
Universal Sequence Preconditioning

Annie Marsden *

Elad Hazan^{*†}

Abstract

We study the problem of preconditioning in sequential prediction. From the theoretical lens of linear dynamical systems, we show that convolving the input sequence corresponds to applying a polynomial to the hidden transition matrix. Building on this insight, we propose a universal preconditioning method that convolves the input with coefficients from orthogonal polynomials such as Chebyshev or Legendre. We prove that this approach reduces regret for two distinct prediction algorithms and yields the first ever sublinear and hidden-dimension-independent regret bounds (up to logarithmic factors) that hold for systems with marginally stable and asymmetric transition matrices. Finally, extensive synthetic and real-world experiments show that this simple preconditioning strategy improves the performance of a diverse range of algorithms, including recurrent neural networks, and generalizes to signals beyond linear dynamical systems.

1 Introduction

In sequence prediction the goal of the learner is to predict the next token accurately according to a specified loss function, such as the mean square error or cross-entropy. This fundamental problem in machine learning has gained increased importance with the rise of large language models, which perform sequence prediction on tokens using cross entropy. The focus of this paper is *preconditioning*, i.e. modifying the target sequence to make it easier to learn. A classic example is **differencing**, introduced by Box and Jenkins in the 1970s [14], which transforms observations y_1, y_2, \dots into successive differences,

$$y_1 - y_0, y_2 - y_1, \dots, y_t - y_{t-1}, \dots$$

It is widely acknowledged that learning this sequence can be “easier” than learning the original sequence for a large number of modalities. In this work we seek a more general framework for sequence preconditioning that captures the same intuition behind differencing and extends it to a broader class of transformations. The question we ask is

What is the general form of sequence preconditioning that enables provably accurate learning?

We address this question by introducing a preconditioning method which takes in n fixed coefficients c_0, \dots, c_n and converts the sequence of observations y_1, \dots, y_t, \dots to the sequence of convolved observations¹

$$c_0 y_0, c_0 y_1 + c_1 y_0, \dots, \sum_{i=0}^n c_i y_{t-i}, \dots$$

From an information-theoretic perspective, approaches of this kind seem futile—predicting y_t or $\sum_i c_i y_{t-i}$ seems equally hard in an adversarial setting. Yet we show that when the data arises from

^{*}Google Deepmind

[†]Princeton University

¹This recovers differencing when $n = 2$, $c_0 = 1$, and $c_1 = -1$.

a linear dynamical system (LDS), there exists a *universal* form of preconditioning that provably improves learnability, independent of the specific system.

In LDS setting, we show that preconditioning significantly strengthens existing prediction methods, leading to new regret bounds. Here, preconditioning has an elegant effect: the preconditioning filter forms coefficients of an n degree polynomial, and the hidden system transition matrix is evaluated on this polynomial—potentially shrinking the domain. In this setting, shrinking the learnable domain is akin to making the problem “easier to learn”, a relationship that is formalized by [28]. This allows us to prove the first dimension-independent sublinear regret bounds for asymmetric linear dynamical systems that are marginally stable.

1.1 Our results

Our main contribution is *Universal Sequence Preconditioning*, a novel method of sequence preconditioning which convolves the target sequence with the coefficients of the n -th monic Chebyshev polynomial. We give a more general form of preconditioning, allowing arbitrary user-specified coefficients, in Algorithm 1 and an online version in Algorithm 4 (Appendix C). We analyze the effect of Universal Sequence Preconditioning on two canonical sequence prediction algorithms in the online setting: (1) convex regression and (2) spectral filtering. In either case, the results are impressive—yielding the first known sublinear regret bounds as compared to the optimal ground-truth predictor that are simultaneously (1) applicable to marginally stable systems, (2) independent of hidden dimension (up to logarithmic factors), and (3) applicable to systems whose transition matrix is asymmetric² (see Table 1).

Algorithm 1 General Sequence Preconditioning (Offline Version)

```

1: Training
2: Input: training data  $(\mathbf{u}_{1:T}^{1:N}, \mathbf{y}_{1:T}^{1:N})$  where  $(\mathbf{u}_t^i, \mathbf{y}_t^i)$  is the  $t$ -th input/output pair in the  $i$ -th sequence;
   coefficients  $c_{0:n}$ ; prediction algorithm  $\mathcal{A}$ .
3: Assert  $c_0 = 1$ .
4: for  $i = 1$  to  $N$  do
5:    $\mathbf{y}_{1:T}^{\text{preconditioned}, i} \leftarrow \text{convolution}(\mathbf{y}_{1:T}^i, c_{0:n})$             $\triangleleft \mathbf{y}_t^{\text{preconditioned}, i} = \mathbf{y}_t^i + \sum_{j=1}^n c_j \mathbf{y}_{t-j}^i$ 
6: end for
7: Train  $\mathcal{A}$  on preconditioned data  $(\mathbf{u}_{1:T}^{1:N}, \mathbf{y}_{1:T}^{\text{preconditioned}, 1:N})$ .
8: Test Time
9: for  $t = 1$  to  $T$  do
10:  Receive  $\mathbf{u}_t$ .
11:  Predict  $\hat{\mathbf{y}}_t \leftarrow \mathcal{A}(\mathbf{u}_{1:t}, \mathbf{y}_{1:(t-1)}) - \sum_{i=1}^n c_i \mathbf{y}_{t-i}$ .
12:  Receive  $\mathbf{y}_t$ .
13: end for

```

First, applying USP to standard convex regression results in regret $\tilde{O}(T^{-2/13})$, which holds simultaneously across the three settings above and remains dimension-independent. For comparison, a naive analysis of regression yields a vacuous regret bound of $O(T^{5/2})$ on marginally stable systems. Second, combining USP with a variant of *spectral filtering* [29] that uses novel filters, the algorithm is able learn a broader class of linear dynamical systems—in particular systems whose hidden transition matrix may be asymmetric. The enhanced method achieves regret $\tilde{O}(T^{-3/13})$, the best known rate under the joint conditions of conditions (1)-(3) discussed above: marginal stability, dimension independence, and asymmetry.

Both results require that the transition matrix eigenvalues have imaginary parts bounded by $O(1/\log T)$ —a near-tight condition for achieving dimension-free regret. Further discussion on this appears in Appendix B.

Empirical results in Section 4 demonstrate that USP consistently improves performance across diverse algorithms—including regression, spectral filtering, and neural networks—and across data types extending beyond linear dynamical systems.

²Our results only hold for asymmetric matrices whose eigenvalues have imaginary component bounded above by $O(1/\log(T))$. This is somewhat tight, see Section B.

Method	Marginally stable	d_{hidden} -free	Asymmetric
Sys-ID	×	×	✓
Regression (open-loop)	×	✓	✓
Regression (closed-loop)	✓	×	✓
Spectral Filtering	✓	✓	×
USP + Regression	✓	✓	~ ✓
USP + Spectral Filtering	✓	✓	~ ✓

Table 1: Comparison of methods for learning LDS. USP extends learning to mildly asymmetric matrices with bounded complex eigenvalues.

1.2 Intuition for Universal Sequence Preconditioning

We now give some brief intuition for the result. Linear dynamical systems (LDS) are perhaps the most basic and well studied dynamical systems in engineering and control science. Given input vectors $\mathbf{u}_1, \dots, \mathbf{u}_T \in \mathbb{C}^{d_{\text{in}}}$, the system generates a sequence of output vectors $\mathbf{y}_1, \dots, \mathbf{y}_T \in \mathbb{C}^{d_{\text{out}}}$ according to the law

$$\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{B}\mathbf{u}_t, \quad \mathbf{y}_t = \mathbf{C}\mathbf{x}_t + \mathbf{D}\mathbf{u}_t, \quad (1)$$

where $\mathbf{x}_0, \dots, \mathbf{x}_T \in \mathbb{C}^{d_{\text{hidden}}}$ is a sequence of hidden states and $(\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D})$ are matrices which parameterize the LDS. We assume w.l.o.g. that $\mathbf{D} = 0$. We can factor out the hidden state \mathbf{x}_t so that the observation at time t is

$$\mathbf{y}_t = \sum_{s=1}^t \mathbf{C}\mathbf{A}^{t-s}\mathbf{B}\mathbf{u}_s.$$

Given coefficients $c_{0:n} = (c_0, \dots, c_n)$ let

$$p_n^c(x) \stackrel{\text{def}}{=} \sum_{i=0}^n c_i x^{n-i}. \quad (2)$$

Consider a “preconditioned” target at time t to be a linear combination of $\mathbf{y}_{t:t-n}$ with coefficients $c_{0:n}$. A key insight is the following identity,

$$\sum_{i=0}^n c_i \mathbf{y}_{t-i} = \sum_{s=0}^{n-1} \left(\sum_{i=0}^s c_i \mathbf{C}\mathbf{A}^{s-i}\mathbf{B} \right) \mathbf{u}_{t-s} + \sum_{s=0}^{t-n-1} \mathbf{C}p_n^c(\mathbf{A})\mathbf{A}^s\mathbf{B}\mathbf{u}_{t-n-s}.$$

If we take $c_0 = 1$ (i.e. a monic polynomial), we can re-write \mathbf{y}_t as

$$\mathbf{y}_t = \underbrace{-\sum_{i=1}^n c_i \mathbf{y}_{t-i}}_{\aleph_0} + \underbrace{\sum_{s=0}^{n-1} \sum_{i=0}^s c_i \mathbf{C}\mathbf{A}^{s-i}\mathbf{B}\mathbf{u}_{t-s}}_{\aleph_1} + \underbrace{\sum_{s=0}^{t-n-1} \mathbf{C}p_n^c(\mathbf{A})\mathbf{A}^s\mathbf{B}\mathbf{u}_{t-n-s}}_{\aleph_2}. \quad (3)$$

This expression highlights our approach as a balance of three terms:

\aleph_0 The **universal preconditioning** term: it depends only on the coefficients $c_{0:n}$ and not on any learning algorithm.

\aleph_1 A term learnable via convex relaxation and regression, for example by denoting

$$\mathbf{Q}_s^{\text{learned}} = \sum_{i=0}^s c_i \mathbf{C}\mathbf{A}^{s-i}\mathbf{B}.$$

The diameter of the coefficient \mathbf{Q}_s depends on the magnitude of the coefficients $c_{0:n}$.

\aleph_2 The residual term with polynomial $p_n^c(\mathbf{A})$. By a careful choice of coefficients $c_{0:n}$, we can force this term to be very small.

The main insight we derive from this expression is the inherent tension between two terms \aleph_1, \aleph_2 . The polynomial $p_n^c(x)$ and its coefficients c_0, \dots, c_n control two competing effects:

1. The preconditioning coefficients grow larger with the degree n of the polynomial and the magnitude of the coefficients c_i . A higher degree polynomial and larger coefficients increase the diameter of the search space over the preconditioning coefficients, and therefore increase the regret bound stemming from the \aleph_1 component learning.
2. On the other hand, a larger search space can allow a broader class of polynomials $p_n(\cdot)$ which can better control of the magnitude of $p_n(\mathbf{A})$, and therefore reduce the search space of the \aleph_2 component.

What choice of polynomial is best? This work considers the Chebyshev polynomial. The reason is the following property of the n -th monic Chebyshev polynomial:

$$\max_{\lambda \in [-1, 1]} |p_n(\lambda)| \leq 2^{-(n-1)}.$$

As an example, consider any LDS whose hidden transition matrix \mathbf{A} is diagonalizable and has eigenvalues in $[-1, 1]$ ³. By the above property, observe that $\|p_n(\mathbf{A})\|_\infty \leq 2 \cdot 2^{-n}$, therefore shrinking the \aleph_2 term at a rate exponential with the number of preconditioning coefficients. We pause to remark on the *universality* of this choice of polynomial. Indeed, one could instead have chosen the preconditioning coefficients to depend on \mathbf{A} so that $p_n^c(\cdot)$ is the characteristic polynomial of \mathbf{A} . By the Cayley-Hamilton theorem, $p_n(\mathbf{A}) = 0$. This means that \aleph_2 term is canceled out completely. However this would have required knowledge of the spectrum of \mathbf{A} . The Chebyshev polynomial, on the other hand, is agnostic to the particular hidden transition matrix. Moreover, even if the spectrum of \mathbf{A} were known, choosing the preconditioning coefficients to form the characteristic polynomial would result in an algorithm which must learn hidden dimension many parameters, which is prohibitive. Instead, the degree of the Chebyshev polynomial must only grow logarithmically with the hidden dimension.

1.3 Related work

Our manuscript is technically involved and incorporates linear dynamical systems, spectral filtering, complex Chebyshev and Legendre polynomials, Hankel and Toeplitz matrix eigendecay, Gaussian quadrature and other techniques. The related work is thus expansive, and due to space limitations we give a detailed treatment in Appendix [A](#).

Preconditioning in the context of time series analysis has roots in the classical work of Box and Jenkins [\[14\]](#). In their foundational text they propose differencing as a method for making the time series stationary, and thus amenable to statistical learning techniques such as ARMA (auto-regressive moving average) [\[10\]](#). The differencing operator can be applied numerous times, and for different lags, giving rise to the ARIMA family of forecasting models. Identifying the order of an ARIMA model, and in particular the types of differencing needed to make a series stationary, is a hard problem. This is a special case of the problem we consider: differencing corresponds to certain coefficients of preconditioning the time series, whereas we consider arbitrary coefficients.

For a thorough introduction to modern control theory and exposition on open loop / closed loop predictors, learning via regression, and spectral filtering, see [\[28\]](#).

The fundamental problem of learning in linear dynamical systems has been studied for many decades, and we highlight several key approaches below:

1. System identification refers to the method of recovering $\mathbf{A}, \mathbf{B}, \mathbf{C}$ from the data. This is a non-convex problem and while many methods have been considered in this setting, they depend polynomially on the hidden dimension.
2. The (auto) regression method predicts according to $\hat{\mathbf{y}}_t = \sum_{i=1}^h \mathbf{M}_i \mathbf{u}_{t-i}$. The coefficients \mathbf{M}_i can be learned using convex regression. The downside of this approach is that if the spectral radius of \mathbf{A} is $1 - \delta$, it can be seen that $\sim \frac{1}{\delta}$ terms are needed.
3. The regression method can be further enhanced with “closed loop” components, that regress on prior observations $\mathbf{y}_{t-1:1}$. It can be shown using the Cayley-Hamilton theorem that using this method, d_h components are needed to learn the system, where d_h is the *hidden* dimension of \mathbf{A} .

³This work considers a broader class of hidden transition matrices.

4. Filtering involves recovering the state x_t from observations. While Kalman filtering is optimal under specific noise conditions, it generally fails in the presence of marginal stability and adversarial noise.
5. Finally, spectral filtering combines the advantages of all methods above. It is an efficient method, its complexity does **not** depend on the hidden dimension, and works for marginally stable systems. However, spectral filtering requires \mathbf{A} to be symmetric, or diagonalizable under the real numbers.

2 Main Results

In this section we formally state our main algorithms and theorems. We show that the Universal Sequence Preconditioning method provides significantly improved regret bounds for learning linear dynamical systems than previously known when used in conjunction with two distinct methods. The first method is simple convex regression, and the second is spectral filtering.

Both algorithms allow for learning in the case of marginally stable linear dynamical systems and allow for certain asymmetric transition matrices of arbitrary high hidden dimension. The regret bounds are free of the hidden dimension (up to logarithmic factors) – which significantly extends the state of the art.

2.1 Universal Sequence Preconditioning Applied to Regression

Algorithm 2 is an instantiation of Algorithm 4 for the method of convex regression. We set the preconditioning coefficients to be the coefficients of the n -th degree (monic) Chebyshev polynomial.

Algorithm 2 Universal Sequence Preconditioning for Regression

- 1: Input: initial parameter \mathbf{Q}^0 ; preconditioning coefficients $c_{0:n}$ from the n -th degree (monic) Chebyshev polynomial; convex constraints $\mathcal{K} = \{(\mathbf{Q}_0, \dots, \mathbf{Q}_{n-1}) \text{ s.t. } \|\mathbf{Q}_j\| \leq C_{\text{domain}} \|c\|_1\}$
 - 2: Assert that $c_0 = 1$.
 - 3: **for** $t = 1$ to T **do**
 - 4: Receive \mathbf{u}_t .
 - 5: Predict $\hat{\mathbf{y}}_t(\mathbf{Q}^t) = -\sum_{i=1}^n c_i \mathbf{y}_{t-i} + \sum_{j=0}^n \mathbf{Q}_j^t \mathbf{u}_{t-j}$.
 - 6: Observe true output \mathbf{y}_t and suffer loss $\ell_t(\mathbf{Q}^t) = \|\hat{\mathbf{y}}_t(\mathbf{Q}^t) - \mathbf{y}_t\|_1$.
 - 7: Update and project:

$$\mathbf{Q}^{t+1} \leftarrow \text{proj}_{\mathcal{K}}(\mathbf{Q}^t - \eta_t \nabla_{\mathbf{Q}} \ell_t(\mathbf{Q}^t)).$$
 - 8: **end for**
-

Theorem 2.1 shows that vanishing loss compared to the optimal ground-truth predictor, at a rate that is independent of the hidden dimension of the system.

Theorem 2.1. *Let $\{\mathbf{u}_t\}_{t=1}^T \in \mathbb{C}^{d_{in}}$ be any sequence of inputs which satisfy $\|\mathbf{u}_t\|_2 \leq 1$ and let $\{\mathbf{y}_t\}_{t=1}^T \in \mathbb{C}^{d_{out}}$ be the corresponding output coming from some linear dynamical system $(\mathbf{A}, \mathbf{B}, \mathbf{C})$ as defined per Eq. 1. Let \mathbf{P} diagonalize \mathbf{A} (note \mathbf{P} exists w.l.o.g.) and let $\kappa = \|\mathbf{P}\| \|\mathbf{P}^{-1}\|$. Assume that $\|\mathbf{B}\| \|\mathbf{C}\| \kappa \leq C_{\text{domain}}$. Let $\lambda_1, \dots, \lambda_{d_h}$ denote the spectrum of \mathbf{A} . If*

$$\max_{j \in [d_h]} |\arg(\lambda_j)| \leq 1/(32 \log_2(2T^3/d_{out}))^2$$

then the predictions $\hat{\mathbf{y}}_1, \dots, \hat{\mathbf{y}}_T$ from Algorithm 2 where the preconditioning coefficients $c_{0:n}$ are chosen to be the coefficients of the n -th monic Chebyshev polynomial satisfy

$$\frac{1}{T} \sum_{t=1}^T \|\hat{\mathbf{y}}_t - \mathbf{y}_t\|_1 \leq \tilde{O}\left(\frac{\|\mathbf{B}\| \|\mathbf{C}\| \kappa \sqrt{d_{out}}}{T^{2/13}}\right),$$

where $\tilde{O}(\cdot)$ hides polylogarithmic factors in T .

The proof of Theorem 2.1 is in Appendix D. For a simple baseline comparison, the regret achieved (via the same proof technique) by the vanilla regression algorithm without preconditioning is $O(C_{\text{domain}} \sqrt{d_{out}} T^{5/2})$ which is not sublinear in T .

