

Structured state-space models are deep Wiener models

Fabio Bonassi*, Carl Andersson*, Per Mattsson*,
Thomas B. Schön*

* Department of Information Technology, Uppsala University,
75105 Uppsala, Sweden. E-mail: name.surname@it.uu.se

Abstract: The goal of this paper is to provide a system identification-friendly introduction to the Structured State-space Models (SSMs). These models have become recently popular in the machine learning community since, owing to their parallelizability, they can be efficiently and scalably trained to tackle extremely long sequence classification and regression problems. Interestingly, SSMs appear as an effective way to learn deep Wiener models, which allows us to reframe SSMs as an extension of a model class commonly used in system identification. To stimulate a fruitful exchange of ideas between the machine learning and system identification communities, we deem it useful to summarize the recent contributions on the topic in a structured and accessible form. At last, we highlight future research directions for which this community could provide impactful contributions.

Copyright © 2024 The Authors. This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0/>)

Keywords: Structured State-space models; system identification; deep learning.

1. INTRODUCTION

Recent years have been characterized by a remarkable research interest towards deep learning tools for data-driven control. An interesting use case is that of nonlinear system identification. Over the years, a variety of (increasingly complex) Neural Network (NN) architectures have indeed been proposed for nonlinear system identification, ranging from gated Recurrent NNs (RNNs) (Bonassi et al., 2022) to Transformer NNs (TNNs) (Sun and Wei, 2022). These models have proven to work well in many challenging identification problems, enabling the synthesis of accurate model predictive control laws, e.g. Lanzetti et al. (2019).

Despite their modeling power, these architectures are known to suffer from computational efficiency problems at training. On the one hand, RNNs are inherently sequential models, implying they have to be iteratively unrolled over the time axis T of training sequences (Bianchi et al., 2017). On the other hand, TNNs are plagued by a quadratic scaling issue ($\mathcal{O}(T^2)$), thus calling for tailored architectures like Longformers, which are however still untapped in the nonlinear system identification realm. Unfortunately, these problems have limited the use of RNNs and TNNs for long-term nonlinear system identification to those cases where enough computational budget is available.

What is more, both RNNs and TNNs are still fairly unexplored when it comes to their system-theoretical properties, such as Incremental Input-to-State Stability (δ ISS) (Bonassi et al., 2022), which are however essential for utilizing these models for control design (Bonassi et al., 2024). Most of the research in this area is tailored to specific

RNNs, see Miller and Hardt (2019); Bonassi et al. (2022), and mainly involves stability promotion via regularization.

To address these problems, Gu et al. (2021) proposed a **Structured State-space Model (SSM)** architecture named S4, which consists of multiple layers composed by LTI discrete-time systems followed by a nonlinear function. The term “structured” stems from the fact that this LTI system is given a specific structure to improve the architecture’s modeling performances while also reducing the computational cost at training (Yu et al., 2018). Nonlinear state-space models are not new, see Marconato et al. (2013), yet their adoption has been hampered by their crucial reliance on the model structure and on the initialization method of learnable parameters. The contribution of the S4 approach towards SSMs has therefore been that of providing (i) a novel, intrinsically stable, parametrization of the LTI system obtained by discretizing a continuous-time Diagonal Plus-Low Rank (DPLR) system, (ii) a new strategy towards the parameters’ initialization problem, (iii) a computationally efficient approach to simulate (and train) these models over extremely long sequences, and (iv) an empirical proof of the state-of-the-art performances of these models in long-term sequence learning problems.

Motivated by these appealing features, many works have continued to build on the S4 architecture. For example, Gupta et al. (2022) and Gu et al. (2022) have explored the benefits entailed by stricter SSM structures, namely the parametrization via diagonal continuous-time systems (S4D), and by simpler initialization strategies. Smith et al. (2022) have explored a novel, and somewhat more computationally efficient, simulation method for diagonal continuous-time parametrizations, named S5. Orvieto et al. (2023) recently investigated the parametrization of the LTI subsystems directly in the discrete time domain, resulting in the Linear Recurrent Unit (LRU) architecture.

* This research was financially supported by *Kjell och Märta Beijer Foundation* and by the project *Deep probabilistic regression – new models and learning algorithms* (contract number: 2021-04301), funded by the Swedish Research Council.

Contribution Despite the appealing results achieved by SSMs in the long-range arena benchmarks sequence classification problems, their use for nonlinear system identification is still unexplored. With this paper, we want to change that by making the following contributions. First of all we show that it is possible to interpret SSMs as *deep Wiener models*, i.e. model structures where several Wiener models are interconnected in series. An interesting note here is that even though the Wiener models have been extremely popular within system identification — see e.g. [Schoukens and Tiels \(2017\)](#) and references therein — their structure has been limited to “single-layer” or parallel architectures ([Wills and Ninness, 2012](#)). Our second contribution is to dissect the recent developments on SSMs and explain them in terms of their structure and parameterization, and to clearly separate this from their initialization, simulation, and training strategies. The presentation in the paper is also done using the language commonly used in the system identification community in order to speed up the use of these tools within this area.

