

Meta-Learning of Neural State-Space Models Using Data From Similar Systems

Chakrabarty, Ankush; Wichern, Gordon; Laughman, Christopher R.

TR2023-087 July 11, 2023

Abstract

Deep neural state-space models (SSMs) provide a powerful tool for modeling dynamical systems solely using operational data. Typically, neural SSMs are trained using data collected from the actual system under consideration, despite the likely existence of operational data from similar systems which have previously been deployed in the field. In this paper, we propose the use of model-agnostic meta-learning (MAML) for constructing deep encoder network-based SSMs, by leveraging a combination of archived data from similar systems (used to meta-train offline) and limited data from the actual system (used for rapid online adaptation). We demonstrate using a numerical example that meta-learning can result in more accurate neural SSM models than supervised- or transfer-learning, despite few adaptation steps and limited online data. Additionally, we show that by carefully partitioning and adapting the encoder layers while fixing the state-transition operator, we can achieve comparable performance to MAML while reducing online adaptation complexity.

World Congress of the International Federation of Automatic Control (IFAC) 2023

Meta-Learning of Neural State-Space Models Using Data From Similar Systems

Ankush Chakrabarty* Gordon Wichern*
Christopher R. Laughman*

* *Mitsubishi Electric Research Laboratories, Cambridge, MA 02139, USA. (Corresponding author e-mail: achakrabarty@ieee.org).*

Abstract: Deep neural state-space models (SSMs) provide a powerful tool for modeling dynamical systems solely using operational data. Typically, neural SSMs are trained using data collected from the actual system under consideration, despite the likely existence of operational data from similar systems which have previously been deployed in the field. In this paper, we propose the use of model-agnostic meta-learning (MAML) for constructing deep encoder network-based SSMs, by leveraging a combination of archived data from similar systems (used to meta-train offline) and limited data from the actual system (used for rapid online adaptation). We demonstrate using a numerical example that meta-learning can result in more accurate neural SSM models than supervised- or transfer-learning, despite few adaptation steps and limited online data. Additionally, we show that by carefully partitioning and adapting the encoder layers while fixing the state-transition operator, we can achieve comparable performance to MAML while reducing online adaptation complexity.

Keywords: Learning, modeling, data-driven methods, Koopman operator, system identification.

1. INTRODUCTION

Data-driven system identification is often a necessary step for model-based design of control systems. While many data-driven modeling frameworks have been demonstrated to be effective, the class of models that contain a state-space description at their core have typically been easiest to integrate with model-based control and estimation algorithms, e.g., model predictive control or Kalman filtering.

Early implementations of neural state-space models (SSMs) employed shallow recurrent layers and were dependent on linearization to obtain linear representations (Zamarreño and Vega, 1998) or linear-parameter-varying system representations (Bao et al., 2020). Recent advancements in deep neural networks have enabled embedding SSMs into the neural architecture explicitly without post-hoc operations (Forgione et al., 2022), and therefore the SSM description can be learned directly during training; see Legaard et al. (2021) for a recent survey. For instance, unmodeled dynamics remaining after procuring a physics-informed prior model can be represented using neural SSMs (Forgione and Piga, 2020, 2021), and additional control-oriented structure can be embedded during training (Skomski et al., 2021). Another interesting direction of research has led to the development of autoencoder-based SSMs, where the neural architecture comprises an encoder that transforms the ambient state-space to a (usually high-dimensional) latent space, a decoder that inverse-transforms a latent state to the corresponding ambient state, and a linear SSM in the latent space that satisfactorily approximates the system’s underlying dynamics (Masti and Bemporad, 2021; Iacob et al., 2021; Bertalan et al., 2019). Even without the decoder, deep

encoder networks have proven useful for neural state-space modeling (Beintema et al., 2021). An argument for the effectiveness of autoencoder-based approaches is based on Koopman operator theory (Koopman and van Neumann, 1932), which posits that a nonlinear system (under some mild assumptions) can be lifted to an infinite-dimensional latent space where the state-transition is linear; an autoencoder allows a finite-dimensional, therefore tractable, approximation of the Koopman lifting/lowering transformations (Lusch et al., 2018).

