved the final manuscript.

Use of large language models. Claude (Anthropic) assisted with literature searches, LaTeX editing, prose refinement, and MATLAB code development. ChatGPT (OpenAI) assisted with manuscript revision and checks of the mathematical presentation and references. The authors reviewed and edited the material and take responsibility for the final manuscript.

Competing interests. The authors declare no financial or non-financial competing interests relevant to this work.

Data and code availability. The simulation code and numerical results supporting this study will be made publicly available through an open-access data archive.

References

[1] Hwang, S.: Frequency domain system identification of helicopter rotor dynamics incorporating models with time periodic coefficients. PhD thesis, University of Maryland, College Park (1997)

[2] Mollerstedt, E., Bernhardsson, B.: A harmonic transfer function model for a diode converter train. In: 2000 IEEE Power Engineering Society Winter Meeting. Conference Proceedings (Cat. No. 00CH37077), vol. 2, pp. 957–962 (2000). https://doi.org/10.1109/PESW.2000.850062 . IEEE

[3] Revzen, S., Guckenheimer, J.M.: Finding the dimension of slow dynamics in a rhythmic system. Journal of The Royal Society Interface **9**(70), 957–971 (2012) https://doi.org/10.1098/rsif.2011.0431

[4] Uyanik, I., Ankarali, M.M., Cowan, N.J., Saranli, U., Morgul, O.: Identification of a vertical hopping robot model via harmonic transfer functions. Transactions of the Institute of Measurement and Control **38**(5), 501–511 (2016) https://doi.org/10.1177/0142331215583327

[5] Van Dooren, P., Sreedhar, J.: When is a periodic discrete-time system equivalent to a time-invariant one? Linear Algebra and its Applications **212**, 131–151 (1994) https://doi.org/10.1016/0024-3795(94)90400-6

[6] Sreedhar, J., Van Dooren, P.: Discrete-time periodic systems: a Floquet approach. In: Proceedings of the Conference on Information Sciences and Systems, pp. 194–195 (1993)

[7] Sandberg, H.: On Floquet-Fourier realizations of linear time-periodic impulse

responses. In: Proceedings of the 45th IEEE Conference on Decision and Control, pp. 1411–1416 (2006). https://doi.org/10.1109/CDC.2006.377638 . IEEE

[8] Bittanti, S., Colaneri, P.: Periodic Systems: Filtering and Control. Springer, London (2009). https://doi.org/10.1007/978-1-84800-911-0

[9] Verhaegen, M., Yu, X.: A class of subspace model identification algorithms to identify periodically and arbitrarily time-varying systems. Automatica **31**(2), 201–216 (1995) https://doi.org/10.1016/0005-1098(94)00091-V

[10] Uyanik, I., Saranli, U., Ankarali, M.M., Cowan, N.J., Morgul, O.: Frequency-domain subspace identification of linear time-periodic (LTP) systems. IEEE Transactions on Automatic Control **64**(6), 2529–2536 (2019) https://doi.org/10.1109/TAC.2018.2867360

[11] Yin, M., Iannelli, A., Smith, R.S.: Subspace identification of linear time-periodic systems with periodic inputs. IEEE Control Systems Letters **5**(1), 145–150 (2020) https://doi.org/10.1109/LCSYS.2020.3000950

[12] Yilmaz, O., Goksuluk, E.C., Uyanik, I.: Sistem tanılama yöntemlerinin zayıf elektrik balıklarının hareket dinamikleri üzerinden karşılaştırılması. Gazi Üniversitesi Mühendislik Mimarlık Fakültesi Dergisi **40**(2), 877–888 (2025) https://doi.org/10.17341/gazimmfd.1384350

[13] Carvalho, M., Soares, C., Lourenço, P., Ventura, R.: COSMIC: fast closed-form identification from large-scale data for LTV systems. arXiv preprint arXiv:2112.04355 (2021). https://doi.org/10.48550/arXiv.2112.04355

[14] Zayyani, H., Habibi, Z., Salman, M., Rashdan, M.: Robust Markovian-Brownian adaptive filter for block sparse system identification in time-varying conditions with BIBO stability analysis for mean-square convergence. Signal, Image and Video Processing **20**(7), 392 (2026) https://doi.org/10.1007/s11760-026-05449-5

[15] Naeimi-Sadigh, A.: Enhanced quaternion recursive least squares: adaptive prediction in stationary and non-stationary environments. Signal, Image and Video Processing **20**(7), 416 (2026) https://doi.org/10.1007/s11760-026-05471-7

[16] Janjanam, L., Saha, S.K., Kar, R., Mandal, D.: Wiener model-based system identification using moth flame optimised Kalman filter algorithm. Signal, Image and Video Processing **16**(5), 1425–1433 (2022) https://doi.org/10.1007/s11760-021-02096-w

[17] Bojanczyk, A.W., Golub, G.H., Van Dooren, P.: Periodic Schur decomposition: algorithms and applications. In: Advanced Signal Processing Algorithms, Architectures, and Implementations III, vol. 1770, pp. 31–42 (1992). https://doi.org/10.1117/12.130915 . SPIE

[18] Varga, A.: Computational issues for linear periodic systems: paradigms, algorithms, open problems. International Journal of Control **86**(7), 1227–1239 (2013) https://doi.org/10.1080/00207179.2013.773088

[19] Yilmaz, O., Uyanik, I.: Floquet-based periodic matrix reconstruction in discrete-time LTP systems. In: 2026 34th Signal Processing and Communications Applications Conference (SIU), pp. 1–4 (2026). https://doi.org/10.1109/SIU71813.2026.11637048

[20] Higham, N.J., Lin, L.: An improved Schur–Padé algorithm for fractional powers of a matrix and their Fréchet derivatives. SIAM Journal on Matrix Analysis and Applications **34**(3), 1341–1360 (2013) https://doi.org/10.1137/130906118

[21] Bini, D.A., Higham, N.J., Meini, B.: Algorithms for the matrix pth root. Numerical Algorithms **39**(4), 349–378 (2005) https://doi.org/10.1007/s11075-004-6709-8

[22] Rosalin, Patnaik, A., Nanda, S., Rout, D.K.: A bias-compensated nlms algorithm based on arctangent framework for system identification. Signal, Image and Video Processing **18**(4), 3595–3601 (2024) https://doi.org/10.

1007/s11760-024-03024-4

[23] Jeong, J.J., Lee, D.-Y.: Unbiased normalized least-mean-square algorithm with input and output noise. Signal, Image and Video Processing **19**(9), 708 (2025) https://doi.org/10.1007/s11760-025-04334-x

[24] Ljung, L.: System Identification: Theory for the User, 2nd edn. Prentice Hall PTR, Upper Saddle River, NJ (1999)

[25] Sert, B., Catalbas, B., Uyanik, I.: Identification and control of a linear time-periodic test bench using harmonic transfer functions and LQR controllers. In: 2023 9th International Conference on Control, Decision and Information Technologies (CoDIT), pp. 2366–2371 (2023). https://doi.org/10.1109/CoDIT58514.2023.10284205 . IEEE