Notation The imaginary unit is denoted by $i = \sqrt{-1}$. Given a vector v , we denote by v' its real transpose. For a time-dependent vector, the discrete-time index k is reported as a subscript, e.g., v_k . Moreover, we denote by $v_{a:b}$ (where $a \leq b$) the sequence $v_{a:b} = (v_a, v_{a+1}, \dots, v_b)$. For this sequence, we indicate by $\text{cat}(v_{a:b})$ the concatenation of its elements, i.e. $\text{cat}(v_{a:b}) = [v'_a, \dots, v'_b]'$, and by $\text{uncat}(\cdot)$ its inverse operation returning a sequence of vectors given their concatenation. Given a complex matrix A , we let $\text{conj}(A)$ be its element-wise complex conjugate and A^* be its Hermitian transpose. Diagonal matrices may be defined via the diag operator, as $A = \text{diag}(a_1, \dots, a_n)$.

2. STRUCTURED STATE-SPACE MODELS

Consider the model depicted in Figure 1, which consists of L Wiener systems interconnected in series. Each of these layers is here referred to as **Structured State-space Layer (SSL)**. Their interconnection results in an SSM, which can be interpreted as a specific configuration of a deep Wiener system. We let the generic ℓ -th SSL ($\ell \in \{1, \dots, L\}$) be represented by a discrete-time state-space model

$$\text{SSL}_\ell : \begin{cases} x_{k+1} = Ax_k + Bu_k, \\ \eta_k = Cx_k + Du_k, \\ y_k = \sigma(\eta_k) + Fu_k, \end{cases} \quad (1)$$

where, for compactness, the layer index is omitted. System (1) is characterized by the input vector $u \in \mathbb{R}^{n_u}$, the intermediate vector $\eta \in \mathbb{R}^{n_y}$, the output vector $y \in \mathbb{R}^{n_y}$, and the complex-valued state vector $x \in \mathbb{C}^{n_x}$. The SSL is parametrized by the matrices $\{A, B, C, D, F\}$. The output transformation $\sigma(\cdot)$ can be any nonlinear, Lipschitz-continuous activation function, such as the tanh, ELU, or Swish, see [Ramachandran et al. \(2017\)](#). In what follows, we aim to provide an overview of the possible structure, parametrization, initialization, and simulation strategies for this SSL.

Remark 1. When a deep SSM is considered ($L > 1$), each layer is parametrized and initialized independently from the others. The simulation is carried out iteratively over the set of layers, meaning that the output sequence of the ℓ -th layer is used as input of the layer $\ell + 1$.

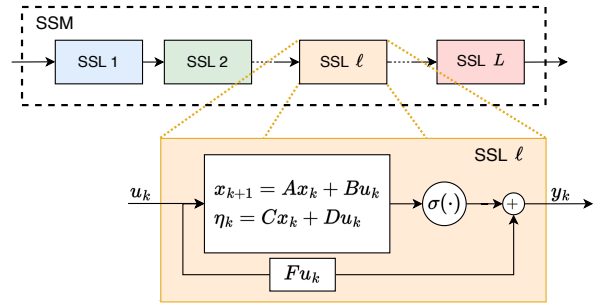


Fig. 1. Schematic of a Structured State-space Model.

We start by observing that, because both the input and output vector of the SSL are real-valued, the system matrices A , B , and C of (1) might be re-written as

$$A = \begin{bmatrix} \tilde{A} & \\ & \text{conj}(\tilde{A}) \end{bmatrix}, \quad B = \begin{bmatrix} \tilde{B} & \\ & \text{conj}(\tilde{B}) \end{bmatrix}, \quad (2)$$

$$C = [\tilde{C} \quad \text{conj}(\tilde{C})],$$

where $\tilde{A} \in \mathbb{C}^{n_\lambda \times n_\lambda}$, $\tilde{B} \in \mathbb{C}^{n_\lambda \times n_u}$, $\tilde{C} \in \mathbb{C}^{n_y \times n_\lambda}$, while $D \in \mathbb{R}^{n_y \times n_u}$, with $n_\lambda = \frac{n_x}{2}$. Matrix $F \in \mathbb{R}^{n_y \times n_u}$ can be interpreted as a skip connection, fixed to I if $n_u = n_y$. Note that if the SSLs are parametrized to be structurally Schur stable, the stability of the SSM is guaranteed.

Proposition 2. If each layer $\ell \in \{1, \dots, L\}$ is parametrized by a Schur-stable matrix \tilde{A} , then the SSM is Incrementally Input-to-State Stable (δ ISS) ([Bonassi et al., 2022, Def. 4](#)).

