

Physics Guided RNNs for Modeling Dynamical Systems: A Case Study in Simulating Lake Temperature Profiles

Xiaowei Jia^{1*}, Jared Willard^{1*}, Anuj Karpatne², Jordan Read³,
Jacob Zwart³, Michael Steinbach¹, Vipin Kumar¹

¹Department of Computer Science and Engineering, University of Minnesota

²Department of Computer Science and Engineering, Virginia Tech

³U.S. Geological Survey

¹{jiaxx221,willa099,stei0062,kumar001}@umn.edu, ²karpatne@vt.edu, ³{jread,jzwart}@usgs.gov

Abstract

This paper proposes a physics-guided recurrent neural network model (PGRNN) that combines RNNs and physics-based models to leverage their complementary strengths and improve the modeling of physical processes. Specifically, we show that a PGRNN can improve prediction accuracy over that of physical models, while generating outputs consistent with physical laws, and achieving good generalizability. Standard RNNs, even when producing superior prediction accuracy, often produce physically inconsistent results and lack generalizability. We further enhance this approach by using a pre-training method that leverages the simulated data from a physics-based model to address the scarcity of observed data. Although we present and evaluate this methodology in the context of modeling the dynamics of temperature in lakes, it is applicable more widely to a range of scientific and engineering disciplines where mechanistic (also known as process-based) models are used, e.g., power engineering, climate science, materials science, computational chemistry, and biomedicine.

1 Introduction

Physics-based models of dynamical systems are often used to study engineering and environmental systems. Despite their extensive use, these models have several well-known limitations due to simplified representations of the physical processes being modeled or challenges in selecting appropriate parameters. Given rapid data growth due to advances in sensor technologies, there is a tremendous opportunity to systematically advance modeling in these domains by using machine learning (ML) methods. However, capturing this opportunity is contingent on a paradigm shift in data-intensive scientific discovery since the black box use of ML often leads to serious false discoveries in scientific applications [9, 12]. This paper presents a novel methodology for combining physics-based models with state-of-the-art deep learning methods to leverage their complementary strengths. Although we present and evaluate this methodology in the context of modeling the dynamics

of temperature in lakes, it is applicable more widely to a range of scientific and engineering disciplines where mechanistic (also known as process-based) models are used, e.g., power engineering, climate science, materials science, computational chemistry, and biomedicine.

Even though physics-based models are based on known laws that govern relationships between input and output variables, they often rely on a large number of unknown parameters. These parameters must be estimated (or calibrated) from observed data, which is often scarce. A standard approach for model calibration is to intelligently search the space of parameter combinations and choose parameter combinations that result in the best performance on training data. This approach is computationally expensive as well as highly prone to over-fitting. Another limitation is that the majority of physics-based models implement approximate forms of physical relationships, either due to incomplete knowledge of certain processes or for practical computing purposes. Such approximations introduce additional parameters to these models that must be calibrated from data, making the process of model calibration even more challenging. These issues individually or in combination can introduce biases and thus result in unacceptable performance. The limitations of physics-based models cut across discipline boundaries and are well known in the scientific community [11].

ML models, which have found tremendous success in several commercial applications where large-scale data is available, e.g., computer vision, and natural language processing, are increasingly being considered as promising alternatives to physics-based models by the scientific community. However, direct application of black-box ML models to a scientific problem encounters three major challenges: 1. State of the art (SOA) ML models that are powerful enough to effectively represent spatial and temporal processes inherent in physical systems can often perform better than traditional em-

*These authors have equal contribution.

pirical models (e.g., regression-based models) used by the science communities as an alternative to physics-based models. However, they require a lot of training data, which is scarce in most practical settings. 2. Empirical models (including the SOA ML models) simply identify statistical relations between inputs and the system variables of interest (e.g., the temperature profile of the lake) without taking in to account any physical laws (e.g., conservation of energy or mass) and thus can produce results that are inconsistent with physical laws. 3. Relationships produced by empirical models can at best be valid only for the set of forcing variable combinations present in the training data and are unable to generalize to scenarios unseen in the training data. For example, a ML model trained on a water body for today's climate may not be accurate for future warmer climate scenarios. As an alternative approach to both physics-based and empirical models, we present Physics-Guided Recurrent Neural Network models (PGRNN) as a general framework for modeling physical phenomena with potential applications for many disciplines. PGRNN incorporates physics into the ML model by generalizing the loss function to include physical laws as the third component beyond the traditional notions of error and model complexity [9].

Our proposed Physics-Guided Recurrent Neural Networks model (PGRNN) is developed in the context of lake water temperature modeling at a daily scale. The temperature of water in a lake is known to be an ecological master factor [13] that controls the growth, survival, and reproduction of fish [19]. Warming water temperatures can increase the occurrence of aquatic invasive species [16, 20], which may displace fish and native aquatic organisms, and result in more harmful algal blooms (HABs) [4]. Understanding temperature change and the resulting biotic winners and losers is timely science that can also be directly applied to inform priority action for natural resources.

The PGRNN model has a number of novel aspects. This model contains two parallel recurrent structures - a standard RNN flow and an energy flow to be able to capture the variation of energy balance over time. While the standard RNN flow models the temporal dependencies that better fit observed data, the energy flow aims to regularize the temporal progression of the model in a physically consistent fashion.

