

# On the Parameterization and Initialization of Diagonal State Space Models

Anonymous

## Abstract

State space models (SSM) have recently been shown to be very effective as a deep learning layer as a promising alternative to sequence models such as RNNs, CNNs, or Transformers. The first version to show this potential was the S4 model, which is particularly effective on tasks involving long-range dependencies by using a prescribed state matrix called the HiPPO matrix. While this has an interpretable mathematical mechanism for modeling long dependencies, it introduces a custom representation and algorithm that can be difficult to implement. On the other hand, a recent variant of S4 called DSS showed that restricting the state matrix to be fully diagonal can still preserve the performance of the original model when using a specific initialization based on approximating S4’s matrix. This work seeks to systematically understand how to parameterize and initialize such diagonal state space models. While it follows from classical results that almost all SSMs have an equivalent diagonal form, we show that the initialization is critical for performance. We explain why DSS works mathematically, by showing that the diagonal restriction of S4’s matrix surprisingly recovers the same kernel in the limit of infinite state dimension. We also systematically describe various design choices in parameterizing and computing diagonal SSMs, and perform a controlled empirical study ablating the effects of these choices. Our final model S4D is a simple diagonal version of S4 whose kernel computation requires just 2 lines of code and performs comparably to S4 in almost all settings, with state-of-the-art results for image, audio, and medical time-series domains, and averaging 85% on the Long Range Arena benchmark.

## 1 Introduction

A core class of models in modern deep learning are sequence models, which are parameterized mappings operating on arbitrary sequences of inputs. Recent approaches based on state space models (SSMs) have outperformed traditional deep sequence models such as recurrent neural networks (RNNs), convolutional neural networks (CNNs), and Transformers, in both computational efficiency and modeling ability. In particular, the S4 model displayed strong results on a range of sequence modeling tasks, especially on long sequences [9]. Its ability to model long-range dependencies arises from being defined with a particular state matrix called the “HiPPO matrix” [6], which allows S4 to be viewed as a convolutional model that decomposes an input onto an orthogonal system of smooth basis functions[1].

However, beyond its theoretical interpretation, actually computing S4 as a deep learning model requires a sophisticated algorithm with many linear algebraic techniques that are difficult to understand and implement. These techniques were necessitated by parameterizing its state matrix as a **diagonal plus low-rank** (DPLR) matrix, which is necessary to capture HiPPO matrices. A natural question is whether simplifications of this parameterization and algorithm are possible. In particular, removing the low-rank term would result in a **diagonal state space model** (DSSM) that is dramatically simpler to implement and understand.

Although it is known that almost all SSMs have an equivalent diagonal form—and therefore (complex) DSSMs are fully expressive algebraically—they may not represent all SSMs numerically, and finding a good initialization is critical. Gu et al. [9] showed that it is difficult to find a performant diagonal SSM, and that many alternative parameterizations of the state matrix – including by random diagonal matrices – are much less effective empirically, which motivated the necessity of the more complicated HiPPO matrix. However, recently Gupta [10] made the empirical observation that a variant of S4 using a *particular diagonal matrix* is nearly as effective as the original S4 method. This matrix is based on the original HiPPO matrix and is

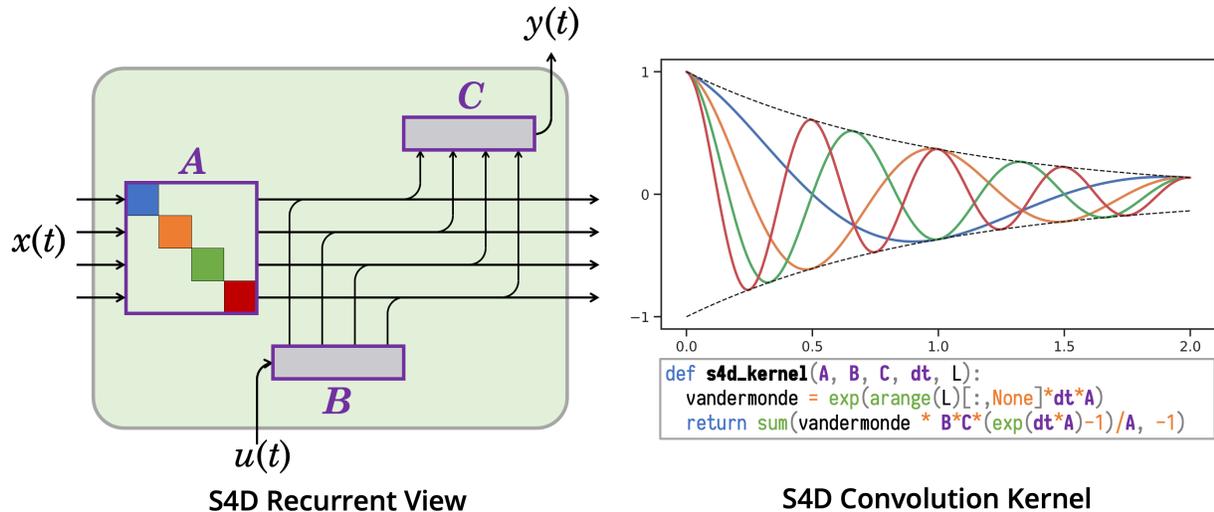


Figure 1: S4D is a diagonal SSM which inherits the strengths of S4 while being much simpler. (Left) The diagonal structure allows it to be viewed as a collection of 1-dimensional SSMs. (Right) As a convolutional model, S4D has a simple interpretable convolution kernel which can be implemented in two lines of code. Colors denote independent 1-D SSMs; purple denotes trainable parameters.

defined by simply chopping off the low-rank term in the DPLR representation.

The discovery of performant diagonal state matrices opens up new possibilities for simplifying deep state space models, and consolidating models such as S4 and DSS to understand and improve them. First, the strongest version of DSS *computes* the SSM with a complex-valued softmax that complicates the algorithm, and is actually less efficient than S4. Additionally, DSS and S4 differ in several auxiliary aspects of *parameterizing* SSMs that can conflate performance effects, making it more difficult to isolate the core effects of diagonal versus DPLR state matrices. Most importantly, DSS relies on *initializing* the state matrix to a particular approximation of S4’s HiPPO matrix. While S4’s matrix has a mathematical interpretation for addressing long-range dependencies, the efficacy of the diagonal approximation to it remains theoretically unexplained.

In this work, we seek to systematically understand how to train diagonal SSMs. We introduce the **S4D** method, a diagonal SSM which combines the best of S4’s *computation* and *parameterization* and DSS’s *initialization*, resulting in a method that is extremely simple, theoretically principled, and empirically effective.

- First, we describe S4D, a simple method outlined by S4 for computing diagonal instead of DPLR matrices, which is based on **Vandermonde matrix multiplication** and is even simpler and more efficient than the DSS. Outside of the core state matrix, we categorize different representations of the other components of SSMs, introducing flexible design choices that capture both S4 and DSS and allow different ways of parameterizing SSMs to be systematically compared (Section 3).
- We provide a new mathematical analysis of DSS’s initialization, showing that the diagonal approximation of the original HiPPO matrix surprisingly produces the same dynamics as S4 when the state size goes to infinity. We propose even simpler variants of diagonal SSMs using different initializations of the state matrix (Section 4).
- We perform a controlled study of these various design choices across many domains, tasks, and sequence lengths, and additionally compare diagonal (S4D) versus DPLR (S4) variants. Our best S4D methods are competitive with S4 on most settings, with near state-of-the-art results on image, audio, and medical time series benchmarks, and achieving **85%** on the Long Range Arena benchmark (Section 5).

## 2 Background

**Continuous State Spaces Models** S4 investigated state space models (1) that are parameterized maps on signals  $u(t) \mapsto y(t)$ . These SSMs are linear time-invariant systems that can be represented either as a linear ODE (equation (1)) or convolution (equation (2)).

$$\begin{aligned} x'(t) &= \mathbf{A}x(t) + \mathbf{B}u(t) & K(t) &= \mathbf{C}e^{t\mathbf{A}}\mathbf{B} \\ y(t) &= \mathbf{C}x(t) & y(t) &= (K * u)(t) \end{aligned} \quad (1) \quad (2)$$

Here the parameters are the state matrix  $\mathbf{A} \in \mathbb{C}^{N \times N}$  and other matrices  $\mathbf{B} \in \mathbb{C}^{N \times 1}$ ,  $\mathbf{C} \in \mathbb{C}^{1 \times N}$ . In the case of DSSMs where  $\mathbf{A}$  is diagonal, and we will overload notation so that  $\mathbf{A}_n, \mathbf{B}_n, \mathbf{C}_n$  denotes the entries of the parameters.

An intuitive way to view the convolution kernel (2) is to interpret it as a linear combination (controlled by  $\mathbf{C}$ ) of **basis kernels**  $K_n(t)$  (controlled by  $\mathbf{A}, \mathbf{B}$ )

$$K(t) = \sum_{n=0}^{N-1} \mathbf{C}_n K_n(t) \quad K_n(t) := \mathbf{e}_n^\top e^{t\mathbf{A}}\mathbf{B} \quad (3)$$

We denote this basis as  $K_{\mathbf{A}, \mathbf{B}}(t) = e^{t\mathbf{A}}\mathbf{B}$  if necessary to disambiguate; note that it is a vector of  $N$  functions. In the case of DSSMs, each function  $K_n(t)$  is just  $e^{t\mathbf{A}_n}\mathbf{B}_n$ .

**S4: Structured State Spaces** As a deep learning model, SSMs have many elegant properties with concrete empirical and computational benefits [8]. For example, the convolutional form (2) can be converted into a temporal recurrence that is substantially faster for autoregressive applications [5].

However, making SSMs effective required overcoming two key challenges: choosing appropriate values for the matrices, and computing the kernel (2) efficiently.