Extensively, neural SSMs have been constructed using data from the target system under consideration. However, in practice, one likely has access to data for a range of similar (not necessarily identical) systems that contain information which could prove useful for speeding up the construction of a neural SSM for the target system. Herein, we show that there is potential in using data from similar systems for neural SSM modeling, and provide a meta-learning approach for tractably obtaining such a model. To the best of our knowledge, the existing literature does not contain a tractable solution to this open problem at this time, although the reduction of model estimation error by using data from linear systems within a prescribed ball has been proven in Xin et al. (2022). Though not for SSMs, meta-learning has been proposed for optimization (Zhan et al., 2022; Chakrabarty, 2022), adaptive control (Richards et al., 2021), and receding horizon control (Arcari et al., 2020; Muthirayan and Khargonekar, 2020).

The main contribution of this work is to propose a meta-learning approach for neural state-space modeling, via deep encoder networks, using data obtained from a range of similar systems. In meta-learning, a (often, deep) neural

network is trained on a variety of similar ‘source’ system models so that it can make accurate predictions for a given target system model, with only a few data points from the target system and a small number of gradient-based adaptation steps. Concretely, we learn encoder weights and state/output matrices from a variety of source systems’ sensor data. Consequently, even with a small amount of the target system’s data, the deep encoder network can be quickly adapted online to obtain a set of encoder weights and state/output matrices representing the target system dynamics. This meta-learned neural SSM is shown via a numerical example to result in higher predictive accuracy than a neural SSM trained solely using target system data, or even a neural SSM trained on the entire source plus target data. We employ model agnostic meta-learning (MAML) algorithms, which are trained by solving a bi-level optimization problem (Finn et al., 2017), wherein the outer loop extracts task-independent features across a set of source tasks, and the inner loop adapts to a specific model with a few iterations and limited data.

2. PRELIMINARIES

We consider a family of parameterized discrete-time nonlinear systems of the form

$$x_{t+1} = f(x_t, \theta_f), \quad (1a)$$

$$y_t = h(x_t, \theta_h) \quad (1b)$$

where $x_t \in \mathbb{R}^{n_x}$ denotes the state of the system at time $t \in \mathbb{N}$ with x_0 being the initial state, $y \in \mathbb{R}^{n_y}$ denotes the measured outputs, f denotes the unknown dynamics, h the unknown output function, and $\theta := [\theta_f, \theta_h] \in \mathbb{R}^{n_\theta}$ denotes a vector of unknown model parameters.

Since f and θ are unknown, our objective is to construct a neural state-space model that can replicate the dynamics of a query system of the form (1) parameterized by θ^* , which is also unknown. Let $Y(\theta^*, T)$ denote a trajectory of outputs generated by the query system over a time range $[0, T]$, where T is small. This implies that the query system dataset

$$\mathcal{D}_{\text{query}} \triangleq Y(\theta^*, T) \equiv Y^*$$

has limited size. One could use this query data to construct a neural state-space model (SSM) of the form

$$z_t = f_{\text{enc}}(Y_{t-H:t}) \quad (2a)$$

$$z_{t+1} = A_z z_t, \quad (2b)$$

$$\hat{y}_t = C_z z_t, \quad (2c)$$

which involves optimizing the weights of an encoder network $f_{\text{enc}}(\cdot)$, along with the elements of the linear decoders A_z and C_z ; a schematic diagram of the proposed neural SSM architecture is shown in Fig. 1. The input to the encoder is a window of length $H \in \mathbb{N}$ containing past measurements denoted

$$Y_{t-H:t} \triangleq \{y_{t-H}, y_{t-H+1}, \dots, y_{t-1}\},$$

and the latent state learned by the encoder net is given by $z \in \mathbb{R}^{n_z}$, where $n_z \in \mathbb{N}$ is a hyperparameter. The estimated output is given by \hat{y}_t , which is computed by the decoder C_z .

We reiterate that training the neural SSM involves computing weights of f_{enc} and the linear state-space layers A_z , C_z . This is performed by minimizing a multi-step

prediction loss as follows. For an input $Y_{t-H:t}$, the latent encoding z_t is computed using (2a), after which, for a prediction horizon of H_p , we recursively compute $z_{t+1}, \dots, z_{t+H_p}$ using (2b). Subsequently, we can compute

$$\hat{Y}_{t:t+H_p-1} = \{\hat{y}_t, \hat{y}_{t+1}, \dots, \hat{y}_{t+H_p-1}\}$$

using (2c). Since training is done offline, the output y_t is available, with which one can construct $Y_{t:t+H_p-1}$. Then, the multi-step predictions can be evaluated via a mean-squared-error (MSE) loss function

$$\ell_{\text{SSM}} = \frac{1}{H_p} \|Y_{t:t+H_p-1} - \hat{Y}_{t:t+H_p-1}\|_2^2 \quad (3)$$

that can be minimized using batched data and stochastic gradient descent methods. Often, for better numerical performance, A_z may be regularized with \mathcal{L}_1 (for sparsity) and/or \mathcal{L}_2 (for stability) norms.