To further improve the learning performance with the scarcity of observed data, we propose a pre-training method that utilizes the simulated data generated by physics-based models. While the simulated data are not an accurate reflection of the observed data, this pre-training algorithm has the potential to produce a better initialized status for the learning model and thus

requires less observed data to fine-tune model parameters. Finally, we show that the proposed PGRNN model has the flexibility to incorporate additional physical constraints that are involved in specific applications. For example, in the lake temperature simulation problem, predicted values of the temperature at different depths should be such that denser water is at a lower depth, which is known as the density-depth constraint [10].

We evaluate the proposed method in a reasonably large real-world system, Lake Mendota in Wisconsin. This lake is chosen for evaluation, as it is one of the most extensively studied lake systems and plenty of observed data is available to evaluate the performance of any new approach. We show that the modeling of energy conservation can successfully improve the learning performance and the generalization capacity. Moreover, the results confirm that the pre-training method can help achieve a reasonable performance even with a small amount of observed data. Finally, we show that after incorporating the density-depth constraint, the PGRNN model can produce both highly accurate and physically meaningful predictions.

2 Problem Formulation

To fully capture the temperature change in a lake system, we are interested in simulating the temperature of water at each depth d , and on each date, t .

Specifically, we consider the physical variables governing the dynamics of lake temperature at every depth and time-step as the set of input drivers, $X = \{x_{d,t}\}$. These chosen features are known to be the primary drivers of lake thermodynamics [7]. This includes meteorological recordings at the surface of water such as the amount of solar radiation, wind speed, air temperature, etc. Given the input drivers, we aim to predict water temperature $Y = \{y_{d,t}\}$. In particular, $y_{d,t}$ denotes the temperature at depth d and at time step t .

3 Preliminaries

3.1 General Lake Model (GLM) The physics-based GLM captures a variety of physical processes governing the dynamics of water temperature in a lake, including the heating of the water surface due to incoming short-wave radiation, the attenuation of radiation beneath the water surface, the mixing of layers with varying energies at different depths, and the loss of heat from the surface of the lake via evaporation or long-wave radiation (shown in Fig. 1). We use GLM as our preferred choice of physics-based model for lake temperature modeling due to its model performance and wide use among the lake modeling community.

The GLM has a number of parameters (e.g., parameters related to vertical mixing, wind sheltering, and

water clarity) that are often calibrated specifically to individual lakes if training data are available. The basic calibration method is to run the model for combinations of parameter values and select the parameter set that minimizes model error. This calibration process can be both labor- and computationally-intensive, and if executed without expert knowledge of parameter meaning and acceptable values, can create model formulations that perform poorly when evaluated against test data. Furthermore, calibration processes applied even in the ideal case of complete and error-free observational are limited by simplifications and rigid formulations of parameters in these physics-based models.

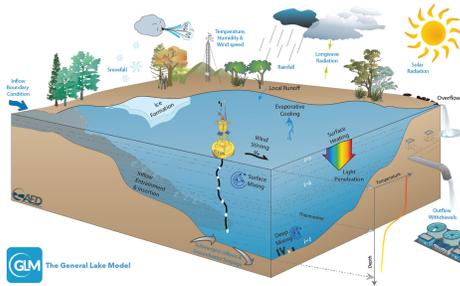


Figure 1: A pictorial description of the physical processes simulated by the General Lake Model [6]. These processes govern the dynamics of temperature in a lake.

3.2 Long-Short Term Memory Networks In essence, the LSTM model defines a transition relationship for hidden representation h_t through an LSTM cell, which combines the input features x_t at each time step and the inherited information from previous time steps.

Each LSTM cell contains a cell state c_t , which serves as a memory and allows reserving information from the past. Specifically, the LSTM first generates a candidate cell state \tilde{c}_t by combining x_t and h_{t-1} , as:

$$(3.1) \quad \tilde{c}_t = \tanh(W_h^c h_{t-1} + W_x^c x_t).$$

Then the LSTM generates a forget gate f_t , an input gate g_t , and an output gate via sigmoid function, as:

$$(3.2) \quad \begin{aligned} f_t &= \sigma(W_h^f h_{t-1} + W_x^f x_t), \\ g_t &= \sigma(W_h^g h_{t-1} + W_x^g x_t), \\ o_t &= \sigma(W_h^o h_{t-1} + W_x^o x_t). \end{aligned}$$

The forget gate is used to filter the information inherited from c^{t-1} , and the input gate is used to filter the candidate cell state at t . Then we compute the new cell state and the hidden representation as:

$$(3.3) \quad \begin{aligned} c_t &= f_t \otimes c_{t-1} + g_t \otimes \tilde{c}_t, \\ h_t &= o_t \otimes \tanh(c_t), \end{aligned}$$

where \otimes denotes the entry-wise product.

As we wish to conduct regression for continuous values, we generate the predicted temperature \hat{y}_t at each time step t via a linear combination of hidden units, as:

$$(3.4) \quad \hat{y}_t = W_y h_t.$$

We apply the LSTM model for each depth separately.

Given the true observation $y_{d,t}$ at every time step and at every depth, our training loss is defined as:

$$(3.5) \quad \mathcal{L}_{\text{RNN}} = \frac{1}{N_d T} \sum_t \sum_d (y_{d,t} - \hat{y}_{d,t})^2,$$

where N_d is the total number of different depths, and T is the number of time steps.