First, Gu et al. [8] showed that naive instantiations of the SSM do not perform well, and instead relied on a particular (real-valued) matrix  $\mathbf{A}$  called the HiPPO-LegS matrix (4).<sup>1</sup> These matrices were derived so that the basis kernels  $K_n(t)$  have closed-form formulas  $L_n(e^{-t})$ , where  $L_n(t)$  are normalized Legendre polynomials. Consequently, the SSM has a mathematical interpretation of decomposing the input signal  $u(t)$  onto a set of infinitely-long basis functions that are orthogonal respect to an exponentially-decaying measure, giving it long-range modeling abilities [1].

Second, S4 introduced a particular parameterization that decomposed this  $\mathbf{A}$  matrix into the sum of a normal and rank-1 matrix (5), which can be unitarily conjugated into a (complex) diagonal plus rank-1 matrix. Leveraging this structured form, they then introduced a sophisticated algorithm for efficiently computing the convolution kernel (2) for state matrices that are **diagonal plus low-rank (DPLR)**.

$$\begin{aligned} \mathbf{A}_{nk} &= - \begin{cases} (2n+1)^{\frac{1}{2}}(2k+1)^{\frac{1}{2}} & n > k \\ n+1 & n = k \\ 0 & n < k \end{cases} & \mathbf{A}_{nk}^{(N)} &= - \begin{cases} (n+\frac{1}{2})^{1/2}(k+\frac{1}{2})^{1/2} & n > k \\ \frac{1}{2} & n = k \\ (n+\frac{1}{2})^{1/2}(k+\frac{1}{2})^{1/2} & n < k \end{cases} \\ \mathbf{B}_n &= (2n+1)^{\frac{1}{2}} \quad \mathbf{P}_n = (n+1/2)^{\frac{1}{2}} & \mathbf{A} &= \mathbf{A}^{(N)} - \mathbf{P}\mathbf{P}^\top, \quad \mathbf{A}^{(D)} := \text{eig}(\mathbf{A}^{(N)}) \\ & \text{(HiPPO-LegS matrix used in S4)} & & \text{(Normal / diagonal plus low-rank form)} \end{aligned} \quad (4) \quad (5)$$

**DSS: Diagonal State Spaces** S4 was originally motivated by searching for a *diagonal state matrix*, which would be even more structured and result in very simple computation of the SSM. However, the HiPPO-LegS matrix cannot be stably transformed into diagonal form [9, Lemma 3.2], and they were unable to find any diagonal matrices that worked well, resulting in the DPLR formulation.

<sup>1</sup>HiPPO also specifies formulas for  $\mathbf{B}$ , but the state matrix  $\mathbf{A}$  is more important. There are many other HiPPO instantiations besides LegS, but HiPPO-LegS is the main one that S4 uses and the term ‘‘HiPPO matrix’’ without the suffix refers to this one.

Gupta [10] made the surprising empirical observation that simply removing the low-rank portion of the DPLR form of the HiPPO-LegS matrix results in a diagonal matrix that performs comparably to the original S4 method. More precisely, their initialization is the diagonal matrix  $\mathbf{A}^{(D)}$ , the diagonalization of  $\mathbf{A}^{(N)}$  in (5). They termed  $\mathbf{A}^{(N)}$  the *skew-HiPPO* matrix, which we will also call the *normal-HiPPO* matrix. To be more specific and disambiguate these variants, we will also call  $\mathbf{A}^{(N)}$  the HiPPO-LegS-N or HiPPO-N matrix and  $\mathbf{A}^{(D)}$  the HiPPO-LegS-D or HiPPO-D matrix.

In addition to this initialization, they proposed a method for computing diagonal state spaces. Beyond these two core differences, several aspects of their parameterization differ from S4’s.

In Sections 3 and 4, we systematically study the components of DSS: We categorize different ways to parameterize and compute the diagonal state space, and explain the theoretical interpretation of this particular diagonal  $\mathbf{A}$  matrix.

Because there are several different concrete matrices with different naming conventions, this table summarizes these special matrices and ways to refer to them.

Matrix	Full Name	Alternate Names
$\mathbf{A}$	HiPPO-LegS	HiPPO matrix, LegS matrix
$\mathbf{A}^{(N)}$	HiPPO-LegS-N	HiPPO-N, skew-HiPPO, normal-HiPPO
$\mathbf{A}^{(D)}$	HiPPO-LegS-D	HiPPO-D, diagonal-HiPPO

### 3 Parameterizing Diagonal State Spaces

We describe various choices for the computation and parameterization of diagonal state spaces. Our categorization of these choices leads to simple variants of the core method. Both DSS and our proposed S4D can be described using a combination of these factors (Section 3.4).

#### 3.1 Discretization

The true continuous-time SSM computes  $y(t) = (K * u)(t) = \int_0^\infty \mathbf{C}e^{s\mathbf{A}}\mathbf{B}u(t-s) ds$ .

In discrete time, we view an input sequence  $u_0, u_1, \dots$  as uniformly-spaced samples from an underlying function  $u(t)$  and must approximate this integral. Standard methods for doing so that preserve the convolutional structure of the model exist. The first step is to discretize the parameters. Two simple choices that have been used in prior work include

$$\begin{aligned}
 \text{(Bilinear)} \quad \bar{\mathbf{A}} &= (\mathbf{I} - \Delta/2\mathbf{A})^{-1}(\mathbf{I} + \Delta/2\mathbf{A}) & \text{(ZOH)} \quad \bar{\mathbf{A}} &= \exp(\Delta\mathbf{A}) \\
 \bar{\mathbf{B}} &= (\mathbf{I} - \Delta/2\mathbf{A})^{-1} \cdot \Delta\mathbf{B} & \bar{\mathbf{B}} &= (\Delta\mathbf{A})^{-1}(\exp(\Delta \cdot \mathbf{A}) - \mathbf{I}) \cdot \Delta\mathbf{B}.
 \end{aligned}$$

With these methods, the discrete-time SSM output is just

$$y = u * \bar{\mathbf{K}} \quad \text{where } \bar{\mathbf{K}} = (\mathbf{C}\bar{\mathbf{B}}, \mathbf{C}\bar{\mathbf{A}}\bar{\mathbf{B}}, \dots, \mathbf{C}\bar{\mathbf{A}}^{L-1}\bar{\mathbf{B}}). \quad (6)$$

These integration rules have both been used in prior works (e.g. LMU and DSS use ZOH [10, 23] while S4 and its predecessors use bilinear [6, 8, 9]).

In Section 5, we show that there is little empirical difference between them. However, we note that there is a curious phenomenon where the bilinear transform actually perfectly smooths out the kernel used in DSS to match the S4 kernel (Section 4 Fig. 2d). We additionally note that numerical integration is a rich and well-studied topic and more stable methods of approximating the convolutional integral may exist. For example, it is well-known that simple rules like the Trapezoid rule [15] can dramatically reduce numerical integration error when the function has bounded second derivative.

### 3.2 Convolution Kernel

The main computational difficulty of the original S4 model is computing the convolution kernel  $\overline{\mathbf{K}}$ . This is extremely slow for general state matrices  $\mathbf{A}$ , and S4 introduced a complicated algorithm for DPLR state matrices. When  $\overline{\mathbf{A}}$  is diagonal, the computation is nearly trivial. Note that

$$\overline{\mathbf{K}}_\ell = \sum_{n=0}^{N-1} \mathbf{C}_n \overline{\mathbf{A}}_n^\ell \overline{\mathbf{B}}_n \implies \overline{\mathbf{K}} = (\overline{\mathbf{B}}^\top \circ \mathbf{C}) \cdot \mathcal{V}_L(\overline{\mathbf{A}}) \quad \text{where } \mathcal{V}_L(\overline{\mathbf{A}})_{n,\ell} = \overline{\mathbf{A}}_n^\ell \quad (7)$$

where  $\circ$  is Hadamard product,  $\cdot$  is matrix multiplication, and  $\mathcal{V}$  is known as a **Vandermonde matrix**. Unpacking this a little more, we can write  $\overline{\mathbf{K}}$  as the following matrix-vector multiplication.

$$\overline{\mathbf{K}} = [\overline{\mathbf{B}}_0 \mathbf{C}_0 \quad \dots \quad \overline{\mathbf{B}}_{N-1} \mathbf{C}_{N-1}] \begin{bmatrix} 1 & \overline{\mathbf{A}}_0 & \overline{\mathbf{A}}_0^2 & \dots & \overline{\mathbf{A}}_0^{L-1} \\ 1 & \overline{\mathbf{A}}_1 & \overline{\mathbf{A}}_1^2 & \dots & \overline{\mathbf{A}}_1^{L-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \overline{\mathbf{A}}_{N-1} & \overline{\mathbf{A}}_{N-1}^2 & \dots & \overline{\mathbf{A}}_{N-1}^{L-1} \end{bmatrix}$$

**Time and Space Complexity** The naive way to compute (7) is by materializing the Vandermonde matrix  $\mathcal{V}_L(\overline{\mathbf{A}})$  and performing a matrix multiplication, which requires  $O(NL)$  time and space.

However, Vandermonde matrices are well-studied and theoretically the multiplication can be computed in  $\tilde{O}(N+L)$  operations and  $O(N+L)$  space. In fact, Vandermonde matrices are closely related to Cauchy matrices, which are the computational core of S4’s DPLR algorithm, and have identical complexity [14].

**Proposition 1.** *The time and space complexity of computing the kernel of diagonal SSMs is equal to that of computing DPLR SSMs.*

We note that on modern parallelizable hardware such as GPUs, a simple fast algorithm is to compute (7) with naive summation (using  $O(NL)$  operations), but without materializing the Vandermonde matrix (using  $O(N+L)$  space). Just as with S4, this may require implementing a custom kernel in some modern deep learning frameworks such as PyTorch to achieve the space savings.

### 3.3 Parameterization

The next question is how to represent the parameters  $\mathbf{A}, \mathbf{B}, \mathbf{C}$ .