Remark 1. Figure 1 illustrates multi-step prediction using a neural SSM with the shaded (fictitious) prediction layer. We emphasize that predicted inputs are not fed back into f_{enc} , but predicted forward using only linear decoders A_z and C_z without additional inputs. Thus, we can consider A_z and C_z as a type of recurrent neural network, whose initial state is provided by lifting via f_{enc} . ◀




Fig. 1. Neural architecture of deep encoder state-space model with an illustration of how to predict using the state-transition operator. The shaded layers are fictitious, and presented for illustrative purposes.

Such neural SSMs have been studied extensively, and various sources have reported their excellent predictive capabilities. However, training the neural SSM using only a limited quantity of target data can lead to poor predictive performance. Our objective in this paper is to use model-agnostic meta-learning (MAML) to learn a generalized neural SSM using data obtained from similar systems to the target system. We can then quickly adapt the generalized neural SSM to the target system despite the scarcity of target data.

To this end, we assume that we have access to a source dataset that consists of state trajectories generated by systems of the form (1) parameterized by different θ vectors, each of which are assumed to be realizations of a distribution Θ that also admits θ^* . Then the source dataset is represented by

$$\mathcal{D}_{\text{source}} = \{Y^k\}_{k=1}^{N_s} \triangleq \{Y(\theta_k, T_k)\}_{k=1}^{N_s},$$

where $\theta_k \sim \Theta$ for each $k = 1, \dots, N_s$.

Our **objective** in this paper is to utilize the source dataset to meta-learn a neural SSM representation offline, and adapt this neural SSM to yield an accurate predictive model for the target system with only a few online iterations and limited target system data. This is common in practical applications where the amount of data received

from a newly deployed target system is typically far less than previously archived customer/user data on similar (source) systems.

3. META-LEARNED NEURAL SSM

3.1 The MAML Algorithm

MAML (Finn et al., 2017) is one of the most well-known and widely used meta-learning (i.e., learning to learn) algorithms. The goal of MAML is to learn a reusable set of model parameters that can be quickly fine-tuned at inference time based on a small amount of adaptation data. MAML achieves this using a nested training scheme where an inner loop fine-tunes a common set of initial model parameters based on a small amount of task-specific adaptation data, and an outer loop that updates the initial set of model parameters across a mini-batch of different tasks. This procedure promotes learning of model parameters that can quickly adapt to new tasks.

An overview of our approach applying MAML for trajectory prediction is shown in Figure 2. Perhaps the most important aspect of applying meta-learning algorithms such as MAML is data partitioning. Formally, we are given a dataset of source systems $\mathcal{D}_{\text{source}}$, comprising trajectory data Y^i for $i = 1, \dots, N_s$ generated by the system (1) with parameters θ_i . We then partition the data for task Y^i into context set \mathcal{C}_{Y^i} for inner-loop updates, and target set \mathcal{T}_{Y^i} to evaluate the loss function for the model parameters adapted in the inner-loop. If we define ω as the set of neural network parameters to represent the mapping from (2), the inner-loop MAML update is

$$\omega_m^i = \omega_{m-1}^i - \beta_{\text{in}} \nabla_{\omega_{m-1}^i} \ell_{\text{SSM}}(\mathcal{C}_{Y^i}; \omega_{m-1}^i) \quad (4)$$

where m is the number of inner-loop updates, β_{in} the inner-loop learning rate, and $\ell_{\text{SSM}}(\mathcal{C}_{Y^i}; \omega_{m-1}^i)$ is the loss function from (3) evaluated on the context set with the neural weights computed after $m - 1$ inner-loop updates. As shown in Figure 2, the inner-loop updates are performed individually using the context set for each trajectory in the batch, while the target sets for all trajectories in the batch are used in the outer-loop. The outer-loop optimization step is typically written as

$$\omega = \omega - \beta_{\text{out}} \nabla_{\omega} \sum_{b=1}^B \ell_{\text{SSM}}(\mathcal{T}_{Y^b}; \omega_m^b) \quad (5)$$

where B is the number of tasks (i.e., trajectories) in a training mini-batch, β_{out} is the outer-loop learning rate, and $\ell(\mathcal{T}_{Y^b}; \omega_m^b)$ is the loss function on the target set \mathcal{T}_{Y^b} after m inner-loop iterations (4) have been completed. By updating the model parameters in the outer-loop as in (5) across B tasks, we promote a parameter set ω that can be quickly adapted at inference time. As shown in the bottom of Figure 2 at inference-time we only perform m inner-loop updates (4) for a query task Y^* and evaluate model performance using updated parameters ω_m^* .