Proof. Owing to the Schur stability of \tilde{A} and Lipschitz continuity of $\sigma(\cdot)$, the SSL can be easily proven (globally) δ ISS. The SSM is then the series of δ ISS systems and is therefore δ ISS ([Angeli, 2002](#)). \square

While (2) ensures that the eigenvalues come in complex-conjugate pairs, both in the linear ([Yu et al., 2018](#)) and the nonlinear realm ([Marconato et al., 2013](#)) it is well known that state-space models call for additional structure to reduce the number of learnable parameters and hence achieve improved modeling performances and computational efficiency. Besides, the initialization of these parameters is also paramount to achieving satisfactory performances. In the following, the strategies proposed in recent literature to address these problems are discussed.

3. DISCRETE-TIME SSL PARAMETRIZATIONS

3.1 Discrete-time diagonal parametrization

One intuitive approach is that of parametrizing the SSL as a discrete-time complex-valued diagonal system, as proposed by [Orvieto et al. \(2023\)](#). In particular, \tilde{A} can be parametrized as a Schur-stable diagonal matrix

$$\tilde{A} = \tilde{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_{n_\lambda}). \quad (3a)$$

Each eigenvalue $j \in \{1, \dots, n_\lambda\}$ is, in turn, parametrized by the modulus $-\exp(\mu_j) \in (0, 1)$ and the phase $\exp(\theta_j)$,

$$\lambda_j = \exp\left(-\exp(\mu_j) + i \exp(\theta_j)\right). \quad (3b)$$

Note that since $|\lambda_j| < 1$ this parametrization guarantees that the SSM is structurally δ ISS (Proposition 2). A relevant design choice advocated by [Orvieto et al. \(2023\)](#) is that of reparametrizing \tilde{B} as

$$\tilde{B} = \text{diag}(\gamma_1, \dots, \gamma_{n_\lambda}) \tilde{B}, \quad (4a)$$

where the normalization factor is defined as

$$\gamma_j = (1 - |\lambda_j|^2)^{\frac{1}{2}}, \quad \forall j \in \{1, \dots, n_\lambda\}. \quad (4b)$$

This ensures that the rows of \tilde{B} are normalized so that white noise inputs yield state trajectories with the same energy content as of the input. Therefore, the set of learnable parameters of this SSL parametrization reads

$$\Theta = \left\{ \{\mu_j, \theta_j\}_{j \in \{1, \dots, n_\lambda\}}, \tilde{B}, \tilde{C}, D, F \right\}. \quad (5)$$

3.2 Initialization strategy

Orvieto et al. (2023) propose to initialize the dynamics matrix randomly inside a portion of the circular crown lying within the unit circle. Letting $0 \leq r < \bar{r} < 1$ be the minimum and maximum modulus of each eigenvalue, respectively, and $0 \leq \underline{\theta} < \bar{\theta} < \pi$ be the minimum and maximum phase,

$$\begin{aligned} \mu_j &\sim \mathcal{U}[\log(-\log(\bar{r})), \log(-\log(r))], \\ \theta_j &\sim \mathcal{U}[\log(\log(\underline{\theta})), \log(\log(\bar{\theta}))], \end{aligned} \quad (6)$$

where $\mathcal{U}[a, b]$ denotes the uniform distribution with support $[a, b]$. The complex matrices \tilde{B} and \tilde{C} , and the real matrices D and F ¹ can be initialized with the Xavier initialization method (Kumar, 2017). That is, they are sampled from a Normal distribution whose variance is scaled by a factor proportional to the number of columns.

4. CONTINUOUS-TIME REPARAMETRIZATIONS

Another approach is that of parametrizing the SSL via a discretized continuous-time system, as proposed by Gu et al. (2021) in the context of S4 models. To this end, we let the SSL be parametrized as

$$\begin{cases} \dot{x}(t) = A_c x(t) + B_c u(t), \\ \eta(t) = C_c x(t), \\ y(t) = \sigma(\eta(t)) + F u(t), \end{cases} \quad (7)$$

where — similarly to (2) — $\{A_c, B_c, C_c\}$ are structured in terms of the blocks $\{\tilde{A}_c, \tilde{B}_c, \tilde{C}_c\}$

$$\begin{aligned} A_c &= \begin{bmatrix} \Gamma \tilde{A}_c & \\ & \Gamma \text{conj}(\tilde{A}_c) \end{bmatrix}, \quad B_c = \begin{bmatrix} \Gamma \tilde{B}_c & \\ \Gamma \text{conj}(\tilde{B}_c) \end{bmatrix}, \\ C_c &= [\tilde{C}_c \text{ conj}(\tilde{C}_c)], \end{aligned} \quad (8)$$

with $\tilde{A}_c \in \mathbb{C}^{n_\lambda \times n_\lambda}$, $\tilde{B}_c \in \mathbb{C}^{n_\lambda \times n_u}$, $\tilde{C}_c \in \mathbb{C}^{n_y \times n_\lambda}$. The factor Γ in (8) is a learnable timescale parameter, which can either be a real and positive scalar (Gu et al., 2021) or a real-valued, positive-definite diagonal matrix (Smith et al., 2022), i.e., $\Gamma = \text{diag}(\gamma_1, \dots, \gamma_{n_\lambda})$.

