Modeling Lake Water Temperature Using Physics-Guided Neural Networks

K.BHASKAR RAO[1], TSSN SRINIVAS[2]
ASSOC.PROFESSOR[1], ASST.PROFESSOR[2]
Department of BS & H
BHIMAVARAM INSTITUTE OF ENGINEERING & TECHNOLOGY

Abstract

This paper introduces a novel framework for learning data science models by using the scientific knowledge encoded in physics-based models. This framework, termed as physics- guided neural network (PGNN), leverages the output of physics-based model simulations along with observational features to generate predictions using a neural network architecture. Further, we present a novel class of learning objective for training neural networks, which ensures that the model predictions not only show lower errors on the training data but are also *consistent* with the known physics. We illustrate the effectiveness of PGNN for the problem of lake temperature modeling, where physical relationships between the temperature, density, and depth of water are used in the learning of neural network model parameters. By using scientific knowledge to guide the construction and learning of neural networks, we are able to show that the proposed framework ensures better generalizability as well as physical consistency of results.

1 Introduction

Physics-based models, which are founded on core sci- entific principles, strive to advance our understanding of the physical world by learning explainable relation- ships between input and output variables. These models can range from solving closed-form equations (e.g. us- ing Navier–Stokes equation for studying laminar flow) to running computational simulations of dynamical sys- tems (e.g. the use of numerical models in climate sci- ence, hydrology, and turbulence modeling). For exam- ple, a number of physics-based models use parameter- ized forms of approximations for representing complex physical processes that are either not fully understood or cannot be solved using computationally tractable methods. Calibrating the parameters in physics-based models is a challenging task because of the combinato- rial nature of the search space. In particular, this can result in the learning of over-complex models that lead to incorrect insights even if they appear interpretable at

a first glance. For example, these and other challenges in modeling hydrological processes using state-of-the- art physics-based models were the subject of a seriesof debate papers in Water Resources Research (WRR)[5, 10, 14].

In contrast, data science methods, that have found tremendous success in several commercial applications where Internet-scale data is available (e.g., natural lan- guage processing, object tracking, and most recently, autonomous driving), are being increasingly anticipated to produce similar accomplishments in scientific do- mains [4, 8, 22]. To capture this excitement, some have even referred to the rise of data science in scientific do- mains as “the end of theory” [1], the idea being that, the increasingly large amounts of data makes it possi- ble to build actionable models without using scientific theories. However, in the absence of adequate informa- tion about the physical mechanisms of real-world pro- cesses, data science approaches are prone to false dis- coveries and could even exhibit serious inconsistencies with known physics. This is because scientific problems often involve complex spaces of hypotheses with non- stationary relationships among the variables that are difficult to capture solely from the data.

Physics-guided data science (PGDS) is an emerging paradigm that aims to leverage the wealth of physical knowledge for improving the effectiveness of data sci- ence models in enabling scientific discovery [9]. Oneof the central goals of PGDS is to ensure the learn-ing of *physically consistent models*, by seamlessly blend-ing physical knowledge in data science methods. Tradi- tional learning frameworks in data science are founded on statistical principles for favoring *simpler* models, e.g., the principle of bias-variance trade-off [3]. While the trade-off between reducing bias and variance is at the heart of a number of machine learning algorithms [23, 3, 24], in scientific applications, another source of in- formation becomes available for ensuring generalizabil- ity, which is the available scientific knowledge. By prun- ing candidate models that are inconsistent with known physics, we can significantly reduce the search space andvariance of models possibly without adversely affecting
their bias. A learning algorithm can then be focusedon the space of physically consistent models, leadingto generalizable and scientifically interpretable models. This is a fundamental objective of PGDS—to include physical *consistency* as a critical

component of model performance along with training accuracy and modelcomplexity. In particular, ensuring physical consistency of a model can be at least as important as improving predictive performance, to safeguard against the possi- bility of learning spurious patterns purely from the data,especially in problems that are of a critical nature and are associated with high risks.

