Pure and Physics-encoded Spatiotemporal Deep Learning for Climate-Vegetation Dynamics

Qianqiu Longyang
Kansas Geological Survey
The University of Kansas
Lawrence, KS 66047
qlongyang@ku.edu

Ruijie Zeng
School of Sustainable Engineering and the Built Environment
Arizona State University
Tempe, AZ 85281
Ruijie.Zeng.1@asu.edu

Abstract

Vegetation is a central hub of water, energy, and carbon exchanges, making its accurate spatiotemporal modeling essential for projecting climate impacts. Existing models for vegetation dynamics often suffer from physical simplifications or limited treatment of spatiotemporal interactions. We present a spatially distributed framework for daily-scale climate-vegetation dynamics, comparing: (1) a Long Short-Term Memory (LSTM) **baseline** with flattened spatial inputs; (2) a **pure** Convolutional LSTM (ConvLSTM) that captures spatial heterogeneity and temporal dependencies while implicitly representing ecohydrological states; and (3) a **physics-encoded** ConvLSTM that serves as a bias-corrector of physically simulated Leaf Area Index (LAI). ConvLSTM architectures outperform the LSTM baseline, and the physics-encoded variant underscores the promise of combining physical knowledge with data-driven models. This framework supports more reliable vegetation projections and highlights the potential of closer collaboration between AI researchers and Earth system scientists to develop trustworthy tools for climate adaptation and mitigation.

1 Introduction

Vegetation regulates atmosphere-land exchanges of water and energy [7], drives the global carbon cycle [18], and supports ecosystem adaptation to a changing climate [1]. Accurate representation of vegetation dynamics and their interactions with climate conditions is therefore essential for reliable projections of water resources, ecosystem productivity, and carbon balance. Intensifying climate variability and extremes drive shifts in vegetation patterns across seasonal to interannual timescales.

However, most process-based land surface models (LSM) and large-scale hydrological models often struggle to represent such variability. Many either (1) prescribe multi-year average monthly vegetation climatology (e.g., Noah LSM [4]), which fails to capture interannual responses, or (2) employ simplified dynamic vegetation modules (e.g., Noah-MP [17]), which suffer from structural errors and parametric uncertainties. These limitations restrict the ability to resolve fine-scale spatiotemporal responses of vegetation to climate drivers, which can propagate through the modeling chain, leading to biased projections under future climate scenarios.

Recent studies have applied deep learning to predict vegetation dynamics [6, 21, 12, 22, 5, 11, 10, 2, 15]. While promising, many of these efforts process spatial inputs in a reduced form (e.g., flattening grids for LSTM networks), which may overlook the full spatiotemporal interactions between environmental drivers and vegetation responses. Fully distributed applications of spatiotemporal deep learning in vegetation dynamics modeling remain limited and are mostly designed for short-term forecasting that incorporates previous vegetation states as inputs. By contrast, climate-vegetation

Tackling Climate Change with Machine Learning: workshop at NeurIPS 2025.

modeling driven solely by climate variables, which can support broader examinations of climate-vegetation interactions and future vegetation projections, has received limited attention.

In this paper, we develop a fully climate-driven, spatially distributed deep learning framework for daily-scale vegetation dynamics. Our main contributions are: (1) introducing two modeling strategies, a pure deep learning approach and a physics-encoded variant that applies deep learning to correct biases in physically simulated vegetation dynamics; (2) evaluating these against an LSTM network baseline in which grid-based spatial inputs are flattened at each time step. This work demonstrates the potential of advanced spatiotemporal neural networks for vegetation modeling and presents a feasible framework for combining data-driven and physics-based approaches to support climate adaptation and mitigation strategies.

2 Data Description

The Upper Colorado River Basin (UCRB, Figure 1) is selected as the testbed due to its diverse hydroclimatic gradients, complex vegetation dynamics, and critical importance for water resources amid increasing drought and climate extremes. This makes it well suited for evaluating the proposed climate-vegetation modeling framework. Although the study focuses on the UCRB, all datasets can be retrieved across the contiguous United States (CONUS), enabling scalable application of the proposed models to larger regional or continental scales.

