

STRUCTURED STATE SPACE SEQUENCE MODELS TRANSFER LEARNING BETWEEN CIRCULAR RESTRICTED THREE-BODY PROBLEM ORBIT FAMILIES

Hunter Quebedeaux*, and Tarek A. Elgohary†

Machine learning has long been plagued with the reliance on large datasets and computational centers. The deployment of machine learning algorithms in cislunar space remains difficult because of these constraints. This work explores a structured state space sequence model, Mamba, and its ability to perform transfer learning between circular restricted three-body problem orbit families. The results show that Mamba learns and performs time series forecasting at a much higher accuracy compared to LSTMs. Learning the same datasets, Mamba can approximate the time series with less learnable parameters for a marginal increase in computational complexity at forward evaluation time.

INTRODUCTION

Dynamical system propagation is the foundation of many applications in astrodynamics. Uncertainty quantification, orbit determination, and guidance, navigation, and control can require millions of propagations.¹ In the domain of cislunar space, performing these propagations can be a challenge in either onboard processing or in the transferring of numerical integration results to a satellite in orbit. With the recent proliferation of machine learning, propagation tasks have been accelerated, albeit with limited accuracy than what is required in sensitive dynamical systems like the circular restricted three-body problem.²⁻⁵ While machine learning accuracy has improved in recent years with the advent of specialized computer hardware, large computation clusters, and large amounts of training data, onboard machine learning autonomy remains an incredible challenge due to the limitations in current satellite hardware, which has a maximum of 50W of available power, and DRAM of a maximum of 1GB.⁶ Coupling the challenge of current space-grade hardware with increased complexity of machine learning algorithms will yield the inability to deploy such algorithms in fully autonomous regimes. Thus, increased interest in both machine learning and cislunar space has revealed a gap in the literature for onboard, intelligent, and autonomous algorithms that require smaller model sizes and datasets. Two avenues to explore low model sizes and low dataset sizes for optimal learning can either be in foundational model research, or alternative learning processes.

In the solution of dynamical propagation, recurrent, date-driven network structures are useful in their ability to process and extract sequential data across multiple time steps. The recurrent structure lends itself well to sequential data approximation of dynamical system trajectories. Long short term memory (LSTMs) are a type of recurrent network structure used in time dependent learning; this

*Graduate Research Assistant, Mechanical and Aerospace Engineering, University of Central Florida, 12760 Pegasus Drive, Orlando, FL.

†Associate Professor, Mechanical and Aerospace Engineering, University of Central Florida, 12760 Pegasus Drive, Orlando, FL.

network structure is based on the use of an internal hidden state that acts as the short-term memory, updated based on the the input sequence and the hidden state's previous and current state.




Figure 1: LSTM Model Definition

For an input sequence $x = (x_1, x_2, \dots, x_T)$, the LSTM layer generates the hidden vector sequence $h = (h_1, h_2, \dots, h_T)$ and output $y = (y_1, y_2, \dots, y_T)$ of the same length by iterating through discrete time points measured at the same frequency, Eq. (1).

$$\begin{aligned} h_t &= \tanh(W_{xh}x_t + W_{hh}h_{t-1} + b_h) \\ y_t &= W_{hy}h_t + b_y \end{aligned} \quad (1)$$

While LSTMs are the conventional choice for time series data, the LSTM is known to degrade for time series beyond certain lengths, making long term forecasting challenging.⁷ Changes to the conventional model choice is one area where improvements to time series learning could be made.

Transfer learning is a learning method used in machine learning in which a neural network is trained on a task with some source dataset, and then this network is retrained on a completely new target dataset related to the original problem, but outside of the scope of the source dataset. By leveraging the use of a network trained on a related problem, retraining on some target domain can greatly reduce the the amount of data or training computation to achieve the same prediction accuracy as the initial trained model.⁸ Transfer learning is an incredibly useful learning technique that has potential benefits in engineering tasks where design requirements may vary slightly per mission, and knowledge of related design requirements can still be useful; this exact design philosophy has been demonstrated to be effective in resource reduction in the entry, descent, and landing problem.⁹ There are several types of transfer learning tasks, such as model finetuning or domain adaptation. In the case of time series forecasting there are several areas in where transfer learning can be beneficial such as when source and target tasks share some similarities, like common dynamics. However, time series forecasting introduces additional challenges because of the temporal dependencies networks make during training or sampling rates of training data. While unproven, we conjecture that the use of S4 models utilize the state space model internal to the network, creating a more robust temporal dependency mapping compared to standard recurrent machine learning models.