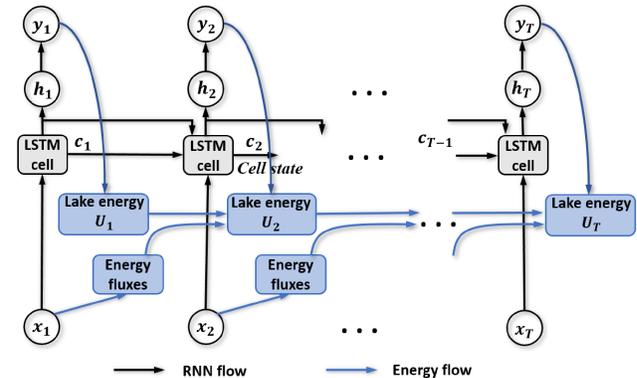


Figure 2: The flow of the PGRNN model.

4 Method

4.1 Energy conservation over time We integrate energy conservation flow into the recurrent process, as shown in Fig. 2. While the recurrent flow in the standard RNN can capture data dependencies across time, the modeling of energy flow ensures that the change of lake environment and predicted temperature conforms to the law of energy conservation.

Traditional LSTM models utilize the LSTM cell to implicitly encode useful information at each time step and pass it to the next time step. In contrast, the energy flow in PGRNN explicitly captures the key factor that leads to temperature change in dynamical systems - the heat energy fluxes that are transferred from one time to the next. In this way, the PGRNN provides a structure within to predict what causes the temporal variation in dynamical systems. Further, note that even though different lake systems have different data distributions, they all conform to the same universal law of energy conservation. Therefore, by complying with the universal law of energy conservation, PGRNN has a better chance at learning generalizable patterns that apply across multiple types of lakes.

In Fig. 3, we show the major incoming and outgoing heat fluxes that impact the lake energy. The incoming heat fluxes include terrestrial long-wave radiation and incoming short-wave radiation. The lake loses heat mainly through the outward fluxes of back radiation (R_{LWout}), sensible heat fluxes (H), and latent evaporative heat fluxes (E). In this work, we ignore the smaller flux terms such as sediment heat flux and advected energy from surface inflows and groundwater.

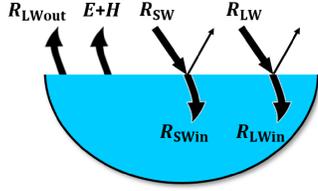


Figure 3: The heat energy fluxes that are modeled in PGRNN. For short-wave radiation (R_{SW}) and long-wave radiation (R_{LW}), a portion of the energy is reflected by the lake surface.

The balance of these components results in a change in the thermal energy (U_t) of the lake. Specifically, the lake will gain energy if the incoming heat fluxes are more than outgoing heat fluxes, and the lake will lose energy if there are more outgoing heat fluxes than incoming heat fluxes. The consistency between lake energy U_t and energy fluxes can be expressed as:

$$(4.6) \quad \begin{aligned} \Delta U_t = & R_{SW}(1 - \alpha_{SW}) + R_{LW_{in}}(1 - \alpha_{LW}) \\ & - R_{LW_{out}} - E - H, \end{aligned}$$

where $\Delta U_t = U_{t+1} - U_t$, α_{SW} is the short-wave albedo (the fraction of short-wave energy reflected by the lake surface) and α_{LW} is the long-wave albedo. In our implementation, we set α_{SW} to 0.07 and α_{LW} to 0.03 which are generally accepted values for lakes from previous scientific studies [7]. All energy components are in Wm^{-2} .

Lake Mendota experiences ice cover during winter months, but here we consider the energy conservation only for ice-free periods since the lake exhibits drastically different reflectance and energy loss dynamics when covered in ice and snow, and the modeling of ice and snow was considered out of scope for this study. As such, we define the loss function term for energy conservation and combine this with the training objective of standard LSTM model in the following equation.

$$(4.7) \quad \begin{aligned} \mathcal{L} = & \mathcal{L}_{RNN} + \lambda_{EC} \mathcal{L}_{EC}, \\ \mathcal{L}_{EC} = & \frac{1}{T_{\text{ice-free}}} \sum_{t \in \text{ice-free}} \text{ReLU}(|\Delta U_t - \mathcal{F}| - \tau_{EC}), \end{aligned}$$

where \mathcal{F} represents the sum of heat fluxes at the right hand side of Eq. 4.6, τ_{EC} is the threshold for the loss of energy conservation. This threshold is introduced because physical processes can be affected by unknown less important factors which are not included in the model, or by observation errors in the meteorological data. The function $\text{ReLU}(\cdot)$ is adopted such that only the difference larger than the threshold is counted towards the penalty. In our implementation, the threshold is set as the largest value of $|\Delta U_t - \mathcal{F}|$ in the GLM model for daily averages. The hyper-parameter λ_{EC} controls the balance between the loss of the standard RNN and the energy conservation loss.

Estimation of Heat Fluxes and Lake Thermal Energy:

Terrestrial long-wave ($R_{LW_{in}}$) radiation is emitted from the atmosphere, and depends on prevailing local conditions like air temperature and cloud cover. Incoming short-wave radiation (R_{SW}) is affected mainly by latitude (solar angle), time of year, and cloud cover. Both factors are included in the input drivers X .