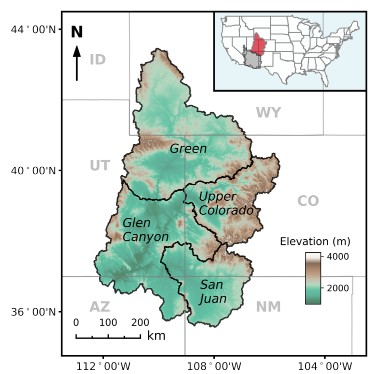

Figure 1: Location of Upper Colorado River Basin (UCRB) and subbasins.

Daily climate variables, including precipitation, near-surface air temperature, surface pressure, specific humidity, surface downward longwave and shortwave radiation, and wind speed, are obtained from NLDAS-2 [20] at 1/8° spatial resolution across CONUS. Vegetation dynamics are represented by the GLOBMAP LAI dataset [13], which offers continuous global coverage at 0.073° resolution, with half-monthly observations for 1981–2000 and 8-day intervals from 2001 onward. The satellite LAI data is resampled to match the climate forcing resolution. The dataset is split temporally for model development: 1981–2001 for training, 2002–2008 for validation, and 2009–2016 for testing.

3 Methods

We employ a one-layer Convolutional Long Short-Term Memory (ConvLSTM) model [19] with a hidden size of 64 to simulate daily LAI dynamics from climate forcing variables. From a deep learning perspective, this task can be framed as a video sequence-to-sequence problem, where climate variable images serve as input frames and the corresponding LAI images are output frames.

Model inputs consist of daily climate variables, organized as spatial grids clipped to 63×54 images. The output is the predicted LAI map at the corresponding temporal resolution.

Two model configurations are implemented (Figure 2): (1) pure ConvLSTM, trained solely on climate forcings to predict LAI; and (2) physics-encoded ConvLSTM variant, which integrates LAI simulated from a process-based vegetation model (combining the Farquhar photosynthesis model [8, 3] and the Jarvis stomatal conductance model [9]) as an additional input channel, with ConvLSTM serving as a bias-corrector to adjust the physically simulated LAI. Key parameters of the physical module are learnable, enabling joint optimization of process-based and data-driven components.

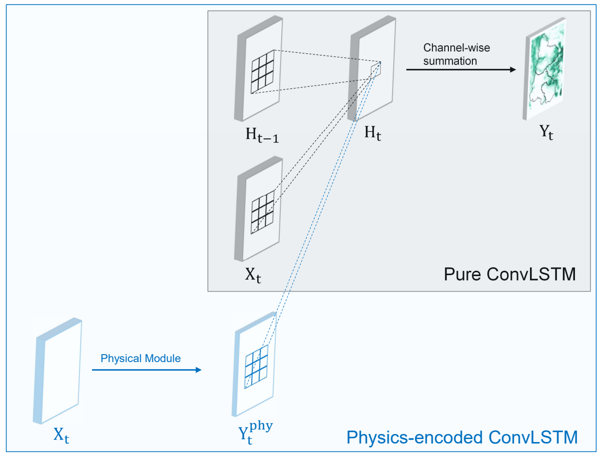

Figure 2: Architecture of our proposed framework showing the pure ConvLSTM and physics-encoded ConvLSTM for climate-vegetation dynamics modeling.

For the pure ConvLSTM, all inputs are normalized prior to training; in the physics-encoded variant, normalization is applied internally after the physical computation. All models are trained using the AdamW [14] optimizer and early stopping is applied to prevent overfitting.