2.2 Universal Sequence Preconditioning Applied to Spectral Filtering

Our second main result is the application of Universal Sequence Preconditioning to the spectral filtering algorithm [29]. In addition to applying USP to spectral filtering, we also propose a novel spectral filtering basis. Both changes to the vanilla spectral filtering algorithm are necessary to extend its sublinear regret bounds to the case of underlying systems with asymmetric hidden transition matrices. First we define two key quantities. Given horizon T and $\alpha \in \mathbb{C}$ let

$$\mu_T(\alpha) \stackrel{\text{def}}{=} [1 \quad \alpha \quad \dots \quad \alpha^{T-1}]^\top, \quad (4)$$

and

$$\mathbf{Z}_T \stackrel{\text{def}}{=} \int_{\alpha \in \mathbb{C}_\beta} \mu_T(\alpha) \mu_T(\bar{\alpha})^\top d\alpha, \quad (5)$$

where $\bar{\alpha} \in \mathbb{C}$ denotes the complex conjugate. The novel spectral filters are the eigenvectors of \mathbf{Z}_{T-n-1} , which we denote as $\phi_1, \dots, \phi_{T-n-1}$. Note that in the standard spectral filtering literature, the spectral filtering matrix is an integral over the real line and does not involve the complex conjugate. Our new matrix has an entirely different structure and although it looks quite similar, it surprisingly upends the proof techniques to ensure exponential spectral decay, a critical property for the method. We therefore introduce a novel proof of spectral decay in Theorem E.6 found in Section E.1.1.

Algorithm 3 Universal Sequence Preconditioning for Spectral Filtering

- 1: Input: initial $\mathbf{Q}_{1:n}^1, \mathbf{M}_{1:k}^1$, horizon T , convex constraints $\mathcal{K} = \{(\mathbf{Q}_0, \dots, \mathbf{Q}_{n-1}, \mathbf{M}_1, \dots, \mathbf{M}_k \text{ s.t. } \|\mathbf{Q}_j\| \leq R_Q \text{ and } \|\mathbf{M}_j\| \leq R_M\}$, parameter n , coefficients $c_{1:n}$.
 - 2: Let $p_n^c(x) = c_0 x^n + c_1 x^{n-1} + \dots + c_n$.
 - 3: Let ϕ_1, \dots, ϕ_n be the top n eigenvectors of \mathbf{Z}_{T-n-1} .
 - 4: Assert $c_0 = 1$.
 - 5: **for** $t = 1$ to T **do**
 - 6: Let $\tilde{\mathbf{u}}_{t-n-1:1}$ be $\mathbf{u}_{t-n-1:1}$ padded with zeros so it has dimension $T - n - 1 \times d_{\text{in}}$.
 - 7: Predict $\hat{\mathbf{y}}_t(\mathbf{Q}^t, \mathbf{M}^t) = -\sum_{i=1}^n c_i \mathbf{y}_{t-i} + \sum_{j=0}^n \mathbf{Q}_j^t \mathbf{u}_{t-j} + \frac{1}{\sqrt{T}} \sum_{j=1}^k \mathbf{M}_j^t \phi_j^\top \tilde{\mathbf{u}}_{t-n-1:1}$.
 - 8: Observe true \mathbf{y}_t , define loss $\ell_t(\hat{\mathbf{y}}_t) = \|\hat{\mathbf{y}}_t(\mathbf{Q}^t, \mathbf{M}^t) - \mathbf{y}_t\|_1$.
 - 9: Update and project: $(\mathbf{Q}^{t+1}, \mathbf{M}^{t+1}) = \text{proj}_{\mathcal{K}}(\mathbf{Q}^t, \mathbf{M}^t) - \eta_t \nabla \ell_t(\mathbf{Q}^t, \mathbf{M}^t)$
 - 10: **end for**
-

Theorem 2.2. Let $\{\mathbf{u}_t\}_{t=1}^T \in \mathcal{R}^{d_{\text{in}}}$ be any sequence of inputs which satisfy $\|\mathbf{u}_t\|_2 \leq 1$ and let $\{\mathbf{y}_t\}_{t=1}^T$ be the corresponding output coming from some linear dynamical system $(\mathbf{A}, \mathbf{B}, \mathbf{C})$ as defined per Eq. 1. Let \mathbf{P} diagonalize \mathbf{A} (note \mathbf{P} exists w.l.o.g.) and let $\kappa = \|\mathbf{P}\| \|\mathbf{P}^{-1}\|$. If R_Q, R_M , and k are chosen to be sufficiently large and

$$\max_{j \in [d_h]} |\arg(\lambda_j)| \leq 1/(32 \log_2(2T^3/d_{\text{out}}))^2,$$

then the predictions $\hat{\mathbf{y}}_1, \dots, \hat{\mathbf{y}}_T$ from Algorithm 3 where the preconditioning coefficients $\mathbf{c}_{0:n}$ are chosen to be the coefficients of the n -th monic Chebyshev polynomial satisfy

$$\frac{1}{T} \sum_{t=1}^T \|\hat{\mathbf{y}}_t - \mathbf{y}_t\|_1 \leq \tilde{O} \left(\frac{\|\mathbf{C}\| \|\mathbf{B}\| \kappa \sqrt{d_{\text{out}}}}{T^{3/13}} \right).$$

For comparison, vanilla spectral filtering without Universal Sequence Preconditioning requires that $\max_{j \in [d_h]} |\arg(\lambda_j)| \leq 1/T$, which is essentially the same as requiring that the matrix be symmetric since the horizon T can become arbitrarily large.

3 Proof Overview

In this section we give a high level overview of the proofs for Theorem 2.1 and Theorem 2.2. We start by recalling the intuition for Universal Sequence Preconditioning developed in Section 1.2 which

shows that if $\{\mathbf{y}_t\}_{t=1}^T$ evolves as a linear dynamical system parameterized by matrices $(\mathbf{A}, \mathbf{B}, \mathbf{C})$ with inputs $\{\mathbf{u}_t\}_{t=1}^T$ then by Equation 3,

$$\mathbf{y}_t = - \underbrace{\sum_{i=1}^n c_i \mathbf{y}_{t-i}}_{\aleph_0} + \underbrace{\sum_{s=0}^{n-1} \sum_{i=0}^s c_i \mathbf{C} \mathbf{A}^{s-i} \mathbf{B} \mathbf{u}_{t-s}}_{\aleph_1} + \underbrace{\sum_{s=0}^{t-n-1} \mathbf{C} p_n^c(\mathbf{A}) \mathbf{A}^s \mathbf{B} \mathbf{u}_{t-n-s}}_{\aleph_2}.$$

Recall that \aleph_0 is the universal preconditioning component, \aleph_1 is the term that can easily be learned by convex relaxation and regression, and \aleph_2 is the critical term that contains $p_n^c(\mathbf{A})$.

Both Theorem 2.1 and Theorem 2.2 use the standard result from online convex optimization (Theorem 3.1 from [26]) that online gradient descent over convex domain \mathcal{K} achieves regret $\frac{3}{2}GD\sqrt{T}$ as compared to the best point in \mathcal{K} , where D denotes the diameter of \mathcal{K} and G denotes the maximum gradient norm.

Regression: Proof of Theorem 2.1 In the case of regression, the domain is chosen so that \aleph_2 may be learned and the proof proceeds by bounding the diameter of such a domain and its corresponding maximum gradient norm to get regret $Cn^2\sqrt{d_{\text{out}}}\|c\|_1\sqrt{T}$ for a universal constant $C > 0$ which depends on the norms of matrices \mathbf{B} and \mathbf{C} from the underlying system. Then \aleph_3 is treated as an un-learnable error term. Let $\lambda(\mathbf{A})$ denote the set of eigenvalues of \mathbf{A} . By the simple magnitude bound of

$$\left\| \sum_{s=0}^{t-n-1} \mathbf{C} p_n(\mathbf{A}) \mathbf{A}^s \mathbf{B} \mathbf{u}_{t-n-s} \right\| \leq \max_{\lambda \in \lambda(\mathbf{A})} |p_n(\lambda)| \cdot T \cdot \|\mathbf{C}\| \cdot \|\mathbf{B}\|,$$

the error of ignoring this term can be very small if $\max_{\lambda \in \lambda(\mathbf{A})} |p_n(\lambda)|$ is small. In the proof of Theorem 2.1 in Appendix D we show that the regret for a generic polynomial p_n^c defined by coefficients $c_{0:n}$ is

$$\sum_{t=1}^T \|\mathbf{y}_t - \hat{\mathbf{y}}_t\|_1 \leq \underbrace{Cn^2\sqrt{d_{\text{out}}}\|c\|_1\sqrt{T}}_{\text{Regret from learning } \aleph_2} + \underbrace{C \max_{\lambda \in \mathcal{D}} |p_n^c(\lambda)| T^2}_{\text{Unlearnable Error Term}},$$

where \mathcal{D} is the region where \mathbf{A} is allowed to have eigenvalues (see Theorem D.1). Therefore, to get sublinear regret, we must choose a polynomial which has bounded ℓ_1 norm of its coefficients, while also exhibits very small infinity norm on the domain of \mathbf{A} 's eigenvalues.

Spectral Filtering: Proof of Theorem 2.2 In the case of spectral filtering, the domain is chosen so that both \aleph_2 and \aleph_3 may be learned. Because spectral filtering learns \aleph_3 , it is able to accumulate less error and hence achieves a better regret bound of $O(T^{-3/13})$ as compared to regression's $O(T^{-2/13})$. At a high level, the proof proceeds by exploiting the fact that $p_n(\mathbf{A})$ shrinks the size of the learnable domain. However this is not enough, in order to extend the result to systems where \mathbf{A} may have complex eigenvalues, the spectral filters must be eigenvalues of a new matrix, defined in Eq. 5, whose domain of integration includes the possibly complex eigenvalues of \mathbf{A} . To get the dimension-independent regret bounds enjoyed by spectral filtering in this new setting where complex eigenvalues may occur, the exponential decay of \mathbf{Z}_T from Eq. 5 must be established. This is nontrivial and requires several novel techniques inspired by [12]. The details are in Appendix E.1.1.

Theorem 2.2 gives the main guarantee for the spectral filtering algorithm, which states that Algorithm 3 instantiated with some choice of polynomial $p_n^c(\cdot)$ achieves regret

$$\tilde{O} \left(\left(n\|c\|_1 + T^{7/6} \max_{\alpha \in \mathbb{C}_\beta} |p_n^c(\alpha)| \right) (n+k) \sqrt{d_{\text{out}}} \sqrt{T} \right).$$

Both this theorem, as well as our new guarantee for convex regression, leads us to the following question: **Is there a universal choice of polynomial $p_n(x)$, where n is independent of hidden dimension, which guarantees sublinear regret?**

3.1 Using the Chebyshev Polynomial over the Complex Plane

For the real line, the answer to this question is known to be positive using the Chebyshev polynomials of the first kind. In general, the n^{th} (monic) Chebyshev polynomial $M_n(x)$ satisfies $\max_{x \in [-1, 1]} |M_n(x)| \leq 2^{-(n-1)}$. However, we are interested in a more general question over the complex plane. Since we care about linear dynamical systems that evolve according to a general asymmetric matrix, we need to extend our analysis to \mathbb{C}_β . This is a nontrivial extension since, in general, functions that are bounded on the real line can grow exponentially on the complex plane. Indeed, $2^{n-1}M_n(x) = \cos(n \arccos(x))$ and while $\cos(x)$ is bounded within $[-1, 1]$ for any $x \in \mathcal{R}$, over the complex numbers we have $\cos(z) = \frac{1}{2}(e^{iz} + e^{-iz})$, which is unbounded. Thus, we analyze the Chebyshev polynomial on the complex plane and provide the following bound.

Lemma 3.1. Let $z \in \mathbb{C}$ be some complex number with magnitude $|\alpha| \leq 1$. Let $M_n(\cdot)$ denote the n -th monic Chebyshev polynomial. If $|\arg(z)| \leq 1/64n^2$, then $|M_n(z)| \leq 1/2^{n-2}$.

We provide the proof in Appendix F. We also must analyze the magnitude of the coefficients of the Chebyshev polynomial, which can grow exponentially with n . We provide the following result.

Lemma 3.2. Let $M_n(\cdot)$ have coefficients c_0, \dots, c_n . Then $\max_{k=0, \dots, n} |c_k| \leq 2^{0.3n}$.

The proof of Lemma 3.2 is in Appendix F. Together, these two lemmas are the fundamental building block for universal sequence preconditioning and for obtaining our new regret bounds.

4 Experimental Evaluation

We empirically validate that convolutional preconditioning with Chebyshev or Legendre coefficients yields significant online regret improvements across various learning algorithms and data types. Below we summarize our data generation, algorithm variants, hyperparameter tuning, and evaluation metrics.

4.1 Synthetic Data Generation

We generate $N = 200$ sequences of length $T = 2000$ via three mechanisms: (i) a noisy linear dynamical system, (ii) a noisy nonlinear dynamical system, and (iii) a noisy deep RNN. Inputs $\mathbf{u}_{1:T} \sim \mathcal{N}(0, I)$.

Linear Dynamical System. Sample $(\mathbf{A}, \mathbf{B}, \mathbf{C})$ with $\mathbf{A} \in \mathcal{R}^{300 \times 300}$ having eigenvalues $\{z_j\}$ drawn uniformly in the complex plane subject to $\text{Im}(z_j) \leq \tau_{\text{thresh}}$ and $L \leq |z_j| \leq U$, and $\mathbf{B}, \mathbf{C} \in \mathcal{R}^{300}$. Then

$$\mathbf{x}_t = \mathbf{A} \mathbf{x}_{t-1} + \mathbf{B} \mathbf{u}_t, \quad \mathbf{y}_t = \mathbf{C} \mathbf{x}_t + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, \sigma^2 I).$$

Nonlinear Dynamical System. Similarly sample $(\mathbf{A}_1, \mathbf{B}_1, \mathbf{C})$ and $(\mathbf{A}_2, \mathbf{B}_2)$ with $\mathbf{A}_i \in \mathcal{R}^{10 \times 10}$, $\mathbf{B}_i, \mathbf{C} \in \mathcal{R}^{10}$. Then

$$\mathbf{x}_t^{(0)} = \mathbf{A}_1 \mathbf{x}_{t-1} + \mathbf{B}_1 \mathbf{u}_t, \quad \mathbf{x}_t^{(1)} = \sigma(\mathbf{x}_t^{(0)}), \quad \mathbf{x}_t = \mathbf{A}_2 \mathbf{x}_t^{(1)} + \mathbf{B}_2 \mathbf{u}_t, \quad \mathbf{y}_t = \mathbf{C} \mathbf{x}_t + \epsilon_t.$$

Deep RNN. We randomly initialize a sparse 10-layer stack of LSTMs with hidden dimension 100 and ReLU nonlinear activations. Given $\mathbf{u}_{1:T}$ we use this network to generate $\mathbf{y}_{1:T}$.

4.2 Algorithms and Preconditioning Variants

We evaluate the following methods: (1) **Regression** (Alg. 2), (2) **Spectral Filtering** (Alg. 3), (3) **DNN Predictor**: n -layer LSTM with dims $[d_1, \dots, d_n]$, ReLU.

Each method is applied with one of:

1. *Baseline*: no preconditioning
2. *Chebyshev*: $\mathbf{c}_{0:n}$ are the coefficients for the n th-Chebyshev polynomial. Note that when $n = 2$ we have $\mathbf{c}_0 = 1$ and $\mathbf{c}_1 = -1$ and therefore this is the method of *differencing* discussed in the introduction.

3. *Legendre*: $c_{0:n}$ are the coefficients for the n th-Legendre polynomial
4. *Learned*: $c_{0:n}$ is a parameter learned jointly with the model parameters

We test polynomial degrees $n \in \{2, 5, 10, 20\}$. This choice of degrees shows a rough picture of the impact of n .

Hyperparameter Tuning. To ensure fair comparison, for each algorithm and conditioning \mathbf{c} variant we perform a grid search over learning rates $\eta \in \{10^{-3}, 10^{-2}, 10^{-1}\}$, selecting the one minimizing average regret across the N sequences. In the case of the learned coefficients, we sweep over the 9 pairs of learning rates $(\eta_{\text{model}}, \eta_{\text{coefficients}}) \in \{10^{-3}, 10^{-2}, 10^{-1}\} \times \{10^{-3}, 10^{-2}, 10^{-1}\}$.

4.3 Results

Tables 2-4 report the mean \pm std of the absolute error over the final 200 predictions, averaged across 200 runs. In the linear and nonlinear cases we train a 2-layer DNN (dims (64, 128)); for RNN-generated data we match the 10-layer (100-dim) generator.

Key observations:

- Preconditioning drastically reduces baseline errors for all algorithms and data types.
- Chebyshev and Legendre yield nearly identical gains.
- For Chebyshev and Legendre, once the degree is higher than 5–10 the performance degrades since $\|\mathbf{c}\|_1$ gets very large (see our Lemma 3.2 which shows that these coefficients grow exponentially fast).
- Improvements decay as the complex threshold τ_{thresh} increases, consistent with our theoretical results which must bound $\text{Im}(z_j)$.
- Learned coefficients excel with regression and spectral filtering but destabilize the DNN on nonlinear and RNN-generated data.

Setting	Baseline	Chebyshev			Legendre			Learned			
		Deg. 2	Deg. 5	Deg. 10	Deg. 2	Deg. 5	Deg. 10	Deg. 2	Deg. 5	Deg. 10	Deg. 20
Regression											
$\tau_{\text{thresh}} = 0.01$	0.74 \pm 0.28	0.25 \pm 0.09	0.15 \pm 0.07	0.77 \pm 0.31	0.36 \pm 0.13	0.14 \pm 0.06	0.64 \pm 0.26	0.52 \pm 0.19	0.27 \pm 0.11	0.17 \pm 0.07	0.24 \pm 0.09
$\tau_{\text{thresh}} = 0.1$	1.92 \pm 0.81	0.84 \pm 0.27	0.66 \pm 0.18	1.90 \pm 0.67	1.10 \pm 0.40	0.63 \pm 0.17	1.66 \pm 0.58	1.34 \pm 0.43	0.57 \pm 0.14	0.55 \pm 0.14	0.56 \pm 0.14 s
$\tau_{\text{thresh}} = 0.9$	2.47 \pm 0.89	1.59 \pm 0.56	2.18 \pm 0.79	2.68 \pm 0.48	1.64 \pm 0.58	1.94 \pm 0.70	2.63 \pm 0.45	1.73 \pm 0.59	0.83 \pm 0.25	0.68 \pm 0.27	0.63 \pm 0.26
Spectral Filtering											
$\tau_{\text{thresh}} = 0.01$	5.94 \pm 3.37	1.72 \pm 0.95	0.69 \pm 0.38	3.25 \pm 1.79	2.78 \pm 1.56	0.66 \pm 0.36	2.74 \pm 1.51	1.99 \pm 0.94	0.54 \pm 0.29	0.55 \pm 0.25	0.61 \pm 0.25
$\tau_{\text{thresh}} = 0.1$	0.89 \pm 0.34	0.42 \pm 0.11	0.34 \pm 0.07	0.86 \pm 0.28	0.54 \pm 0.17	0.33 \pm 0.07	0.76 \pm 0.24	0.69 \pm 0.28	0.31 \pm 0.06	0.37 \pm 0.08	0.45 \pm 0.28
$\tau_{\text{thresh}} = 0.9$	10.17 \pm 8.80	9.87 \pm 8.90	12.66 \pm 8.18	32.90 \pm 22.02	9.42 \pm 8.39	11.53 \pm 7.46	28.83 \pm 19.24	7.93 \pm 4.42	6.31 \pm 4.19	5.73 \pm 3.80	5.86 \pm 4.02
2-layer DNN											
$\tau_{\text{thresh}} = 0.01$	4.49 \pm 2.02	2.31 \pm 1.18	2.62 \pm 1.52	10.36 \pm 6.05	2.79 \pm 1.34	2.35 \pm 1.35	8.92 \pm 5.20	2.89 \pm 1.24	1.56 \pm 0.64	0.79 \pm 0.25	0.40 \pm 0.16
$\tau_{\text{thresh}} = 0.1$	9.41 \pm 7.34	2.66 \pm 1.76	1.52 \pm 0.72	4.65 \pm 3.03	4.24 \pm 3.06	1.44 \pm 0.71	4.03 \pm 2.64	6.54 \pm 4.32	3.22 \pm 1.93	1.59 \pm 0.95	0.80 \pm 0.48
$\tau_{\text{thresh}} = 0.9$	2.45 \pm 1.31	2.24 \pm 1.34	3.49 \pm 2.27	11.48 \pm 8.15	2.17 \pm 1.25	3.10 \pm 2.00	9.97 \pm 7.05	1.29 \pm 0.66	0.71 \pm 0.34	0.43 \pm 0.18	0.24 \pm 0.11

Table 2: Linear dynamical system data (detailed in Sec. 4.1) across varying complex threshold τ_{thres} .