Remark 2. When partitioning trajectory data Y^i into context set \mathcal{C}_{Y^i} used for inner-loop fine-tuning and target set \mathcal{T}_{Y^i} used to evaluate the fine-tuning, one may use different approaches for training and inference. At inference time, we will typically use the first several points of the observed trajectory as the context set, and any subsequent points

Algorithm 1 Meta-training SSM with MAML/ANIL

Require: $\omega \leftarrow$ weights of neural SSM
Require: $\mathcal{D}_{\text{source}} \leftarrow$ source task dataset
Require: $\omega_{\text{in}} \subseteq \omega$ \triangleright ANIL: $\omega_{\text{in}} \subset \omega$, MAML: $\omega_{\text{in}} = \omega$
Require: $\beta_{\text{in}}, \beta_{\text{out}}, M$ \triangleright learning rates and # iters

- 1: Randomly initialize ω
- 2: **while** not done **do** \triangleright outer-loop
- 3: Sample batch $\{Y^b\}_{k=1}^B$ from $\mathcal{D}_{\text{source}}$
- 4: **for** $b = 1$ to B **do** \triangleright inner-loop
- 5: Partition Y^b into \mathcal{C}_{Y^b} and \mathcal{T}_{Y^b}
- 6: $\omega_0^b \leftarrow \omega$ \triangleright copy current weights
- 7: **for** $m = 1$ to M **do** \triangleright adaptation steps
- 8: **for** $\omega_l \in \omega_{\text{in}}$ **do** \triangleright step through layers
- 9: $(\omega_l)_m^b \leftarrow$ update using (6)
- 10: **end for**
- 11: **end for**
- 12: **end for**
- 13: $\omega \leftarrow$ update using (5)
- 14: **end while**
- 15: Return $\omega_{\infty} \leftarrow$ final trained weights

Algorithm 2 SSM inference with MAML/ANIL

Require: $\omega_{\infty} \leftarrow$ weights of meta-trained neural SSM
Require: $\mathcal{D}_{\text{query}} \leftarrow$ query task dataset
Require: $\omega_{\text{in}} \subseteq \omega$ \triangleright ANIL: $\omega_{\text{in}} \subset \omega$, MAML: $\omega_{\text{in}} = \omega$
Require: β_{in}, M \triangleright learning rates and # iters

- 1: $\mathcal{C}_{Y^*} \leftarrow$ all available data in $\mathcal{D}_{\text{query}}$
- 2: $\omega_0^* \leftarrow \omega_{\infty}$ \triangleright use meta-trained weights
- 3: **for** $m = 1$ to M **do** \triangleright online adaptation
- 4: **for** $\omega_l \in \omega_{\text{in}}$ **do**
- 5: $(\omega_l)_m^* \leftarrow$ update using (6)
- 6: **end for**
- 7: **end for**
- 8: Use ω_M^* for SSM predictions

as the target set to simulate the real-world scenario, where we first fine-tune our meta learned model and then use it. At training time, we randomly sample $|\mathcal{C}_{Y^i}|$ consecutive points from anywhere in the trajectory as the context set, and $|\mathcal{T}_{Y^i}|$ consecutive points from anywhere in trajectory (not necessarily after the context set) as the target set. We take this approach during training to ensure that the learned model does not always expect the context set to depend on initial conditions, while target set is in steady-state regions for example. In our experiments we use $|\mathcal{C}_{Y^i}| = |\mathcal{T}_{Y^i}| = 12$. \blacktriangleleft

3.2 Partial network meta-learning with ANIL

The nested training scheme in MAML can lead to a very difficult optimization problem (Antoniou et al., 2019), and updating all model parameters in the inner-loop may be unnecessary (Raghu et al., 2019). For this reason, the almost no inner-loop (ANIL) algorithm was proposed in Raghu et al. (2019), which updates only the parameters in the last layer of the network in the inner-loop (4), while parameters for all layers are updated in the outer-loop (5). The intuition being that the outer-loop update promotes extraction of low-level features that are reusable across tasks, and the inner-loop promotes rapid learning of a final task specific layer.