Under this parametrization, the discrete-time system matrices $\{A, B, C, D\}$ can finally be expressed in terms of their continuous-time counterparts, $\{A_c, B_c, C_c\}$, by defining the sampling time τ and a discretization method, such as forward Euler, the bilinear transform, or Zero-Order Hold (ZOH). As shown in Section 3.1, an additional structure can now be imposed on \tilde{A}_c to mitigate the overparametrization, thus making the training procedure more scalable, and to enforce the structural stability of the SSL and hence, owing to Proposition 2, the stability of the SSM. The possible structures of \tilde{A}_c are now discussed.

¹ If $n_u = n_y$ the skip connection F is usually fixed to the identity.

4.1 Continuous-time diagonal parametrization

One strategy, advocated by Gu et al. (2022) and Smith et al. (2022), is that of parametrizing \tilde{A}_c as a diagonal matrix, i.e. $\tilde{A}_c = \tilde{\Lambda}_c$, where

$$\tilde{\Lambda}_c = \text{diag}(\lambda_1, \dots, \lambda_{n_\lambda}). \quad (9a)$$

Each eigenvalue $j \in \{1, \dots, n_\lambda\}$ is, in turn, parametrized by the logarithm of its real and imaginary part, denoted as α_j^{re} and α_j^{im} , respectively. That is,

$$\lambda_j = -\exp(\alpha_j^{\text{re}}) + i \exp(\alpha_j^{\text{im}}). \quad (9b)$$

Note that, because $\Re(\lambda_j) < 0$, this parametrization structurally guarantees the asymptotic stability of the SSL, and hence the δ ISS of the SSM (Proposition 2). Overall, the set of learnable parameters is

$$\Theta = \left\{ \{\alpha_j^{\text{re}}, \alpha_j^{\text{im}}\}_{j \in \{1, \dots, n_\lambda\}}, \tilde{B}_c, \tilde{C}_c, F, \Gamma \right\}. \quad (10)$$

As a side note, it is worth pointing out that the eigenvalues can also be parametrized in terms of the modulus and phase instead, as in (3b).

4.2 Continuous-time DPLR parametrization

The slightly more general Diagonal Plus Low-Rank (DPLR, Gu et al. (2021)) structure has also been proposed. According to this strategy, \tilde{A}_c is parametrized as

$$\tilde{A}_c = \tilde{\Lambda}_c - \tilde{P} \tilde{Q}^*, \quad (11)$$

where $\tilde{\Lambda}_c$ is a diagonal component defined as

$$\tilde{\Lambda}_c = \text{diag}(\lambda_1, \dots, \lambda_{n_\lambda}) \quad (12a)$$

and $\tilde{P}, \tilde{Q} \in \mathbb{C}^{n_\lambda \times n_r}$ yield a low rank component, with $n_r \ll n_\lambda$. In order to guarantee the structural stability of the SSL, and hence the δ ISS of the SSM, one can take $\tilde{P} = \tilde{Q}$ (see Gu et al. (2021)) and

$$\lambda_j = -\phi(\alpha_j^{\text{re}}) + i \alpha_j^{\text{im}}, \quad (12b)$$

with $\phi(\cdot)$ being any strictly-positive function, applied element-wise, e.g. $\phi(\alpha) = \max(0, \alpha) + \varepsilon$, with $\varepsilon > 0$. This ensures the negative definiteness of \tilde{A}_c . The set of learnable parameters, under the DPLR parametrization, is

$$\Theta = \left\{ \{\alpha_j^{\text{re}}, \alpha_j^{\text{im}}\}_{j \in \{1, \dots, n_\lambda\}}, \tilde{P}, \tilde{Q}, \tilde{B}_c, \tilde{C}_c, F, \Gamma \right\}. \quad (13)$$

4.3 Initialization strategies

In the case of continuous-time DPLR parametrizations Gu et al. (2021) proposed to resort to the so-called HiPPO framework (Gu et al., 2020) to initialize the learnable matrices \tilde{B}_c , $\tilde{\Lambda}_c$, \tilde{P} , and \tilde{Q} . In particular, being it associated with long-term memory properties, the HiPPO matrix is regarded as a suitable initialization for (11). Such initialization is carried out by (i) building the HiPPO-LegS matrix, (ii) computing its normal and low-rank components, and (iii) applying the eigenprojection, which yields the diagonal matrix $\tilde{\Lambda}_c$ and the projected low-rank components \tilde{P} and \tilde{Q} . The HiPPO framework also provides an initialization for the input matrix \tilde{B}_c . For more details, the interested reader is referred to (Gu et al., 2020) or Appendix A of (Bonassi et al., 2023). The complex matrix \tilde{C}_c and the real matrix F are randomly initialized, e.g., with the Xavier method.