Setting	Baseline	Chebyshev			Legendre			Learned			
		Deg. 2	Deg. 5	Deg. 10	Deg. 2	Deg. 5	Deg. 10	Deg. 2	Deg. 5	Deg. 10	Deg. 20
Spectral Filtering											
$\tau_{\text{thresh}} = 0.01$	153.8 \pm 15.7	43.7 \pm 19.4	0.92 \pm 0.31	0.26\pm 0.15	78.4 \pm 34.7	2.82 \pm 1.17	0.34\pm 0.19	1.04 \pm 0.29	0.07 \pm 0.03	0.05\pm 0.03	0.07 \pm 0.03
$\tau_{\text{thresh}} = 0.01$	124.1\pm 68.4	33.4 \pm 17.3	2.02 \pm 1.23	0.36\pm 0.31	56.7 \pm 30.5	3.81 \pm 1.67	0.35\pm 0.26	2.85 \pm 1.13	0.04 \pm 0.01	0.02\pm 0.01	0.07 \pm 0.02
$\tau_{\text{thresh}} = 0.9$	165.5\pm 84.5	43.41 \pm 16.39	1.27\pm 0.37	1.90 \pm 0.80	76.9 \pm 29.1	3.32 \pm 1.09	1.64\pm 0.69	1.79 \pm 0.61	0.10\pm 0.04	0.12 \pm 0.05	0.17 \pm 0.07
2-layer DNN											
$\tau_{\text{thresh}} = 0.01$	10.46\pm 5.10	3.45 \pm 1.19	0.19\pm 0.16	0.43 \pm 0.22	3.55 \pm 1.56	0.18\pm 0.14	0.37 \pm 0.19	45.40 \pm 11.35	25.34 \pm 9.06	12.73 \pm 4.60	6.43\pm 2.32
$\tau_{\text{thresh}} = 0.1$	4.20\pm 1.10	3.38 \pm 0.79	0.14\pm 0.05	0.41 \pm 0.16	3.42 \pm 0.79	0.25\pm 0.14	0.36 \pm 0.13	67.70 \pm 28.40	31.99 \pm 11.78	15.95 \pm 5.87	8.01\pm 2.95
$\tau_{\text{thresh}} = 0.9$	6.72\pm 2.35	2.45 \pm 1.12	0.08\pm 0.03	0.28 \pm 0.15	3.35 \pm 0.98	0.20\pm 0.07	0.24 \pm 0.13	58.65 \pm 17.09	28.74 \pm 8.07	14.39 \pm 4.06	7.26\pm 2.05

Table 3: Nonlinear data (detailed in Sec. 4.1) across varying complex threshold τ_{thres} .

Setting	Baseline	Chebyshev			Legendre			Learned			
		Deg. 2	Deg. 5	Deg. 10	Deg. 2	Deg. 5	Deg. 10	Deg. 2	Deg. 5	Deg. 10	Deg. 20
10-layer DNN	0.54 ± 0.23	0.29 ± 0.10	0.08 ± 0.03	0.13 ± 0.05	0.37 ± 0.14	0.09 ± 0.03	0.12 ± 0.04	1.49 ± 0.93	2.13 ± 1.05	1.04 ± 0.51	0.5 ± 0.24

Table 4: Performance (average absolute error of the last 200 predictions) of a 10-layer DNN (detailed in Sec. 4.2) on data generated from the same model (detailed in Sec. 4.1).

4.4 ETTh1 Dataset

To evaluate whether our proposed preconditioning approach generalizes to real-world time series, we conduct experiments on the well-established **ETTh1** dataset from the Electricity Transformer Temperature (ETT) benchmark [41]. The ETTh1 dataset consists of continuous hourly measurements of load and oil temperature collected from electricity transformers and has been used in several recent works [41, 39, 35, 24, 25, 40, 34].

We study the effect of preconditioning on a 10-layer LSTM with hidden dimension 100 per layer using the Adam optimizer. We set the horizon to be $T = 5000$ and we sweep over a broader range of learning rates $\eta \in \{10^{-j}\}_{j=0,1,2,3,4,5}$. As before we consider (i) no preconditioning (baseline), (ii) fixed Chebyshev coefficients, (iii) fixed Legendre coefficients, and (iv) coefficients learned jointly with model parameters.

As seen in Figure 1, preconditioning with Chebyshev and Legendre for degree 5 the best performance after only the first 1000 iterations, while the performance of jointly learning the coefficients is worse at this stage. The performance of all three preconditioning methods are roughly on par with each other by 2500 iterations and by the full horizon $T = 5000$, jointly learning the coefficients results in the best average prediction error.




Figure 1: Absolute prediction error on final 200 predictions averaged over 10 independent runs for 10-layer LSTM with layer dimension 100 using Adam optimizer and sweeping over learning rates for each run.

5 Discussion

There are many settings in machine learning where universal, rather than learned, rules have proven very efficient. For example, physical laws of motion can be learned directly from observation data. However, Newton’s laws of motion succinctly crystallize very general phenomenon, and have proven very useful for large scale physics simulation engines. Similarly, in the theory of mathematical optimization, adaptive gradient methods have revolutionized deep learning. Their derivation as a consequence of regularization in online regret minimization is particularly simple [21], and thousands of research papers have not dramatically improved the initial basic ideas. These optimizers are, at the very least, a great way to initialize learned optimizers [38].

By analogy, our thesis in this paper is that universal preconditioning based on the solid theory of dynamical systems can be applicable to many domains or, at the very least, an initialization for other learning methods.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: We claim a rigorous new method for sequence preconditioning that improves upon the classical differencing method in theory and practice.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: Throughout we demonstrate understanding of the limitations of our approach, such as restricting our results to linear dynamical systems with bounded imaginary part of the transition matrix.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: Our main theoretical results are Theorems 2.1 and 2.2 in which the assumptions made are stated. We give only proof overview in main paper, but complete one in appendix.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Yes, we describe our algorithms completely in 2 and 3. Our code repository will also be open sourced.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: A link to the code will be provided in the final version. However, we omit the GitHub repository with all of our code at the moment in order not to violated the double blind policy and keep the paper anonymous.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Experiment details are expansively explained in section 4 and the appendix.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: all our graphs include error bars and statistical significance.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification:

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: The paper conforms simply in every respect with the NeurIPS Code of Ethics as explained.

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: [NA]

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: [NA]

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: All previous work is fully credited and cited throughout.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: Our new assets are well documented and the documentation is provided.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: [NA]

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: [NA]

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: [NA]

References

- [1] Yasin Abbasi-Yadkori, Peter Bartlett, and Varun Kanade. Tracking adversarial targets. In *International Conference on Machine Learning*, pages 369–377. PMLR, 2014.
- [2] Yasin Abbasi-Yadkori and Csaba Szepesvári. Regret bounds for the adaptive control of linear quadratic systems. In *Proceedings of the 24th Annual Conference on Learning Theory*, pages 1–26, 2011.
- [3] Milton Abramowitz and Irene A Stegun. *Handbook of mathematical functions with formulas, graphs, and mathematical tables*, volume 55. US Government printing office, 1948.
- [4] Naman Agarwal, Brian Bullins, Elad Hazan, Sham Kakade, and Karan Singh. Online control with adversarial disturbances. In *International Conference on Machine Learning*, pages 111–119. PMLR, 2019.
- [5] Naman Agarwal, Xinyi Chen, Evan Dogariu, Vlad Feinberg, Daniel Suo, Peter Bartlett, and Elad Hazan. Futurefill: Fast generation from convolutional sequence models. *arXiv preprint arXiv:2410.03766*, 2024.
- [6] Naman Agarwal, Elad Hazan, and Karan Singh. Logarithmic regret for online control. In *Advances in Neural Information Processing Systems*, pages 10175–10184, 2019.
- [7] Naman Agarwal, Daniel Suo, Xinyi Chen, and Elad Hazan. Spectral state space models. *arXiv preprint arXiv:2312.06837*, 2023.
- [8] Lars Valerian Ahlfors and Lars V Ahlfors. *Complex analysis*, volume 3. McGraw-Hill New York, 1979.
- [9] Oren Anava, Elad Hazan, and Shie Mannor. Online learning for adversaries with memory: price of past mistakes. *Advances in Neural Information Processing Systems*, 28, 2015.
- [10] Oren Anava, Elad Hazan, Shie Mannor, and Ohad Shamir. Online learning for time series prediction. In *Conference on learning theory*, pages 172–184. PMLR, 2013.
- [11] Ainesh Bakshi, Allen Liu, Ankur Moitra, and Morris Yau. A new approach to learning linear dynamical systems. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 335–348, 2023.
- [12] Bernhard Beckermann and Alex Townsend. On the singular values of matrices with displacement structure. *SIAM Journal on Matrix Analysis and Applications*, 38(4):1227–1248, 2017.
- [13] Dimitri P Bertsekas. *Dynamic Programming and Optimal Control*. Athena Scientific, 2007.
- [14] George EP Box and Gwilym M Jenkins. Time series analysis. forecasting and control. *Holden-Day Series in Time Series Analysis*, 1976.
- [15] Nicolo Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.
- [16] Xinyi Chen and Elad Hazan. Black-box control for linear dynamical systems. In *Conference on Learning Theory*, pages 1114–1143. PMLR, 2021.
- [17] Alon Cohen, Avinatan Hasidim, Tomer Koren, Nevena Lazic, Yishay Mansour, and Kunal Talwar. Online linear quadratic control. In *International Conference on Machine Learning*, pages 1029–1038, 2018.
- [18] Alon Cohen, Tomer Koren, and Yishay Mansour. Learning linear-quadratic regulators efficiently with only \sqrt{T} regret. In *International Conference on Machine Learning*, pages 1300–1309, 2019.
- [19] Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. Regret bounds for robust adaptive control of the linear quadratic regulator. In *Advances in Neural Information Processing Systems*, pages 4188–4197, 2018.

- [20] E. B. Dimitrov and N. Nikolov. Sharp bounds for the extreme zeros of classical orthogonal polynomials. *Journal of Approximation Theory*, 162(10):1793–1804, 2010.
- [21] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of machine learning research*, 12(7), 2011.
- [22] Maryam Fazel, Rong Ge, Sham Kakade, and Mehran Mesbahi. Global convergence of policy gradient methods for the linear quadratic regulator. In *International Conference on Machine Learning*, pages 1467–1476, 2018.
- [23] Udaya Ghai, Holden Lee, Karan Singh, Cyril Zhang, and Yi Zhang. No-regret prediction in marginally stable systems. In *Conference on Learning Theory*, pages 1714–1757. PMLR, 2020.
- [24] Albert Gu, Karan Goel, and Christopher Re. Efficiently modeling long sequences with structured state spaces. In *International Conference on Learning Representations*, 2022.
- [25] Rohit Gupta and et al. Simplifying long-range modeling with structured state space models. *arXiv preprint arXiv:2208.04933*, 2022.
- [26] Elad Hazan. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.
- [27] Elad Hazan, Holden Lee, Karan Singh, Cyril Zhang, and Yi Zhang. Spectral filtering for general linear dynamical systems. In *Advances in Neural Information Processing Systems*, pages 1680–1689, 2018.
- [28] Elad Hazan and Karan Singh. Introduction to online nonstochastic control. *arXiv preprint arXiv:2211.09619*, 2022.
- [29] Elad Hazan, Karan Singh, and Cyril Zhang. Learning linear dynamical systems via spectral filtering. In *Advances in Neural Information Processing Systems*, pages 6702–6712, 2017.
- [30] Rudolph Emil Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82.1:35–45, 1960.
- [31] CJ Knight and ACR Newbery. Trigonometric and gaussian quadrature. *MATHEMATICS of computation*, pages 575–581, 1970.
- [32] Mark Kozdoba, Jakub Marecek, Tigran Tchakian, and Shie Mannor. On-line learning of linear dynamical systems: Exponential forgetting in kalman filters. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 4098–4105, 2019.
- [33] Horia Mania, Stephen Tu, and Benjamin Recht. Certainty equivalence is efficient for linear quadratic control. In *Advances in Neural Information Processing Systems*, pages 10154–10164, 2019.
- [34] An Nguyen, Yifan Chen, and Jian Li. Preconditioning state-space models for stable long-horizon forecasting. *arXiv preprint arXiv:2403.07890*, 2024.
- [35] Yixuan Nie, Nina Narodytska, and Ankit Patel. A time series is worth 64 words: Long-term forecasting with transformers. *arXiv preprint arXiv:2209.13466*, 2022.
- [36] Max Simchowitz, Karan Singh, and Elad Hazan. Improper learning for non-stochastic control. In *Conference on Learning Theory*, pages 3320–3436. PMLR, 2020.
- [37] Josef Stoer, Roland Bulirsch, R Bartels, Walter Gautschi, and Christoph Witzgall. *Introduction to numerical analysis*, volume 1993. Springer, 1980.
- [38] Olga Wichrowska, Niru Maheswaranathan, Matthew W Hoffman, Sergio Gomez Colmenarejo, Misha Denil, Nando Freitas, and Jascha Sohl-Dickstein. Learned optimizers that scale and generalize. In *International conference on machine learning*, pages 3751–3760. PMLR, 2017.
- [39] Haixu Wu, Jianmin Xu, Jie Wang, and Mingsheng Long. Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting. In *Advances in Neural Information Processing Systems*, 2021.

- [40] Ailing Zeng, Muxi Chen, Lei Zhang, and Qiang Xu. Are transformers effective for time series forecasting? In *AAAI Conference on Artificial Intelligence*, 2023.
- [41] Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hongyun Xiong, Wancai Zhang, and Yan Li. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2021.
- [42] Kemin Zhou, John C. Doyle, and Keith Glover. *Robust and Optimal Control*. Prentice-Hall, Inc., USA, 1996.

Contents

1 Introduction	1
1.1 Our results	2
1.2 Intuition for Universal Sequence Preconditioning	3
1.3 Related work	4
2 Main Results	5
2.1 Universal Sequence Preconditioning Applied to Regression	5
2.2 Universal Sequence Preconditioning Applied to Spectral Filtering	6
3 Proof Overview	6
3.1 Using the Chebyshev Polynomial over the Complex Plane	8
4 Experimental Evaluation	8
4.1 Synthetic Data Generation	8
4.2 Algorithms and Preconditioning Variants	8
4.3 Results	9
4.4 ETTh1 Dataset	10
5 Discussion	10
A Related work	18
B Memory Capacity of Linear Dynamical Systems	19
C Online Version of Preconditioning	20
D Proof of Convolutional Preconditioned Regression Performance Theorem 2.1	20
E Proof of Theorem 2.2	23
E.1 Proving Approximation Lemma E.2	27
E.1.1 Proof of the Spectral Filtering Property Lemma E.4	29
F Chebyshev Polynomials Evaluated in the Complex Plane	36
G Plots from Experiments	38
G.1 Experiments with Linear Dynamical System Data	38
G.2 Nonlinear Data	40
G.3 Data from a DNN	41

A Related work

Preconditioning for sequence prediction. Preconditioning in the context of time series analysis has roots in the classical work of Box and Jenkins [14]. In this foundational text they propose differencing as a method for making the time series stationary, and thus amenable to statistical learning techniques such as ARMA (auto-regressive moving average) [10]. The differencing operator can be applied numerous times, and for different lags, giving rise to the ARIMA family of forecasting models.

Identifying the order of an ARIMA model, and in particular the types of differencing needed to make a series stationary, is a hard problem. This is a special case of the problem we consider: differencing corresponds to certain coefficients of preconditioning the time series, whereas we consider arbitrary coefficients.

Background on control of linear dynamical systems. Linear dynamical systems are the most fundamental and basic model in control theory, and have been studied for more than a century. For a thorough introduction, see the texts [13, 42, 28].

A rigorous proof that the Cayley-Hamilton theorem implies that d_h learned closed-loop components are sufficient to learn any LDS is given in [7, 27].

The seminal work of Kalman on state-space representation and filtering [30] allows one to learn any LDS with hidden-dimension parameters under stochastic and generative assumptions. Closed-loop auto-regressive learning subsumes Kalman filtering in the presence of adversarial noise, see e.g. [32]. [23] provide a method to learn a marginally stable LDS in the presence of bounded adversarial noise and asymmetric matrices, however their regret bound depends on the hidden dimension of the system. More recently, [11] use tensor and moment-based methods to learn a LDS with stochastic noise without dependence on the system’s condition number. However, their algorithmic complexity still scales polynomially with the hidden dimension.

In this work we consider regret in the context of *learning* linear dynamical systems. This is related to, but different from, *control* of the systems. We survey regret minimization for control next.

Regret for classical control models. The first works addressing control in the machine learning community assume either no perturbation in the dynamics at all, or i.i.d. Gaussian perturbations. Much of this work has considered obtaining low regret in the online LQR setting [2, 19, 33, 18] where a fully-observed linear dynamic system is driven by i.i.d. Gaussian noise via $x_{t+1} = \mathbf{A}x_t + \mathbf{B}u_t + w_t$, and the learner incurs constant quadratic state and input cost $\ell(x, u) = \frac{1}{2}x^\top \mathbf{Q}x + \frac{1}{2}u^\top \mathbf{R}u$. The optimal policy for this setting is well-approximated by a *state feedback controller* $u_t = Kx_t$, where K is the solution to the Discrete Algebraic Riccati Equation (DARE), and thus regret amounts to competing with this controller. Recent algorithms [33, 18] attain \sqrt{T} regret for this setting, with polynomial runtime and polynomial regret dependence on relevant problem parameters. A parallel line of work by [17] establishes \sqrt{T} regret in a variant of online LQR where the system is known to the learner, noise is stochastic, but an adversary selects quadratic loss functions ℓ_t at each time t . Again, the regret is measured with respect to a best-in-hindsight state feedback controller.

Provable control in the Gaussian noise setting via the policy gradient method was studied in [22]. Other relevant work from the machine learning literature includes tracking adversarial targets [1].