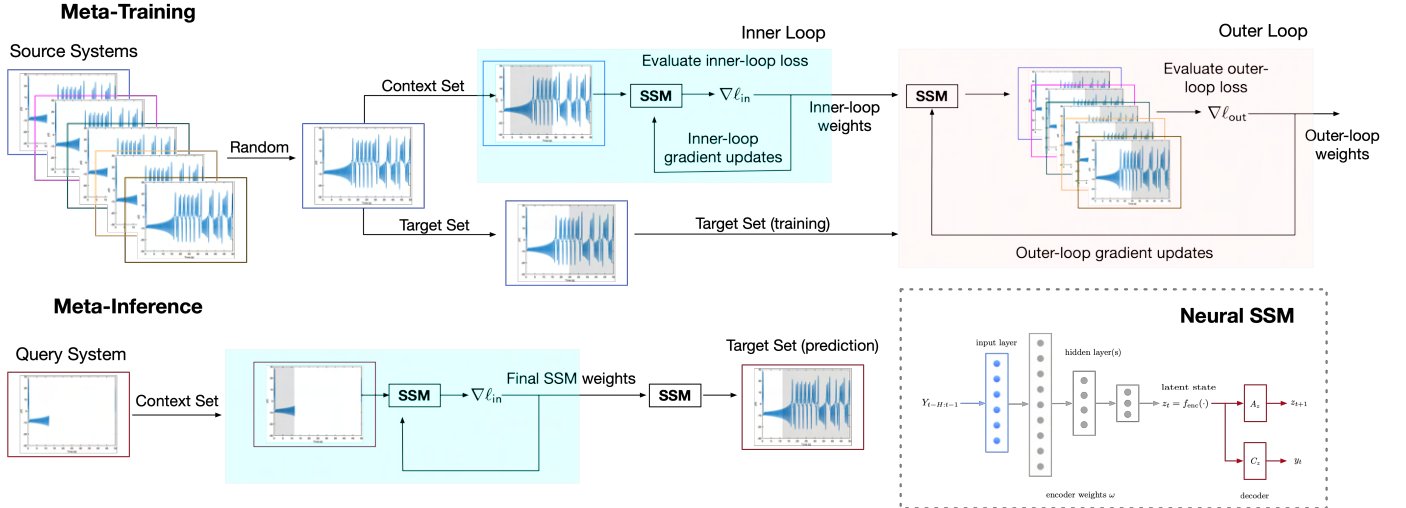


Fig. 2. Schematic diagram of the overall meta-learning process.

In the few-shot image classification problem commonly studied in meta-learning research, the last layer is referred to as a “classifier” and the earlier layers that are not updated in the inner-loop as the “feature extractor.” However, for the neural SSM architecture studied in this work, it is unclear whether the encoder weights (2a) or the state-space model parameters (2b) and (2c) should be common across tasks or fine-tuned in the inner-loop, so we explore both possibilities. Formally, if $\omega = (\omega_1, \dots, \omega_L)$, where ω_l are the network parameters for layer l and L is the total number of layers in the network, we then define $\omega_{\text{in}} \subseteq \omega$ as the subset of parameters updated in the inner-loop, and the inner-loop update from (4) becomes

$$(\omega_l)_m^i = \begin{cases} \omega_l, & \omega_l \notin \omega_{\text{in}} \\ (\omega_l)_{m-1}^i - \beta_{\text{in}} \nabla \ell_{\text{SSM}}(\mathcal{C}_{Y^i}; (\omega_l)_{m-1}^i), & \omega_l \in \omega_{\text{in}} \end{cases} \quad (6)$$

The outer-loop of ANIL remains unchanged from MAML, and we note that when $\omega_{\text{in}} = \omega$ ANIL and MAML are identical. We summarize the meta-training regime using source tasks via ANIL/MAML in Algorithm 1, and explain how to update the weights for the query system in Algorithm 2.