Gu et al. (2022) proposed a similar strategy to initialize the learnable parameters of the continuous-time diagonal parametrizations (10). They suggested to initialize $\tilde{\Lambda}_c$ with the diagonal component of the normalized HiPPO-LegS matrix, discarding the low-rank terms \tilde{P} and \tilde{Q} .

As an alternative, we point out that a strategy similar to the one described in Section 3.2 could be also applied to continuous-time diagonal parametrizations. Given a range for the modulus of eigenvalues, $0 < \underline{r} < \bar{r} < \frac{\pi}{\tau}$, and for their phase, $\frac{\pi}{2} < \underline{\theta} < \bar{\theta} \leq \pi$, one can sample

$$\lambda_j = r \exp(i\theta), \quad r \sim \mathcal{U}[\underline{r}, \bar{r}], \quad \theta \sim \mathcal{U}[\underline{\theta}, \bar{\theta}], \quad (14a)$$

and initialize α_j^{re} and α_j^{im} as

$$\alpha_j^{\text{re}} = \log(-\Re(\lambda_j)), \quad \alpha_j^{\text{im}} = \log(\Im(\lambda_j)). \quad (14b)$$

Under (14), Γ can be initialized to the identity matrix. The complex-valued matrices \tilde{B}_c and \tilde{C}_c , and the real-valued matrix F can be randomly initialized.

4.4 Comments on the parametrization strategies

For clarity purposes, the parameterizations have been introduced here in reverse chronological order. The continuous-time DPLR parametrization of S4 has indeed been the one that sparked a renewed research interest in SSMs (Gu et al., 2021). This strategy of parametrizing SSLs in the continuous-time domain has been proposed for two reasons. Firstly, it enables the HiPPO framework to be used for the weights' initialization. Second, it allows simulation of the learned SSMs with a sampling time potentially different from that of the training data, while conventional (discrete-time) RNNs would call for a new training.

We point out, however, that the choice of parametrizing SSLs in the continuous-time domain implies the *aliasing* problem, usually overlooked in the SSMs literature. Since (7) is discretized with the sampling time τ of the available training data, one should ensure that (i) the eigenvalues are not initialized beyond the Nyquist-Shannon bandwidth $\frac{\pi}{\tau}$ and (ii) the eigenvalues remains within this bandwidth throughout the model's training. Eigenvalues beyond such bandwidth likely lead to reduced modeling capabilities and poor performances in extrapolation especially when different sampling times are adopted. In this regard, we note that the HiPPO-based initialization strategies proposed by Gu et al. (2021), Gu et al. (2022), and Smith et al. (2022) do not, in general, meet these requirements. It is indeed well-known that the HiPPO matrix's eigenvalues grow quickly with n_λ . To mitigate this issue, one should initialize the timescale parameter Γ so that $\tilde{\Lambda}_c$ has eigenvalues within the Nyquist bandwidth.²

5. COMPUTATIONAL EFFICIENCY OF SSM

Because SSMs are learned with a simulation error minimization strategy, it is paramount to compute the output trajectory of the model — given the input sequence $u_{0:T}$ — as efficiently as possible. As discussed in Remark 1, an SSM can be simulated by sequentially evaluating its SSLs. In what follows, we therefore focus on how the generic SSL (1) can be efficiently simulated.

² In the SSM literature Γ is often randomly sampled from a distribution which is uniform in its log-space (Smith et al., 2022).

Since the SSL's dynamics are linear and asymptotically stable, and the static nonlinearity only affects the output variable, the state trajectory can be computed via the truncated Lagrange equation

$$x_{t+1} \approx \underbrace{[A^{R-1}B \ A^{R-2}B \ \dots \ AB \ B]}_{\mathfrak{B}_R} \text{cat}(u_{t-R+1:t}), \quad (15)$$

with the truncation $R \leq T$ being sufficiently large to ensure that the spectral radius of A^R is negligible. In the SSM literature, (15) is known as the *convolutional form*, as one can retrieve the state trajectory by convolving the input sequence $u_{0:T}$ with the *filter* \mathfrak{B}_R ,

$$x_{1:T+1} = \text{uncat}(\mathfrak{B}_R \odot \text{cat}(u_{0:T})), \quad (16a)$$

$$\eta_t = Cx_t + Du_t, \quad \forall t \in \{0, \dots, T\}.$$

The output $y_{0:T}$ can then be straightforwardly computed from $\eta_{0:T}$ by applying a static output transformation,

$$y_t = \sigma(\eta_t) + Fu_t, \quad \forall t \in \{0, \dots, T\}. \quad (16b)$$

The convolutional form (16) enjoys a noteworthy computational efficiency since, as discussed in the remainder of this section, both the materialization of the filter \mathfrak{B}_R and the convolution itself can be easily parallelized. This allows the computational cost of simulating the SSL over a T -long input sequence to scale as $\mathcal{O}(T \log T)$ in place of the $\mathcal{O}(T^2)$ entailed by the iterative, sequential, application of (1) over the time axis (Smith et al., 2022). The algorithms allowing such a cheap computation of (16a) are now outlined.