Building on past foundational model research for time series forecasting, this work explores Mamba, a state space based machine learning model in its ability to approximate exact dynamical

systems compared to conventional data-driven recurrent networks. In addition to this foundational model research, we investigate alternative learning processes; namely the ability for Mamba to transfer learn in time series forecasting applications. With these two machine learning choices, we demonstrate that the use of state-space based machine learning models improves numerical accuracy when generating a prediction, as well as improved transfer learning abilities, ultimately reducing the total amount of data when learning between different initial conditions for systems governed by the same dynamical systems. In the following sections, we introduce the Mamba network and specific transfer learning methods examined. The experiments conducted showcase the ability for Mamba to transfer learn compared to contemporary LSTM networks through a model finetuning approach.

METHODOLOGY

Structured State Space Sequence Models & Mamba

An interesting class of machine learning models that has been recently introduced is the structured state space sequence (S4) models, utilizing a linear state space model internal to the network. S4 models are a class of sequence models aimed to address the issue of scaling during long sequences while capturing long dependencies for vision, language, and audio. S4 models utilize a state space model to create a black-box representation in a sequence model, where the matrices $\mathbf{A}$, $\mathbf{B}$, $\mathbf{C}$, and $\mathbf{D}$ are learned via gradient descent.

$$\begin{aligned}\dot{x} &= \mathbf{A}x + \mathbf{B}u \\ y &= \mathbf{C}x + \mathbf{D}u\end{aligned}\tag{2}$$

One model of this class is the Mamba model. Mamba was originally designed for natural language processing tasks, but has recently been demonstrated to have increased performance in operations involving the training of neural operators, multivariate time series forecasting, and time series forecasting of exact dynamical systems in limited data training scenarios.^{10–12} Mamba augments the model with a context-dependent selective scan mechanisms as well as the incorporation of a gated multilayer perceptron (MLP). This class of selective state space model is pictured in Figure 2.




Figure 2: Single Layer Mamba Block

Among the state-of-the-art sequential models long short-term memory (LSTMs), and transformers are the most popular. Traditional recurrent models like LSTMs have been proven to be robust in time series tasks and have been applied to solve a wide variety of tasks adjacent to dynamical systems such as approximating the Sun-Earth three-body problem.¹³ The transformer primarily excels within the natural language processing (NLP) domain and has been implemented into applications

like time series forecasting and classification, but their effectiveness in this domain is questionable, with some work suggesting that transformers are not the best choice for simpler time series forecasting problems.¹⁴ In their original design, Mamba networks aimed to replace transformers with a less memory intensive, less computationally complex model.¹⁰

Transfer Learning

Transfer learning can imply a variety of learning techniques, but in this work we will examine inductive transfer learning, as defined by the taxonomy of Pan and Yang.¹⁵ Inductive transfer learning refers to learning between target and source domain on labeled data, and learning between different tasks. Within this learning technique, there are a variety of subcategories of learning that can affect the training efficiency and effectiveness. In this work, we examine inductive transfer learning in the full-model finetuning sense. In full-model finetuning, machine learning models are trained on a specific domain, and that fully trained network is used to learn on some other domain, Figure 3. This is in contrast to feature-extraction transfer learning, which refers to training a network on a specific domain, freezing the upstream layers or weights of the network, and retraining on the downstream unfrozen layers or weights, pictured in Figure 4.