4. SIMULATION RESULTS

In this section, we validate our meta-learned neural SSM on a family of parameter-uncertain unforced van der Pol oscillators, where each oscillator is given by

$$\dot{x}_1 = x_2, \quad \dot{x}_2 = \theta x_2(1 - x_1^2) - x_1, \quad y = x. \quad (7)$$

4.1 Data collection and implementation details

The source and target systems are generated by sampling θ uniformly from the (unknown) range $\Theta = \mathcal{U}([0.5, 2])$, which induces a broad range of dynamics since θ is effectively a damping parameter. In particular, we simulate until $T = 20$ and collect output data for $N_s = 200$ source systems with unique $\theta \sim \Theta$ values for each system, with a sampling period of 0.01 s. The target system $\theta^* = 1.572$, which is also unknown. For each of the source systems, the initial state is randomly sampled from $[-1, 1]^2$, while the target system initial state is fixed at $[1, -0.5]^T$; the

final time is sampled randomly from $T \sim \mathcal{U}([10, 40])$ s, with the sampling period kept constant for all systems. The past window length $H = 10$ and predictive window length $H_p = 5$.

The deep encoder network (see Fig. 1) we use as the neural SSM has an encoder consisting of an input layer that takes x , and passes it through 5 hidden layers, each of which has 128 neurons, to an output layer that generates a latent variable of dimension $n_z = 128$. This latent variable is updated using a 128×128 linear layer with no bias, and the output is computed using a 128×2 linear layer with no bias. The entire network uses rectified linear units (ReLU) for activation, and the weights are initialized using Xavier initialization. The deep neural network is implemented entirely in PyTorch (Paszke et al., 2019). For MAML and ANIL implementation, we adapt the open-source PyTorch-based `learn2learn` toolbox (Arnold et al., 2020). The batch size $B = 32$, with the number of inner-loop iterations fixed at $m = 10$, and the total number of meta-training epochs is set to 10^4 . The step-sizes for MAML are $\beta_{\text{in}} = 0.01$ and $\beta_{\text{out}} = 0.001$. Online, the MAML is allowed 40 adaptation steps.

4.2 Meta-learning vs. supervised and transfer-learning

To judge the effectiveness of the proposed MAML-SSM, we compare against a few baselines, all of which have the same architecture as MAML-SSM but are trained differently, using a single training loop rather than an inner-outer bilevel loop as in MAML. These include:

- **SSM**: which is trained by supervised learning using only the target system data: this is the classical one-training-loop approach for learning neural SSMs. The SSM is allowed $10 \times$ more steps to train
- **All-NoAdapt-SSM**: which is training by supervised learning using all the source system and target system data, without allowing online adaptation steps.
- **Xfer-SSM**, which is trained by transfer learning; that is, supervised learning on all the source system data, and allowing a few online adaptation steps exploiting the target system data. The same number of adaptation steps are allowed for MAML-SSM and Xfer-SSM.




Fig. 3. Comparison of MAML-SSM with baselines. (*upper, middle*) State x_1 and x_2 of the oscillator. (*lower*) Comparison of sum-squared-error (SSE).

To provide some intuition, the first baseline SSM is selected to understand whether MAML-SSM can enable learning from similar systems, or whether using data from other systems negatively affects predictive quality. The second baseline All-NoAdapt-SSM is selected to study whether supervised learning would be enough to learn a good predictive model if it was provided all the data available, both from source and target systems. The third baseline Xfer-SSM is to test whether meta-learning can outperform transfer-learning in few-shot, data-poor situations.

The performance of MAML-SSM and the baselines is illustrated in Fig. 3. At online inference, the context set is generated by using data from the first 400 time steps, and the target set to be predicted by the neural SSM is the next 3000 time steps. The top and middle subplots show the evolution of the states of the system, with the dashed line denoting the true query system states. The lowest subplot shows the evolution of the sum-squared-error¹ between the true state and the predicted state over the 3000 time steps. It is immediately clear that MAML-SSM incurs the smallest prediction SSE and outperforms the baseline, with the transfer learning Xfer-SSM performing second best with double the SSE, closely followed by SSM, and with greedily using all the data as in All-NoAdapt-SSM exhibiting significantly worse performance uniformly across time. From the x_1 subplot, we can reason that SSM most likely overfits the context data, and therefore, it exhibits excellent predictive accuracy in within-training-set points of time, with marked deterioration at out-of-set points. Conversely, All-NoAdapt-SSM severely underfits the data, which is plausible, as it inherently tries to find a set of neural SSM weights to fit the ‘average’ dynamics induced by the source and query systems, rather than adapting to the query system. As expected, meta- and transfer-learning exhibit good performance, but our proposed MAML-SSM outperforms Xfer-SSM. The paucity of query system data and the few number of adaptation