5.1 Convolution via Parallel Scan

As proposed by Smith et al. (2022), the task of scalably computing (16a) can be addressed via Parallel Scan (Blelloch, 1990). This approach stems from the idea that intermediate blocks of \mathfrak{B}_R can be computed in parallel and combined, allowing to scale as $\mathcal{O}(T \log_2 T)$ times the complexity of the product $A \cdot A$. While such a product in the particular case of diagonal parametrization scales as $\mathcal{O}(n_\lambda)$, for non-diagonal matrices its complexity is $\mathcal{O}(n_\lambda^3)$, thus rapidly increasing with the state dimension. Thus, while the parallel scan approach can, in principle, be applied to any parametrization, it has primarily been utilized for diagonal ones, see Orvieto et al. (2023). For more details on how parallel scan is defined and implemented the reader is addressed to Smith et al. (2022), Appendix H.

Remark 3. Given the structure (2) of the SSL (1), one can instead simulate the subsystem $\tilde{x}_{k+1} = \tilde{A}\tilde{x}_k + \tilde{B}u_k$ and $\tilde{\eta}_k = \tilde{C}\tilde{x}_k + Du_k$, and then compute the output as $y_k = \sigma(2\Re(\tilde{\eta}_k)) + Fu_k$. This allows for further reduction of the computational burden without any approximations.

5.2 Convolution via FFT

For SSLs parametrized by the DPLR structure described in Section 4.2, it is more efficient to carry out the convolution (16a) in the frequency domain. Operating in this domain, DPLR-parametrized SSLs can be simulated as efficiently as in diagonal parametrizations (Gu et al., 2021).

Let us consider, for the purpose of explanation, the case $n_u = n_y = 1$.³ Then, this approach consists in (i) computing the Fast Fourier Transform (FFT) of the input

³ This approach has been generalized to the case $n_u > 1$ and/or $n_y > 1$ by applying it separately on each input-output pair.

signal $\mathcal{U}(\omega) = \text{FFT}(u_{0:T})$, (ii) multiplying it (frequency-wise) with the Fourier transform $\tilde{\mathcal{H}}(\omega)$ of the discretized linear subsystem defined by the triplet $(\tilde{A}_c, \tilde{B}_c, \tilde{C}_c)$, (iii) computing the inverse FFT (iFFT), and (iv) applying the nonlinear output transformation. That is,

$$\tilde{\eta}_{0:T} = \text{iFFT}(\tilde{\mathcal{H}}(\omega) \cdot \mathcal{U}(\omega)), \quad (17a)$$

$$\eta_t = 2\Re(\tilde{\eta}_t), \quad t \in \{0, \dots, T\}, \quad (17b)$$

$$y_t = \sigma(\eta_t) + F u_t, \quad t \in \{0, \dots, T\}. \quad (17c)$$

Note that both the FFT and iFFT operations scale with complexity $\mathcal{O}(T \log T)$. Gu et al. (2021) have shown that, for continuous-time DPLR-parametrized SSLs (like S4), the Fourier transform $\tilde{\mathcal{H}}(\omega)$ can be efficiently evaluated by resorting to black-box Cauchy kernels. The interested reader may see Appendix B of the extended version (Bonassi et al., 2023) for a derivation using this paper’s notation. With this approach, one issue to pay attention to is frequency leakage due to the input’s nonperiodicity.

6. NUMERICAL EXAMPLE

To preliminary quantify the modeling capabilities of the described SSM architectures, we considered the well-known Silverbox benchmark Wigren and Schoukens (2013), for which a public dataset is available. The Silverbox is an electronic circuit that mimics the input-output behavior of a damped mechanical system with a nonlinear elastic coefficient, and it has often been used for benchmarking, e.g., Wiener models (Tiels, 2015). The training and validation data consist of ten experiments in which a multisine input with random frequency components is applied to the apparatus. Each experiment features 8192 input-output datapoints, collected with a sampling rate of 610.35 Hz. The validation dataset has been constructed by extracting 76 subsequences (of length $T = 512$ steps) from a random experiment. From the remaining nine experiments 684 training subsequences have been extracted.

The Silverbox dataset also contains an independent *test dataset* that can be used to assess the accuracy of the identified models. This test dataset consists of 40500 samples collected by exciting the system with a filtered Gaussian noise having linearly increasing mean value. As noted by Tiels (2015), the test dataset is characterized by two regions. The first 25000 time-steps allow quantification of the model’s accuracy in interpolation, i.e., within operating regions well explored by the training and validation data. The successive time-steps allow, instead, enable the assessment of the model’s extrapolation performances in operating regions not explored in the training and validation datasets, particularly in terms of the input signal amplitude. For this reason, the modeling performances scored by the identified SSMs—measured by RMSE [mV] and FIT index [%]—are reported both on the overall test dataset and limitedly to the interpolatory regime.