**Parameterization of  $\mathbf{A}$ .** Note that the kernel  $K(t) = \mathbf{C}e^{t\mathbf{A}}\mathbf{B}$  blows up to  $\infty$  as  $t \rightarrow \infty$  if  $\mathbf{A}$  has any eigenvalues with positive real part. Goel et al. [5] found that this is a serious constraint that affects the stability of the model, especially when using the SSM as an autoregressive generative model. They propose to force the real part of  $\mathbf{A}$  to be negative, also known as the left-half plane condition in classical controls, by parameterizing the real part inside an exponential function  $\mathbf{A} = -\exp(\mathbf{A}_{Re}) + i \cdot \mathbf{A}_{Im}$ .

We note that instead of exp, any activation function can be used as long as its range is bounded on one side, such as ReLU, softplus, etc. The original DSS does not constrain the real part of  $\mathbf{A}$ , which is sufficient for simple tasks involving fixed-length sequences, but could become unstable in other settings.

**Parameterization of  $\mathbf{B}, \mathbf{C}$ .** Another choice in the parameterization is how to represent  $\mathbf{B}$  and  $\mathbf{C}$ . Note that the computation of the final discrete convolution kernel  $\overline{\mathbf{K}}$  depends only on the elementwise product  $\mathbf{B} \circ \mathbf{C}$  (equation (7)). Therefore DSS chose to parameterize this product directly, which they call  $\mathbf{W}$ , instead of  $\mathbf{B}$  and  $\mathbf{C}$  individually.

However, we observe that this is equivalent to keeping independent  $\mathbf{B}$  and  $\mathbf{C}$ , and simply freezing  $\mathbf{B} = \mathbf{1}$  while training  $\mathbf{C}$ . Therefore, just as S4 has separate parameters  $\mathbf{A}, \mathbf{B}$ , and  $\mathbf{C}$  and uses a fixed initialization

for  $\mathbf{A}$  and  $\mathbf{B}$ , S4D also proposes separate  $\mathbf{A}$ ,  $\mathbf{B}$ , and  $\mathbf{C}$  and uses fixed initializations for  $\mathbf{A}$  (discussed in Section 4) and  $\mathbf{B}$  (set to  $\mathbf{1}$ ). Then the difference between S4D and DSS is simply that DSS does not train  $\mathbf{B}$ . In our ablations, we show that training  $\mathbf{B}$  gives a minor but consistent improvement in performance.

As described in [1], S4 initializes  $\mathbf{C}$  randomly with standard deviation 1 (in contrast to standard deep learning initializations, which scale with the dimension e.g.  $N^{-\frac{1}{2}}$ ), which is variance-preserving for S4’s ( $\mathbf{A}$ ,  $\mathbf{B}$ ) as a consequence of the HiPPO theory. Because it turns out that the diagonal approximation to HiPPO has similar theoretical properties, we retain this initialization in the diagonal case.

### 3.4 S4D: the Diagonal Version of S4

A key component of our exposition is disentangling the various choices possible in representing and computing state space models. With this categorization, different choices can be mixed and matched to define variants of the core method. Table 1 compares S4, DSS, and S4D, which have a core structure and kernel computation, but have various choices of other aspects of the parameterization.

Table 1: (**Parameterization choices for Structured SSMS.**) Aside from the core structure of  $\mathbf{A}$  and the computation of its convolution kernel, SSMS have several design choices which are consolidated in S4D.

Method	Structure	Kernel Computation	Discretization	Constraint $\Re(\mathbf{A})$	Trainable $\mathbf{B}$	Initialization of $\mathbf{A}$
<b>S4</b>	DPLR	Cauchy	Bilinear	exp	Yes	HiPPO
<b>DSS</b>	diagonal	softmax	ZOH	id (none)	No	HiPPO-D
<b>S4D</b>	diagonal	Vandermonde	either	exp / ReLU	optional	various

**Comparison to S4 and DSS.** We will define the base version of S4D to match the parameterization of S4 (i.e. bilinear discretization,  $\Re(\mathbf{A})$  parameterized with exp, trainable  $\mathbf{B}$ , and HiPPO-D initialization), but many other variants are possible. Note that unlike DSS, the output of S4D would be *exactly the same as masking out the low-rank component of S4’s DPLR representation*. Thus comparing S4D vs. S4 is a comparison of diagonal vs. DPLR representations of  $\mathbf{A}$  while controlling all other factors. In our empirical study in Section 5, we systematically ablate the effects of each of these components.

We elaborate more on the comparisons between S4, DSS, and S4D below.

**Kernel computation.** The original S4 work briefly considered the diagonal case as motivation [9, Section 3.1], and explicitly mentioned the connection to Vandermonde products and the computational complexity of diagonal SSMS. However, their focus was the more complex DPLR representation because it is difficult to find a performant diagonal state matrix. Compared to S4, we fleshed out details of the Vandermonde connection and its computational complexity, which matches that of S4.

On the other hand, DSS empirically found an effective diagonal state matrix, but introduced a more complicated method based on a **complex softmax** for computing it. Compared to S4D, this softmax essentially *normalizes* by the row-sums of the Vandermonde matrix, so we also refer to this distinction as “softmax normalization”. This makes the kernel more complicated than necessary, and has a few concrete drawbacks. First, the row-normalization effectively makes the model dependent on a particular sequence length  $L$ , and special logic is required to handle different sequence lengths. Second, it does not expose the optimal computational complexity of the method, and the original version of DSS in fact uses  $O(N)$  more memory than S4 and S4D do.<sup>2</sup>

**Discretization.** S4D disentangles the discretization method from the kernel computation (equation (7)), so that any discretization can be used, whereas previous methods required a specific discretization. For

<sup>2</sup>An early version of DSS claimed that it did not require a custom kernel while S4 does, but this is because of its extra memory usage. The PyTorch implementation of S4 has an optional custom CUDA kernel primarily to save this factor of  $N$  in space.

example, DSS requires the zero-order hold (ZOH) discretization because the exp term in the ZOH formula lends itself to be computed with a softmax. On the other hand, when  $\mathbf{A}$  is not diagonal, ZOH involves a matrix exponential which can be slower to compute, so S4 uses the bilinear discretization which can be computed efficiently for DPLR matrices.

**Eigenvalue constraint.** All methods can enforce any constraint on the eigenvalues of  $\mathbf{A}$ . While DSS found that letting them be unconstrained has slightly better performance, our experiments find that the difference is negligible and we recommend constraining negative real part of  $\mathbf{A}$  as is standard practice in control systems. This ensures stability even in unbounded autoregressive settings.

**The full model.** The entire S4D method is very straightforward to implement, requiring just a few lines of code each for the parameterization and initialization, kernel computation, and full forward pass (Listing 1). This minimal model maps an input sequence of length  $L$  to an output of the same length; given multiple input channels, independent S4D layers are broadcast over them. Other details such as the initialization of  $\Delta$  are the same as in S4 and DSS.

---

**Listing 1** Full Numpy example of the parameterization and computation of a 1-dimensional S4D-Lin model

---

```
def parameters(N, dt_min=1e-3, dt_max=1e-1):
    # Initialization
    log_dt = np.random.rand() * (np.log(dt_max)-np.log(dt_min)) + np.log(dt_min) # Geometrically uniform timescale
    A = -0.5 + 1j * np.pi * np.arange(N) # S4D-Lin initialization
    B = np.ones(N)
    C = np.random.randn(N) + 1j * np.random.randn(N) # Variance preserving initialization
    return log_dt, np.log(-A.real), A.imag, B, C

def kernel(L, log_dt, log_A_real, A_imag, B, C):
    # Discretization (e.g. bilinear transform)
    dt, A = np.exp(log_dt), -np.exp(log_a_real) + 1j * a_imag
    dA, dB = (1+dt*A/2) / (1-dt*A/2), dt*B / (1-dt*A/2)

    # Computation (Vandermonde matrix multiplication - can be optimized)
    return (B*C) @ (dA[:, None] ** np.arange(L))

def forward(u, parameters):
    L = u.shape[-1]
    K = kernel(L, *parameters)
    K_f, u_f = np.fft.fft(K, n=2*L), np.fft.fft(input, n=2*L) # Convolve y = u * K using FFT
    return np.fft.ifft(K_f*u_f, n=2*L)[..., :L]
```

---

Finally, note that different combinations of parameterization choices can lead to slightly different implementations of the kernel; Fig. 1 illustrates the S4D kernel with ZOH discretization which can be simplified even further.

## 4 Initialization of Diagonal State Matrices

The critical question remains: which diagonal state matrices  $\mathbf{A}$  are actually effective? We comment on the limitations of DSSMs, and then provide three instantiations of S4D that perform well empirically.

**Expressivity and Limitations of Diagonal SSMs.** We first present a simplified view on the expressivity of DSSMs mentioned by [10]. First, it is well-known that almost all matrices diagonalize over the complex plane. Therefore it is critical to use complex-valued matrices in order to use DSSMs.

**Proposition 2.** *The set  $\mathcal{D} \subset \mathbb{C}^{N \times N}$  of diagonalizable matrices is dense in  $\mathbb{C}^{N \times N}$ , and has full measure (i.e. its complement has measure 0).*

It is also well known that the state space  $(\mathbf{A}, \mathbf{B}, \mathbf{C})$  is exactly equivalent to (i.e. expresses the same map  $u \mapsto y$ ) the state space  $(\mathbf{V}^{-1}\mathbf{A}\mathbf{V}, \mathbf{V}^{-1}\mathbf{B}, \mathbf{C}\mathbf{V})$ , known in the SSM literature as a state space transformation. Therefore Proposition 2 says that (almost) all SSMs are equivalent to a DSSM.

However, we emphasize that Proposition 2 is about *expressivity* which does not guarantee strong performance of a trained model after optimization. For example, Gu et al. [9] and Gupta [10] show that when  $\mathbf{A}$  is parameterized as a dense real matrix or diagonal complex matrix, which are both fully expressive, it performs much worse than S4 if randomly initialized.

Second, Proposition 2 does not take into account numerical representations of data, which was the original reason S4 required a low-rank correction term instead of a pure diagonalization. In Section 5.2, we also show that two different initializations with the *same spectrum* (i.e., are equivalent to the same diagonal  $\mathbf{A}$ ) can have very different performance.