¹ The sum-squared-error (SSE) between two signals $\chi, \chi' \in \mathbb{R}^{d \times T}$ is given by $\sum_{t=1}^T \|\chi_t - \chi'_t\|_2^2$.

iterations require the trained learner (before adaptation) to have a set of neural weights that can rapidly adapt to the query system: since this is explicitly what MAML is trained for, the MAML-SSM has the ability to rapidly learn a good predictive model for the given example. In contrast, the transfer learning approach (before adaptation) is not explicitly trained keeping the adaptation in mind, therefore the set of initial weights is likely tailored to generating good predictions for the family of source systems, and the few-shot nature of the adaptation is not enough to enable Xfer-SSM to learn a model as accurate as MAML-SSM; this is especially notable from the x_1 plot, where transfer-learning completely fails to capture the positive half of the oscillatory dynamics.

4.3 Full vs. partial network meta-learning

As described in Section 3.2, ANIL can be used to partially meta-learn neural SSMs by splitting the neural SSM into a base-learner (layers with parameters $\omega_l \notin \omega_{in}$) and a meta-learner (layers with parameters $\omega_l \in \omega_{in}$), wherein only the meta-learner is adapted online; the base-learner is fixed. An interesting question for neural SSMs based on deep encoder networks is whether the base-learner should be the encoder layers or the state and output linear layers, as they have clearly different roles to play in the neural SSM. To answer this question, we propose two implementations: ANIL-SSM, where the base-learner is A_z, C_z , with f_{enc} set to be the meta-learner; and, ANIL-SSM-R which has f_{enc} as the base-learner and meta-learns A_z, C_z .

Table 1. Comparison of SSE. Full (MAML) vs. partial (ANIL) meta-learning of neural SSM.

EXPERIMENT	MEDIAN SSE		
Size/Steps	MAML-SSM	ANIL-SSM	ANIL-SSM-R
200/10	1.4×10^3	4.6×10^3	6.8×10^3
200/40	1.6×10^3	2.0×10^3	6.9×10^3
200/100	1.1×10^3	2.1×10^3	1.0×10^4
500/10	1.3×10^3	1.7×10^3	6.3×10^3
500/40	2.3×10^3	1.9×10^3	4.1×10^3
500/100	1.6×10^2	6.0×10^1	1.5×10^3
1000/10	9.1×10^2	8.5×10^2	5.5×10^3
1000/40	2.2×10^1	2.2×10^1	2.6×10^3
1000/100	3.4×10^0	3.1×10^0	5.6×10^2

Table 1 compiles the median SSE values obtained over 100 query runs with MAML-SSM, ANIL-SSM, and ANIL-SSM-R from unique and randomly selected initial conditions. The leftmost column of the table indicates the number of context data points used for inference, and the number of online adaptation steps. It is easy to deduce that for small-size context sets, MAML outperforms both ANIL variants, whereas when the context set is larger, and more adaptation steps are allowed, ANIL-SSM performs slightly better than MAML-SSM. The most marked improvement is in the ‘500/100’ case, indicating that ANIL has the potential to surpass MAML, albeit with more context data and adaptation steps. We see that ANIL-SSM-R is uniformly worse, often by orders of magnitude, than its competitors, indicating that rapidly adapting the encoding/lifting transformation with A_z and C_z being reused online without alteration allows the deep encoder network to be more expressive over a wider range of dynamics than reusing the encoding/lifting transformation. This can be

explained by realizing that f_{enc} is a deep network with nonlinear activation functions rather than a single linear layer, and therefore it can induce a much richer family of transformations than A_z and C_z , allowing for greater expressivity in modeling nonlinear dynamics.

5. CONCLUSIONS

In this paper, we proposed a methodology for designing neural state-space models using limited data from the query system to be modeled, relying more on data from similar systems. We provide a deep learning-based framework called meta-learning for tractably learning from similar systems offline, and adapting to the query system online in a few-shot manner; we also demonstrate the potential of the proposed approach using a numerical example. Interestingly, we show that meta-learning can outperform transfer learning in data-poor settings when the number of online adaptation steps is limited. We also describe how to meta-learn a subset of the layers of a neural SSM, and report that meta-learning lifting transformations is more advantageous than meta-learning the state-transition operator itself. In future work, we will investigate alterations to the MAML algorithm for faster training convergence, and apply these techniques to the modeling of energy systems.