In this paper, we introduce a novel framework of knowledge discovery in scientific problems that combines the power of deep learning with physics- based models, termed as physics-guided neural networks (PGNN). In particular, we present an approach to lever- age the output of physics-based model simulations alongwith the observational features for making predictions using a neural network architecture. Further, we intro- duce a novel class of physics-based learning objective fortraining neural networks, which ensures that the learnednetworks not only admit to lower errors on the training data set but also produce outputs that are consistent with our scientific understanding of the physical world. The use of *physical consistency* as a learning objective isaimed at ensuring better generalizability of the learned network, thus serving as the third pillar of generaliza- tion performance estimate alongside training accuracy and model complexity as follows:

$$\text{Performance} \propto \text{Accuracy} + \text{Simplicity} + \textit{Consistency.}$$

To demonstrate the framework of PGNN, we consider the illustrative problem of modeling the temperature ofwater in a lake at varying depths and times, using me- teorological observations as well as physics-based modelsimulations. For this problem, we exploit an interest- ing relationship between the temperature, density, and depth of water at any given time in a lake to construct aphysics-based learning objective for training neural net- works. While the methodological descriptions of PGNN in this paper are centered around the problem of lake temperature modeling, similar formulations of PGNNcan be explored in a wide range of scientific disciplines involving physics-based models.

The remainder of this paper is organized as follows. Section 2 briefly describes the illustrative problem of lake temperature modeling that is the focus of this paper. Section 3 presents the proposed framework of PGNN. Section 4 discusses experimental results while Section 5 provides concluding remarks Lake Temperature Modeling

We demonstrate the framework of PGNN for the illus- trative problem of lake temperature modeling. The tem-perature of water in a lake is governed by a variety of physical processes pictorially shown in Figure 1, e.g., theheating of the water surface due to incoming shortwave radiation from the sun, the attenuation of radiation be- neath the surface and the mixing of layers with varying energies at different depths, and the dissipation of heat from the surface of the lake via evaporation or longwave radiation.

Knowledge of these physical processes can help us model the dynamics of water temperature in a lake,which is known to be an ecological "master factor" [12] that controls the growth, survival, and reproduction of fish (e.g., [20]). Warming water temperatures can in- crease the occurrence of aquatic invasive species [17, 21], which may displace fish and native aquatic organisms, and result in more harmful algal blooms (HABs) [6, 15]. Understanding temperature change and the resulting bi-otic winners and losers is timely science that can also be directly applied to inform priority action for natu- ral resources. Accurate water temperatures (observed or modeled) are critical to understanding contemporary change, and for predicting future thermal for economi- cally valuable fish.

Since observational data of water temperature at broad spatial scales is incomplete (or non-existent insome regions) high-quality temperature modeling is nec-essary. Of particular interest is the problem of model- ing the temperature of water, $Y_{d,t}$, at a given depth[1],d, and on a certain time, t. This problem is referredto as 1D-modeling of temperature (depth being the sin-gle dimension). A number of physics-based models have

.

been developed that make use of input drivers availableat every depth and time-step, $\tilde{X}_{d,t}$, to produce model es-
timates of temperature at every depth and time, $Y^{Phy}_{d,t}$. These models have a number of parameters (e.g., pa- rameters related to vertical mixing, wind energy inputs, and water clarity) whose values can be set to default val- ues or custom-calibrated for each lake if some training data is available. The basic idea behind these calibra- tion steps is to run the model for each possible com-bination of parameter values and select the one thathas maximum agreement with the observations (lowestRMSE). Because this step of custom-calibrating is both

based model. In particular, we can complement the deficiencies of a physics-based model in the PGNN framework, by learning features extracted as complex combinations of input drivers and physics-based model outputs. We adopt a basic multi-layer perceptron architecture to regress the temperature, $Y_{d,t}$, on any given depth and time, using the input attributes, $X_{d,t}$. For a fully-connected network with L hidden layers, this amounts to the following modeling equations relating the input attributes on a data instance, **x**, to its targetprediction, $\hat{y}$:

labor- and computation-intensive, there is a trade-off between increasing the accuracy of the model and ex- panding the feasability of study to a large number of

Physics-guided Neural Network

Our proposed framework of PGNN uses the scientific knowledge contained in physics-based models in two different ways: (a) by ingesting the output of physics- based models in the neural network framework, and

(b) by using a novel physics-based learning objective to ensure the learning of physically consistent predictions, as described in the following.