4 Results

We compare three models: (1) a one-layer LSTM baseline with flattened spatial inputs (with hidden size of 64), (2) the pure ConvLSTM, and (3) the physics-encoded ConvLSTM. All models are trained, validated, and tested against remotely sensed LAI. Model performance is evaluated using the Nash-Sutcliffe Efficiency (NSE) [16], calculated for each grid cell along the temporal dimension to assess how well the models capture sub-monthly scale temporal dynamics (Figure 3). NSE is a normalized statistic that measures the relative magnitude of residual variance compared to the variance of the observed data. Higher NSE values correspond to better predictive performance and a value of 1 indicates perfect agreement between predictions and observations.

The pure ConvLSTM achieves the highest predictive skill, with spatial NSE maps showing consistently superior performance in diverse vegetation types (Figure 3b) and a median NSE of 0.867, compared to 0.839 for the physics-encoded ConvLSTM and 0.787 for the LSTM baseline. It reproduces observed LAI spatial patterns and temporal dynamics with high fidelity (Figures 3c-d). This likely reflects its ability to implicitly represent ecohydrological state variables (e.g., soil moisture, carbon pool) in its hidden states.

The physics-encoded ConvLSTM also outperforms the LSTM baseline across diverse vegetation types and reduces the seasonal phase shifts in LAI seen in the baseline (Figure 3d). It demonstrates particular advantages in certain water-limited regions (e.g., shrublands during the spring and summer of 2009; see Figure 3d), where purely data-driven models tend to produce unrealistic fluctuations. The incorporation of physical constraints appears to enhance stability under complex vegetation-climate interactions, indicating potential for more reliable application in climate projection scenarios, where extrapolation beyond the historical training range is often unavoidable.

Nevertheless, the physics-encoded ConvLSTM does not match the pure ConvLSTM in overall accuracy. This may be due to the trade-off between flexibility and constraint: the model must jointly optimize both ConvLSTM weights and physical parameters, and errors or simplifications in the physical model can propagate into the hybrid system. Additionally, the optimization process is more complex because the physical parameters must be learned from scratch, which can slow convergence. Even when initialized with plant functional type (PFT)-based parameter sets, uncertainties arising

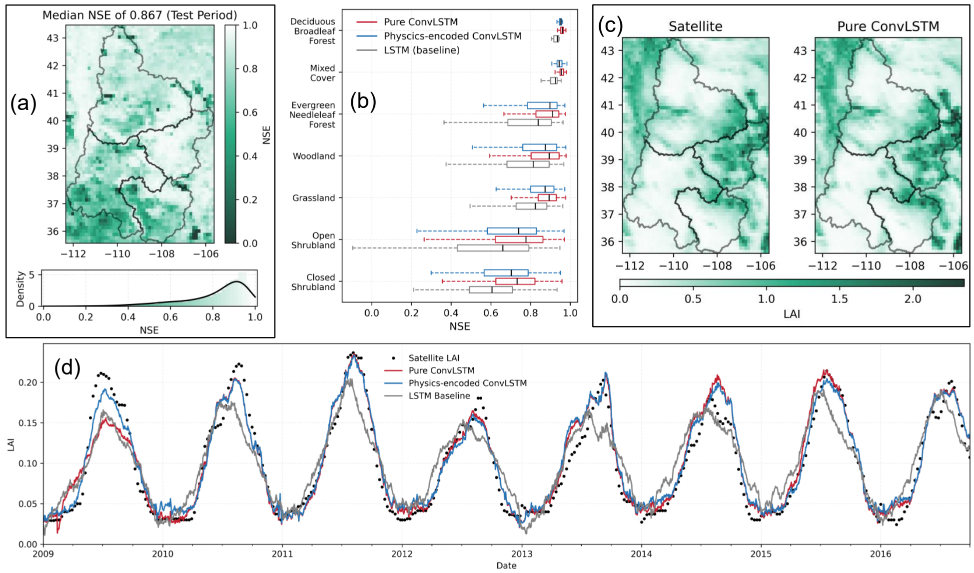

Figure 3: (a) Spatial distribution and probability density of NSE for the pure ConvLSTM over the study area during the test period (2009–2016); (b) Boxplots of NSE for the LSTM baseline, pure ConvLSTM, and physics-encoded ConvLSTM across dominant vegetation types; (c) Example spatial comparison between pure ConvLSTM-simulated LAI and satellite-derived LAI on May 9, 2014; (d) Example comparisons of simulated vegetation dynamics (Open Shrubland) from the LSTM baseline, pure ConvLSTM, and physics-encoded ConvLSTM during the test period.