As for the outgoing energy fluxes, we estimate E , H , and $R_{LW_{out}}$ separately using the input drivers and modeled surface temperature.

The sensible heat flux and latent evaporative heat flux can be computed based on the previous study [7]:

$$(4.8) \quad E = -\rho_a C_E \nu \kappa_{10} \frac{\omega}{p} (e_s - e_a),$$

$$H = -\rho_a c_a C_H \kappa_{10} (T_s - T_a),$$

where C_H is the bulk aerodynamic coefficients for sensible heat transfer, and C_E the bulk aerodynamic coefficients for latent heat transfer. Both coefficients are estimated from Hicks' collection of ocean and lake data [5]. The coefficient ω is the ratio of molecular mass of water to molecular mass of dry air ($=0.622$), ν the latent heat of vaporization ($=2.453 \times 10^6$), and c_a the specific heat capacity of air ($=1005$). The variable T_a is the air temperature, and κ_{10} the wind speed (m/s) above the lake referenced to 10m height. Both these variables are included or can be derived from input drivers. T_s is the surface water temperature in degrees Kelvin obtained through the feed-forward process. The air density ρ_a is computed as $\rho_a = 0.348(1+r)/(1+1.61r)p/T_a$, where p is air pressure (hPa) and r is the water vapour mixing ratio (both derived from input drivers). The vapour pressure (e_s and e_a) is calculated by the linear formula from Tabata [23]:

$$(4.9) \quad \begin{aligned} e_s = & 10^{(9.28603523 \frac{2322.37885}{T_s + 273.15})}, \\ e_a = & (S_{RH} RH / 100) e_s, \end{aligned}$$

where S_{RH} is the relative humidity scaling factor ($=1$, obtained through calibrating the GLM model) and RH is the relative humidity (included in input drivers).

The back radiation $R_{LW_{out}}$ is estimated as:

$$(4.10) \quad R_{LW_{out}} = \epsilon_s \delta T_s^4,$$

where ϵ_s is the emissivity of the water surface ($=0.97$), and δ is the Stefan-Boltzmann constant ($=5.6697 \times 10^{-8} \text{Wm}^{-2}\text{K}^{-4}$).

On the left-hand side of Eq. 4.6, the total thermal energy of the lake at time t is:

$$(4.11) \quad U_t = c_w \sum_d a_d y_{d,t} \rho_{d,t} \partial z_d,$$

where c_w is the specific heat of water ($4186 \text{J kg}^{-1}\text{C}^{-1}$), a_d the cross-sectional area of the water column (m^2), and ∂z_d the thickness (m) of the layer at depth d . In this work, we simulate water temperature for every 0.5m and thus we set $\partial z_d=0.5$. The computation of U_t requires the output of temperature $y_{d,t}$ through feed-forward

process for all the depths, as well as cross-sectional depth area a_d .

Note that the modeling of energy flow using the procedure described above does not require any input of true labels/observations. According to Eqs. 4.8-4.11, the heat fluxes and lake energy are computed using only input drivers and predicted temperature. In light of these observations, we can extend this model to incorporate the energy conservation for other systems which have only a few labeled data points. This enforces predicted temperatures in unmonitored systems to also follow the universal law of energy conservation.

4.2 Pre-training using Physical Simulations In many environmental systems, observed data is limited due to the extensive labor required to deploy sensors and process the data. Therefore, the learning model needs to be effectively trained using only small amount of observed data. For scientific problems, it is also critical to build generalizable models that can be applied to different systems. For example, a model built on one lake should be able to generalize to another lake which has limited observed data. However, traditional data science models are known to suffer from the data heterogeneity problem where training data and testing data are from different distributions [15]. Hence, the RNN-based model trained with limited observed data is very likely to lead to unsatisfactory performance.

To address the issue of limited observed data, We pretrain the PGRNN model using the simulated data produced by a simple physics-based model. In particular, given the input drivers, we can run the GLM to predict temperature at every depth and at every day. These simulated temperature data by GLM conform to underlying physical laws used to build the GLM. Hence, the pre-training using these simulated data can result in a physically consistent initialized model. When applying the pretrained model to a new system, we fine-tune the model using limited observed data. In our experiments, we show that this pre-training method can achieve high accuracy given very few observed data.

4.3 Density-depth Constraint While the modeling of energy flow allows the incorporation of the most important physical process that controls water temperature dynamics, other physical constraints relevant to lake temperature exist. The incorporation of these additional constraints can help guide the model to make predictions that are consistent with real-world physics. To demonstrate the capacity of the PGRNN model to incorporate these constraints, we consider an illustrative example for adding density-depth constraint as follows.