6.1 Identification results

The training procedure was implemented in PyTorch 2.1 and is described in more detail in the accompanying code⁴ or in the Appendix C of the extended version (Bonassi et al., 2023). The following SSM configurations have been

Table 1. Performance of the identified SSMs

Model	First 25000 steps		Full	
	RMSE	FIT	RMSE	FIT
S4 ($L = 4, n_\lambda = 10$) ⁵	0.81	97.60	4.73	95.49
S5 ($L = 4, n_\lambda = 10$) ⁵	0.73	97.78	3.56	96.48
S5R ($L = 4, n_\lambda = 10$)	0.77	97.62	4.58	95.92
LRU ($L = 4, n_\lambda = 10$)	0.73	97.71	4.18	96.37
TCNN (Andersson et al., 2019)	0.75	-	4.9	-
LSTM (Andersson et al., 2019)	0.31	-	4.0	-
BLA (Tiels, 2015)	-	-	13.7	-
Wiener (Tiels, 2015)	1.9	-	9.2	-
Grey-box NARX (Ljung et al., 2004)	-	-	0.3	-

considered for identification. Additional details about their structures, initialization, and training hyperparameters are reported in Bonassi et al. (2023), Appendix D.

S4 (Gu et al., 2021) — SSM whose layers are parametrized in the continuous-time domain (7)-(8) by a DLPR-structured state matrix (11) initialized via HiPPO.

S5 (Smith et al., 2022) — SSM whose layers are parametrized by continuous-time diagonal systems (7)-(9) initialized via HiPPO.

S5R — SSM whose layers are parametrized by continuous-time diagonal systems (7)-(9). Similar to S5, but initialized by random sampling of the eigenvalues, (14).

LRU (Orvieto et al., 2023) — SSM whose layers are parametrized by discrete-time diagonal systems (2)-(4).

In Table 1 the performance metrics scored by the best identified SSM are reported, and they are compared to some of those reported in the literature⁶. Note that, although these SSMs are more accurate than traditional Wiener models, they are still in line with those achieved by other deep learning models like TCNNs and LSTMs, while featuring significantly less learnable parameters⁷ and high parallelizability of the training procedure.

At last, let us point out that training competitive SSMs requires, in general, a careful selection of the architecture hyperparameters. These not only include the number of layers and the state and output size of each layer, but also a suitable initialization strategy for the selected parametrization. The computational efficiency of these models and their reduced parameter footprint come at the cost of a slightly increased architecture design effort. Further testing of SSMs for nonlinear system identification is thus advisable to establish empirical design criteria.

7. CONCLUSIONS AND RESEARCH DIRECTIONS

Structured State-space Models represent an interesting approach to identifying deep Wiener models. In this paper we summarized recent developments of SSMs in the machine learning community, attempting to disentangle their structure, parameterization, initialization, and simulation aspects. The goal was to make these models more accessi-

⁵ At least one eigenvalue falls beyond the Nyquist frequency.

⁶ For detailed accounts of the literature, the reader is referred to Andersson et al. (2019) and Maroli et al. (2019).

⁷ A diagonal SSL parametrized in continuous-time has $n_\lambda(3 + 2n_y + 2n_u) + n_y n_u$ learnable parameters, compared to the $n_x(4 + 3n_x + n_u + n_y) + n_y$ learnable parameters of an LSTM layer.

⁴ Source code: <https://github.com/bonassifabio/SSM-sysid>

ble for the system identification community, to stimulate further investigation of these models for identification. We believe there is scope for interesting future developments—possible research directions are summarized below.

Minimal structures Structures with fewer learnable parameters should be considered, as long as they ensure the model's stability and allow for an efficient simulation. In this context, it would be advisable to evaluate whether parameterizations in the transfer function domain, similar to dynoNet Forgiione and Piga (2021), yield less over-parametrized and more effective models.

Data-driven initializations Novel strategies could exploit training data to improve and tailor the initialization of learnable parameters. For Wiener models, for example, initialization based on best linear approximators are known to boost the identified model performances.

Discrete- or continuous-time parametrizations? Despite the revived interest in SSMs originated from continuous-time parameterizations, Orvieto et al. (2023) questioned their real need, showing how similar performances can be achieved also with discrete-time parameterizations. This issue should also be further investigated in light of the aliasing problem discussed in Section 4.4.

Extensive benchmarking The performances of SSMs for nonlinear system identification should be assessed on a broader set of benchmarks, possibly featuring long-term time dependencies, for which SSMs were proposed.