The non-stochastic control problem. The most accepted and influential control model stemming from the machine learning community was established in [4], who obtain \sqrt{T} -regret in the more general and challenging setting where both the Lipschitz loss function and the perturbations are adversarially chosen. The key insight behind this result is combining an improper controller parametrization known as disturbance-based control with online convex optimization with memory due to [9]. Follow-up work by [6] achieves logarithmic pseudo-regret for strongly convex, adversarially selected losses and well-conditioned stochastic noise. Further extensions were made for linear control with partial observation [36], system identification with adversarial noise [16], and many more settings surveyed in [28].

Online learning and online convex optimization. We make extensive use of techniques from the field of online learning and regret minimization in games [15, 26]. Following previous work from

machine learning, we consider regret minimization in sequence prediction, where the underlying sequence follows a linear dynamical system.

Spectral filtering for learning linear dynamical systems. The spectral filtering technique was put forth in [29] for learning symmetric linear dynamical systems. In [27], the technique was extended for more general dynamical systems using closed-loop regression; however, this required hidden-dimension parameters and polynomial dependence on the approximation guarantee. Spectral filtering techniques were recently used in non-linear sequence prediction, notably in the context of large language models, albeit with no theoretical guarantees [7]. As convolutional models, these methods are attractive for sequence prediction due to faster generation and inference as compared to attention-based models [5].

While several methods exist that can learn in the presence of asymmetric transition matrices [30, 11, 23], their performance depends on hidden dimension. On the other hand, spectral filtering methods [29] achieve regret which is independent of hidden dimension, even for marginally stable systems. However, these spectral filtering methods were limited to systems with symmetric transition matrices. In contrast, real-world dynamical systems often have asymmetric transition matrices with large hidden dimension, necessitating a more general approach. In this paper, we provide such an approach by extending the theory of spectral filtering to handle asymmetric systems, as long as the complex component of their eigenvalues is bounded.

In this paper we dramatically improve the spectral filtering technique and broaden its applicability in two major aspects: First, for general asymmetric linear dynamical systems we remove the dependence on the hidden dimension. Second, we improve the dependence of the number of learned parameters from polynomial to logarithmic. \square

B Memory Capacity of Linear Dynamical Systems

The hidden dimension d_h , which is the dimension of the transition matrix \mathbf{A} , plays a significant role in the expressive power of LDS. One of the most important features of the hidden dimension is that an LDS can memorize and recall inputs from up to d_{hidden} iterations in the past.

This can be seen with the system where \mathbf{B}, \mathbf{C} are identity, and \mathbf{A} is given by the permutation matrix

$$\mathbf{A}_{d_{\text{hidden}}}^{\text{perm}} = \begin{bmatrix} 0 & 0 & \cdots & 0 & 1 \\ 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 0 & 0 \\ 0 & 0 & \cdots & 1 & 0 \end{bmatrix},$$

which implements the memory operator $\mathbf{y}_t = \mathbf{u}_{t-d_h}$. Observe that any method which uses fewer than d_{hidden} parameters will fail to implement this memory operator and therefore, for general linear dynamical systems, d_{hidden} parameters are *necessary*.

Seemingly, this contradicts our promised results, which allows for learning a general LDS without hidden dimension dependence. The explanation is in the spectrum of the system. Notice that the eigenvalues of the permutation matrix \mathbf{A} above are the d_{hidden} roots of unity given by

$$\lambda_1, \dots, \lambda_{d_{\text{hidden}}} \in \left\{ e^{2\pi i \frac{k}{d}}, k = 1, 2, \dots, d_{\text{hidden}} \right\}.$$

Note that these eigenvalues have complex component as large as $1 - 1/d_{\text{hidden}}$.

Although in general a LDS can express signals with d_{hidden} memory, and thus intuitively might require d_{hidden} parameters, there are notable special cases that allow for efficient learning, i.e. learning the LDS with far fewer parameters. A notable case is that of spectral filtering, which allows efficient learning of a *symmetric LDS* with poly-logarithmic (in the desired accuracy ϵ) number of parameters. The significance of a symmetric transition matrix \mathbf{A} is that it can be diagonalized over the real numbers.

The natural question that arises is **which asymmetric matrices can be learned by spectral filtering efficiently, and which characterization of their spectrum allows for efficient learning?**

The answer we offer is surprisingly broad. For a LDS with transition matrix \mathbf{A} , let $\lambda_1, \dots, \lambda_d$ be its complex eigenvalues. We show that we can learn up to ε accuracy any LDS for which the largest eigenvalue has imaginary part bounded by $\frac{1}{\text{poly log } \frac{1}{\varepsilon}}$. We note that the spectral radius can be arbitrarily close to, or even equal to, one. The only restriction is on the complex part, which is mildly constrained as a logarithmic function of ε .

As per the permutation matrix example, this dependence is necessary and nearly tight - if the imaginary component of the eigenvalues of \mathbf{A} becomes large, any learning method requires parameterization that depends on the hidden dimension of the system.

C Online Version of Preconditioning

Algorithm 4 Universal Sequence Preconditioning (Online Version)

- 1: Input: sequence prediction model f_θ with initial parameter θ^0 ; loss function $\ell(\cdot)$; preconditioning coefficients $\mathbf{c}_{1:n}$.
- 2: **for** $t = 1$ to T **do**
- 3: Receive \mathbf{u}_t
- 4: Predict

$$\hat{\mathbf{y}}_t = - \sum_{j=1}^n \mathbf{c}_j \mathbf{y}_{t-j} + f_{\theta^t}(\mathbf{u}_{1:t}, \mathbf{y}_{1:t-1})$$

- 5: Observe true output \mathbf{y}_t and suffer loss $\ell_t(\hat{\mathbf{y}}_t, \mathbf{y}_t)$.
- 6: Update via projected gradient descent

$$\theta^{t+1} \leftarrow \text{proj}_{\mathcal{K}}(\theta^t - \eta_t \nabla_{\theta} \ell_t(\hat{\mathbf{y}}_t, \mathbf{y}_t))$$

- 7: **end for**
-

D Proof of Convolutional Preconditioned Regression Performance

Theorem 2.1

We will prove Theorem 2.1 first by proving the general result of Algorithm 2 for any choice of preconditioning coefficients $\mathbf{c}_{0:n}$. Then we will apply the Chebyshev coefficients to the result to get Theorem 2.1. For convenience, we restate Theorem 2.1 here.

Theorem D.1 (General Form of Theorem 2.1). *Let $\{\mathbf{u}_t\}_{t=1}^T \in \mathcal{R}^{d_m}$ be any sequence of inputs which satisfy $\|\mathbf{u}_t\|_2 \leq 1$ and let $\{\mathbf{y}_t\}_{t=1}^T$ be the corresponding output coming from some linear dynamical system $(\mathbf{A}, \mathbf{B}, \mathbf{C})$ as defined per Eq. 1. Let \mathbf{P} diagonalize \mathbf{A} (note \mathbf{P} exists w.l.o.g.) and let $\kappa = \|\mathbf{P}\| \|\mathbf{P}^{-1}\|$. Assume that $\|\mathbf{B}\| \|\mathbf{C}\| \kappa \leq C_{\text{domain}}$. Assume that $\|\mathbf{B}\| \|\mathbf{C}\| \kappa \leq C_{\text{domain}}$. Then the predictions $\hat{\mathbf{y}}_1, \dots, \hat{\mathbf{y}}_T$ from Algorithm 2 satisfy*

$$\sum_{t=1}^T \|\hat{\mathbf{y}}_t - \mathbf{y}_t\|_1 \leq C_{\text{domain}} \left(\frac{3}{2} n^2 \sqrt{d_{\text{out}}} \|\mathbf{c}\|_1 \sqrt{T} + \max_{\lambda \in \lambda(\mathbf{A})} |p_n^c(\lambda)| T^2 \right).$$

Proof of Theorem D.1 For the remainder of the proof we will denote

$$\hat{\mathbf{y}}_t(\mathbf{Q}) = - \sum_{i=1}^n \mathbf{c}_i \mathbf{y}_{t-i} + \sum_{j=0}^n \mathbf{Q}_j \mathbf{u}_{t-j},$$

so that Algorithm 2 outputs $\hat{\mathbf{y}}_t = \hat{\mathbf{y}}_t(\mathbf{Q}^t)$. Recall we define the domain

$$\mathcal{K} = \{(\mathbf{Q}_0, \dots, \mathbf{Q}_{n-1}) \text{ s.t. } \|\mathbf{Q}_j\| \leq C_{\text{domain}} n \|\mathbf{c}\|_1\}.$$

For convenience, we will use the shorthand \mathbf{Q} to refer to $(\mathbf{Q}_0, \dots, \mathbf{Q}_{n-1})$. First we prove that the regret of Algorithm 2 as compared to the best $\mathbf{Q}^* \in \mathcal{K}$ in hindsight is small. Then we prove that the best $\mathbf{Q}^* \in \mathcal{K}$ in hindsight achieves small prediction error. Let

$$G = \max_{t \in [T]} \|\nabla_{\mathbf{Q}} \ell_t(\mathbf{Q}^t)\|,$$

and let

$$D = \max_{\mathbf{Q}^1, \mathbf{Q}^2 \in \mathcal{K}} \|\mathbf{Q}^1 - \mathbf{Q}^2\|.$$

By Theorem 3.1 from [26],

$$\sum_{t=1}^T \ell_t(\mathbf{Q}^t) - \min_{\mathbf{Q}^* \in \mathcal{K}} \sum_{t=1}^T \ell_t(\mathbf{Q}^*) \leq \frac{3}{2} G D \sqrt{T}.$$

Therefore it remains to bound D and G . First we bound D . By definition of \mathcal{K} ,

$$\begin{aligned} D &= \max_{\mathbf{Q}^1, \mathbf{Q}^2 \in \mathcal{K}} \|\mathbf{Q}^1 - \mathbf{Q}^2\| \\ &= \max_{(\mathbf{Q}_0^1, \dots, \mathbf{Q}_{n-1}^1), (\mathbf{Q}_0^2, \dots, \mathbf{Q}_{n-1}^2) \in \mathcal{K}} \|(\mathbf{Q}_0^1, \dots, \mathbf{Q}_{n-1}^1) - (\mathbf{Q}_0^2, \dots, \mathbf{Q}_{n-1}^2)\| \\ &\leq \sum_{j=0}^{n-1} \|\mathbf{Q}_j^1 - \mathbf{Q}_j^2\| \\ &\leq 2C_{\text{domain}} n^2 \|\mathbf{c}\|_1. \end{aligned}$$

Next we bound the gradient. Recall that

$$\begin{aligned} \ell_t(\mathbf{Q}) &= \|\hat{\mathbf{y}}_t(\mathbf{Q}) - \mathbf{y}_t\|_1 \\ &= \left\| - \sum_{i=1}^n \mathbf{c}_i \mathbf{y}_{t-i} + \sum_{j=0}^n \mathbf{Q}_j \mathbf{u}_{t-j} - \mathbf{y}_t \right\|_1. \end{aligned}$$

Note that, in general, $\nabla_{\mathbf{A}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_1 = \text{sign}(\mathbf{A}\mathbf{x} - \mathbf{b})\mathbf{x}^\top$. Since $\|\mathbf{x}\mathbf{y}^\top\|_F \leq \|\mathbf{x}\|_2 \|\mathbf{y}\|_2$ we have

$$\|\nabla_{\mathbf{A}} f(\mathbf{A})\|_F \leq \sqrt{d} \|\mathbf{x}\|_2,$$

where d is the dimension of \mathbf{b} . Using this and the assumption that for any $t \in [T]$, $\|\mathbf{u}_t\|_2 \leq 1$, we have

$$\|\nabla_{\mathbf{Q}_i} \ell_t(\mathbf{Q})\|_F \leq \sqrt{d_{\text{out}}} \|\mathbf{u}_{t-i}\|_2 \leq \sqrt{d_{\text{out}}}.$$

Therefore,

$$G = \max_{t \in [T]} \|\nabla_{\mathbf{Q}} \ell_t(\mathbf{Q}^t)\|_F \leq n \sqrt{d_{\text{out}}}.$$

Thus we have a final regret bound

$$\sum_{t=1}^T \ell_t(\mathbf{Q}^t) - \min_{\mathbf{Q}^* \in \mathcal{K}} \sum_{t=1}^T \ell_t(\mathbf{Q}^*) \leq \frac{3}{2} C_{\text{domain}} n^2 \sqrt{d_{\text{out}}} \|\mathbf{c}\|_1 \sqrt{T}.$$

Next we show that if $(\mathbf{u}_{1:T}, \mathbf{y}_{1:T})$ is a linear dynamical system parameterized by $(\mathbf{A}, \mathbf{B}, \mathbf{C})$, then for any $t \in [T]$,

$$\min_{\mathbf{Q}^* \in \mathcal{K}} \ell_t(\mathbf{Q}^*) \leq \|\mathbf{C}\| \|\mathbf{B}\| \cdot \max_{\lambda \in \lambda(\mathbf{A})} |p_n^c(\lambda)| \cdot T,$$

where p_n^c denotes the polynomial

$$p_n^c(x) \stackrel{\text{def}}{=} \sum_{i=0}^n \mathbf{c}_i x^{n-i},$$

and $\lambda(\mathbf{A})$ denotes the set of eigenvalues of \mathbf{A} . Indeed, if $(\mathbf{u}_{1:T}, \mathbf{y}_{1:T})$ is a linear dynamical system parameterized by $(\mathbf{A}, \mathbf{B}, \mathbf{C})$ then

$$\mathbf{y}_t = \sum_{s=1}^t \mathbf{C} \mathbf{A}^{t-s} \mathbf{B} \mathbf{u}_s.$$

With some linear algebra we get that convolving $\mathbf{y}_{t:t-n}$ with coefficients $\mathbf{c}_{0:n}$ results in

$$\sum_{i=0}^n \mathbf{c}_i \mathbf{y}_{t-i} = \sum_{s=0}^{n-1} \sum_{i=0}^s \mathbf{c}_i \mathbf{C} \mathbf{A}^{s-i} \mathbf{B} \mathbf{u}_{t-s} + \sum_{s=0}^{t-n-1} \mathbf{C} p_n^c(\mathbf{A}) \mathbf{A}^s \mathbf{B} \mathbf{u}_{t-n-s},$$

or equivalently (since $\mathbf{c}_0 = 1$ due to the assertion in Algorithm 2),

$$\mathbf{y}_t = - \sum_{i=1}^n \mathbf{c}_i \mathbf{y}_{t-i} + \sum_{s=0}^{n-1} \sum_{i=0}^s \mathbf{c}_i \mathbf{C} \mathbf{A}^{s-i} \mathbf{B} \mathbf{u}_{t-s} + \sum_{s=0}^{t-n-1} \mathbf{C} p_n^{\mathbf{c}}(\mathbf{A}) \mathbf{A}^s \mathbf{B} \mathbf{u}_{t-n-s}.$$

Set $\hat{\mathbf{Q}}_s = \sum_{i=0}^s \mathbf{c}_i \mathbf{C} \mathbf{A}^{s-i} \mathbf{B}$ and set $\hat{\mathbf{Q}} = (\hat{\mathbf{Q}}_0, \dots, \hat{\mathbf{Q}}_{n-1})$. Since we assumed $C_{\text{domain}} \geq \|\mathbf{C}\| \|\mathbf{B}\|$,

$$\|\hat{\mathbf{Q}}_i\| \leq \sum_{j=0}^i |\mathbf{c}_j| \|\mathbf{C} \mathbf{A}^{i-j} \mathbf{B}\| \leq \|\mathbf{C}\| \|\mathbf{B}\| \sum_{j=0}^i |\mathbf{c}_j| \leq C_{\text{domain}} n \|\mathbf{c}\|_1.$$

Therefore $\hat{\mathbf{Q}} \in \mathcal{K}$. Moreover,

$$\begin{aligned} \hat{\mathbf{y}}_t(\hat{\mathbf{Q}}) - \mathbf{y}_t &= \left(- \sum_{i=1}^n \mathbf{c}_i \mathbf{y}_{t-i} + \sum_{j=0}^{n-1} \hat{\mathbf{Q}}_j \mathbf{u}_{t-j} \right) \\ &\quad - \left(- \sum_{i=1}^n \mathbf{c}_i \mathbf{y}_{t-i} + \sum_{s=0}^{n-1} \sum_{i=0}^s \mathbf{c}_i \mathbf{C} \mathbf{A}^{s-i} \mathbf{B} \mathbf{u}_{t-s} + \sum_{s=0}^{t-n-1} \mathbf{C} p_n^{\mathbf{c}}(\mathbf{A}) \mathbf{A}^s \mathbf{B} \mathbf{u}_{t-n-s} \right) \\ &= \sum_{s=0}^{t-n-1} \mathbf{C} p_n^{\mathbf{c}}(\mathbf{A}) \mathbf{A}^s \mathbf{B} \mathbf{u}_{t-n-s}. \end{aligned}$$

Therefore, $\|\hat{\mathbf{y}}_t(\hat{\mathbf{Q}}) - \mathbf{y}_t\|_1 = \|\sum_{s=0}^{t-n-1} \mathbf{C} p_n^{\mathbf{c}}(\mathbf{A}) \mathbf{A}^s \mathbf{B} \mathbf{u}_{t-n-s}\|_1$. Let \mathbf{A} be diagonalized by some \mathbf{P} so that

$$\mathbf{A} = \mathbf{P} \mathbf{D} \mathbf{P}^{-1},$$

where \mathbf{D} is the diagonalization of \mathbf{A} and let $\kappa = \|\mathbf{P}\| \|\mathbf{P}^{-1}\|$, note that we can assume this w.l.o.g. since the set of diagonalizable matrices over the complex numbers is dense and therefore if \mathbf{A} is not diagonalizable we may apply an arbitrarily small perturbation to it. Then since $\max_{j \in [d_h]} |\lambda_j(\mathbf{A})| \leq 1$,

$$\begin{aligned} \|\mathbf{C} p_n^{\mathbf{c}}(\mathbf{A}) \mathbf{A}^j \mathbf{B}\| &= \|\mathbf{C} \mathbf{P} p_n^{\mathbf{c}}(\mathbf{D}) \mathbf{D}^j \mathbf{P}^{-1} \mathbf{B}\| \\ &\leq \max_{k \in [d_h]} |p_n^{\mathbf{c}}(\mathbf{D}_{kk})| \cdot \|\mathbf{C}\| \|\mathbf{P}\| \|\mathbf{P}^{-1}\| \|\mathbf{B}\| \\ &\leq \max_{\lambda \in \lambda(\mathbf{A})} |p_n^{\mathbf{c}}(\lambda)| \cdot \|\mathbf{C}\| \|\mathbf{B}\| \kappa. \end{aligned}$$

Thus,

$$\|\hat{\mathbf{y}}_t(\hat{\mathbf{Q}}) - \mathbf{y}_t\|_1 \leq \|\mathbf{C}\| \|\mathbf{B}\| \kappa \cdot \max_{\lambda \in \lambda(\mathbf{A})} |p_n^{\mathbf{c}}(\lambda)| \cdot T,$$

and so (recalling that we showed $\hat{\mathbf{Q}} \in \mathcal{K}$),

$$\begin{aligned} \min_{\mathbf{Q}^* \in \mathcal{K}} \sum_{t=1}^T \ell_t(\mathbf{Q}^*) &\leq \sum_{t=1}^T \ell_t(\hat{\mathbf{Q}}) \\ &= \sum_{t=1}^T \|\hat{\mathbf{y}}_t(\hat{\mathbf{Q}}) - \mathbf{y}_t\|_1 \\ &\leq \|\mathbf{C}\| \|\mathbf{B}\| \kappa \cdot \max_{\lambda \in \lambda(\mathbf{A})} |p_n^{\mathbf{c}}(\lambda)| \cdot T^2. \end{aligned}$$

Since $C_{\text{domain}} \geq \|\mathbf{C}\| \|\mathbf{B}\| \kappa$ we conclude,

$$\sum_{t=1}^T \ell_t(\mathbf{Q}^t) \leq C_{\text{domain}} \left(\frac{3}{2} n^2 \sqrt{d_{\text{out}}} \|\mathbf{c}\|_1 \sqrt{T} + \max_{\lambda \in \lambda(\mathbf{A})} |p_n^{\mathbf{c}}(\lambda)| T^2 \right).$$

■

Next we choose \mathbf{c} to be the coefficients of the n -th monic Chebyshev polynomial to get the original theorem, Theorem 2.1 restated here for convenience.