It is known that the density of water monotonically

increases with depth and thus can be used as constraints on the outputs of PGRNN. We first transform the predicted temperature Y into the density values ρ according to the following known physical equation [14]:

$$(4.12) \quad \rho = 1000 \times \left(1 - \frac{(Y + 288.9414) \times (Y - 3.9863)^2}{508929.2 \times (Y + 68.12963)}\right).$$

Then, we add an extra penalty for violation of density-depth relationship. Specifically, on any pair of consecutive depths d and $d + 1$, if $\rho_{d,t}$ is larger than $\rho_{d+1,t}$, then this is considered as a violation to the density-depth relation. In this way, we define the loss of density-depth constraint as:

$$(4.13) \quad \Delta\rho_{d,t} = \rho_{d,t} - \rho_{d+1,t},$$

$$\mathcal{L}_{DC} = \frac{1}{T(N_d - 1)} \sum_t \sum_d \text{ReLU}(\Delta\rho_{d,t}),$$

where $\text{ReLU}(\cdot)$ is used to ensure that only pairs with inverse density values are counted towards the penalty.

Combining this with Eq. 4.7, the complete training objective becomes:

$$(4.14) \quad \mathcal{L} = \mathcal{L}_{RNN} + \lambda_{EC}\mathcal{L}_{EC} + \lambda_{DC}\mathcal{L}_{DC},$$

where λ_{DC} is the hyper-parameter to control the penalty for violating the density-depth constraint.

5 Experiment

Our dataset was collected from Lake Mendota in Wisconsin, USA. This lake system is reasonably large (~ 40 km² in area) and displays sufficient dynamics in the temperature profiles over time. Observations of lake temperature were collated from a variety of sources, including the North Temperate Lakes Long-Term Ecological Research Program and a web resource that collates data from federal and state agencies, academic monitoring campaigns, and citizen data [17]. These temperature observations vary in their distribution across depths and time. There are certain days when observations are available at multiple depths while only a few or no observations are available on some other days.

The input drivers that describe prevailing meteorological conditions are available on a continuous daily basis from April 02, 1980 to December 30, 2014. Specifically, we use a set of 10 drivers as input variables, which include short-wave and long-wave radiation, air temperature, relative humidity, wind speed, frozen and snowing indicators, etc. In contrast, observational data for training and testing the models is not uniform, as measurements were made at varying temporal and spatial (depth) resolutions. In total, 13,158 observations were used for the study period.

5.1 Performance for prediction and generalization We use the observed data from April 02, 1980 to October 31, 1991 and from June 01, 2003 to December 30, 2014 as training data (in total 8,037 observations).

Table 1: Performance of RNN, PGRNN, as well as pretrained models from Lake Mendota (RNN_p and $PGRNN_p$) and Florida ($RNN_{\bar{p}}$ and $PGRNN_{\bar{p}}$) using different amounts of observed data.

Method	2%	20%	100%
RNN	2.311(± 0.240)	1.531(± 0.083)	1.489(± 0.091)
PGRNN	2.156(± 0.178)	1.484(± 0.102)	1.466(± 0.063)
RNN_p	1.650(± 0.169)	1.417(± 0.113)	1.385(± 0.080)
$PGRNN_p$	1.592(± 0.175)	1.409(± 0.103)	1.377(± 0.076)
$RNN_{\bar{p}}$	1.930(± 0.173)	1.476(± 0.096)	1.398(± 0.069)
$PGRNN_{\bar{p}}$	1.759(± 0.147)	1.470(± 0.091)	1.394(± 0.071)

Then we apply the trained model to predict the temperature at different depths for the period from November 01, 1991 to May 31, 2003 (5,121 observations).

To give a sense of the overall performance of the many variations of the PGRNN approach we evaluate in this section, we provide the following statistics for the state-of-the-art GLM model. The uncalibrated GLM model achieves the RMSE of 2.950 in the test period. The GLM model can also be fine-tuned to fit each lake system by optimizing the parameter set to minimize model error [21]. If we use the same 2% training data to optimize parameters in the GLM model, it can reach the test RMSE of 2.645. If we fine-tune it using 100% training data, it will reach the test RMSE of 2.253.

To verify that energy conservation and pre-training helps improve the prediction and generalization, we compare RNN and PGRNN in terms of their prediction RMSE. Here we do not include the basic neural network since it produces RMSE around 1.8, which is far higher than standard RNN.

To test whether each model can perform well using reduced observed data, we randomly select different proportion of data from the training period. For example, to select 20% of training data, we remove every observation in our training period with 0.8 probability. The test data stays the same regardless of training data selection. We repeat each test 10 times and report the mean RMSE and standard deviation.

According to Table 1 (rows 1-2), we can observe that PGRNN consistently outperforms standard RNN. The gap is especially obvious when using smaller subsets of observed data (e.g., 2% data). PGRNN can reach reasonable accuracy using a small amount of observed data because the law of energy conservation regularizes the model to retain physical consistency.

We can see that PGRNN using only 2% observed data can outperform the fully calibrated GLM model (using 100% data). The PGRNN model tuned using 100% data and the pretrained models can achieve much lower RMSE than the GLM model.

We also verify that the pre-training can indeed improve prediction accuracy and generalizability of the

model. A basic premise of pre-training our models is that GLM simulations, though imperfect, provide a synthetic realization of physical responses of a lake to a given set of meteorological drivers. Hence, pre-training a neural network using GLM simulations allows the network to emulate a synthetic but physical phenomena. Our hypothesis is that such a pretrained model requires fewer labeled samples to achieve good generalization performance, even if the GLM simulations do not match with the observations. To test this hypothesis, we conduct an experiment where we generate GLM simulations with input drivers from Lake Mendota. These simulations have been created using a GLM model with generic parameter values that are not calibrated for Lake Mendota, resulting in large errors in modeled temperature profiles with respect to the real observations on Lake Mendota (RMSE=2.950). Nevertheless, these simulated data are physically consistent and by using them for pre-training, we can demonstrate the power of our ML models to work with limited observed data while leveraging the physics inherent in the physical models.