**S4D-LegS.** The HiPPO-LegS matrix has DPLR representation  $\mathbf{A}^{(D)} - \mathbf{P}\mathbf{P}^\top$ , and Gupta [10] showed that simply approximating it with  $\mathbf{A}^{(D)}$  works quite well (5). Our first result is providing a clean mathematical interpretation of this method. Theorem 3 shows a surprising fact that does not hold in general for DPLR matrices (Appendix A.1), and arises out of the special structure of this particular matrix.

**Theorem 3.** *Let  $\mathbf{A} = \mathbf{A}^{(N)} - \mathbf{P}\mathbf{P}^\top$  and  $\mathbf{B}$  be the HiPPO-LegS matrices, and  $K_{\mathbf{A}, \mathbf{B}}(t)$  be its basis (Fig. 1). As the state size  $N \rightarrow \infty$ , the SSM basis  $K_{\mathbf{A}^{(N)}, \mathbf{B}/2}(t)$  limits to  $K_{\mathbf{A}, \mathbf{B}}(t)$  (Fig. 2).*

(Note that these are *unitarily* equivalent to  $\mathbf{A}^{(D)}$ , which preserves the stability and timescale [1] of the system.)

We define **S4D-LegS** to be the S4D method for this choice of diagonal  $\mathbf{A} = \mathbf{A}^{(D)}$ . Theorem 3 explains the empirical results in [10] whereby this system performed quite close to S4, but was usually slightly worse. This is because DSS is a variant of S4D-LegS, which by Theorem 3 is a noisy approximation to S4-LegS. Fig. 2 illustrates this result, and also shows a curious phenomenon involving different discretization rules that is open for future work.

**S4D-Inv.** To further simplify S4D-LegS, we analyze the structure of  $\mathbf{A}^{(D)} = \text{diag}(\mathbf{A})$  in more detail. The real part is easy to understand, which follows from the analysis in [9]:

**Proposition 4.**  $\Re(\mathbf{A}) = -\frac{1}{2}\mathbf{1}$

Let the imaginary part be sorted, i.e.  $\Im(\mathbf{A})_n$  is the  $n$ -th largest (positive) imaginary component. We empirically deduced the following conjecture for the asymptotics of the imaginary part.

**Conjecture 5.** *As  $N \rightarrow \infty$ ,  $\Im(\mathbf{A})_0 \rightarrow \frac{1}{\pi}N^2 + c$  where  $c \approx 0.5236$  is a constant. For a fixed  $N$ , the other eigenvalues satisfy an inverse scaling in  $n$ :  $\Im(\mathbf{A})_n = \Theta(n^{-1})$ .*

Fig. 3 empirically supports this conjecture. Based on Conjecture 5, we propose the initialization S4D-Inv to use the following inverse-law diagonal matrix which closely approximates S4D-LegS.

$$\text{(S4D-Inv)} \quad \mathbf{A}_n = -\frac{1}{2} + i\frac{N}{\pi} \left( \frac{N}{2n+1} - 1 \right) \quad (8) \quad \text{(S4D-Lin)} \quad \mathbf{A}_n = -\frac{1}{2} + i\pi n \quad (9)$$

**S4D-Lin.** While S4D-Inv can be seen as an approximation to the original S4-LegS, we propose an even simpler scaling law for the imaginary parts that can be seen as an approximation of S4-FouT ([0]), where the imaginary parts are simply the Fourier series frequencies (i.e. matches the diagonal part of the DPLR form of S4-FouT). Fig. 1 (Right) illustrates the S4D-Lin basis  $e^{t\mathbf{A}}\mathbf{B}$ , which are simply damped Fourier basis functions.

**General Diagonal SSM Basis Functions.** The empirical study in Section 5 performs many ablations of different diagonal initializations, showing that many natural variants of the proposed methods do not perform

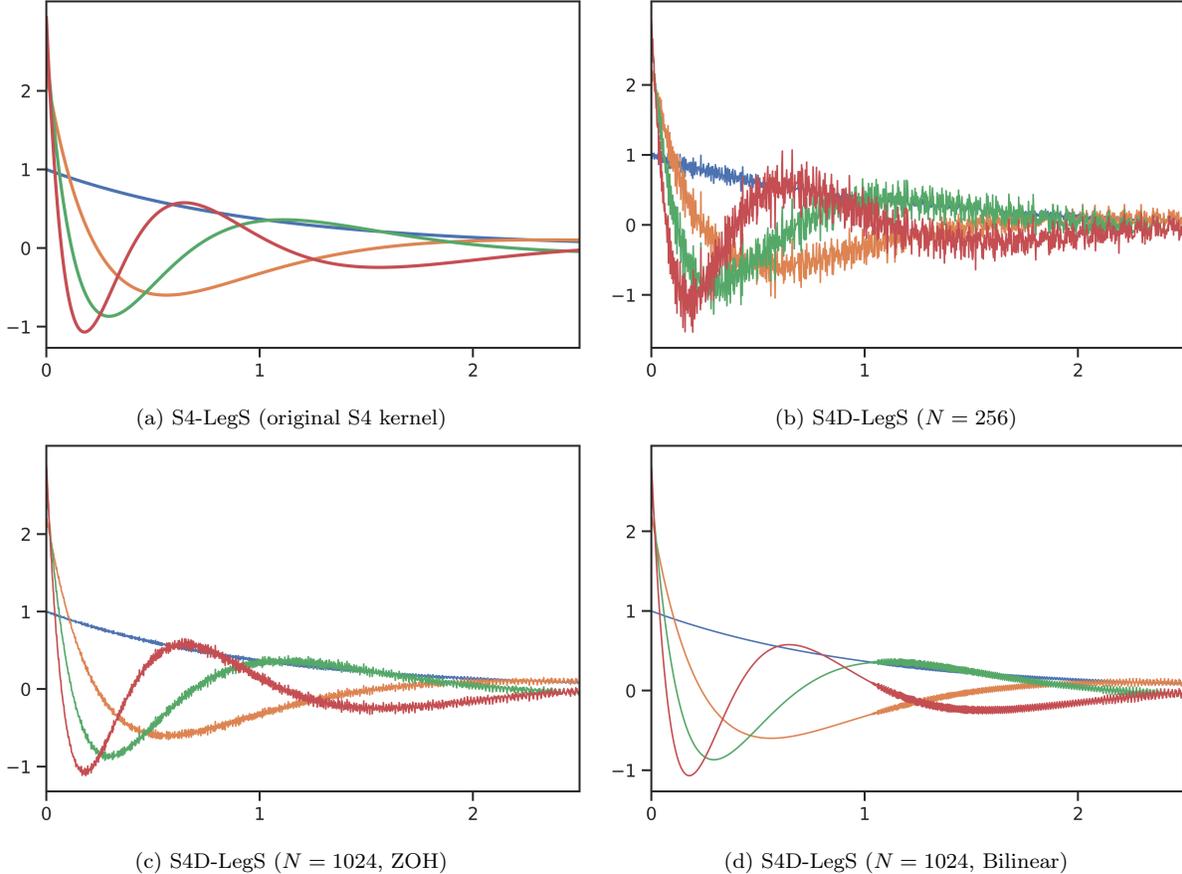


Figure 2: **(Visualization of Theorem 3)**. (a) The particular  $(\mathbf{A}, \mathbf{B})$  matrix chosen in S4 results in smooth basis functions  $e^{t\mathbf{A}}\mathbf{B}$  with a closed form formula in terms of Legendre polynomials. By the HiPPO theory, convolving against these functions has a mathematical interpretation as orthogonalizing against an exponentially-decaying measure. (b, c) By special properties of this state matrix, removing the low-rank term of its NPLR representation produces the same basis functions as  $N \rightarrow \infty$ , explaining the empirical effectiveness of DSS. (c) Curiously, the bilinear transform instead of ZOH smooths out the kernel to exactly match S4-LegS as  $N$  grows.

as well. The overall guiding principles for the diagonal state matrix  $\mathbf{A}$  are twofold, which can be seen from the closed form of the basis functions  $K_n(t) = e^{t\mathbf{A}_n}\mathbf{B}_n$  (Eq. (3)).

First, the real part of  $\mathbf{A}_n$  controls the decay rate of the function.  $\mathbf{A}_n = -\frac{1}{2}$  is a good default that bounds the basis functions by the envelope  $e^{-\frac{t}{2}}$ , giving a constant timescale (Fig. 1 (Right)).

Second, the imaginary part of  $\mathbf{A}_n$  controls the oscillating frequencies of the basis function. Critically, these need to be spread out, which explains why random initializations of  $\mathbf{A}$  do not perform well. The S4D-Inv and S4D-Lin use simple asymptotics for these imaginary components that provide interpretable methods. We believe that alternative initializations that have different mathematical interpretations may exist, which is an interesting question for future work.

## 5 Experiments

Our experimental study shows that S4D has strong performance in a wide variety of domains and tasks, including the well-studied Long Range Arena (LRA) benchmark where the best S4D variant is competitive with S4 on all tasks and significantly outperforms all non-SSM baselines.

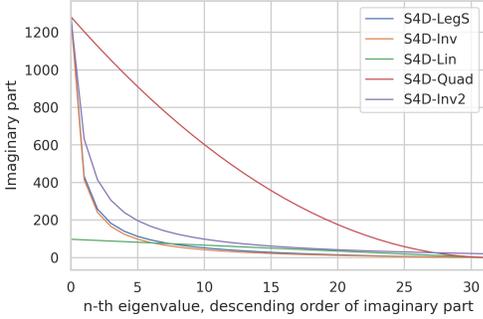


Figure 3: (**S4D eigenvalues.**) All S4D methods have eigenvalues  $-\frac{1}{2} + \lambda_n i$ . S4D-LegS theoretically approximates dynamics of the original (non-diagonal) S4 (Blue), and has eigenvalues following an inverse law  $\lambda_n \propto n^{-1}$  (Orange). The precise law is important: other scaling laws with the same range, including an inverse law with different constant (Purple) and a quadratic law (Red), perform empirically worse (Section 5.2). A very different linear law based on Fourier frequencies also performs well (Green).

We begin with controlled ablations of the various representations of diagonal state space models.