Figure 3: LSTM Finetuning Learning (Learnable in Red)




Figure 4: LSTM Feature Extraction Learning (Learnable in Red)

TRANSFER LEARNING BETWEEN ORBIT FAMILIES

Transfer learning requires the use of a source domain that is initially learned on for some number of training epochs, and a target domain which typically learned for a lesser amount of epochs or data. In this study, we examined that former; transferring knowledge of a source domain to a target domain and learning in the lesser epoch count. The chosen dynamical system to transfer between is the circular restricted three body problem (CR3BP), because of the unique challenges the cislunar domain possess when placing satellite and spacecraft due to the highly nonlinear and chaotic dynamics. Consider the nondimensional form of the CR3BP with the defining nondimensional constraints with the states normalized by $DU = 389703$ km and $TU = 382981$ seconds:

$$\begin{aligned}
 \ddot{x} - 2\dot{y} &= x - (1 - \mu)\frac{x + \mu}{r_1^3} - \mu\frac{x - 1 + \mu}{r_2^3}, & \mu &= \frac{m_2}{m_1 + m_2} \\
 \ddot{y} + 2\dot{x} &= y - (1 - \mu)\frac{y}{r_1^3} - \mu\frac{y}{r_2^3}, & \rho_1 &= \sqrt{(x - x_1)^2 + y^2 + z^2} \\
 \ddot{z} &= -\left(\frac{1 - \mu}{r_1^3} + \frac{\mu}{r_2^3}\right)z, & \rho_2 &= \sqrt{(x - x_2)^2 + y^2 + z^2} \\
 & & x_1 &= -\mu \\
 & & x_2 &= 1 - \mu
 \end{aligned} \tag{3}$$

For the following problems, all data was generated using initial conditions in the NASA JPL three-body problem database.¹⁶ Each trajectory was generated using an adaptive Runge-Kutta 8(7) integrator with a relative tolerance of 10^{-9} and an absolute tolerance of 10^{-12} , separated by a time delta of $\Delta t = 0.001TU$, and is considered the true trajectory. The Mamba network was chosen to contain a hidden state size of 16, with the LSTM chosen to have a hidden size of 30. The source domain was generated using this numerical integrator, and trained on a single orbit period for 5 training epochs. The Mamba and LSTM models are then transferred with all parameters still trainable, for only 2 training epochs, again training in the new target domain generated using the same numerical integrator and trained on for 1 orbit period to then predict the next 4 periods into the future.

Two error metrics were considered in this problem. Firstly, we considered the percent error of the predicted trajectory to the true trajectory

$$\% \text{ Error} = \frac{\mathbf{y}_{truth}(t) - \mathbf{y}_{pred}(t)}{\mathbf{y}_{truth}(t)} \tag{4}$$

then, in the CR3BP, the only known conserved quantity is Jacobi's integral, calculated in the non-dimensional form as

$$J_c = (x^2 + y^2) + 2\left(\frac{1 - \mu}{\rho_1} + \frac{\mu}{\rho_2}\right) - v^2 \tag{5}$$

will act as another metric for how well a network will be able to learn a trajectory.

Long Period to Short Period

The first tested transfer learning problem is a knowledge transfer of a trajectory in the long period family to the short period family, both centered around L4. The source domain orbit selected from the NASA JPL three-body problem database was ID number 751, with the target orbit of ID number 755. Figures 5 and 6 plot the orbit prediction of the source and target domains generated by a Mamba and LSTM network. Figures 7 and 8 plot the Jacobi constant over the predicted trajectories, and Figures 9 - 12 plot the respective percent errors for the source and target domains for the Mamba and LSTM predictions.




Figure 5: Long Period Source Domain Phase Space




Figure 6: Short Period Target Domain Phase Space




Figure 7: Long Period Source Domain Jacobi Constant




Figure 8: Short Period Target Domain Jacobi Constant




Figure 9: Long Period Source Domain Percent Error for Mamba Prediction




Figure 10: Short Period Target Domain Percent Error for Mamba Prediction




Figure 11: Long Period Source Domain Percent Error for LSTM Prediction




Figure 12: Short Period Target Domain Percent Error for LSTM Prediction

Long Period to Dragonfly South

The next tested transfer learning problem is a knowledge transfer of a trajectory in the long period family to the dragonfly south family. The source domain orbit selected from the NASA JPL three-body problem database was a long period orbit centered around L4 with ID number 751, and the target orbit was selected as a dragonfly south with ID number 71. Figures 13 and 14 plot the orbit prediction of the source and target domains generated by a Mamba and LSTM network. Figures 15 and 16 plot the Jacobi constant over the predicted trajectories, and Figures 17 - 20 plot the respective percent errors for the source and target domains for the Mamba and LSTM predictions.