We pre-train RNN and PGRNN using such simulated data and we report their performance when fine-tuned with true observations in Table 1 (rows 3-4). The comparisons between RNN and RNN_p and between PGRNN and $PGRNN_p$ show that the pre-training can significantly improve the performance. The improvement is especially obvious given small amount of observed data. Moreover, we find that the training RNN and PGRNN model commonly takes 150-200 epochs to converge while the training for RNN_p and $PGRNN_p$ only takes 30-50 epochs to converge. This demonstrates that pre-training can provide a better initialized status for the learning model.

To assess the generalization ability of ML models in input conditions different from what we have observed previously, we conduct a different experiment. Here we generate GLM simulations for a synthetic lake with input drivers from Florida (which are very different from the typically much colder conditions in Wisconsin) and then pretrain RNN and PGRNN using the simulated data by GLM based on these input drivers. We show the performance of pretrained models ($RNN_{\bar{p}}$ and $PGRNN_{\bar{p}}$) in Table 1 (rows 5-6).

While $RNN_{\bar{p}}$ and $PGRNN_{\bar{p}}$ trained using these input drivers and simulated data in Florida have very poor performance when directly applied to Lake Mendota (RMSE=9.010 for $RNN_{\bar{p}}$ and 8.657 for $PGRNN_{\bar{p}}$), once they are refined using observed data from Lake Mendota, they get much better. Here we note that $PGRNN_{\bar{p}}$ improves with a much greater amount than $RNN_{\bar{p}}$, which shows the increased generalization power obtained by having a model with physics built into it.

GLM has rarely been tested on more than a handful of years worth of observation data for a single lake. Even when trained and tested on only 2 years of observations, which should produce lower RMSE, GLM model performance is still worse than our PGRNN model performance [1]. Our PGRNN model performance for lake water temperature over multiple decades is unprecedented and will be a valuable tool for water resource managers and lake scientists.

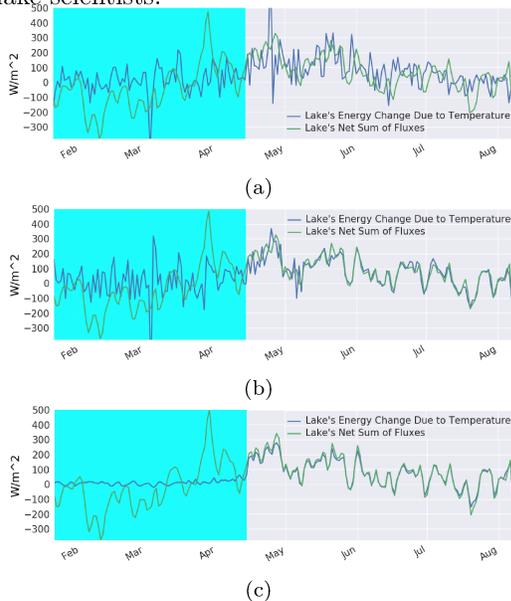


Figure 4: The sum of heat fluxes and the lake energy change generated by (a) RNN, (b) PGRNN, and (c) the generic GLM, from January 21, 1989 to August 09, 1989. The blue part on the left indicates the frozen period (where we do not apply energy conservation).

Heat fluxes and lake energy change To visualize how PGRNN contributes to a physically consistent solution, we testify whether the gap between incoming and outgoing heat energy fluxes matches the lake energy change over time. Specifically, we train RNN and PGRNN using observed data from the first ten years. Then, we show the curves for the gap between incoming and outgoing heat fluxes and the change of lake energy over time for a certain test period (Fig. 4). These two curves should be well aligned (in the ice-free period) if the learning model follows the energy conservation.

It can be seen that the two curves generated by standard RNN have significant differences with each other, which indicates the physical inconsistency over time. In contrast, by modeling the energy flow, PGRNN can generate energy fluxes that to a greater degree match the lake energy change. Because our energy conservation approach has ignored minor terms (e.g., energy fluxes through sediment heating or advection), it is expected that these curves do not completely overlap.

Nevertheless, this close agreement confirms that the PGRNN model is guided by the universal law of energy conservation over time.

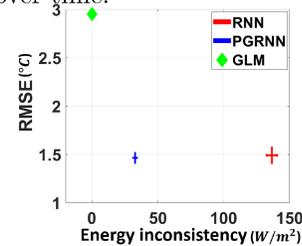


Figure 5: The performance of RNN, PGRNN and generic GLM by RMSE and energy inconsistency.

In Fig. 4 (c), it can be seen that the two curves produced by the generic GLM model are even better aligned than the the curves obtained from the PGRNN model. In real-world systems, the temperature data can be noisy or affected by other unknown factors. Hence, we adopt the soft Lagrangian regularization (Eq. 4.7), which may not lead to the perfect match between heat fluxes and lake energy change.