- In Section 5.1, we compare the different methods of parameterizing and computing a diagonal state space model (Section 3).
- In Section 5.2, we compare our proposed initializations of the critical  $\mathbf{A}$  matrix and perform several ablations showing that simple variants can substantially degrade performance, underscoring the importance of choosing  $\mathbf{A}$  carefully (Section 4).
- In Section 5.3, we compare our proposed S4D methods against the original S4 method (and the variants proposed in [1]).

**Methodology and Datasets.** In order to study the effects of different S4 and S4D variants in a controlled setting, we propose the following protocol. We focus on three datasets covering a varied range of data modalities (image pixels, biosignal time series, audio waveforms), sequence lengths (1K, 4K, 16K), and tasks (classification and regression with bidirectional and causal models).

- **Sequential CIFAR (sCIFAR).** CIFAR-10 images are flattened into a sequence of length 1024, and a bidirectional sequence model is used to perform 10-way classification.
- **BIDMC Vital Signs.** EKG and PPG signals of length 4000 are used to predict respiratory rate (RR), heart rate (HR), and blood oxygen saturation (SpO2). We focus on SpO2 in this study.
- **Speech Commands (SC).**<sup>3</sup> A 1-second raw audio waveform comprising 16000 samples is used for 35-way spoken word classification. We use an autoregressive (AR) model to vary the setting; this causal setting more closely imitates autoregressive speech generation, where SSMs have shown recent promise [5].

We fix a simple architecture and training protocol that works generically. The architecture has 4 layers and hidden dimension  $H = 128$ , resulting in  $\sim 100K$  parameters. All results are averaged over multiple seeds (full protocol and results including std. reported in Appendix B).

## 5.1 Parameterization, Computation, Discretization

Given the same diagonal SSM matrices  $\mathbf{A}, \mathbf{B}$ , there are many variants of how to parameterize the matrices and compute the SSM kernel described in Section 3. We ablate the different choices described in Table 1. Results are in Table 2, and show that:

- (i) Computing the model with a softmax instead of Vandermonde product does not make much difference
- (ii) Training  $\mathbf{B}$  is consistently slightly better
- (iii) Different discretizations (Section 3.1) do not make a noticeable difference
- (iv) Unrestricting the real part of  $\mathbf{A}$  (Section 3.3) may be slightly better

Trainable B	Method	sCIFAR	SC (AR)	BIDMC (SpO2)
No	Softmax	85.04	89.80	0.1299
No	Vandermonde	84.78	89.62	0.1355
Yes	Softmax	85.37	90.06	0.1170
Yes	Vandermonde	85.37	90.34	0.1274

Discretization	Real part of A	sCIFAR	SC (AR)	BIDMC (SpO2)
Bilinear	exp	85.20	89.52	0.1193
	ReLU	85.06	90.22	0.1172
	-	85.35	90.58	0.1102
ZOH	exp	85.02	89.93	0.1303
	ReLU	84.98	90.03	0.1232
	-	85.15	90.19	0.1289

Table 2: Ablations of different parameterizations of diagonal SSMS using S4D-Inv. (*Left*) trainability and computation; (*Right*) discretization and parameterization.

These ablations show that for a fixed initialization  $(\mathbf{A}, \mathbf{B})$ , different aspects of parameterizing SSMS make little difference overall. This justifies the parameterization and algorithm S4D uses (Section 3.4), which preserves the choices of the original S4 model and is simpler than DSS. For the remaining of the experiments in Section 5.2 and Section 5.3, we fix the S4D parameterization and algorithm described in Section 3. Note that this computes exactly the same kernel as the original S4 algorithm when the low-rank portion is set to 0, allowing controlled comparisons of the critical state matrix  $\mathbf{A}$  for the remainder of this section.

## 5.2 S4D Initialization Ablations

The original S4 model proposed a specific formula for the  $\mathbf{A}$  matrix, and the first diagonal version [10] used a specific matrix based on it. Our new proposed variants S4D-Inv and S4D-Lin also define precise formulas for the initialization of the  $\mathbf{A}$  matrix (8). This raises the question of whether the initialization of the  $\mathbf{A}$  still needs to be so precise, despite the large simplifications from the original version. We perform several natural ablations on these initializations, showing that even simple variations of the precise formula can degrade performance.

**Imaginary part scaling factor.** The scaling rules for the imaginary parts of S4D-Inv and S4D-Lin are simple polynomial laws, but how is the constant factor chosen and how important is it? These constants are based on approximations to HiPPO methods (e.g. Conjecture 5). Note that the range of imaginary components for S4D-Inv and S4D-Lin are quite different (Fig. 3); the largest imaginary part is  $\frac{N^2}{\pi}$  for S4D-Inv and  $\pi N$  for S4D-Lin.

We consider scaling all imaginary parts by a constant factor of 0.01 or 100.0 to investigate whether the constant matters. Note that this preserves the overall shape of the basis functions (Fig. 1) and simply changes the frequencies, and it is not obvious that this should degrade performance. However, both changes substantially reduce the performance of both S4D methods in all settings.

**Randomly initialized imaginary part.** Next, we consider choosing the imaginary parts randomly. For S4D-Inv, we keep the real parts equal to  $-\frac{1}{2}$  and set each imaginary component to

$$\mathbf{A}_n = -\frac{1}{2} + i\frac{N}{\pi} \left( \frac{N}{2u+1} - 1 \right) \quad u \sim N \cdot \mathcal{U}[0, 1] \quad (10)$$

Note that when  $u$  is equally spaced in  $[0, 1]$  instead of uniformly random, this exactly recovers S4D-Inv (8), so this is a sensible random approximation to it.

Similarly, we consider a variant of S4D-Lin

$$\mathbf{A}_n = -\frac{1}{2} + i\pi u N \quad u \sim N \cdot \mathcal{U}[0, 1] \quad (11)$$

that is equal to equation (9) when  $u$  is equally spaced instead of random.

<sup>3</sup>We note that a line of prior work including S4 [9, 12, 16] all used a smaller 10-class subset of SC, so our results on the full dataset are not directly comparable.

Table 3a (*Random Imag*) shows that this small change causes minor degradation in performance. We additionally note that the randomly initialized imaginary ablation can be interpreted as follows. Fig. 3 shows the asymptotics of the imaginary parts of SSM matrices, where the imaginary parts of the eigenvalues correspond to  $y$ -values corresponding to uniformly spaced nodes on the  $x$ -axis. This ablation then replaces the uniform spacing on the  $x$ -axis with uniformly random  $x$  values.

**Randomly initialized real part.** We consider initializing the real part of each eigenvalue as  $-\mathcal{U}[0, 1]$  instead of fixing them to  $-\frac{1}{2}$ . Table 3a(Left, *Random Real*) shows that this also causes minor but consistent degradation in performance on the ablation datasets. Finally, we also consider randomizing both real and imaginary parts, which degrades performance even further.

Table 3: (*Initialization and Trainability ablations*)

Ablation	sCIFAR	SC (AR)	BIDMC
	<b>S4D-Lin</b>	<b>85.12</b>	<b>90.66</b>
Scale 0.01	-7.27	-1.92	+0.040
Scale 100	-7.91	-4.04	+0.077
Random Imag	-0.42	-3.08	-0.001
Random Real	-0.73	-0.87	+0.011
Random Both	-1.28	-5.88	+0.007

Ablation	sCIFAR	SC (AR)	BIDMC
	<b>S4D-Inv</b>	<b>84.79</b>	<b>90.27</b>
Scale 0.01	-5.03	-0.08	+0.028
Scale 100	-7.77	-52.31	+0.034
Random Imag	-0.29	-0.52	+0.010
Random Real	0.12	-2.18	+0.032
Random Both	-1.55	-0.55	+0.024
S4D-Inv2	-2.62	-39.84	+0.005
S4D-Quad	-1.83	-0.62	+0.024
S4D-Random	-6.32	-1.95	+0.034
S4D-Real	-5.45	-10.17	+0.066

Frozen (A, B)	sCIFAR		SC (AR)		BIDMC
	Acc (first)	Acc (best)	Acc (first)	Acc (best)	RMSE (best)
<b>S4-LegS</b>	53.63	86.19	33.87	85.33	0.1049
<b>S4-LegT</b>	54.76	86.30	8.77	57.35	0.1106
<b>S4-FouT</b>	55.28	86.05	9.27	69.57	0.1072
<b>S4-LegS+FouT</b>	54.38	86.53	34.06	83.37	0.0887
<b>S4D-LegS</b>	50.87	84.81	22.76	77.18	0.0960
<b>S4D-Inv</b>	53.19	84.40	18.49	76.53	0.0995
<b>S4D-Lin</b>	51.75	84.96	19.09	75.58	0.0935

Trainable (A, B)	sCIFAR		SC (AR)		BIDMC
	Acc (first)	Acc (best)	Acc (first)	Acc (best)	RMSE (best)
<b>S4-LegS</b>	54.23	86.29	62.19	90.68	0.1033
<b>S4-LegT</b>	55.16	86.12	55.86	90.42	0.1146
<b>S4-FouT</b>	55.89	85.93	60.56	90.83	0.1136
<b>S4-LegS+FouT</b>	55.00	86.18	61.76	91.01	0.0970
<b>S4D-LegS</b>	50.41	85.64	47.54	88.47	0.1148
<b>S4D-Inv</b>	53.42	84.59	45.73	89.69	0.1132
<b>S4D-Lin</b>	52.23	85.75	47.68	89.56	0.1032

(a) Ablations of the initialization of the diagonal  $\mathbf{A}$  matrix in S4D. Very simple changes that largely preserve the structure of the diagonal eigenvalues all degrade performance.

(b) Results for all S4 and S4D methods on the ablation datasets, when the  $\mathbf{A}$  and  $\mathbf{B}$  matrices are either frozen (*Top*) or trained (*Bottom*). Diagonal state matrices are highly competitive with full DPLR versions, achieving strong results on all datasets.