Figure 13: Long Period Source Domain Phase Space




Figure 14: Dragonfly South Target Domain Phase Space




Figure 15: Long Period Source Domain Jacobi Constant




Figure 16: Dragonfly South Target Domain Jacobi Constant




Figure 17: Long Period Source Domain Percent Error for Mamba Prediction




Figure 18: Dragonfly South Target Domain Percent Error for Mamba Prediction




Figure 19: Long Period Source Domain Percent Error for LSTM Prediction




Figure 20: Dragonfly South Target Domain Percent Error for LSTM Prediction

Butterfly North to Short Period

Finally, the last tested transfer learning problem is a knowledge transfer of a trajectory in the butterfly north family to the short period family. The source domain orbit selected from the NASA JPL three-body problem database was a butterfly north orbit with ID number 404, and the target orbit was selected as a short period orbit centered around L4 with ID number 755. Figures 21 and 22 plot the orbit prediction of the source and target domains generated by a Mamba and LSTM network. Figures 23 and 24 plot the Jacobi constant over the predicted trajectories, and Figures 25 - 28 plot the respective percent errors for the source and target domains for the Mamba and LSTM predictions.




Figure 21: Butterfly North Source Domain Phase Space




Figure 22: Short Period Target Domain Phase Space




Figure 23: Butterfly North Source Domain Jacobi Constant




Figure 24: Short Period Target Domain Jacobi Constant




Figure 25: Butterfly North Source Domain Percent Error for Mamba Prediction




Figure 26: Short Period Target Domain Percent Error for Mamba Prediction




Figure 27: Butterfly North Source Domain Percent Error for LSTM Prediction




Figure 28: Short Period Target Domain Percent Error for LSTM Prediction

DISCUSSION

These results show increased accuracy of prediction when transfer learning with a Mamba network compared to an LSTM for time series forecasting. In all problems, the Mamba network was able to accurately predict the source domain and predict the trajectory in the target domain when knowledge was transferred. In comparison, the LSTM was unable to accurately transfer knowledge to all target domains in a meaningful way; all predicted orbits in the target domain do not adhere to the conservation of Jacobi's integral, and visual inspection of the predicted orbits shows no meaningful learning. The most surprising result comes from the transfer learning between the long period and dragonfly south orbit. With the specific dragonfly south orbit, the periodic solution becomes unstable around two periods of propagation, with the trajectory breaking out periodic dragonfly orbit. With only knowledge of one orbit period of the long period orbit, and knowledge of only the periodic portion of the trajectory, Mamba is able to predict the chaotic part of the trajectory, with

the LSTM fully unable to meaningfully predict any part of the dragonfly orbit. In addition, Mamba can perform these propagations with a reduction in network parameter size with only a marginal increase in time complexity, as depicted in Table 1.

Network Model	Total Parameters	Memory (bytes)	Time Complexity
Mamba	906	7248	$\mathcal{O}(n \log n)$
LSTM	4746	37968	$\mathcal{O}(n)$

Table 1: Parameter Statistics Per Network Structure

Based on these results, we conjecture the Mamba networks have better transfer learning capabilities, further supporting the claim that Mamba networks learn time series governed by exactly known dynamical system better than other recurrent network structures.¹² Incredibly, full-model finetuning as examined in this work can be susceptible to catastrophic forgetting, as demonstrated by the LSTM in the Long Period to Dragonfly South transfer learning example. Typically, the power of transfer learning can be used in feature-extraction transfer learning; this is the act of freezing early layers or weights of a pretrained neural network. Being “higher” in the chain of calculation, these layers or weights are already trained to process the inputs features of the network model; this allows all other layers or weights of the network to learn on the new dataset. The benefits to learning using the feature-extraction techniques include more stable learning, but falter in the robustness to distribution shifts in the dataset compared to full model finetuning. While unproven, these results suggest that even in the most difficult learning domains, the LSTM networks demonstrate catastrophic forgetting. On the other hand, the Mamba networks show incredible resilience to changing domains. We conjecture that this is due to the inclusion of an SSM, capturing the true dynamics of the dataset, compared to learning based on an arbitrary network structure. However, more work is needed to substantially prove this. This would involve studying the effects that feature-extraction transfer learning has on the ability for both networks to learn this domain. As well, more work is needed to fully test other cislunar orbit families such as quasiperiodic orbits to fully understand the transfer learning capabilities of Mamba.