REFERENCES

- Andersson, C., Ribeiro, A.H., Tiels, K., Wahlström, N., and Schön, T.B. (2019). Deep convolutional networks in system identification. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, 3670–3676. IEEE.
- Angeli, D. (2002). A lyapunov approach to incremental stability properties. *IEEE Transactions on Automatic Control*, 47(3), 410–421.
- Bianchi, F.M., Maiorino, E., Kampffmeyer, M.C., Rizzi, A., and Jenssen, R. (2017). *Recurrent neural networks for short-term load forecasting: an overview and comparative analysis*. Springer.
- Belloch, G.E. (1990). Prefix sums and their applications. In J.H. Reif (ed.), *Synthesis of parallel algorithms*. Morgan Kaufmann Publishers Inc.
- Bonassi, F., Andersson, C., Mattsson, P., and Schön, T.B. (2023). Structured state-space models are deep wiener models. *arXiv preprint arXiv:2312.06211*. Extended version.
- Bonassi, F., Farina, M., Xie, J., and Scattolini, R. (2022). On Recurrent Neural Networks for learning-based control: recent results and ideas for future developments. *Journal of Process Control*, 114, 92–104.
- Bonassi, F., La Bella, A., Farina, M., and Scattolini, R. (2024). Nonlinear MPC design for incrementally ISS systems with application to GRU networks. *Automatica*, 159, 111381.
- Forgiione, M. and Piga, D. (2021). dynoNet: A neural network architecture for learning dynamical systems. *International Journal of Adaptive Control and Signal Processing*, 35(4), 612–626.
- Gu, A., Dao, T., Ermon, S., Rudra, A., and Ré, C. (2020). Hippo: Recurrent memory with optimal polynomial projections. *Advances in neural information processing systems*, 33, 1474–1487.
- Gu, A., Goel, K., Gupta, A., and Ré, C. (2022). On the parameterization and initialization of diagonal state space models. *Advances in Neural Information Processing Systems*, 35, 35971–35983.
- Gu, A., Goel, K., and Ré, C. (2021). Efficiently modeling long sequences with structured state spaces. *arXiv preprint arXiv:2111.00396*.
- Gupta, A., Gu, A., and Berant, J. (2022). Diagonal state spaces are as effective as structured state spaces. *Advances in Neural Information Processing Systems*, 35, 22982–22994.
- Kumar, S.K. (2017). On weight initialization in deep neural networks. *arXiv preprint arXiv:1704.08863*.
- Lanzetti, N. et al. (2019). Recurrent neural network based MPC for process industries. In *2019 18th European Control Conference (ECC)*, 1005–1010. IEEE.
- Ljung, L., Zhang, Q., Lindskog, P., and Juditski, A. (2004). Estimation of grey box and black box models for nonlinear circuit data. *IFAC Proceedings Volumes*, 37(13), 399–404.
- Marconato, A., Sjöberg, J., Suykens, J.A., and Schoukens, J. (2013). Improved initialization for nonlinear state-space modeling. *IEEE Transactions on instrumentation and Measurement*, 63(4), 972–980.
- Maroli, J.M., Özgüner, Ü., and Redmill, K. (2019). Non-linear system identification using temporal convolutional networks: a silverbox study. *IFAC-PapersOnLine*, 52(29), 186–191.
- Miller, J. and Hardt, M. (2019). Stable recurrent models. In *International Conference on Learning Representations*. ArXiv preprint arXiv:1805.10369.
- Orvieto, A. et al. (2023). Resurrecting recurrent neural networks for long sequences. *arXiv preprint arXiv:2303.06349*.
- Ramachandran, P., Zoph, B., and Le Quoc, V. (2017). Searching for activation functions. *arXiv preprint arXiv:1710.05941*.
- Schoukens, M. and Tiels, K. (2017). Identification of block-oriented nonlinear systems starting from linear approximations: A survey. *Automatica*, 85, 272–292.
- Smith, J.T., Warrington, A., and Linderman, S. (2022). Simplified state space layers for sequence modeling. In *The Eleventh International Conference on Learning Representations*.
- Sun, Y. and Wei, H.L. (2022). Efficient mask attention-based narmax (mab-narmax) model identification. In *2022 27th International Conference on Automation and Computing (ICAC)*, 1–6. IEEE.
- Tiels, K. (2015). *Wiener system identification with generalized orthonormal basis functions*. PhD thesis, Vrije Universiteit Brussel.
- Wigren, T. and Schoukens, J. (2013). Three free data sets for development and benchmarking in nonlinear system identification. In *2013 European control conference (ECC)*, 2933–2938. IEEE.
- Wills, A. and Ninness, B. (2012). Generalised hammerstein-wiener system estimation and a benchmark application. *Control Engineering Practice*, 20(11), 1097–1108.
- Yu, C., Ljung, L., and Verhaegen, M. (2018). Identification of structured state-space models. *Automatica*, 90, 54–61.