**Ablation: Other S4D matrices.** Other simple variants of initializations show that it is not just the range of the eigenvalues but the actual distribution that is important (Fig. 3). Both S4D-Inv2 and S4D-Quad have real part  $-\frac{1}{2}$  and imaginary part satisfying the same maximum value as Conjecture 5. The S4D-Inv2 initialization uses the same formula as S4D-Inv, but replaces a  $2n + 1$  in the denominator with  $n + 1$ . The S4D-Quad initialization uses a polynomial law with power 2 instead of  $-1$  (S4D-Inv) or  $1$  (S4D-Lin).

$$\text{(S4D-Inv2)} \quad \mathbf{A}_n = -\frac{1}{2} + i\frac{N}{\pi} \left( \frac{N}{n+1} - 1 \right) \quad (12)$$

$$\text{(S4D-Quad)} \quad \mathbf{A}_n = \frac{1}{\pi}(1 + 2n)^2 \quad (13)$$

We include two additional methods here that are not based on the proposed S4D-Inv or S4D-Lin methods. First, S4D-Rand uses a randomly initialized diagonal  $\mathbf{A}$ , and validates that it performs poorly, in line with earlier findings [9, 10]. Second, S4D-Real uses a particular real initialization with  $\mathbf{A}_n = -(n + 1)$ . This is the exact same spectrum as the original S4(-LegS) method, which validates that it is not just the diagonalization that matters, highlighting the limitations of Proposition 2.

Table 4: (**Ablation datasets:** Full results with larger models.) For Speech Commands, we show both an autoregressive model as in the ablations, and an unconstrained bidirectional model.

MODEL	sCIFAR	SC		BIDMC		
	TEST	AR	Bi.	HR	RR	SPO2
S4-LegS	<b>91.80</b> (0.43)	<b>93.60</b> (0.13)	96.08 (0.15)	<b>0.332</b> (0.013)	<u>0.247</u> (0.062)	0.090 (0.006)
S4-FouT	<u>91.22</u> (0.25)	91.78 (0.10)	95.42 (0.20)	<u>0.339</u> (0.020)	0.301 (0.030)	<b>0.068</b> (0.003)
S4D-LegS	89.92 (1.69)	<u>93.57</u> (0.09)	95.97 (0.14)	0.367 (0.001)	0.248 (0.036)	0.102 (0.001)
S4D-Inv	90.69 (0.06)	93.40 (0.67)	<u>96.18</u> (0.27)	0.373 (0.024)	0.254 (0.022)	0.110 (0.001)
S4D-Lin	90.42 (0.03)	93.37 (0.05)	<b>96.25</b> (0.03)	0.379 (0.006)	<b>0.226</b> (0.008)	0.114 (0.003)

Table 5: (**Long Range Arena**) Accuracy on full suite of LRA tasks. Hyperparameters in Appendix B.

MODEL	LISTOPS	TEXT	RETRIEVAL	IMAGE	PATHFINDER	PATH-X	AVG
S4-LegS	59.60 (0.07)	86.82 (0.13)	90.90 (0.15)	<u>88.65</u> (0.23)	<u>94.20</u> (0.25)	<b>96.35</b>	<b>86.09</b>
S4-FouT	57.88 (1.90)	86.34 (0.31)	89.66 (0.88)	<b>89.07</b> (0.19)	<b>94.46</b> (0.24)	<b>X</b>	77.90
S4D-LegS	<u>60.47</u> (0.34)	86.18 (0.43)	89.46 (0.14)	88.19 (0.26)	93.06 (1.24)	91.95	84.89
S4D-Inv	60.18 (0.35)	<b>87.34</b> (0.20)	<b>91.09</b> (0.01)	87.83 (0.37)	93.78 (0.25)	<u>92.80</u>	<u>85.50</u>
S4D-Lin	<b>60.52</b> (0.51)	<u>86.97</u> (0.23)	<u>90.96</u> (0.09)	87.93 (0.34)	93.96 (0.60)	<b>X</b>	78.39
S4 (original)	58.35	76.02	87.09	87.26	86.05	88.10	80.48
Transformer	36.37	64.27	57.46	42.44	71.40	<b>X</b>	53.66

### 5.3 Full Comparisons of S4D and S4 Methods

**Trainable  $\mathbf{A}$ ,  $\mathbf{B}$  matrices.** Table 3b shows the performance of all S4D and S4 variants [1] on the ablation datasets. We observe several interesting phenomena:

- (i) Freezing the matrices performs comparably to training them on sCIFAR and BIDMC, but is substantially worse on SC. We hypothesize that this results from  $\Delta$  being poorly initialized for SC, so that at initialization models do not have context over the entire sequence, and training  $\mathbf{A}$  and  $\mathbf{B}$  helps adjust for this. As further evidence, the *finite window methods* S4-LegT and S4-FouT (defined in [1]) have the most limited context and suffer the most when  $\mathbf{A}$  is frozen.
- (ii) The full DPLR versions are often slightly better than the diagonal version throughout the entire training curve. We report the validation accuracy after 1 epoch of training on sCIFAR and SC to illustrate this phenomenon.

**Large models on ablation datasets.** Finally, we relax the strict requirements on model size and regularization for the ablation datasets, and show the performance of S4 and S4D variants on the test sets with a larger model (architecture and training details in Appendix B) when the model size and regularization is simply increased (Table 4). We note that results for each dataset are better than the original S4 model, which was already state-of-the-art on these datasets [8, 9].

**Long Range Arena.** We use the same hyperparameter setting for the state-of-the-art S4 model in [0] on the Long Range Arena benchmark for testing long dependencies in sequence models. S4D variants are highly competitive on all datasets except Path-X, and outperform the S4 variants on several of them. On Path-X using this hyperparameter setting with bidirectional models, only S4D-Inv, our simpler approximation to the original S4-LegS model, achieves above random chance, and has an average of 85% on the full LRA suite, more than 30 points better than the original Transformer [21].

Table 6: (**S4D Path-X Ablations.**) Ablating parameterization choices for models with less than 200K parameters.

<b>S4D</b>	Identity $\Re(\mathbf{A})$	ReLU $\Re(\mathbf{A})$	ZOH disc.	Frozen $\mathbf{B}$	ZOH + softmax	<b>DSS</b>
92.12 (0.34)	92.32 (0.16)	92.29 (0.20)	92.09 (0.08)	91.66 (0.62)	90.92 (0.34)	89.72 (0.33)

**Final parameterization ablations on Path-X.** Finally, we return to the parameterization choices presented in Section 3 and ablated in Section 5.1, and ablate them once more on the difficult Path-X dataset. We use small models of between 150K and 200K parameters (differing only depending on whether  $\mathbf{B}$  is trained). We fix the S4D-LegS initialization (i.e., the diagonal HiPPO initialization (5)).

We start from the base S4D parameterization based on S4: Bilinear discretization, exp  $\Re(\mathbf{A})$ , trainable  $\mathbf{B}$ , and no softmax (Table 1). We ablate each of these choices one at a time for the discretization, constraint on  $\Re(\mathbf{A})$ , trainability of  $\mathbf{B}$ , and normalization. We also consider the combination that defines DSS: ZOH discretization, identity  $\Re(\mathbf{A})$ , frozen  $\mathbf{B}$ , softmax normalization.

Table 6 shows that the default S4 parameterization choices are a strong baseline. As in Section 5.1, we find that most of the other choices do not make much difference:

- (i) letting  $\Re(\mathbf{A})$  be unconstrained has little benefit, and can theoretically cause instabilities, so we do not recommend it,
- (ii) the bilinear vs. ZOH discretizations make no difference,
- (iii) training  $\mathbf{B}$  helps slightly, for a minor increase in parameter count and no change in speed.

Finally, on this task – unlike the easier ablation datasets in Section 5.1 – the softmax normalization of DSS actually hurts, and we do not recommend it.

## 6 Conclusion

State space models based on S4 are a promising family of models for modeling many types of sequential data, with particular strengths for continuous signals and long-range interactions. These models are a large departure from conventional sequence models such as RNNs, CNNs, and Transformers, with many new ideas and moving parts. This work provides a more in-depth exposition for all aspects of working with S4-style models, from their core structures and kernel computation algorithms, to miscellaneous choices in their parameterizations, to new theory and methods for their initialization. We systematically analyzed and ablated each of these components, and provide recommendations for building a state space model that is as simple as possible, while as theoretically principled and empirically effective as S4. We believe that S4D can be a strong generic sequence model for a variety of domains, that opens new directions for state space models theoretically, and is much more practical to understand and implement for practitioners.