Theorem (Restatement of Theorem 2.1). Let $\{\mathbf{u}_t\}_{t=1}^T \in \mathcal{R}^{d_{\text{in}}}$ be any sequence of inputs which satisfy $\|\mathbf{u}_t\|_2 \leq 1$ and let $\{\mathbf{y}_t\}_{t=1}^T$ be the corresponding output coming from some linear dynamical system $(\mathbf{A}, \mathbf{B}, \mathbf{C})$ as defined per Eq. 1. Let \mathbf{P} diagonalize \mathbf{A} (note \mathbf{P} exists w.l.o.g.) and let $\kappa = \|\mathbf{P}\| \|\mathbf{P}^{-1}\|$. Assume that $\|\mathbf{B}\| \|\mathbf{C}\| \kappa \leq C_{\text{domain}}$. Let $\lambda_1, \dots, \lambda_{d_h}$ denote the spectrum of \mathbf{A} . If

$$\max_{j \in [d_h]} |\text{Arg}(\lambda_j)| \leq \left(64 \log_2 \left(\frac{8T^{3/2}}{3\sqrt{d_{\text{out}}}} \right) \right)^{-2},$$

then the predictions $\hat{\mathbf{y}}_1, \dots, \hat{\mathbf{y}}_T$ from Algorithm 2 where the preconditioning coefficients $\mathbf{c}_{0:n}$ are chosen to be the coefficients of the n -th monic Chebyshev polynomial satisfy

$$\frac{1}{T} \sum_{t=1}^T \|\hat{\mathbf{y}}_t - \mathbf{y}_t\|_1 \leq \frac{9C_{\text{domain}} \sqrt{d_{\text{out}}} \log_2^2(3T)}{T^{2/13}}.$$

Proof. From Theorem D.1 we have

$$\sum_{t=1}^T \ell_t(\mathbf{Q}^t) \leq C_{\text{domain}} \left(\frac{3}{2} n^2 \sqrt{d_{\text{out}}} \|\mathbf{c}\|_1 \sqrt{T} + \max_{\lambda \in \lambda(\mathbf{A})} |p_n^c(\lambda)| T^2 \right).$$

By Lemma 3.1 if for any eigenvalue λ of \mathbf{A} , $|\arg(\lambda)| \leq 1/64n^2$ then

$$\max_{\lambda \in \lambda(\mathbf{A})} |p_n^c(\lambda)| \leq \frac{1}{2^{n-2}}.$$

Moreover, by Lemma 3.2, $\|\mathbf{c}\|_1 \leq 2^{0.3n}$. Thus the Chebyshev-conditioned predictor class satisfies

$$\sum_{t=1}^T \ell_t(\mathbf{Q}^t) \leq C_{\text{domain}} \left(\frac{3}{2} n^2 \sqrt{d_{\text{out}}} 2^{0.3n} \sqrt{T} + 2^{-(n-2)} T^2 \right).$$

Picking

$$n = \frac{10}{13} \log_2 \left(\frac{8}{3\sqrt{d_{\text{out}}}} T^{3/2} \right),$$

we get

$$\begin{aligned} \sum_{t=1}^T \ell_t(\mathbf{Q}^t) &\leq 2 \left(\frac{3}{2} \left(\frac{10}{13} \log_2 \left(\frac{8}{3\sqrt{d_{\text{out}}}} T^{3/2} \right) \right)^2 \sqrt{d_{\text{out}}} \right)^{10/13} 4^{3/13} T^{11/13} \\ &\leq 9C_{\text{domain}} \sqrt{d_{\text{out}}} \log_2^2(3T) T^{11/13}. \end{aligned}$$

Dividing both sides by T we get the stated result. ■

E Proof of Theorem 2.2

First, we prove a more general form that holds for any choice of $\mathbf{c}_{0:n}$ and resulting polynomial p_n^c .

Theorem E.1. Let $\{\mathbf{u}_t\}_{t=1}^T \in \mathcal{R}^{d_{\text{in}}}$ be any sequence of inputs which satisfy $\|\mathbf{u}_t\|_2 \leq 1$ and let $\{\mathbf{y}_t\}_{t=1}^T$ be the corresponding output coming from some linear dynamical system $(\mathbf{A}, \mathbf{B}, \mathbf{C})$ as defined per Eq. 1. Let \mathbf{P} diagonalize \mathbf{A} (\mathbf{P} exists w.l.o.g.) and let $\kappa = \|\mathbf{P}\| \|\mathbf{P}^{-1}\|$. Let $B_n = \max_{\alpha \in \mathbb{C}_\beta} |p_n^c(\alpha)|$. Let $C > 0$ and $c > 1$ be universal constants. If

$$\begin{aligned} R_Q &\geq \|\mathbf{C}\| \|\mathbf{B}\| \|\mathbf{c}\|_1 \\ R_M &\geq C \log(T)^{1/6} T^{7/6} \|\mathbf{C}\| \|\mathbf{B}\| \kappa B_n \\ k &\geq 6 \log(T) \log \left(R_M / \sqrt{d_{\text{out}}} \epsilon \right) / \log(c) \end{aligned}$$

then the predictions $\hat{\mathbf{y}}_1, \dots, \hat{\mathbf{y}}_T$ from Algorithm 3 satisfy

$$\sum_{t=1}^T \ell_t(\mathbf{Q}^t, \mathbf{M}^t) \leq 1 + C \|\mathbf{C}\| \|\mathbf{B}\| \left(n \|\mathbf{c}\|_1 + \log^2(T) T^{7/6} \max_{\alpha \in \mathbb{C}_\beta} |p_n^c(\alpha)| \right) (n + k) \sqrt{d_{\text{out}}} \sqrt{T}.$$

Critically, we observe that the guaranteed regret bound does not depend on the hidden dimension of the dynamics matrix \mathbf{A} . While the general version of Algorithm [3](#) is interesting in its own right, we show that by choosing the coefficients of the algorithm to be the Chebyshev coefficients we obtain sublinear absolute error when the signal comes from a linear dynamical system.

Theorem (Detailed Version of Theorem [2.2](#)). Let $\{\mathbf{u}_t\}_{t=1}^T \in \mathcal{R}^{d_{\text{in}}}$ be any sequence of inputs which satisfy $\|\mathbf{u}_t\|_2 \leq 1$ and let $\{\mathbf{y}_t\}_{t=1}^T$ be the corresponding output coming from some linear dynamical system $(\mathbf{A}, \mathbf{B}, \mathbf{C})$ as defined per Eq. [1](#). Let \mathbf{P} diagonalize \mathbf{A} (\mathbf{P} exists w.l.o.g.) and let $\kappa = \|\mathbf{P}\| \|\mathbf{P}^{-1}\|$. Let $B_n = \max_{\alpha \in \mathbb{C}_\beta} |p_n^c(\alpha)|$. Let $C > 0$ and $c > 1$ be universal constants. Suppose that

$$\begin{aligned} R_Q &\geq \|\mathbf{C}\| \|\mathbf{B}\| \|\mathbf{c}\|_1, \\ R_M &\geq C \log(T)^{1/6} T^{7/6} \|\mathbf{C}\| \|\mathbf{B}\| \kappa B_n, \\ k &\geq 6 \log(T) \log(R_M / \sqrt{d_{\text{out}} \epsilon}) / \log(c), \end{aligned}$$

and moreover that for any eigenvalue, λ_j , of \mathbf{A}

$$\max_{j \in [d_h]} |\arg(\lambda_j)| \leq \left(64 \log_2 \left(\frac{8T^{3/2}}{3\sqrt{d_{\text{out}}}} \right) \right)^{-2},$$

then the predictions $\hat{\mathbf{y}}_1, \dots, \hat{\mathbf{y}}_T$ from Algorithm [3](#) where the preconditioning coefficients $\mathbf{c}_{0:n}$ are chosen to be the coefficients of the n -th monic Chebyshev polynomial satisfy

$$\frac{1}{T} \sum_{t=1}^T \|\hat{\mathbf{y}}_t - \mathbf{y}_t\|_1 \leq \tilde{O} \left(\frac{\|\mathbf{C}\| \|\mathbf{B}\| \kappa}{T^{3/13}} \right),$$

where $\tilde{O}(\cdot)$ hides polylogarithmic factors in T .

Proof. From Theorem [D.1](#) we have

$$\sum_{t=1}^T \ell_t(\mathbf{Q}^t, \mathbf{M}^t) \leq 1 + C \|\mathbf{C}\| \|\mathbf{B}\| \left(n \|\mathbf{c}\|_1 + \log^2(T) T^{7/6} \max_{\alpha \in \mathbb{C}_\beta} |p_n^c(\alpha)| \right) (n+k) \sqrt{d_{\text{out}}} \sqrt{T}.$$

By Lemma [3.1](#) if for any eigenvalue λ of \mathbf{A} , $|\arg(\lambda)| \leq 1/64n^2$ then

$$\max_{\lambda \in \lambda(\mathbf{A})} |p_n^c(\lambda)| \leq \frac{1}{2^{n-2}}.$$

Moreover, by Lemma [3.2](#), $\|\mathbf{c}\|_1 \leq 2^{0.3n}$. Thus the Chebyshev-conditioned predictor class satisfies

$$\sum_{t=1}^T \ell_t(\mathbf{Q}^t) \leq 1 + C \|\mathbf{C}\| \|\mathbf{B}\| \left(n^2 2^{0.3n} + \log^2(T) T^{7/6} 2^{-(n-2)} \right) (n+k) \sqrt{d_{\text{out}}} \sqrt{T}.$$

Picking

$$n = \frac{10}{13} \log_2 \left(4 \log^2(T) T^{7/6} \right),$$

we get

$$n^2 2^{0.3n} + \log^2(T) T^{7/6} 2^{-(n-2)} \leq n^2 \left(4 \log^2(T) T^{7/6} \right)^{3/13} \leq 4n^2 \log(T) T^{7/26},$$

and therefore

$$\sum_{t=1}^T \ell_t(\mathbf{Q}^t) \leq 4C \|\mathbf{C}\| \|\mathbf{B}\| n^3 k \sqrt{d_{\text{out}}} \log(T) T^{10/13}.$$

Observing that both n and k are bounded polylogarithmically in T and dividing both sides by T we get the stated result. \blacksquare

Remark on the loss function. The reader may notice we use the ℓ_1 , or absolute loss, rather than Euclidean or other loss functions. All norms are equivalent up to the (output) dimension, and thus learning to predict as well as the best linear dynamical system in hindsight is meaningful in any norm. However, we make this technical choice since it greatly simplifies the regret bounds, and in particular the bound on the gradient norms, which is technically involved. We conjecture that sublinear regret bounds are attainable in other norms as well, and leave it for future work.

We prove Theorem 2.2 on an (equivalent) algorithm, where we rescale the parameter \mathbf{M}_j by \sqrt{T} and the input $\langle \phi_j, \mathbf{u}_{(t-n-i):1} \rangle$ by $1/\sqrt{T}$. We account for this rescaling by increasing the size of the domain for \mathbf{M} by \sqrt{T} . The proof of Theorem 2.2 proceeds in two parts. The first is to show that any linear dynamical signal is well approximated by a predictor of the form in Algorithm 3 $\hat{\mathbf{y}}(\mathbf{Q}, \mathbf{M})$ where $(\mathbf{Q}, \mathbf{M}) \in \mathcal{K}$.

Lemma E.2 (Approximation Lemma). Suppose $\mathbf{y}_{1:T}$ evolves as a linear dynamical system characterized by $(\mathbf{A}, \mathbf{B}, \mathbf{C})$ as in Eq. 1. Consider domain

$$\mathcal{K} = \{(\mathbf{Q}_1, \dots, \mathbf{Q}_n, \mathbf{M}_1, \dots, \mathbf{M}_k) \text{ s.t. } \|\mathbf{Q}_i\| \leq R_Q, \text{ and } \|\mathbf{M}_i\| \leq R_M\}.$$

Let \mathbf{P} diagonalize \mathbf{A} (\mathbf{P} exists w.l.o.g.) and let $\kappa = \|\mathbf{P}\| \|\mathbf{P}^{-1}\|$. Let $B_n = \max_{\alpha \in \mathbb{C}_\beta} |p_n^c(\alpha)|$. Let $C > 0$ and $c > 1$ be universal constants. If

$$\begin{aligned} R_Q &\geq \|\mathbf{C}\| \|\mathbf{B}\| \|\mathbf{c}\|_1 \\ R_M &\geq C \log(T)^{1/6} T^{7/6} \|\mathbf{C}\| \|\mathbf{B}\| \kappa B_n \\ k &\geq 6 \log(T) \log\left(R_M / \sqrt{d_{\text{out}} \epsilon}\right) / \log(c) \end{aligned}$$

then there exists $(\hat{\mathbf{Q}}_1, \dots, \hat{\mathbf{Q}}_n, \hat{\mathbf{M}}_1, \dots, \hat{\mathbf{M}}_k) \in \mathcal{K}$ such that for prediction (as in Algorithm 3)

$$\hat{\mathbf{y}}_t(\hat{\mathbf{Q}}, \hat{\mathbf{M}}) = - \sum_{i=1}^n c_i \mathbf{y}_{t-i} + \sum_{j=0}^n \hat{\mathbf{Q}}_j^t \mathbf{u}_{t-j} + \frac{1}{\sqrt{T}} \sum_{j=1}^k \hat{\mathbf{M}}_j^t \phi_j^\top \tilde{\mathbf{u}}_{t-n-1:1},$$

it holds that

$$\|\hat{\mathbf{y}}_t - \mathbf{y}_t\|_1 \leq \epsilon.$$

Lemma E.2 is the heart of the result and requires several pages of proofs and novel ideas to establish the result. We therefore dedicate Section E.1 to its proof.

The next result provides the regret of Online Gradient Descent when compared to the best $(\mathbf{Q}^*, \mathbf{M}^*) \in \mathcal{K}$.

Lemma E.3 (Online Gradient Descent). Recall the domain \mathcal{K} in Algorithm 3 be

$$\mathcal{K} = \{(\mathbf{Q}_1, \dots, \mathbf{Q}_n, \mathbf{M}_1, \dots, \mathbf{M}_k) \text{ s.t. } \|\mathbf{Q}_i\| \leq R_Q \text{ and } \|\mathbf{M}_i\| \leq R_M\}.$$

Given coefficients $\mathbf{c}_{1:n}$, let $p_n^c(x) = x^n + \mathbf{c}_1 x^{n-1} + \dots + \mathbf{c}_n$ and let r denote the maximum absolute value of the coefficients. The iterates $\hat{\mathbf{y}}_t^A$ output by Algorithm 3 satisfy

$$\sum_{t=1}^T \ell_t(\mathbf{Q}^t, \mathbf{M}^t) - \min_{(\mathbf{Q}^*, \mathbf{M}^*) \in \mathcal{K}} \sum_{t=1}^T \ell_t(\mathbf{Q}^*, \mathbf{M}^*) \leq \frac{3}{2} (nR_Q + kR_M)(n+k) \sqrt{d_{\text{out}}} \sqrt{T}.$$

Proof of Lemma E.3 Let $G = \max_{t \in [T]} \|\nabla_{\mathbf{Q}, \mathbf{M}} \ell_t(\mathbf{Q}^t, \mathbf{M}^t)\|$ and let

$$D = \max_{(\mathbf{Q}^1, \mathbf{M}^1), (\mathbf{Q}^2, \mathbf{M}^2) \in \mathcal{K}} \|(\mathbf{Q}^1, \mathbf{M}^1) - (\mathbf{Q}^2, \mathbf{M}^2)\|.$$

By Theorem 3.1 from [26],

$$\sum_{t=1}^T \ell_t(\mathbf{Q}^t, \mathbf{M}^t) - \min_{(\mathbf{Q}^*, \mathbf{M}^*) \in \mathcal{K}} \sum_{t=1}^T \ell_t(\mathbf{Q}^*, \mathbf{M}^*) \leq \frac{3}{2} G D \sqrt{T}.$$

Therefore it remains to bound G and D . By definition of \mathcal{K} we have

$$D \leq nR_Q + kR_M.$$

For G we compute the subgradient at any \mathbf{Q}_i and \mathbf{M}_i . Note that, in general, $\nabla_{\mathbf{A}} \|\mathbf{Ax} - \mathbf{b}\|_1 = \text{sign}(\mathbf{Ax} - \mathbf{b})\mathbf{x}^\top$. Since $\|\mathbf{xy}^\top\|_F \leq \|\mathbf{x}\|_2 \|\mathbf{y}\|_2$ we have

$$\|\nabla_{\mathbf{A}} f(\mathbf{A})\|_F \leq \sqrt{d} \|\mathbf{x}\|_2,$$

where d is the dimension of \mathbf{b} . Using this and the assumption that for any $t \in [T]$, $\|\mathbf{u}_t\|_2 \leq 1$, we have

$$\begin{aligned} \|\nabla_{\mathbf{M}_j} \ell_t(\mathbf{Q}, \mathbf{M})\| &\leq \sqrt{d_{\text{out}}} \|\phi_j^\top \tilde{\mathbf{u}}_{t-n-1:1} T^{-1/2}\| \\ &\leq \sqrt{d_{\text{out}}} \|\phi_j\|_1 \|\tilde{\mathbf{u}}_{t-n-1:1}\|_\infty T^{-1/2} \\ &\leq \sqrt{d_{\text{out}}}. \quad (\|\tilde{\mathbf{u}}_{t-n-1:1}\|_\infty \leq 1 \text{ and } \|\phi_j\|_1 \leq \sqrt{T} \text{ since } \|\phi_j\|_2 \leq 1) \end{aligned}$$

Next,

$$\|\nabla_{\mathbf{Q}_i} \ell_t(\mathbf{Q})\|_F \leq \sqrt{d_{\text{out}}} \|\mathbf{u}_{t-i}\|_2 \leq \sqrt{d_{\text{out}}}.$$

Therefore,

$$G = \max_{t \in [T]} \|\nabla_{(\mathbf{Q}, \mathbf{M})} \ell_t(\mathbf{Q}^t, \mathbf{M}^t)\|_F \leq (n+k) \sqrt{d_{\text{out}}}.$$

Therefore, we have

$$\sum_{t=1}^T \ell_t(\mathbf{Q}^t, \mathbf{M}^t) - \min_{(\mathbf{Q}^*, \mathbf{M}^*) \in \mathcal{K}} \sum_{t=1}^T \ell_t(\mathbf{Q}^*, \mathbf{M}^*) \leq \frac{3}{2} (nR_Q + kR_M)(n+k) \sqrt{d_{\text{out}}} \sqrt{T}.$$