We summarize the average gap between these two curves in ice-free periods as the energy inconsistency. In Fig. 5, we show the RMSE and the energy inconsistency of RNN, PGRNN and the generic GLM model in the entire test period. Compared with RNN and GLM, PGRNN can significantly reduce both prediction RMSE and energy inconsistency.

5.2 Density-depth relation We now show that the incorporation of density-depth constraint can further improve the prediction. Specifically, from all the pairs of consecutive depths, we compute the fraction of pairs where the model makes physically inconsistent predictions (i.e., the density-depth relationship is violated). We report the average of this fraction over all the test data as the measure of density inconsistency.

According to Fig. 6, we can observe that PGRNN reduces RMSE compared with the standard RNN, but still retains high density inconsistency. In contrast, the incorporation of the density-depth constraint effectively reduces density inconsistency to 0.0021.

In Fig. 7, we show the predicted density values (computed from temperature by Eq. 4.12) by RNN, PGRNN and PGRNN+DC on a certain date when observations are available at most depths. It can be seen that the RNN results in a high prediction error since the predictions stay far away from true observations. Besides, the density values start to decrease below depth 15m, which violates the density-depth constraint. Compared with the standard RNN, the PGRNN improves the prediction accuracy, but some predicted values still violate the density-depth constraint (at depth

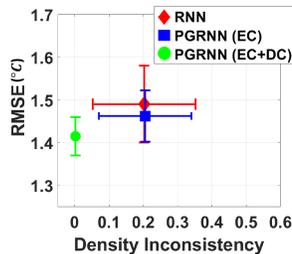


Figure 6: The performance of RNN and PGRNN (EC/EC+DC) by RMSE and density inconsistency.

9-12m). In contrast, the PGRNN+DC model achieves both higher overall accuracy and physically meaningful density predictions.

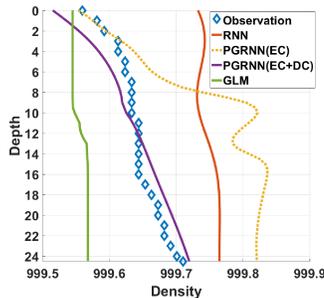


Figure 7: The obtained density values at different depths by different methods and ground-truth observations on May 20, 2002.

5.3 Sensitivity test We conduct sensitivity tests to examine the impact of hyper-parameters (λ_{EC} and λ_{DC}) on the performance (Fig. 8).

As observed from Fig. 8 (a), the performance improves when we increase the weight of energy conservation loss from 0 to 0.02. However, as we adopt extremely large λ_{EC} , the RMSE will be high due to the ignorance of the standard supervised training loss.

In Fig. 8 (b), we can see that the performance is relatively stable as we add weight to the density-depth constraint. However, when the value of λ_{DC} is larger than 3000, the error starts to increase sharply. This is also because the training focuses on generating physically meaningful outputs while failing to fit the observed data.

6 Related Work

The standard and most common approach for incorporating physical knowledge in machine learning models is to use it to guide feature selection or feature construction, but such approaches cannot learn in absence of labels. Recently, several papers have considered the use of constraints (based on physical knowledge) on the output of a machine learning model so that these models can be trained even with unlabeled data by relying

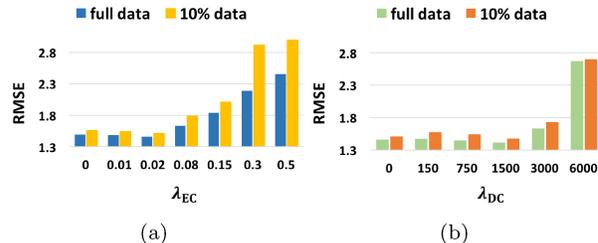


Figure 8: Sensitivity tests: the performance (with full data or 10% data) with respect to (a) the weight of energy conservation λ_{EC} , and (b) the weight of density-depth constraint λ_{DC} .

on physical principles [10, 18, 22]. Another very common approach is residual modeling, where an ML model is used to predict the errors made by a physics-based model [2, 3, 24]. In addition to generating results that may be inconsistent with physical laws, these residual modeling approaches do not have the capability to train ML models with unlabeled data by relying on physical principles. The approach presented in [10] is related to both of these types of approaches, as it uses the output of a physics-based model as an additional feature in a neural network model (making it related to residual modeling approaches), and also makes use of constraints on the output to allow for label free learning. The PGRNN approach presented in this paper is unique in that it provides a powerful framework for modeling spatial and temporal processes. This framework makes it possible to incorporate complex physical laws such as energy and mass conservation (that cannot be enforced by just considering the output of the model) and can also be pre-trained using physics based models. Together these properties enable development of models that are much more generalizable and can learn with little observational data.

7 Acknowledgement

This work was funded by the USGS grant G18AC00352, NSF Awards 1029711 and DTC seed grant. JSR was funded by the Department of the Interior Northeast and North Central Climate Adaptation Science Centers. Access to computing facilities was provided by Minnesota Supercomputing Institute.

8 Conclusion

In this paper, we propose a novel learning model PGRNN that integrates energy conservation and a density-depth constraint into standard recurrent neural networks for monitoring dynamical systems for scientific knowledge discovery. For the prediction of lake temperature dynamics, RNN simulations had better performance (RMSE) than a commonly used physics-based

model with an optimized parameter set [1], but suffered physical inconsistencies such as violations of the conservation of energy and denser water being on top of less dense water. Our PGRNN approach, which used a loss function with physical constraints, reduced or eliminated two different types of physical inconsistencies (energy conservation and depth-density consistency) while also improving the model accuracy.