## References

- [1] Anonymous. How to train your hippo: State space models with generalized basis projections. *arXiv preprint*, 2022.
- [2] Shaojie Bai, J Zico Kolter, and Vladlen Koltun. Trellis networks for sequence modeling. In *The International Conference on Learning Representations (ICLR)*, 2019.
- [3] Yann N Dauphin, Angela Fan, Michael Auli, and David Grangier. Language modeling with gated convolutional networks. In *International conference on machine learning*, pages 933–941. PMLR, 2017.
- [4] N Benjamin Erichson, Omri Azencot, Alejandro Queiruga, Liam Hodgkinson, and Michael W Mahoney. Lipschitz recurrent neural networks. In *International Conference on Learning Representations*, 2021.
- [5] Karan Goel, Albert Gu, Chris Donahue, and Christopher Ré. It’s raw! audio generation with state-space models. *arXiv preprint arXiv:2202.09729*, 2022.
- [6] Albert Gu, Tri Dao, Stefano Ermon, Atri Rudra, and Christopher Ré. Hippo: Recurrent memory with optimal polynomial projections. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [7] Albert Gu, Caglar Gulcehre, Tom Le Paine, Matt Hoffman, and Razvan Pascanu. Improving the gating mechanism of recurrent neural networks. In *The International Conference on Machine Learning (ICML)*, 2020.
- [8] Albert Gu, Isys Johnson, Karan Goel, Khaled Saab, Tri Dao, Atri Rudra, and Christopher Ré. Combining recurrent, convolutional, and continuous-time models with the structured learnable linear state space layer. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [9] Albert Gu, Karan Goel, and Christopher Ré. Efficiently modeling long sequences with structured state spaces. In *The International Conference on Learning Representations (ICLR)*, 2022.
- [10] Ankit Gupta. Diagonal state spaces are as effective as structured state spaces. *arXiv preprint arXiv:2203.14343*, 2022.
- [11] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [12] Patrick Kidger, James Morrill, James Foster, and Terry Lyons. Neural controlled differential equations for irregular time series. *arXiv preprint arXiv:2005.08926*, 2020.
- [13] James Morrill, Cristopher Salvi, Patrick Kidger, James Foster, and Terry Lyons. Neural rough differential equations for long time series. *The International Conference on Machine Learning (ICML)*, 2021.
- [14] Victor Pan. *Structured matrices and polynomials: unified superfast algorithms*. Springer Science & Business Media, 2001.
- [15] Anthony Ralston and Philip Rabinowitz. *A first course in numerical analysis*. Courier Corporation, 2001.
- [16] David W Romero, Anna Kuzina, Erik J Bekkers, Jakub M Tomczak, and Mark Hoogendoorn. Ckconv: Continuous kernel convolution for sequential data. *arXiv preprint arXiv:2102.02611*, 2021.
- [17] David W Romero, Robert-Jan Brintjes, Jakub M Tomczak, Erik J Bekkers, Mark Hoogendoorn, and Jan C van Gemert. Flexconv: Continuous kernel convolutions with differentiable kernel sizes. In *The International Conference on Learning Representations (ICLR)*, 2022.
- [18] T Konstantin Rusch and Siddhartha Mishra. Unicornn: A recurrent model for learning very long time dependencies. *The International Conference on Machine Learning (ICML)*, 2021.
- [19] Noam Shazeer. Glu variants improve transformer. *arXiv preprint arXiv:2002.05202*, 2020.

- [20] Chang Wei Tan, Christoph Bergmeir, Francois Petitjean, and Geoffrey I Webb. Time series extrinsic regression. *Data Mining and Knowledge Discovery*, pages 1–29, 2021. doi: <https://doi.org/10.1007/s10618-021-00745-9>.
- [21] Yi Tay, Mostafa Dehghani, Samira Abnar, Yikang Shen, Dara Bahri, Philip Pham, Jinfeng Rao, Liu Yang, Sebastian Ruder, and Donald Metzler. Long range arena : A benchmark for efficient transformers. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=qVyeW-grC2k>.
- [22] Trieu H Trinh, Andrew M Dai, Minh-Thang Luong, and Quoc V Le. Learning longer-term dependencies in RNNs with auxiliary losses. In *The International Conference on Machine Learning (ICML)*, 2018.
- [23] Aaron Voelker, Ivana Kajić, and Chris Eliasmith. Legendre memory units: Continuous-time representation in recurrent neural networks. In *Advances in Neural Information Processing Systems*, pages 15544–15553, 2019.

# A Method Details

## A.1 Proofs

We prove Theorem 3, and then show why this it is a surprising result that is not true in general to low-rank perturbations of SSMs.

We start with the interpretation of the S4-LegS matrix shown in [1], which corresponds to Fig. 1 (Left).

**Theorem 6.** *Let  $\mathbf{A}, \mathbf{B}, \mathbf{P}$  be the matrices defined in equation (4). The SSM kernels  $K_n(t) = \mathbf{e}_n^\top e^{t\mathbf{A}} \mathbf{B}$  have the closed form formula*

$$K_n(t) = L_n(e^{-t})e^{-t}$$

where  $L_n$  are the Legendre polynomials shifted and scaled to be orthonormal on the interval  $[0, 1]$ .

**Lemma A.1.** *The functions  $L_n(e^{-t})$  are a complete orthonormal basis with respect to the measure  $\omega(t) = e^{-t}$ .*

*Proof.* The polynomials are defined to be orthonormal on  $[0, 1]$ , i.e.

$$\int_0^1 L_n(t)L_m(t) dt = \delta_{n,m}.$$

By the change of variables  $t = e^{-s}$  with  $\frac{dt}{ds} = -e^{-s}$ ,

$$-\int_{-\infty}^0 L_n(e^{-s})L_m(e^{-s})e^{-s} ds = \delta_{n,m} = \int_0^{\infty} L_n(e^{-s})L_m(e^{-s})e^{-s} ds$$

which shows the orthonormality.

Completeness follows from the fact that polynomials are complete. □

*Proof of Theorem 3.* We start with the standard interpretation of SSMs as convolutional systems. The SSM  $x'(t) = \mathbf{A}x(t) + \mathbf{B}u(t)$  is equivalent to the convolution

$$x_n(t) = (u * K_n)(t) = \int_{-\infty}^t u(s)K_n(t-s) ds = \int_0^{\infty} u(t-s)K_n(s) ds$$

for the SSM kernels (equation (3)).

Defining  $u^{(t)}(s) = u(t-s)$ , we can write this as

$$x_n(t) = \langle u^{(t)}, K_n \rangle_\omega$$

where  $\omega(s) = e^{-s}$  and  $\langle p(s), q(s) \rangle_\omega = \int_0^{\infty} p(s)q(s)\omega(s) ds$  is the inner product in the Hilbert space of L2 functions with respect to measure  $\omega$ .

By Theorem 6, the  $K_n$  are a complete orthonormal basis in this Hilbert space. There  $x_n(t)$  represents a decomposition of the function  $u^{(t)}$  with respect to this basis, and can be recovered as a linear combination of these projections

$$u^{(t)} = \sum_{n=0}^{\infty} x_n(t)K_n.$$

Pointwise over the inner times  $s$ ,

$$u^{(t)}(s) = \sum_{n=0}^{\infty} x_n(t)K_n(s).$$

This implies that

$$\begin{aligned}
u(t) &= u^{(t)}(0) = \sum_{n=0}^{\infty} x_n(t) K_n(0) \\
&= \sum_{n=0}^{\infty} x_n(t) L_n(0) = \sum_{n=0}^{\infty} x_n(t) (2n+1)^{\frac{1}{2}} \\
&= \mathbf{B}^\top x(t)
\end{aligned}$$

Intuitively, due to the function reconstruction interpretation of HiPPO [1], we can approximate  $u(t)$  using knowledge in the current state  $x(t)$ . There in the limit  $N \rightarrow \infty$ , the original SSM is equivalent to

$$\begin{aligned}
x'(t) &= \mathbf{A}x(t) + \mathbf{B}u(t) \\
&= \mathbf{A}x(t) + \frac{1}{2}\mathbf{B}u(t) + \frac{1}{2}\mathbf{B}u(t) \\
&= \mathbf{A}x(t) + \frac{1}{2}\mathbf{B}\mathbf{B}^\top x(t) + \frac{1}{2}\mathbf{B}u(t) \\
&= \mathbf{A}x(t) + \mathbf{P}\mathbf{P}^\top x(t) + \frac{1}{2}\mathbf{B}u(t) \\
&= \mathbf{A}^N x(t) + \frac{1}{2}\mathbf{B}u(t)
\end{aligned}$$

□

Finally, we remark that this phenomenon where removing the low-rank correction to a DPLR matrix approximates the original dynamics, is unique to this HiPPO-LegS matrix. We note that if instead of  $\mathbf{P}\mathbf{P}^\top$ , a *random* rank-1 correction is added to the HiPPO-LegS matrix in Theorem 3, the resulting SSM kernels look completely different and in fact diverge rapidly as the magnitude of  $\mathbf{P}$  increases (Fig. 4).

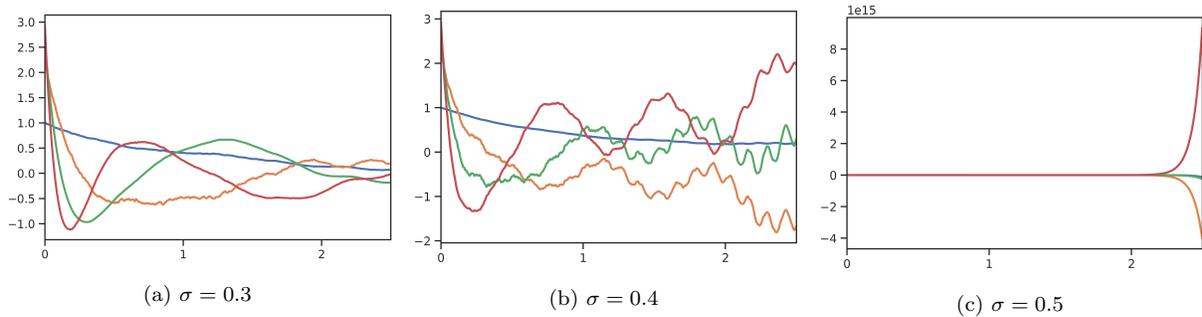


Figure 4: Basis kernels for  $(\mathbf{A} + \mathbf{P}\mathbf{P}^\top, \mathbf{B})$  for HiPPO-LegS  $(\mathbf{A}, \mathbf{B})$  and random i.i.d. Gaussian  $\mathbf{P}$  with varying std  $\sigma$ , illustrating that the SSM basis is very sensitive to low-rank perturbations. Note that the normal-HiPPO matrix  $\mathbf{A}^{(N)} = \mathbf{A} + \mathbf{P}\mathbf{P}^\top$  for  $\mathbf{P}$  with entries of magnitude  $N^{\frac{1}{2}}$  which is far larger, highlighting how unexpected the theoretical result Theorem 3 is.

Similarly, Fig. 5a shows a new S4 variant called S4-FouT that is also DPLR [0], but removing the low-rank component dramatically changes the SSM kernels.

## B Experiment Details

**Ablation datasets training protocol.** The architecture has 4 layers and hidden dimension  $H = 128$ , resulting in around 100K trainable parameters. All results are averaged over 2 or 3 seeds.