REFERENCES

- Antoniou, A., Edwards, H., and Storkey, A. (2019). How to train your MAML. In *Proc. of the Int. Conf. on Learning Representations (ICLR)*.
- Arcari, E., Carron, A., and Zeilinger, M.N. (2020). Meta learning MPC using finite-dimensional Gaussian process approximations. *arXiv preprint arXiv:2008.05984*.
- Arnold, S.M.R., Mahajan, P., Datta, D., Bunner, I., and Zarkias, K.S. (2020). learn2learn: A library for Meta-Learning research. URL <http://arxiv.org/abs/2008.12284>.
- Bao, Y., Velni, J.M., Basina, A., and Shahbakhhti, M. (2020). Identification of state-space linear parameter-varying models using artificial neural networks. *IFAC-PapersOnLine*, 53(2), 5286–5291.
- Beintema, G., Toth, R., and Schoukens, M. (2021). Non-linear state-space identification using deep encoder networks. In *Learning for Dynamics and Control*, 241–250. PMLR.
- Bertalan, T., Dietrich, F., Mezić, I., and Kevrekidis, I.G. (2019). On learning Hamiltonian systems from data. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 29(12), 121107.
- Chakrabarty, A. (2022). Optimizing closed-loop performance with data from similar systems: A Bayesian meta-learning approach. In *Proc. IEEE Conf. Dec. Control (CDC)*, To appear. IEEE.
- Finn, C., Abbeel, P., and Levine, S. (2017). Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, 1126–1135. PMLR.
- Forgione, M., Mejari, M., and Piga, D. (2022). Learning neural state-space models: do we need a state estimator? *arXiv preprint arXiv:2206.12928*.
- Forgione, M. and Piga, D. (2020). Model structures and fitting criteria for system identification with neural networks. In *2020 IEEE 14th International Conference on Application of Information and Communication Technologies (AICT)*, 1–6. IEEE.
- Forgione, 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.
- Iacob, L.C., Beintema, G.I., Schoukens, M., and Tóth, R. (2021). Deep identification of nonlinear systems in Koopman form. In *2021 60th IEEE Conference on Decision and Control (CDC)*, 2288–2293. IEEE.
- Koopman, B.O. and van Neumann, J. (1932). Dynamical systems of continuous spectra. *Proceedings of the National Academy of Sciences*, 18(3), 255–263.
- Legaard, C.M., Schranz, T., Schweiger, G., Drgoña, J., Falay, B., Gomes, C., Iosifidis, A., Abkar, M., and Larsen, P.G. (2021). Constructing neural network-based models for simulating dynamical systems. *arXiv preprint arXiv:2111.01495*.
- Lusch, B., Kutz, J.N., and Brunton, S.L. (2018). Deep learning for universal linear embeddings of nonlinear dynamics. *Nature communications*, 9(1), 1–10.
- Masti, D. and Bemporad, A. (2021). Learning nonlinear state-space models using autoencoders. *Automatica*, 129, 109666.
- Muthirayan, D. and Khargonekar, P.P. (2020). Meta-learning guarantees for online receding horizon learning control. *arXiv preprint arXiv:2010.11327*.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. (2019). PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, 8024–8035. Curran Associates, Inc.
- Raghu, A., Raghu, M., Bengio, S., and Vinyals, O. (2019). Rapid learning or feature reuse? Towards understanding the effectiveness of MAML. In *Proc. of the Int. Conf. on Learning Representations (ICLR)*.
- Richards, S.M., Azizan, N., Slotine, J.E., and Pavone, M. (2021). Adaptive-control-oriented meta-learning for nonlinear systems. *CoRR*, abs/2103.04490. URL <https://arxiv.org/abs/2103.04490>.
- Skomski, E., Drgoña, J., and Tuor, A. (2021). Automating discovery of physics-informed neural state space models via learning and evolution. In *Learning for Dynamics and Control*, 980–991. PMLR.
- Xin, L., Ye, L., Chiu, G., and Sundaram, S. (2022). Identifying the dynamics of a system by leveraging data from similar systems. In *2022 American Control Conference (ACC)*, 818–824. doi: 10.23919/ACC53348.2022.9867413.
- Zamarreño, J.M. and Vega, P. (1998). State-space neural network. Properties and application. *Neural networks*, 11(6), 1099–1112.
- Zhan, S., Wichern, G., Laughman, C., Chong, A., and Chakrabarty, A. (2022). Calibrating building simulation models using multi-source datasets and meta-learned Bayesian optimization. *Energy and Buildings*, 270, 112278.