We also studied the ability of pre-training these models using simulated data to deal with the scarcity of observed data. Using the simulated data from a poorly parameterized physics-based model, we observed an increasing model performance of PGRNN over RNN with fewer observation data used for training. Thus, PGRNN can leverage the strengths of physics-based models while also filling in knowledge gaps by overlaying features learned from data.

The proposed method can also be adjusted to model other important physical laws in dynamical systems, such as the law of mass conservation. Since energy conservation and mass conservation are universal laws in dynamical systems, the proposed PGRNN model can be applied to a variety of scientific problems such as nutrient exchange in lake systems and analysis of crop field production, as well as engineering problems such as auto-vehicle refueling design. Moreover, the proposed model allows incorporation of additional physical constraints specific to different tasks. Therefore, we anticipate this work to be an important stepping-stone towards more innovations of machine learning for scientific knowledge discovery.

An extended version of this paper is available at [8].

References

- [1] Louise C Bruce et al. A multi-lake comparative analysis of the general lake model (glm): Stress-testing across a global observatory network. *Environmental Modelling & Software*, 102:274–291, 2018.
- [2] Urban Forssell and Peter Lindskog. Combining semi-physical and neural network modeling: An example of usefulness. *IFAC Proceedings Volumes*, 1997.
- [3] Franz Hamilton, Alun L Lloyd, and Kevin B Flores. Hybrid modeling and prediction of dynamical systems. *PLoS computational biology*, 13(7):e1005655, 2017.
- [4] Ted D Harris and Jennifer L Graham. Predicting cyanobacterial abundance, microcystin, and geosmin in a eutrophic drinking-water reservoir using a 14-year dataset. *Lake and reservoir management*, 2017.
- [5] BB Hicks. Some evaluations of drag and bulk transfer coefficients over water bodies of different sizes. *Boundary-Layer Meteorology*, 3(2):201–213, 1972.
- [6] MR Hipsey et al. Glm-general lake model: Model overview and user information. 2014.
- [7] Matthew R Hipsey et al. A general lake model (glm 2.4) for linking with high-frequency sensor data from the global lake ecological observatory network (gleon). *Geoscientific Model Development*, 2017.
- [8] Xiaowei Jia et al. Physics-guided rnns for simulating lake temperature profiles. *arXiv preprint*, 2019.
- [9] Anuj Karpatne, Gowtham Atluri, James H Faghmous, Michael Steinbach, Arindam Banerjee, Auroop Ganguly, Shashi Shekhar, Nagiza Samatova, and Vipin Kumar. Theory-guided data science: A new paradigm for scientific discovery from data. *TKDE*, 2017.
- [10] Anuj Karpatne, William Watkins, Jordan Read, and Vipin Kumar. Physics-guided neural networks (pgnn): An application in lake temperature modeling. *arXiv preprint arXiv:1710.11431*, 2017.
- [11] Upmanu Lall. Debates the future of hydrological sciences: A (common) path forward? one water. one world. many climes. many souls. *WRR*, 2014.
- [12] David Lazer et al. The parable of google flu: traps in big data analysis. *Science*, 2014.
- [13] John J Magnuson et al. Temperature as an ecological resource. *American Zoologist*, 19(1):331–343, 1979.
- [14] James L Martin et al. *Hydrodynamics and transport for water quality modeling*. CRC Press, 2018.
- [15] Sinno Jialin Pan, Qiang Yang, et al. A survey on transfer learning. *TKDE*, 2010.
- [16] Frank J Rahel and Julian D Olden. Assessing the effects of climate change on aquatic invasive species. *Conservation biology*, 22(3):521–533, 2008.
- [17] Emily K Read et al. Water quality data for national-scale aquatic research: The water quality portal. *Water Resources Research*, 2017.
- [18] Hongyu Ren, Russell Stewart, Jiaming Song, Volodymyr Kuleshov, and Stefano Ermon. Learning with weak supervision from physics and data-driven constraints. *AI Magazine*, 2018.
- [19] James J Roberts et al. Fragmentation and thermal risks from climate change interact to affect persistence of native trout in the colorado river basin. *Global Change Biology*, 2013.
- [20] James J Roberts et al. Nonnative trout invasions combined with climate change threaten persistence of isolated cutthroat trout populations in the southern rocky mountains. *North American Journal of Fisheries Management*, 2017.
- [21] Patricia A Soranno et al. Lagos-ne: a multi-scaled geospatial and temporal database of lake ecological context and water quality for thousands of us lakes. *GigaScience*, 2017.
- [22] Russell Stewart and Stefano Ermon. Label-free supervision of neural networks with physics and domain knowledge. In *AAAI*, volume 1, pages 1–7, 2017.
- [23] S Tabata. A simple but accurate formula for the saturation vapor pressure over liquid water. *Journal of Applied Meteorology*, 12(8):1410–1411, 1973.
- [24] Tianfang Xu and Albert J Valocchi. Data-driven methods to improve baseflow prediction of a regional groundwater model. *Computers & Geosciences*, 2015.