■

Combining Lemma E.2 and Lemma E.3 proves Theorem 2.2

Proof of Theorem 2.2 By Lemma E.3 the iterates from Algorithm 3 $(\mathbf{Q}^1, \mathbf{M}^1), \dots, (\mathbf{Q}^T, \mathbf{M}^T)$ satisfy

$$\sum_{t=1}^T \ell_t(\mathbf{Q}^t, \mathbf{M}^t) - \min_{(\mathbf{Q}^*, \mathbf{M}^*) \in \mathcal{K}} \sum_{t=1}^T \ell_t(\mathbf{Q}^*, \mathbf{M}^*) \leq \frac{3}{2} (nR_Q + kR_M)(n+k) \sqrt{d_{\text{out}}} \sqrt{T}.$$

Next, if $\mathbf{y}_{1:T}$ comes from a linear dynamical system, by Lemma E.2 if $R_Q \geq \|\mathbf{C}\| \|\mathbf{B}\| \|\mathbf{c}\|_1$, $R_M \geq C \log(T)^{1/6} T^{7/6} \max_{\alpha \in \mathcal{C}_\beta} |p_n^c(\alpha)| \|\mathbf{CB}\|$ and $k \geq C \log^2(T \max_{\alpha \in \mathcal{C}_\beta} |p_n^c(\alpha)| \|\mathbf{CB}\| / (d_{\text{out}} \epsilon))$ (where $C > 0$ is a universal constant) then there exists $(\hat{\mathbf{Q}}, \hat{\mathbf{M}}) \in \mathcal{K}$ such that

$$\sum_{t=1}^T \ell_t(\hat{\mathbf{Q}}, \hat{\mathbf{M}}) \leq \epsilon T.$$

Since $(\hat{\mathbf{Q}}, \hat{\mathbf{M}}) \in \mathcal{K}$ we must also have that

$$\min_{(\mathbf{Q}^*, \mathbf{M}^*) \in \mathcal{K}} \sum_{t=1}^T \ell_t(\mathbf{Q}^*, \mathbf{M}^*) \leq \epsilon T.$$

Therefore,

$$\begin{aligned} \sum_{t=1}^T \ell_t(\mathbf{Q}^t, \mathbf{M}^t) &\leq \min_{(\mathbf{Q}^*, \mathbf{M}^*) \in \mathcal{K}} \sum_{t=1}^T \ell_t(\mathbf{Q}^*, \mathbf{M}^*) + \frac{3}{2} (nR_Q + kR_M)(n+k) \sqrt{d_{\text{out}}} \sqrt{T} \\ &\leq \epsilon T + \frac{3}{2} (nR_Q + kR_M)(n+k) \sqrt{d_{\text{out}}} \sqrt{T}. \end{aligned}$$

Plugging in the values for R_Q and R_M and for k when $\epsilon = 1/T$ we conclude

$$\begin{aligned} \sum_{t=1}^T \ell_t(\mathbf{Q}^t, \mathbf{M}^t) &\leq \epsilon T + \frac{3}{2} \|\mathbf{C}\| \|\mathbf{B}\| \left(n \|\mathbf{c}\|_1 \right. \\ &\quad \left. + C \log^2 \left(T \max_{\alpha \in \mathcal{C}_\beta} |p_n(\alpha)| \|\mathbf{C}\| \|\mathbf{B}\| \right) T^{7/6} \max_{\alpha \in \mathcal{C}_\beta} |p_n^c(\alpha)| \right) (n+k) \sqrt{d_{\text{out}}} \sqrt{T} \\ &\leq 1 + C \|\mathbf{C}\| \|\mathbf{B}\| \left(n \|\mathbf{c}\|_1 + \log^2(T) T^{7/6} \max_{\alpha \in \mathcal{C}_\beta} |p_n^c(\alpha)| \right) (n+k) \sqrt{d_{\text{out}}} \sqrt{T}. \end{aligned}$$

■

E.1 Proving Approximation Lemma E.2

The difficult result to prove is Lemma E.2. As discussed in Section 3, any signal from a linear dynamical system can be broken down into two parts: a linear autoregressive term and a spectral filtering term (see Eq (3)). The “spectral filtering term” is $\sum_{s=0}^{t-n-1} \mathbf{C} p_n(\mathbf{A}) \mathbf{A}^s \mathbf{B} \mathbf{u}_{t-n-s}$. We call it as such because its structure is amenable to spectral filtering. Indeed, we construct a set of spectral filters that satisfy the following.

Lemma E.4. Let $\mu(\alpha)$ be as defined in Eq (4) and \mathbf{Z} be as defined in Eq (5) and let $\phi_1, \dots, \phi_{T-n}$ be the eigenvectors of \mathbf{Z}_{T-n} . There is a universal constant $C > 0$ such that

$$\max_{\alpha \in \mathbb{C}_\beta} |\mu(\alpha)^\top \phi_j| \leq C \log(T)^{1/6} T^{2/3} e^{-j/6 \log T}.$$

The proof of Lemma E.4 is inspired by previous spectral filtering literature but requires several key and nontrivial adaptations to incorporate the polynomial $p_n^c(\cdot)$, to argue that the new spectral filtering matrix \mathbf{Z} defined over a complex domain exhibits exponential eigenvalue decay, and to generally extend the analysis to the case where the spectrum of \mathbf{A} may lie in the complex plane. We prove both lemmas in the subsequent subsections. With Lemma E.4 in hand, we may prove Lemma E.2

Proof of Lemma E.2 Suppose $\mathbf{y}_{1:T}$ evolves as a linear dynamical system characterized by $(\mathbf{A}, \mathbf{B}, \mathbf{C})$. Then given inputs $\mathbf{u}_{1:t}$,

$$\mathbf{y}_t = \sum_{s=1}^t \mathbf{C} \mathbf{A}^{t-s} \mathbf{B} \mathbf{u}_s.$$

With some linear algebra we get that convolving $\mathbf{y}_{t:t-n}$ with coefficients $\mathbf{c}_{0:n}$ results in

$$\sum_{i=0}^n \mathbf{c}_i \mathbf{y}_{t-i} = \sum_{s=0}^{n-1} \sum_{i=0}^s \mathbf{c}_i \mathbf{C} \mathbf{A}^{s-i} \mathbf{B} \mathbf{u}_{t-s} + \sum_{s=0}^{t-n-1} \mathbf{C} p_n^c(\mathbf{A}) \mathbf{A}^s \mathbf{B} \mathbf{u}_{t-n-s},$$

or equivalently (since $\mathbf{c}_0 = 1$ due to the assertion in Algorithm 2),

$$\mathbf{y}_t = - \sum_{i=1}^n \mathbf{c}_i \mathbf{y}_{t-i} + \sum_{s=0}^{n-1} \sum_{i=0}^s \mathbf{c}_i \mathbf{C} \mathbf{A}^{s-i} \mathbf{B} \mathbf{u}_{t-s} + \sum_{s=0}^{t-n-1} \mathbf{C} p_n^c(\mathbf{A}) \mathbf{A}^s \mathbf{B} \mathbf{u}_{t-n-s}. \quad (6)$$

Set $\hat{\mathbf{Q}}_s = \sum_{i=0}^s \mathbf{c}_i \mathbf{C} \mathbf{A}^{s-i} \mathbf{B}$ and set $\hat{\mathbf{Q}} = (\hat{\mathbf{Q}}_0, \dots, \hat{\mathbf{Q}}_{n-1})$. Then

$$\|\hat{\mathbf{Q}}_i\| \leq \sum_{j=0}^i |\mathbf{c}_j| \|\mathbf{C} \mathbf{A}^{i-j} \mathbf{B}\| \leq \|\mathbf{C}\| \|\mathbf{B}\| \sum_{j=0}^i |\mathbf{c}_j| \leq \|\mathbf{C}\| \|\mathbf{B}\| \|\mathbf{c}\|_1.$$

Next we turn our attention to the spectral filtering parameters. Recall our definition of $\mu_{p_n}(\alpha)$,

$$\mu(\alpha) \stackrel{\text{def}}{=} [1 \quad \alpha \quad \dots \quad \alpha^{t-n-1}]^\top$$

and

$$\mathbf{u}_{t:1} \stackrel{\text{def}}{=} [\mathbf{u}_t \quad \mathbf{u}_{t-1} \quad \dots \quad \mathbf{u}_1]^\top. \quad (7)$$

Let \mathbf{A} be diagonalized by some \mathbf{P} so that

$$\mathbf{A} = \mathbf{P} \mathbf{D} \mathbf{P}^{-1},$$

where \mathbf{D} is the diagonalization of \mathbf{A} , note that we can assume this w.l.o.g. since the set of diagonalizable matrices over the complex numbers is dense and therefore if \mathbf{A} is not diagonalizable we may apply an arbitrarily small perturbation to it. Recall the notation from Algorithm 3 that $\hat{\mathbf{u}}_{t:1}$ is $\mathbf{u}_{t:1}$

padded so that it has shape $(T - n - 1) \times d_{\text{in}}$. Then,

$$\begin{aligned}
\sum_{s=0}^{t-n-1} \mathbf{C} p_n^c(\mathbf{A}) \mathbf{A}^s \mathbf{B} \mathbf{u}_{t-n-s} &= \sum_{s=0}^{t-n-1} \mathbf{C} \mathbf{P} p_n^c(\mathbf{D}) \mathbf{D}^s \mathbf{P}^{-1} \mathbf{B} \mathbf{u}_{t-n-s} \\
&= \sum_{s=0}^{t-n-1} \mathbf{C} \mathbf{P} \left(\sum_{m=1}^{d_h} p_n^c(\alpha_m) \alpha_m^s \mathbf{e}_m \mathbf{e}_m^\top \right) \mathbf{P}^{-1} \mathbf{B} \mathbf{u}_{t-n-s} \\
&= \sum_{m=1}^{d_h} \sum_{s=0}^{t-n-1} \mathbf{C} \mathbf{P} \mathbf{e}_m \mathbf{e}_m^\top \mathbf{P}^{-1} \mathbf{B} p_n^c(\alpha_m) \alpha_m^s \mathbf{u}_{t-n-s} \\
&= \sum_{m=1}^{d_h} \mathbf{C} \mathbf{P} \mathbf{e}_m \mathbf{e}_m^\top \mathbf{P}^{-1} \mathbf{B} \mu_{p_n}(\alpha_m)^\top \tilde{\mathbf{u}}_{t-n-1:1} \\
&= \sum_{m=1}^{d_h} \mathbf{C} \mathbf{P} \mathbf{e}_m \mathbf{e}_m^\top \mathbf{P}^{-1} \mathbf{B} \mu_{p_n}(\alpha_m)^\top \left(\sum_{j=1}^{T-n} \phi_j \phi_j^\top \right) \tilde{\mathbf{u}}_{t-n-1:1} \\
&\quad \text{(Orthonormality of the filters)} \\
&= \sum_{j=1}^{T-n} \left(\sum_{m=1}^{d_h} \mathbf{C} \mathbf{P} \mathbf{e}_m \mathbf{e}_m^\top \mathbf{P}^{-1} \mathbf{B} \mu_{p_n}(\alpha_m)^\top \phi_j \right) \phi_j^\top \tilde{\mathbf{u}}_{t-n-1:1}.
\end{aligned}$$

Therefore defining

$$\hat{\mathbf{M}}_j \stackrel{\text{def}}{=} T^{1/2} \sum_{m=1}^{d_h} \mathbf{C} \mathbf{P} \mathbf{e}_m \mathbf{e}_m^\top \mathbf{P}^{-1} \mathbf{B} \mu_{p_n}(\alpha_m)^\top \phi_j,$$

we have

$$\sum_{s=0}^{t-n-1} \mathbf{C} p_n^c(\mathbf{A}) \mathbf{A}^s \mathbf{B} \mathbf{u}_{t-n-s} = \sum_{j=1}^{T-n} \hat{\mathbf{M}}_j \frac{\phi_j^\top \tilde{\mathbf{u}}_{t-n-1:1}}{\sqrt{T}}.$$

Next we bound the norm of $\hat{\mathbf{M}}_j$. Let \mathbf{S} be the diagonal matrix with entries $\mathbf{S}_{mm} = \mu_{p_n}(\alpha_m)^\top \phi_j$. Note that $\hat{\mathbf{M}}_j = \mathbf{C} \mathbf{P} \mathbf{S} \mathbf{P}^{-1} \mathbf{B}$. Recalling that $\mu_{p_n^c}(\alpha) = p_n^c(\alpha) \mu(\alpha)$,

$$\begin{aligned}
\|\hat{\mathbf{M}}_j\| &= \|T^{1/2} \mathbf{C} \mathbf{P} \mathbf{S} \mathbf{P}^{-1} \mathbf{B}\| \\
&\leq T^{1/2} \max_{m \in [d_h]} |p_n^c(\alpha_m)| \cdot \max_{m \in [d_h]} |\mu(\alpha_m)^\top \phi_j| \cdot \|\mathbf{C}\| \|\mathbf{P}\| \|\mathbf{P}^{-1}\| \|\mathbf{B}\|.
\end{aligned}$$

By Lemma E.4, there are universal constants $C > 0$ and $c > 1$ such that

$$\max_{\alpha \in \mathbb{C}_\beta} |\mu(\alpha)^\top \phi_j| \leq C \log(T)^{1/6} T^{2/3} c^{-j/6 \log T}.$$

Therefore,

$$\|\hat{\mathbf{M}}_j\| \leq C \log(T)^{1/6} T^{7/6} \max_{\alpha \in \mathbb{C}_\beta} |p_n^c(\alpha)| \cdot \|\mathbf{C}\| \|\mathbf{P}\| \|\mathbf{P}^{-1}\| \|\mathbf{B}\| \cdot c^{-j/6 \log T}.$$

Therefore $(\hat{\mathbf{Q}}_0, \dots, \hat{\mathbf{Q}}_{n-1}, \hat{\mathbf{M}}_1, \dots, \hat{\mathbf{M}}_k) \in \mathcal{K}$ as long as $R_Q \geq \|\mathbf{C}\| \|\mathbf{B}\| \|\mathbf{c}\|_1$ and $R_M \geq C \log(T)^{1/6} T^{7/6} \max_{\alpha \in \mathbb{C}_\beta} |p_n^c(\alpha)| \|\mathbf{C}\| \|\mathbf{P}\| \|\mathbf{P}^{-1}\| \|\mathbf{B}\|$. Therefore, recalling Eq. 6, we have shown,

$$\mathbf{y}_t = - \sum_{i=1}^n \mathbf{c}_i \mathbf{y}_{t-i} + \sum_{s=0}^{n-1} \sum_{i=0}^s \hat{\mathbf{Q}}_s \mathbf{u}_{t-s} + \sum_{j=1}^{T-n} \hat{\mathbf{M}}_j \frac{\phi_j^\top \tilde{\mathbf{u}}_{t-n-1:1}}{\sqrt{T}}.$$

Observe that $\hat{\mathbf{y}}(\hat{\mathbf{Q}}, \hat{\mathbf{M}})$ is nearly identical to the above, but truncates the terms involving $\hat{\mathbf{M}}_j$ after $j > k$, where k is a fixed parameter input to Algorithm 3. Therefore,

$$\hat{\mathbf{y}}_t(\hat{\mathbf{Q}}, \hat{\mathbf{M}}) - \mathbf{y}_t = \sum_{j=k+1}^{T-n} \hat{\mathbf{M}}_j \frac{\phi_j^\top \tilde{\mathbf{u}}_{t-n-1:1}}{\sqrt{T}}.$$

Therefore, to bound the error we have

$$\begin{aligned}
\|\hat{\mathbf{y}}_t(\hat{\mathbf{Q}}, \hat{\mathbf{M}}) - \mathbf{y}_t\|_2 &= \left\| \sum_{j=k+1}^{T-n} \mathbf{M}_j \phi_j^\top \mathbf{u}_{(t-n-1):1} T^{-1/2} \right\|_2 \\
&\leq \sum_{j=k+1}^{T-n} \|\hat{\mathbf{M}}_j\| \|\phi_j^\top \mathbf{u}_{(t-n-1):1} T^{-1/2}\|_2 \\
&\leq \sum_{j=k+1}^{T-n} \|\hat{\mathbf{M}}_j\| \|\phi_j^\top\|_1 \|\mathbf{u}_{(t-n-1):1}\|_\infty T^{-1/2} \\
&\leq \sum_{j=k+1}^{T-n} \|\hat{\mathbf{M}}_j\| \quad (\|\phi_j\|_2 = 1 \text{ and therefore } \|\phi_j\|_1 \leq \sqrt{T}) \\
&\leq T R_M \\
&\leq \epsilon / \sqrt{d_{\text{out}}}. \quad (\text{For } k \geq 6 \log(T) \log(R_M / \sqrt{d_{\text{out}}}) / \log(c).)
\end{aligned}$$

Finally,

$$\|\hat{\mathbf{y}}_t(\hat{\mathbf{Q}}, \hat{\mathbf{M}}) - \mathbf{y}_t\|_1 \leq \sqrt{d_{\text{out}}} \|\hat{\mathbf{y}}_t(\hat{\mathbf{Q}}, \hat{\mathbf{M}}) - \mathbf{y}_t\|_2 \leq \epsilon.$$

■

E.1.1 Proof of the Spectral Filtering Property Lemma E.4

In order to prove Lemma E.4 we require two further helper lemmas. The first is Lemma E.5 which roughly argues that the Lipschitz constant of a function $f : \mathbb{C}_\beta \rightarrow \mathbb{R}$ can be bounded by a polynomial of the expectation of the function on \mathbb{C}_β . The second is Theorem E.6 which argues that the matrix \mathbf{Z} from Eq. 5 that we use to derive the new spectral filters exhibits exponentially fast decay in its spectrum. This is a nontrivial and novel proof that does not follow from previous spectral filtering literature and, as such, we have chosen to use the word ‘‘Theorem’’ rather than ‘‘Lemma’’ to refer to it.

Lemma E.5. Let $L > 0$, $g_{\max} > 0$ and $0 \leq \beta \leq 1$. Define

$$\mathbb{C}_\beta = \{z \in \mathbb{C} : |z| \leq 1, |\arg(z)| \leq \beta\},$$

and let \mathcal{F} be the set of non-negative, L -Lipschitz functions $f : \mathbb{C}_\beta \rightarrow \mathbb{R}$ such that $\max_{z \in \mathbb{C}_\beta} f(z) = g_{\max}$. Then

$$\min_{f \in \mathcal{F}} \int_{\mathbb{C}_\beta} f(z) dz \geq \frac{\beta g_{\max}^3}{24 L^2}.$$

Proof of Lemma E.5 **1. The extremal function.** Fix $f \in \mathcal{F}$ and choose z_\star with $f(z_\star) = g_{\max}$. Among admissible functions we consider only those that decrease as fast as the Lipschitz constraint allows in every radial direction, replacing f by the *extremal cone*. Set

$$h(z) := [g_{\max} - L|z - z_\star|]_+, \quad r := \frac{g_{\max}}{L}.$$

Because $0 \leq h \leq f$ on \mathbb{C}_β , it suffices to lower bound $\int_{\mathbb{C}_\beta} h$.