from inaccurate vegetation classification can introduce unrealistic constraints. This limitation is common in physics-based deep learning for vegetation dynamics, as process-based vegetation models often rely heavily on PFTs, which can reduce model flexibility and overall performance. Future work could address this by introducing dynamic PFT representations to enhance accuracy and adaptability.

5 Conclusions and impacts

This study develops and evaluates a fully climate-driven, spatially distributed deep learning framework for simulating daily vegetation dynamics, comparing a pure ConvLSTM, a physics-encoded ConvLSTM, and an LSTM baseline. The pure ConvLSTM achieves the highest predictive accuracy, capturing fine-scale spatiotemporal variability across diverse vegetation types. The physics-encoded variant, while slightly less accurate than pure ConvLSTM, demonstrates improved stability in water-limited regions and offers greater potential for reliable application to climate projection scenarios, where extrapolation beyond the historical record is inevitable.

From a climate change perspective, accurate simulation of vegetation-climate interactions is essential for projecting future water availability, ecosystem resilience, and carbon cycle feedbacks. Purely data-driven models can excel within historical domains but may struggle under nonstationary climate conditions. Incorporating physical constraints can help mitigate these risks, enabling more robust long-term projections and supporting climate adaptation and mitigation strategies. The proposed framework is broadly applicable wherever gridded climate forcings are available, making it scalable to continental and global applications.

Acknowledgments and Disclosure of Funding

We acknowledge the financial support from the **Advance Water Observatory and Decision Support System (AWODSS)** program through the funded project "A machine-learning based vegetation dynamics to improve hydrologic prediction under changing climate".

References

[1] Jill T Anderson and Bao-Hua Song. Plant adaptation to climate change—where are we? *Journal of Systematics and Evolution*, 58(5):533–545, 2020.

[2] Vitus Benson, Claire Robin, Christian Requena-Mesa, Lazaro Alonso, Nuno Carvalhais, José Cortés, Zhihan Gao, Nora Linscheid, Mélanie Weynants, and Markus Reichstein. Multi-modal learning for geospatial vegetation forecasting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 27788–27799, 2024.

[3] CJ Bernacchi, EL Singsaas, CARLOS Pimentel, AR Portis Jr, and Stephen P Long. Improved temperature response functions for models of rubisco-limited photosynthesis. *Plant, cell & environment*, 24(2):253–259, 2001.

[4] Fei Chen and Jimy Dudhia. Coupling an advanced land surface–hydrology model with the penn state–ncar mm5 modeling system. part i: Model implementation and sensitivity. *Monthly weather review*, 129(4):569–585, 2001.

[5] Zhi-Ting Chen, Hong-Yan Liu, Chong-Yang Xu, Xiu-Chen Wu, Bo-Yi Liang, Jing Cao, and Deliang Chen. Deep learning projects future warming-induced vegetation growth changes under ssp scenarios. *Advances in Climate Change Research*, 13(2):251–257, 2022.

[6] Zhiting Chen, Hongyan Liu, Chongyang Xu, Xiuchen Wu, Boyi Liang, Jing Cao, and Deliang Chen. Modeling vegetation greenness and its climate sensitivity with deep-learning technology. *Ecology and Evolution*, 11(12):7335–7345, 2021.

[7] Gregory Duveiller, Josh Hooker, and Alessandro Cescatti. The mark of vegetation change on earth's surface energy balance. *Nature communications*, 9(1):679, 2018.

[8] Graham D Farquhar, S von von Caemmerer, and Joseph A Berry. A biochemical model of photosynthetic co2 assimilation in leaves of c3 species. *planta*, 149(1):78–90, 1980.