CONCLUSION

Due to the highly nonlinear nature of the system, conventional recurrent network structures like LSTMs have difficulty learning non-planar maneuvers, in regular learning, and in finetuning transfer learning. Past research has demonstrated that Mamba, or generally, SSM based models, learn data governed by differential equations exceptionally well. Expanding this work into the domain of transfer learning, the Mamba network has demonstrated improved full-model fine tuning compared to LSTM networks in transfer learning for time series regression tasks in a variety of CR3BP orbit regimes. With improved state vector regression as proven by the conservation of Jacobi’s constant, SSM based machine learning models can approximate time series in less learnable parameters compared to a conventional LSTM solution for marginal increases in computational complexity.

REFERENCES

- [1] P. Sivasankar, H. Quebedeaux, and T. A. Elgohary, “Application of a state space based neural network model for Uncertainty Propagation in dynamical systems,” *DDAS2024 - 5th International Conference*, 2024.
- [2] H. Quebedeaux, B. G. Lewis Jr, and T. A. Elgohary, “Parallelization Techniques for Quantifying Uncertainty Using Embedded Hardware,” *2022 AAS/Astrodynamics Specialist Conference*, 2022.

- [3] H. Quebedeaux and T. A. Elgohary, "PHYSICS-INFORMED NEURAL NETWORKS FOR ORBIT PROPAGATION IN ASTRODYNAMICS PROBLEMS," *2023 AAS/Astrodynamics Specialist Conference*, 2023.
- [4] H. Quebedeaux, "Investigating Shallow Neural Networks for Orbit Propagation Deployed on Spaceflight-Like Hardware," 2023.
- [5] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *Journal of Computational physics*, Vol. 378, 2019, pp. 686–707.
- [6] B. Yost and S. Weston, "State-of-the-art small spacecraft technology," tech. rep., 2024.
- [7] C. C. Aggarwal *et al.*, *Neural networks and deep learning*, Vol. 10. Springer, 2018.
- [8] K. Weiss, T. M. Khoshgoftaar, and D. Wang, "A survey of transfer learning," *Journal of Big data*, Vol. 3, 2016, pp. 1–40.
- [9] L. Jones and T. A. Elgohary, "AN INVESTIGATION OF TRANSFER LEARNING FOR GUIDANCE AND CONTROL OF MARS ENTRY, DESCENT, AND LANDING MISSIONS," *2024 AAS GN&C Conference*, 2024.
- [10] A. Gu and T. Dao, "Mamba: Linear-Time Sequence Modeling with Selective State Spaces," Dec. 2023.
- [11] S. K. Bhethanabhotla, O. Swelam, J. Siems, D. Salinas, and F. Hutter, "Mamba4Cast: Efficient Zero-Shot Time Series Forecasting with State Space Models," *arXiv preprint arXiv:2410.09385*, 2024.
- [12] H. Quebedeaux and T. A. Elgohary, "Using Mamba for Modeling Dynamical Systems in a Limited Data Scenario," *DDAS2024 - 5th International Conference*, 2024.
- [13] A. K. Singh, V. K. Srivastava, S. Agarwal, and P. Chakraborty, "Sun-Earth system Lagrange point spacecraft orbit computation by numerical method and deep learning technique," *Acta Astronautica*, Vol. 219, 2024, pp. 408–416.
- [14] A. Zeng, M. Chen, L. Zhang, and Q. Xu, "Are transformers effective for time series forecasting?," *Proceedings of the AAAI conference on artificial intelligence*, Vol. 37, 2023, pp. 11121–11128.
- [15] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Transactions on knowledge and data engineering*, Vol. 22, No. 10, 2009, pp. 1345–1359.
- [16] NASA, "Periodic Orbits Tool," NASA Solar System Dynamics, 2024.