2. Geometry around z_\star . It can be seen by Euclidean geometry, the point $z_\star = 0$ minimize the volume of the intersection with \mathbb{C}_β , up to a factor of 4, to make it the extremal function. It is depicted in Figure 2. Notice that the area of the intersection is colored in blue, and it is particularly easy to integrate over.

3. Integrating h .

$$\int_{\mathbb{C}_\beta} h \geq \int_0^\beta \int_0^r (g_{\max} - L\rho) \rho d\rho d\varphi = \frac{\beta g_{\max}^3}{6 L^2}.$$

Since $h \in \mathcal{F}$, and is the extremal function, we obtain the Lemma statement. ■




Figure 2: The origin minimizes the intersection area up to factor $\frac{1}{4}$.

Theorem E.6. *Given horizon T let*

$$\mu(\alpha) \stackrel{\text{def}}{=} [1 \quad \alpha \quad \dots \quad \alpha^{T-1}]^\top,$$

and

$$\mathbf{Z}_T \stackrel{\text{def}}{=} \int_{\alpha \in \mathbb{C}_\beta} \mu(\alpha) \mu(\bar{\alpha})^\top d\alpha,$$

where $\bar{\alpha} \in \mathbb{C}$ denotes the complex conjugate and $\mathbb{C}_\beta = \{z \in \mathbb{C} \text{ s.t. } |z| \leq 1 \text{ and } |\arg(z)| \leq \theta\}$. There exist constants $C > 0$ and $c > 1$, independent of T , such that if $\sigma_j(\mathbf{Z})$ denotes the j -th largest singular value of \mathbf{Z} ,

$$\sigma_j \leq C\beta \cdot c^{-j/\log T} \cdot \log T.$$

This proof is nontrivial and does not follow from previous results on Spectral Filtering due to the fact that it is no longer a Hankel matrix. It also does not immediately follow from general work on spectral decay, for instance [12], since it does not have an obvious displacement structure. We prove its spectral decay borrowing heavily from techniques from [12], but we must introduce a novel Gaussian quadrature argument that applies to the complex plane.

Proof of Theorem E.6 Observe that we can write \mathbf{Z}_{jk} as

$$\mathbf{Z}_{jk} = \int_{z \in \mathbb{C}_\beta} z^j \bar{z}^k dz,$$

or, equivalently

$$\begin{aligned} \mathbf{Z}_{jk} &= \int_{r \in [0,1]} \int_{\theta \in \mathbb{C}_\beta} (re^{i\theta})^j (re^{-i\theta})^k r dr d\theta \\ &= \int_{r \in [0,1]} \int_{\theta \in \mathbb{C}_\beta} r^{j+k+1} \cos(|j-k|\theta) dr d\theta + i \cdot \text{sign}(j-k) \int_{r \in [0,1]} \int_{\theta \in \mathbb{C}_\beta} \sin(|j-k|\theta) dr d\theta. \end{aligned}$$

We will use Gaussian quadrature in what follows. We have two non-standard features that we must manage. First, it is commonly known that Gaussian quadrature with n nodes is exact when integrating a polynomial of degree $2n-1$ or less. However, we need to ensure exactness when integrating a trigonometric function. The second non-standard feature is that we are integrating a 2-dimensional function, ie our domain is 2-dimensional specified by (r, θ) , while standard Gaussian quadrature applies only to 1-dimensional functions.

For the first non-standard feature, [31] shows that Gaussian quadrature is still exact even when integrating trigonometric functions. Specifically, they show that if the domain of integration is some compact set \mathcal{C} , then there exist $n+1$ weights $w_1^{\text{trig}}, \dots, w_n^{\text{trig}}$ and nodes $\theta_1, \dots, \theta_n$ such that for any $f(\theta)$ which can be written in the form

$$f(\theta) = \sum_{\ell=1}^n a_\ell \sin(\ell\theta) + \sum_{\ell=0}^n b_\ell \cos(\ell\theta), \quad (8)$$

we have

$$\int_{\theta \in \mathcal{C}} f(\theta) d\theta = \sum_{\ell=1}^{n+1} w_\ell^{\text{trig}} f(\theta_\ell). \quad (9)$$

Let $f_1^{(j,k)}(r, \theta) = r^{j+k+1} \cos(|j-k|\theta)$ and $f_2^{(j,k)}(r, \theta) = r^{j+k+1} \sin(|j-k|\theta)$. Using the above, observe that both $f_1^{(j,k)}$ and $f_2^{(j,k)}$ can be written as polynomials in r and θ with degree at most $2|j-k| \leq 2T$.

For the second nonstandard feature, i.e. that we are integrating over a 2-dimensional function, we can use Tchakaloff's Theorem, which is an extension of quadrature rules to general Borel measures on \mathbb{R}^d . Given any positive Borel measure on \mathcal{R}^d , μ , which has compact support, Tchakaloff's Theorem guarantees the existence of $n+1$ weights w_1, \dots, w_{n+1} and nodes v_1, \dots, v_{n+1} such that for any polynomial $P(x)$ on \mathcal{R}^d with degree at most n ,

$$\int_{\mathcal{R}^d} P(x) \mu(dx) = \sum_{i=1}^{n+1} w_i P(v_i). \quad (10)$$

Since our measure is compactly supported (we have $dr d\theta$ integrated over a compact domain) we have that there exist weights w_1^2, \dots, w_{2T+1}^2 and nodes $(r_1, \theta_1), \dots, (r_{2T+1}, \theta_{2T})$ such that for any $j, k \leq T$ both $f_1^{(j,k)}(r, \theta)$ and $f_2^{(j,k)}(r, \theta)$ satisfy

$$\int_{r \in [0,1]} \int_{\theta \in \mathbb{C}_\beta} f_i^{(j,k)}(r, \theta) = \sum_{\ell=1}^{2T} w_\ell f_1^{(j,k)}(r_\ell, \theta_\ell).$$

Therefore

$$\begin{aligned} \mathbf{Z}_{jk} &= \sum_{\ell=1}^{2T+1} \left(w_\ell^2 r_\ell^{j+k+1} \cos(|j-k|\theta_\ell) \right) + i \operatorname{sign}(j-k) \left(\sum_{\ell=1}^{n+1} w_\ell^2 r_\ell^{j+k+1} \sin(|j-k|\theta_\ell) \right) \\ &= \sum_{\ell=1}^{2T+1} w_\ell^2 r_\ell^{j+k+1} (\cos((j-k)\theta_\ell) + i \sin((j-k)\theta_\ell)) \\ &= \sum_{\ell=1}^{2T+1} w_\ell^2 r_\ell^{j+k+1} e^{i(j-k)\theta_\ell} \\ &= \sum_{\ell=1}^{2T+1} w_\ell^2 (r_\ell e^{i\theta_\ell})^j (r_\ell e^{-i\theta_\ell})^k \\ &= \sum_{\ell=1}^{2T+1} w_\ell^2 z_\ell^j \bar{z}_\ell^k, \end{aligned}$$

where $z_\ell \stackrel{\text{def}}{=} r_\ell e^{i\theta_\ell}$. We now proceed by using the same proof technique from [12]. Indeed, we have shown that \mathbf{Z} has a so-called ‘‘Fiedler factorization’’,

$$\mathbf{Z} = \mathbf{K} \mathbf{K}^*, \quad (11)$$

where

$$\begin{aligned} \mathbf{w} &= [w_1 \quad \dots \quad w_{2T+1}], \\ \mathbf{z} &= [z_1 \quad \dots \quad z_{2T+1}], \\ \mathbf{K} &= [\mathbf{w} \quad \operatorname{diag}(\mathbf{z})\mathbf{w} \quad \dots \quad (\operatorname{diag}(\mathbf{z}))^{2T-1} \mathbf{w}]. \end{aligned}$$

Indeed,

$$\begin{aligned}
(\mathbf{K}\mathbf{K}^*)_{jk} &= \sum_{\ell=1}^{2T-1} (\mathbf{K})_{j\ell} \mathbf{K}_{\ell k}^* \\
&= \sum_{\ell=1}^{2T-1} (\text{diag}(\mathbf{z})^j \mathbf{w})_{\ell} (\text{diag}(\mathbf{z}^*)^k \mathbf{w})_{\ell} \\
&= \sum_{\ell=1}^{2T-1} (\text{diag}(\mathbf{z})^j \mathbf{w})_{\ell} (\text{diag}(\mathbf{z}^*)^k \mathbf{w})_{\ell} \\
&= \sum_{\ell=1}^{2T-1} z_{\ell}^j w_{\ell} \bar{z}_{\ell}^k w_{\ell} \\
&= \mathbf{Z}_{jk}.
\end{aligned}$$

Let

$$\mathbf{A} \stackrel{\text{def}}{=} \text{diag}(\mathbf{z}) \quad \text{and} \quad \mathbf{Q} \stackrel{\text{def}}{=} \begin{bmatrix} 0 & & & -1 \\ 1 & & & \\ & \ddots & & \\ & & 1 & 0 \end{bmatrix}.$$

Observe that both \mathbf{A} and \mathbf{Q} are normal and

$$\mathbf{A}\mathbf{Z} - \mathbf{Z}\mathbf{Q} = \mathbf{s}\mathbf{e}_{2T-1}^{\top},$$

where $\mathbf{s} \in \mathbb{C}^{2T-1}$ and \mathbf{e}_{2T-1} is the final basis vector $(0, \dots, 0, 1)$. Therefore, by Theorem 2.1 from [12], if E is a set containing the eigenvalues of \mathbf{A} and F is a set containing the eigenvalues of \mathbf{Q} ,

$$\sigma_{j+k}(\mathbf{Z}) \leq Z_k(E, F)\sigma_j(\mathbf{Z}). \quad (12)$$

Therefore, our next task is to separate the spectrum of \mathbf{A} and \mathbf{Q} . Let $n = 2T + 1$. As shown in [12], \mathbf{Q} has eigenvalues,

$$\sigma(\mathbf{Q}) = \left\{ x \in \mathbb{C} : x = \exp(2\pi i(j + \frac{1}{2})/n), 0 \leq j \leq n-1 \right\}.$$

Observe that for any $x \in \sigma(\mathbf{Q})$, $|x| = 1$. We will bound the Zolotarev number $Z(E, F)$ by showing that the eigenvalues of \mathbf{A} , i.e. $\{z_1, \dots, z_{2T+1}\}$ defined above have magnitude bounded away from 1. This is a major departure from what is done in [12] and provides a novel way of showing that matrices with non-Hermitian Fiedler factorizations still exhibit spectral decay. We proceed by showing that $\max_{\ell \in [2T-1]} |r_{\ell}| \leq 1 - \delta_T$ for $\delta_T = c/(T+1)^2$, where c is an absolute constant independent of T . To establish this bound, we use the fact that the quadrature nodes $\{r_{\ell}\}$ are the Gauss–Legendre nodes on the interval $[0, 1]$ with $n = 2T + 1$ points. By definition of the Gauss–Legendre rule, these nodes are obtained by an affine transformation of the zeros of the Legendre polynomial $P_n(x)$ on $[-1, 1]$, namely

$$r_{\ell} = \frac{1 + x_{n,\ell}}{2}, \quad \ell = 1, \dots, n,$$

where $x_{n,1} < \dots < x_{n,n}$ are the ordered zeros of $P_n(x)$. For a proof of this fact see Section 3.6 of [37]. Theorem 3 of [20] provides sharp bounds on the extreme zeros of Jacobi polynomials. Specializing their result to the Legendre case ($\alpha = \beta = 0$), they show that the largest zero $x_{n,n}$ satisfies

$$1 - \frac{3}{4(n+1)^2} < x_{n,n} < 1 - \frac{1}{2(2n+1)^2}.$$

This inequality implies that the largest Gauss–Legendre node lies strictly below 1 by an amount proportional to $1/n^2$. Substituting $r_{\max} = (1 + x_{n,n})/2$ and rearranging gives

$$1 - r_{\max} = \frac{1 - x_{n,n}}{2} \geq \frac{c}{(n+1)^2}$$

for some universal constant $c > 0$. Recalling that $n = 2T + 1$, we therefore have

$$r_{\max} \leq 1 - \frac{c}{(2T+1)^2}.$$

Hence every quadrature node r_ℓ lies strictly inside the unit interval and, consequently, each complex node $z_\ell = r_\ell e^{i\theta_\ell}$ used in the factorization satisfies $|z_\ell| \leq 1 - \delta_T$ with $\delta_T = c/(2T+1)^2$.

We now use this to bound the Zolotarev number between our node set and the spectrum of Q . Let

$$E = \sigma(\mathbf{A}) = \{z_\ell : z_\ell = r_\ell e^{i\theta_\ell}, \ell = 0, \dots, 2T-1\} \quad \text{and} \quad F = \sigma(\mathbf{Q}) = \{e^{2\pi i(j+1/2)/n} : j = 0, \dots, n-1\}.$$

Let

$$\mathbb{D} \stackrel{\text{def}}{=} \{z \in \mathbb{C} : |z| < 1\} \quad \text{the open unit disk, and} \quad \mathbb{T} \stackrel{\text{def}}{=} \{z \in \mathbb{C} : |z| = 1\}$$

its boundary. For $\rho > 0$ let

$$\overline{D}_\rho \stackrel{\text{def}}{=} \{z \in \mathbb{C} : |z| \leq \rho\}.$$

Consider the ‘‘Cayley transform’’

$$\Phi(z) := i \frac{1+z}{1-z}.$$

This Möbius map sends $\mathbb{T} \setminus \{1\}$ onto the real line \mathbb{R} via $\Phi(e^{i\theta}) = \cot(\theta/2)$, and maps the open unit disk \mathbb{D} onto the open upper half-plane, with the identity

$$\text{Im}\Phi(z) = \frac{1-|z|^2}{|1-z|^2}, \quad z \in \mathbb{D}$$

(see, for instance, [8] Chapter 3).

Because F contains $e^{i\theta_j}$ where $\theta_j = (2j+1)\pi/n$ we have

$$\Phi(F) = \left\{ \cot\left(\frac{(2j+1)\pi}{2n}\right) \right\}_{j=0}^{n-1} \subset (-\infty, -1/t] \cup [1/t, \infty), \quad t := \tan\left(\frac{\pi}{2n}\right).$$

Hence the image of F consists of two symmetric subsets of the real line, separated by the gap $(-1/t, 1/t)$.

The image of E under Φ lies strictly above the real axis, since $E \subset \overline{D}_\rho \subset \mathbb{D}$. Using the identity for $\text{Im}\Phi(z)$, we find that for every $z \in E$,

$$\text{Im}\Phi(z) \geq \frac{1-\rho^2}{(1+\rho)^2} =: \eta > 0,$$

and moreover

$$|\Phi(z)| \leq \frac{1+\rho}{1-\rho} =: M.$$

Thus

$$\Phi(E) \subset \{w \in \mathbb{C} : \text{Im}w \geq \eta, |w| \leq M\},$$

a compact ‘‘cap’’ in the upper half-plane bounded away from \mathbb{R} by the horizontal line $\text{Im}w = \eta$.

To map this cap into the unit disk while keeping the real axis invariant, we use the Möbius transformation

$$\Psi_\eta(w) := \frac{w - i\eta}{w + i\eta}.$$

This function maps the line $\text{Im}(w) = \eta$ to \mathbb{T} and the region $\text{Im}(w) > \eta$ to \mathbb{D} . Therefore we have,

$$\Psi_\eta(\Phi(E)) \subset \overline{D}_{\rho'}, \quad \rho' := \max_{|w| \leq M, \text{Im}w \geq \eta} \left| \frac{w - i\eta}{w + i\eta} \right| = \frac{M - \eta}{M + \eta} = \frac{2\rho}{1 + \rho^2} < 1.$$

Observe that $1 - \rho' = \frac{(1-\rho)^2}{1+\rho^2}$. Since Ψ_η maps the real axis to \mathbb{T} , the set $\Psi_\eta(\Phi(F))$ lies on \mathbb{T} and consists of two arcs that are symmetric about 1, separated by the image of the real gap $(-1/t, 1/t)$.

Next, we apply the inverse Cayley transform

$$\Xi(u) := i \frac{1+u}{1-u},$$

which maps $\mathbb{T} \setminus \{1\}$ onto \mathbb{R} (see Chapter 3 of [8]). For any real x , one can compute that

$$\Xi(\Psi_\eta(x)) = i \frac{1 + \frac{x-i\eta}{x+i\eta}}{1 - \frac{x-i\eta}{x+i\eta}} = \frac{x}{\eta}.$$

Hence the real gap $(-1/t, 1/t)$ under this composition becomes the larger gap $(-\gamma, \gamma)$ with $\gamma = \frac{1}{\eta} \cdot \frac{1}{t}$. Next we compose this with the inversion map $y \mapsto 1/y$ which sends $\mathbb{R} \setminus (-\gamma, \gamma)$ to

$$[-1/\gamma, -\gamma] \cup [\gamma, 1/\gamma] = \left[-\frac{\eta}{t}, -\frac{t}{\eta}\right] \cup \left[\frac{t}{\eta}, \frac{\eta}{t}\right].$$

Therefore, the image of F under the overall Möbius transformation $A := y \mapsto 1/\Xi(\Psi_\eta(\Phi(y)))$ is a pair of symmetric real intervals which have the same form of those considered in [12, Lem. 5.1]. The Möbius transformations Ξ and $y \mapsto 1/y$ have real coefficients and therefore map circles and lines symmetric with respect to the real axis to other circles or lines that are likewise symmetric (see [8, Chapter 3]). Hence, the image of $\Psi_\eta(\Phi(E))$ under these maps remains contained in a disk that is symmetric about \mathbb{R} . From the computation of $\rho' = \frac{M-\eta}{M+\eta} = \frac{2\rho}{1+\rho^2}$, we obtain

$$1 - \rho' = \frac{(1 - \rho)^2}{1 + \rho^2}.$$

Thus, after these transformations, the image of E lies inside a disk whose distance from the unit circle is still of order δ_T^2 . For bounding purposes, we may replace this image by the concentric disk $\overline{D}_{\rho'}$, since enlarging the inner set only increases $Z_k(E, F)$.

Because Zolotarev numbers are invariant under composition with Möbius transformations (see [12, Def. 2.3 and Sec. 5]), we may compute

$$Z_k(E, F) = Z_k(A(\Phi(E)), A(\Phi(F))) \leq Z_k(\overline{D}_{\rho'}, [-\eta/t, -t/\eta] \cup [t/\eta, \eta/t]),$$

where we have used monotonicity in the inner set by replacing $A(\Phi(E))$ with the larger disk $\overline{D}_{\rho'}$.