[9] PG Jarvis. The interpretation of the variations in leaf water potential and stomatal conductance found in canopies in the field. *Philosophical Transactions of the Royal Society of London. B, Biological Sciences*, 273(927):593–610, 1976.

[10] Klaus-Rudolf Kladny, Marco Milanta, Oto Mraz, Koen Hufkens, and Benjamin D Stocker. Enhanced prediction of vegetation responses to extreme drought using deep learning and earth observation data. *Ecological Informatics*, 80:102474, 2024.

[11] Guohua Liu, Mirco Migliavacca, Christian Reimers, Basil Kraft, Markus Reichstein, Andrew D Richardson, Lisa Wingate, Nicolas Delpierre, Hui Yang, and Alexander J Winkler. Deepphenomem v1. 0: deep learning modelling of canopy greenness dynamics accounting for multi-variate meteorological memory effects on vegetation phenology. *Geoscientific Model Development*, 17(17):6683–6701, 2024.

[12] Tian Liu, Huaan Jin, Ainong Li, Hongliang Fang, Dandan Wei, Xinyao Xie, and Xi Nan. Estimation of vegetation leaf-area-index dynamics from multiple satellite products through deep-learning method. *Remote sensing*, 14(19):4733, 2022.

[13] Yang Liu, Ronggao Liu, and Jing M Chen. Retrospective retrieval of long-term consistent global leaf area index (1981–2011) from combined avhrr and modis data. *Journal of Geophysical Research: Biogeosciences*, 117(G4), 2012.

[14] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.

[15] Francesco Martinuzzi, Miguel D Mahecha, Gustau Camps-Valls, David Montero, Tristan Williams, and Karin Mora. Learning extreme vegetation response to climate drivers with recurrent neural networks. *Nonlinear Processes in Geophysics*, 31(4):535–557, 2024.

[16] J Eamonn Nash and Jonh V Sutcliffe. River flow forecasting through conceptual models part i—a discussion of principles. *Journal of hydrology*, 10(3):282–290, 1970.

[17] Guo-Yue Niu, Zong-Liang Yang, Kenneth E Mitchell, Fei Chen, Michael B Ek, Michael Barlage, Anil Kumar, Kevin Manning, Dev Niyogi, Enrique Rosero, et al. The community noah land surface model with multiparameterization options (noah-mp): 1. model description and evaluation with local-scale measurements. *Journal of Geophysical Research: Atmospheres*, 116(D12), 2011.

[18] Zongyao Sha, Yongfei Bai, Ruren Li, Hai Lan, Xueliang Zhang, Jonathon Li, Xuefeng Liu, Shujuan Chang, and Yichun Xie. The global carbon sink potential of terrestrial vegetation can be increased substantially by optimal land management. *Communications Earth & Environment*, 3(1):8, 2022.

[19] Xingjian Shi, Zhourong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. *Advances in neural information processing systems*, 28, 2015.

[20] Youlong Xia, Kenneth Mitchell, Michael Ek, Justin Sheffield, Brian Cosgrove, Eric Wood, Lifeng Luo, Charles Alonge, Helin Wei, Jesse Meng, et al. Continental-scale water and energy flux analysis and validation for the north american land data assimilation system project phase 2 (nldas-2): 1. intercomparison and application of model products. *Journal of Geophysical Research: Atmospheres*, 117(D3), 2012.

[21] Wentao Yu, Jing Li, Qinhuo Liu, Jing Zhao, Yadong Dong, Cong Wang, Shangrong Lin, Xinran Zhu, and Hu Zhang. Spatial–temporal prediction of vegetation index with deep recurrent neural networks. *IEEE Geoscience and Remote Sensing Letters*, 19:1–5, 2021.

[22] Pu Zhang, Zhipeng Li, Heyu Zhang, Jie Ding, Xufeng Zhang, Rui Peng, and Yiming Feng. Simulation model of vegetation dynamics by combining static and dynamic data using the gated recurrent unit neural network-based method. *International Journal of Applied Earth Observation and Geoinformation*, 112:102901, 2022.