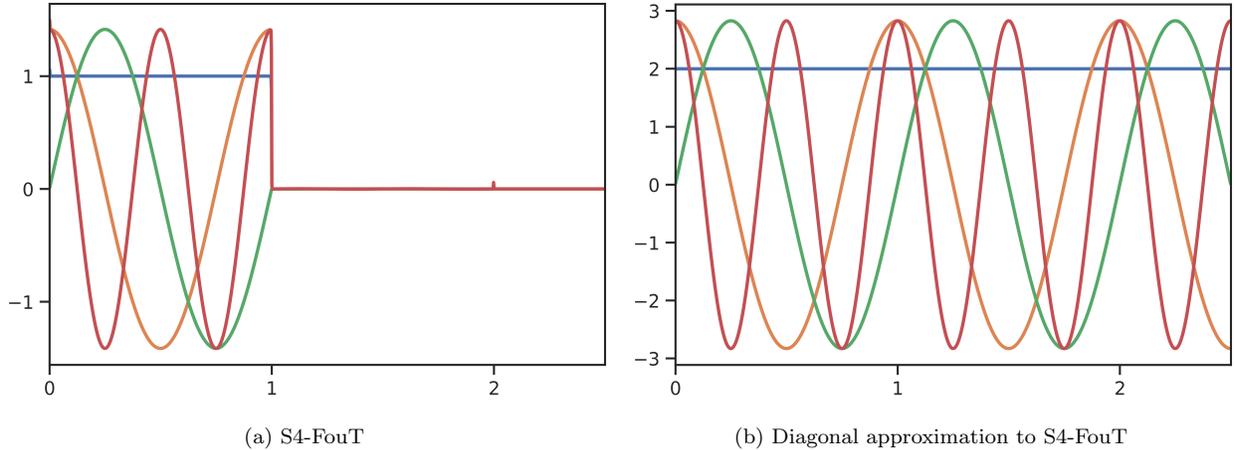


Figure 5: (a) S4-FouT is a new method that produces truncated Fourier basis functions. This captures *sliding Fourier transforms* into a state space. (b) Removing the low-rank term from the FouT matrix (as is the relation between (b) and (a)) does *not* approximate S4-FouT. This diagonal state matrix has real part 0 that produces infinite oscillations and does not perform well empirically.

Table 7: Full results for Table 2 (Left) including standard deviations.

Trainable $\mathbf{B}$	Method	sCIFAR	SC (AR)	BIDMC (SpO2)
No	Softmax	85.04 (0.22)	89.80 (0.21)	0.1299 (0.0048)
No	Vandermonde	84.78 (0.16)	89.62 (0.03)	0.1355 (0.0039)
Yes	Softmax	85.37 (0.43)	90.06 (0.11)	0.1170 (0.0039)
Yes	Vandermonde	85.37 (0.43)	90.34 (0.18)	0.1274 (0.0020)

All models use learning rate 0.004, 0.01 weight decay, and no other regularization or data augmentation. For the classification tasks (sCIFAR and SC), we use a cosine scheduler with 1 epoch warmup and decaying to 0. For the regression task (BIDMC), we use a multistep scheduler following [8, 18].

Reported results are all best validation accuracy, except for the large models in Table 4.

**Full results for parameterization ablations.** Table 7 and Table 8 contain the raw results for Table 2 including standard deviations.

**Full results for large models on ablations datasets.** Table 9 and Table 10 show full results comparing out proposed methods against the best models from the literature.

Table 8: Full results for Table 2 (Right) including standard deviations.

Discretization	Real part	sCIFAR	SC (AR)	BIDMC (SpO2)
Bilinear	Exp	85.20 (0.18)	89.52 (0.01)	0.1193 (0.0069)
Bilinear	-	85.35 (0.27)	90.58 (0.37)	0.1102 (0.0075)
Bilinear	ReLU	85.06 (0.06)	90.22 (0.25)	0.1172 (0.0063)
ZOH	Exp	85.02 (0.24)	89.93 (0.07)	0.1303 (0.0014)
ZOH	-	85.15 (0.13)	90.19 (0.58)	0.1289 (0.0035)
ZOH	ReLU	84.98 (0.72)	90.03 (0.13)	0.1232 (0.0065)

Table 9: (**Sequential CIFAR image classification.**) Table 10: (**BIDMC Vital signs prediction.**) RMSE for predicting respiratory rate (RR), heart rate (HR), and blood oxygen (SpO2).

Model	sCIFAR	Model	HR	RR	SpO2
S4-LegS	<b>91.80</b> (0.43)	S4-LegS	<b>0.332</b> (0.013)	0.247 (0.062)	0.090 (0.006)
S4-FouT	91.22 (0.25)	S4-FouT	<u>0.339</u> (0.020)	0.301 (0.030)	<b>0.068</b> (0.003)
S4-(LegS+FouT)	<u>91.58</u> (0.17)	S4-(LegS+FouT)	0.344 (0.032)	<b>0.163</b> (0.008)	<u>0.080</u> (0.007)
S4D-LegS	89.92 (1.69)	S4D-LegS	0.367 (0.001)	0.248 (0.036)	0.102 (0.001)
S4D-Inv	90.69 (0.06)	S4D-Inv	0.373 (0.024)	0.254 (0.022)	0.110 (0.001)
S4D-Lin	90.42 (0.03)	S4D-Lin	0.379 (0.006)	<u>0.226</u> (0.008)	0.114 (0.003)
Transformer [22]	62.2	UnICORNN [18]	1.39	1.06	0.869
FlexConv [17]	80.82	coRNN [18]	1.81	1.45	-
TrellisNet [2]	73.42	CKConv	2.05	1.214	1.051
LSTM [7, 11]	63.01	NRDE [13]	2.97	1.49	1.29
r-LSTM [22]	72.2	LSTM	10.7	2.28	-
UR-GRU [7]	<u>74.4</u>	Transformer	12.2	2.61	3.02
HiPPO-RNN [6]	61.1	XGBoost [20]	4.72	1.67	1.52
LipschitzRNN [4]	64.2	Random Forest [20]	5.69	1.85	1.74
		Ridge Regress. [20]	17.3	3.86	4.16

**Long Range Arena.** Our Long Range Arena experiments follow the same setup as the original S4 paper with some differences in model architecture and hyperparameters. The main global differences are as follows:

**Bidirectional** The original S4 layer is unidirectional or causal, which is an unnecessary constraint for the classification tasks appearing in LRA. Goel et al. [5] propose a bidirectional version of S4 that simply concatenates two S4 convolution kernels back-to-back. We use this for all tasks.

**GLU feedforward** S4 consists of  $H$  independent 1-dimensional SSMS, each of which are processed by an independent S4 SSM mapping  $(\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D})$ . These outputs are then mixed with a position-wise linear layer, i.e.  $\mathbf{W}y$  for a learned matrix  $\mathbf{W} \in \mathbb{R}^{H \times H}$ . Instead of this linear mapping, we use a GLU activation  $(\mathbf{W}_1 y) \circ \sigma(\mathbf{W}_2 y)$  for  $\mathbf{W}_1, \mathbf{W}_2 \in \mathbb{R}^{H \times H}$  [3]. These have been empirically found to improve linear layers of DNNs in general [19].

**Cosine scheduler** Instead of the plateau scheduler used in [9], we use a cosine annealing learning rate scheduler for all tasks.

**Regularization** Almost all tasks used no dropout and 0.05 weight decay.

**Architecture** Almost all tasks used an architecture with 6 layers,  $H = 256$  hidden units, BatchNorm, pre-norm placement of the normalization layer.

Exceptions to the above rules are described below. Full hyperparameters are in Table 11.

**sCIFAR / LRA Image.** This dataset is grayscale sequential CIFAR-10, and the settings for this task were taken from S4’s hyperparameters on the normal sequential CIFAR-10 task. In particular, this used LayerNorm instead of BatchNorm, a larger number of hidden features  $H$ , post-norm instead of pre-norm, and minor dropout. We note that the choice of normalization and increased  $H$  do not make a significant difference on final performance, still attaining classification accuracy in the high 80’s. Dropout does seem to make a difference.

**BIDMC.** We used a larger state size of  $N = 256$ , since we hypothesized that picking up higher frequency features on this dataset would help. We also used a step scheduler that decayed the LR by 0.5 every 100 epochs, following prior work [8, 18].

Table 11: The values of the best hyperparameters found for all datasets; full models on ablation datasets (Top) and LRA (Bottom). LR is learning rate and WD is weight decay. BN and LN refer to Batch Normalization and Layer Normalization.

	Depth	Features $H$	State Size $N$	Norm	Pre-norm	Dropout	LR	Batch Size	Epochs	WD	$(\Delta_{min}, \Delta_{max})$
<b>sCIFAR</b>	6	512	64	LN	False	0.1	0.01	50	200	0.05	(0.001, 0.1)
<b>SC</b>	6	128	64	BN	True	0	0.01	16	40	0.05	(0.001, 0.1)
<b>BIDMC</b>	6	128	256	LN	True	0	0.01	32	500	0.05	(0.001, 0.1)
<b>ListOps</b>	8	128	64	BN	False	0	0.01	50	40	0.05	(0.001, 0.1)
<b>Text</b>	6	256	64	BN	True	0	0.01	16	32	0.05	(0.001, 0.1)
<b>Retrieval</b>	6	256	64	BN	True	0	0.01	64	20	0.05	(0.001, 0.1)
<b>Image</b>	6	512	64	LN	False	0.1	0.01	50	200	0.05	(0.001, 0.1)
<b>Pathfinder</b>	6	256	64	BN	True	0	0.004	64	200	0.03	(0.001, 0.1)
<b>Path-X</b>	6	256	64	BN	True	0	0.0005	32	50	0.05	(0.0001, 0.01)

**ListOps.** We hypothesized that this task benefits from deeper models, because of the explicit hierarchical nature of the task, so the architecture used here had 8 layers and  $H = 128$  hidden features. However, results are very close with much smaller models. We also found that post-norm generalized better than pre-norm, but results are again close (less than 1% difference).

**PathX.** As described in Anonymous [1], the initialization range for PathX is decreased from  $(\Delta_{min}, \Delta_{max}) = (0.001, 0.1)$  to  $(\Delta_{min}, \Delta_{max}) = (0.0001, 0.01)$ .