For this canonical configuration of a disk inside two symmetric real intervals, the classical Zolotarev estimate—derived via elliptic function theory and stated in [12, Lem. 5.1]—implies that for absolute constants $C, C' > 0$,

$$Z_k \leq \exp\left(-\frac{\pi^2}{2} \frac{k}{\log(C/\tau) + \log(C'/(1 - \rho'))}\right),$$

where $\tau := t/\eta$. Using $t = \tan(\frac{\pi}{2n}) \asymp 1/n$, $\eta = \frac{1-\rho}{1+\rho} \asymp \delta_T$, and $1 - \rho' \asymp \delta_T^2$, we compute

$$\log \frac{1}{\tau} = \log \frac{\eta^{-1}}{t} = \log n + \log(1/\delta_T) + O(1), \quad \log \frac{1}{1 - \rho'} = 2 \log(1/\delta_T) + O(1).$$

Therefore there is a universal $c > 0$ such that

$$Z_k(E, F) \leq \exp\left(-c \frac{k}{\log n + \log(1/\delta_T)}\right) \leq \exp\left(-c \frac{k}{\log(4T^3)}\right).$$

Recalling Eq. [12], the result from [12], we have

$$\sigma_{j+k} \leq c^{-k/\log(T)} \sigma_j.$$

Given any index $m > 2$, either m is odd, $m = 1 + 2k$, or m is even, $m = 2 + 2k$. Then

$$\sigma_m = \sigma_{1+2k} \leq C \cdot c^{-2k/\log T} \cdot \sigma_1,$$

or

$$\sigma_m = \sigma_{2+2k} \leq C \cdot c^{-2k/\log T} \cdot \sigma_2.$$

Either way, we can bound $\max\{\sigma_1, \sigma_2\}$ by $\text{tr}(\mathbf{Z})$. Recall

$$\mathbf{Z}_{jk} = \int_{r \in [0,1]} \int_{\theta \in \mathbb{C}_\beta} (re^{i\theta})^j (re^{-i\theta})^k r dr d\theta.$$

Evaluating this integral for $j = k$ we get

$$\mathbf{Z}_{jj} = \frac{4\beta}{2j+2}.$$

Therefore,

$$\text{tr}(\mathbf{Z}) = 4\beta \sum_{j=1}^T \frac{1}{2(j+1)} \leq 2\beta \log(T).$$

Plugging this in, we conclude that for any $j \in [T]$,

$$\sigma_j \leq C\beta \cdot c^{-j/\log T} \cdot \log T.$$

■

With Lemma E.5 and Theorem E.6 in hand, we are ready to prove Lemma E.4

Proof of Lemma E.4 Let

$$f_j(\alpha) \stackrel{\text{def}}{=} |\phi_j^\top \mu(\alpha)|^2,$$

If $f_j(\alpha)$ is L -Lipschitz, letting $g_{\max} = \max_{\alpha \in \mathbb{C}_\beta} f_j(\alpha)$ by Lemma E.5

$$\int_{\alpha \in \mathbb{C}_\beta} f_j(\alpha) d\alpha \geq \frac{\beta g_{\max}^3}{24L^2},$$

or equivalently,

$$\max_{\alpha \in \mathbb{C}_\beta} f_j(\alpha) \leq \left(\frac{24L^2}{\beta} \int_{\alpha \in \mathbb{C}_\beta} f_j(\alpha) d\alpha \right)^{1/3}.$$

Observe that

$$\begin{aligned} \int_{\alpha \in \mathbb{C}_\beta} f_j(\alpha) d\alpha &= \int_{\alpha \in \mathbb{C}_\beta} |\phi_j^\top \mu(\alpha)|^2 d\alpha \\ &= \phi_j^\top \left(\int_{\alpha \in \mathbb{C}_\beta} \mu(\alpha) \mu(\bar{\alpha})^\top d\alpha \right) \bar{\phi}_j \\ &= \sigma_j. \end{aligned}$$

Therefore conclude we have shown,

$$\max_{\alpha \in \mathbb{C}_\beta} |\mu(\alpha)^\top \phi_j| = \max_{\alpha \in \mathbb{C}_\beta} \sqrt{f_j(\alpha)} \leq \left(\frac{24L^2}{\beta} \sigma_j \right)^{1/6}. \quad (13)$$

The remainder of the proof consists of bounding the Lipschitz constant L and bounding the eigenvalue σ_j . To bound the Lipschitz constant of f_j ,

$$\begin{aligned} L &\leq \max_{\alpha \in S} |f'_j(\alpha)| \\ &= \max_{\alpha \in S} 2|\text{Re}(\phi_j^\top \mu'(\alpha) \cdot \phi_j^\top \mu(\alpha))| \\ &\leq \max_{\alpha \in S} \|\phi_j\|_2^2 \cdot \|\mu(\alpha)\|_2 \cdot \|\mu'(\alpha)\|_2 \\ &= \max_{\alpha \in S} \|\mu(\alpha)\|_2 \cdot \|\mu'(\alpha)\|_2. \end{aligned}$$

We have

$$\begin{aligned} \|\mu(\alpha)\|_2^2 &= \left(\sum_{s=0}^{T-n-1} |\alpha|^{2s} \right) \\ &\leq T. \end{aligned}$$

Next, to bound $\|\mu'(\alpha)\|_2$ we observe,

$$(\mu'(\alpha))_j = \frac{d}{d\alpha} \alpha^j = j\alpha^{j-1} \mathbf{1}_{j>0}.$$

Therefore, $|(\mu'(\alpha))_j| \leq T$ and $\|\mu'(\alpha)\|_2^2 \leq T^3$. Thus,

$$L^2 \leq T^4. \quad (14)$$

By Theorem [E.6](#) there are universal constants $C > 0$ and $c > 1$ such that

$$\sigma_j \leq C\beta \cdot c^{-j/\log T} \cdot \log T.$$

Plugging this in, we get for $C' = 24C$,

$$\max_{\alpha \in \mathbb{C}_\beta} |\mu(\alpha)^\top \phi_j| \leq \left(24CT^4 c^{-j/\log T} \log T \right)^{1/6} = C' \log(T)^{1/6} T^{2/3} c^{-j/6 \log T}.$$

■

F Chebyshev Polynomials Evaluated in the Complex Plane

In this section we let T_n denote the n -th Chebyshev polynomial and let M_n denote the monic form.

Proof of Lemma [3.1](#) We use that $M_n(z) = T_n(z)/2^{n-1}$ and

$$T_n(z) = \cos(n \arccos(z)). \quad (15)$$

If $|\operatorname{Im}(z)| \leq 1/64n^2$ then by Lemma [F.2](#), $\arccos(z) \leq 1/n$. Therefore $n \arccos(z) \leq 1$ and so by Fact [F.1](#),

$$T_n(z) = \cos(n \arccos(z)) \leq 2 \quad (16)$$

Now we turn to the derivative $M'_n(z)$. It's a fact that

$$M'_n(z) = \frac{n}{2^{n-1}} U_{n-1}(z), \quad (17)$$

where U_{n-1} is the Chebyshev polynomial of the second kind. We next use the fact that

$$U_{n-1}(z) = \begin{cases} 2 \sum_{\substack{j=0 \\ j \text{ even}}}^n T_j(z), & n \text{ even}, \\ 2 \sum_{\substack{j=0 \\ j \text{ odd}}}^n T_j(z), & n \text{ odd}. \end{cases}$$

By Eq [\(16\)](#), $|T_j(z)| \leq 2$ for any j and therefore

$$|U_{n-1}(z)| \leq n. \quad (18)$$

Therefore,

$$|M'_n(z)| \leq \frac{n^2}{2^{n-1}}.$$

■

Fact F.1. Let $z \in \mathbb{C}$. Then $|\cos(z)| \leq 2$ whenever $|\operatorname{Im}| \leq 1$.

Proof of Fact [F.1](#)

$$\begin{aligned} |\cos(x + iy)| &= (\cos^2 x \cosh^2 y + \sin^2 x \sinh^2 y)^{1/2} \quad (\text{Uses standard complex cosine identity.}) \\ &= (\cos^2 x + \cos^2 x (\cosh^2 y - 1) + \sin^2 x \sinh^2 y)^{1/2} \\ &= (\cos^2 x + \cos^2 x \sinh^2 y + \sin^2 x \sinh^2 y)^{1/2} \quad (\cosh^2 y - \sinh^2 y = 1) \\ &= (\cos^2 x + \sinh^2 y)^{1/2} \quad (\cos^2 x + \sin^2 x = 1) \\ &\leq (1 + \sinh^2 y)^{1/2} \quad (\sinh^2 y \leq 2 \text{ when } |y| \leq 1.) \\ &\leq 2. \end{aligned}$$

■

Lemma F.2. Let $z \in \mathbb{C}$ with $|z| \leq 1$. Then $|\operatorname{Im}(\arccos(z))| \leq 1/n$ whenever $|\arg(z)| \leq 1/64n^2$.

Proof of Lemma F.2. Let $re^{i\theta} = z$ and assume $|\theta| \leq 1/64n^2$. We use the Taylor series for $\arccos(\cdot)$,

$$\begin{aligned} \arccos(re^{i\theta}) &= \frac{\pi}{2} - \sum_{k=0}^{\infty} a_k (re^{i\theta})^{2k+1} && (\text{For } a_k = \frac{(2k)!}{4^k (k!)^2 (2k+1)}) \\ &= \frac{\pi}{2} - \sum_{k=0}^{\infty} a_k r^{2k+1} e^{i(2k+1)\theta} && (\text{De Moivre's Theorem}) \\ &= \frac{\pi}{2} - \sum_{k=0}^{\infty} a_k r^{2k+1} \cos((2k+1)\theta) - i \sum_{k=0}^{\infty} a_k r^{2k+1} \sin((2k+1)\theta). \\ &&& (e^{i\theta} = \cos \theta + i \sin \theta) \end{aligned}$$

Therefore,

$$\text{Im}(\arccos(re^{i\theta})) = \sum_{k=0}^{\infty} a_k r^{2k+1} \sin((2k+1)\theta).$$

Then

$$\begin{aligned} |\text{Im}(\arccos(re^{i\theta}))| &\leq \sum_{k=0}^{\infty} a_k |r|^{2k+1} |\sin((2k+1)\theta)| \\ &\leq \sum_{k=0}^{\infty} a_k |r|^{2k+1} \min(1, (2k+1)|\theta|) && (|\sin(x)| \leq \min(|x|, 1)) \\ &\leq \sum_{k=0}^{\infty} a_k \min(1, (2k+1)|\arg(z)|) && (|r| \leq 1) \\ &\leq \sum_{k=0}^K a_k (2k+1)|\arg(z)| + \sum_{k=K+1}^{\infty} a_k. && (\text{For any arbitrary } K \geq 0) \end{aligned}$$

Now we bound a_k .

$$\begin{aligned} a_k(2k+1) &= \frac{(2k)!}{4^k (k!)^2} \\ &\leq \frac{\sqrt{2\pi(2k)}(2k/e)^{2k} (1 + \frac{1}{2k})}{4^k \left(\sqrt{2\pi k} (k/e)^k \right)^2} && (\text{Stirling's Formula}) \\ &= \left(1 + \frac{1}{2k} \right) / \sqrt{\pi k} \\ &\leq 1/\sqrt{k}. \end{aligned}$$

Therefore we also have that $a_k \leq 1/k^{3/2}$. Using this (and noting that $a_0 = 1$) we see,

$$\begin{aligned} |\text{Im}(\arccos(z))| &\leq |\arg(z)| \left(1 + \sum_{k=1}^K \frac{1}{\sqrt{k}} \right) + \sum_{k=K}^{\infty} \frac{1}{k^{3/2}} \\ &\leq 4 \left(|\arg(z)| \sqrt{K} + \frac{1}{\sqrt{K}} \right) \\ &\leq \frac{8}{\sqrt{K}} && (\text{For } |\arg(z)| \leq 1/K) \\ &\leq \frac{1}{n} && (\text{For } K \geq 64n^2.) \end{aligned}$$

Therefore, for $|\arg(z)| \leq 1/64n^2$, we have that

$$|\text{Im}(\arccos(z))| \leq 1/n.$$

■

Proof of Lemma 3.2. We bound the coefficients of the Chebyshev polynomial. From Chapter 22 of [3],

$$T_n(x) = \frac{n}{2} \sum_{m=0}^{\lfloor n/2 \rfloor} (-1)^m \frac{(n-m-1)!}{m!(n-2m)!} (2x)^{n-2m}. \quad (19)$$

Therefore

$$M_n(x) = \frac{1}{2^{n-1}} T_n(x) = n \sum_{m=0}^{\lfloor n/2 \rfloor} (-1)^m \frac{(n-m-1)!}{m!(n-2m)!} 2^{-2m} x^{n-2m}.$$

Let $c_m = \frac{(n-m-1)!}{m!(n-2m)!} 2^{-2m}$. Then

$$\begin{aligned} \max_{m=0,\dots,n} c_m &\leq \max_{m=0,\dots,n} \binom{n-m}{m} 4^{-m} \\ &\leq \max_{m=0,\dots,n} \left(\frac{(n-m)e}{4m} \right)^m && \left(\binom{n}{k} \leq (ne/k)^k \right) \\ &\leq \max_{c \in [0,1]} \left(\frac{(1-c)e}{4c} \right)^{cn} && (\text{Letting } m = cn) \\ &\leq 2^{0.3n}. && (\max_{c \in [0,1]} ((1-c)e/4c)^c \leq 2^{0.3}) \end{aligned}$$

■

G Plots from Experiments

The details of the experiments are in Section 4, please refer to that for specifics.

G.1 Experiments with Linear Dynamical System Data

As our theory shows, for Chebyshev polynomials (the same can be shown for Legendre polynomials), the coefficients of the polynomial grow exponentially and therefore the performance gains vanish after the degree is too high. However, learning the optimal coefficients is able to sidestep this issue.

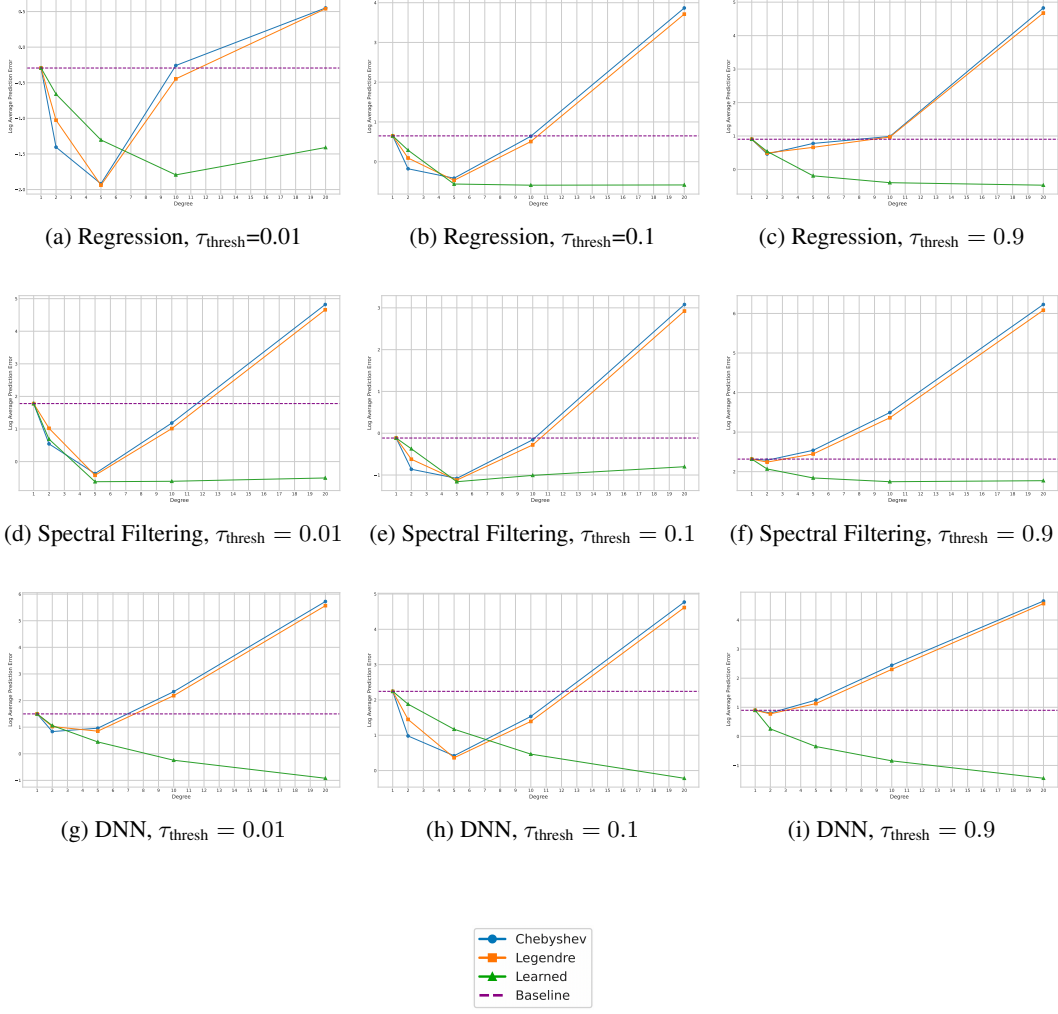


Figure 3: Absolute prediction error averaged over 200 independent runs with data generated from a linear dynamical system with varying complex threshold.

G.2 Nonlinear Data

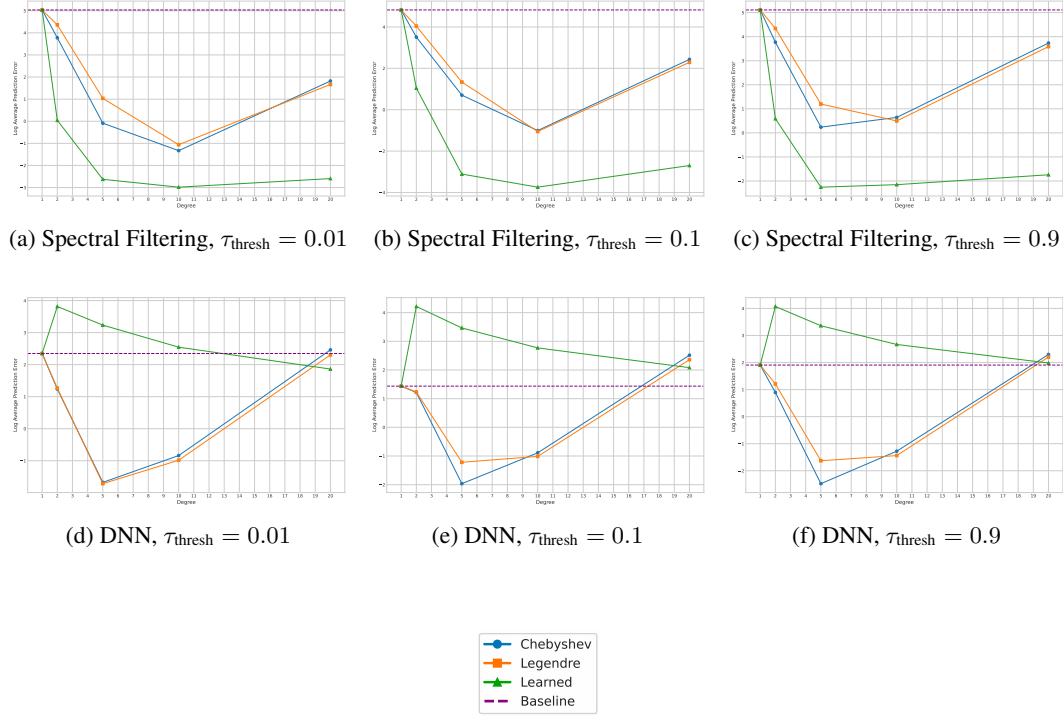


Figure 4: Absolute prediction error averaged over 200 independent runs with data generated from a nonlinear dynamical system with varying complex threshold.

G.3 Data from a DNN

Finally we generate data from a 10-layer sparse neural network which stacks 100-dimensional LSTMs with ReLU nonlinear activations.




Figure 5: Absolute prediction error of a 10-layer DNN model averaged over 200 independent runs.