Ingesting Physics-based Model Simulations Since the meteorological observations of input drivers are only available at the surface of the lake but we need to estimate lake temperature at varying values of depth, we consider depth as another attribute in the

list of input variables, $\tilde{X}_{d,t}$. We further augment this

set of attributes using the simulated model outputs ofa generic physics-based model, Y^{Phy}, resulting in the

where n is the number of training instances.$_{d,t}$

Physics-based Learning Objective: Apart from minimizing training errors, we exploit an inter-esting physical relationship between the temperature, density, and depth of water in a lake, that serves as the basis of our physics-based learning objective used for training PGNN. In the following, we introduce the two key components of this physical relationship and de- scribe our approach for using it to ensure the learningof physically consistent predictions.

Temperature–Density Relationship: The temperature, Y, and density, ρ, of water are non-linearlyrelated to each other according to the following known physical equation [13]:

(3.7)

predictions of a PGNN model, can be used as a mea- sure of physical consistency of the PGNN model. Note that a data science model that is inconsistent with the known physical laws is likely to be learning specific and non-generalizable patterns from the training data (e.g., arising from noise in the training attributes as well as labels), in the pursuit of minimizing its training loss. When the sizes of both the training and test sets are small (as is common in many scientific problems), such subtle forms of overfitting may go unnoticed even af- ter using standard evaluation frameworks (e.g., cross- validation) and conventional regularizers based on sta- tistical notions of model complexity. The learned mod- els, when applied on novel unseen instances that were not adequately represented in the training and test sets,can then result in poor generalization performance, of- ten coming off a surprise [11].

In contrast to conventional learning objectives, since the known laws of physics are assumed to hold equally well for any unseen data instance, ensuring the physical consistency of model outputs as a learning ob- jective in PGNN can help in achieving better generaliza- tion performance even when the training data is small and not fully representative. Additionally, the output of a PGNN model can also be interpreted by a domain expert and ingested in scientific workflows, thus leading to scientific advancements. Note that in our particular problem of lake temperature modeling, even though the neural network is being trained to improve its accuracy on the task of predicting water temperatures, PGNNensures that the temperature predictions also translate

bonding between water molecules)[2]. This function is convex in the range of temperature values that thisis relationship holds, making it possible to compute its gradients and use then in the back-propagation algorithm. Given the temperature predictions of a

PGNN model at a given depth and time, $\hat{Y}_{d,t}$, we can

use Equation 3.7 to compute the corresponding density prediction, $\hat{\rho}_{d,t}$.

Density–Depth Relationship: The density of water monotonically increases with depth as shown inthe example plot of Figure 2(b). Formally, the densityof water at two different depths, d_1 and d_2, on the same time-step, t, are related to each other in the following manner:

to consistent relationships between other physical vari- ables, namely density and depth. Similar forms of rela- tionships can be leveraged in other scientific problems involving multiple physical variables that are related to each other, thus resulting in a wholesome solution to physical problems.

To compute a quantifiable measure of physical con- sistency using the density-depth relationship encoded in Equation 3.8, we consider the pair-wise differences be- tween density predictions at consecutive depths on the same time-step. In particular, we sort the depth values available for a certain time-step, t, in increasing order: $d_1 \ldots d_{n_t}$ (where n_t is the number of obervations on time-step t), and compute the consecutive differences as follows: pronged view of generalization performance (accuracy, simplicity, and physical consistency) introduced in Sec- tion 1. As the final loss function is differentiable almost everywhere, we use the backpropagation algorithm to compute and transmit gradients at the output and hid- den layers.

3 Results

Data We consider the example lake of Mille Lacs in Minnesota, USA, to evaluate and analyze the PGNN models presented in this paper. This is a reasonably large lake (536 km^2) that shows sufficient dynamics in the temperature profiles across depth over time, mak- ing it an interesting test case for analysis. Water tem- perature observations for our study lake were collated from a variety of sources including Minnesota Depart- ment of Natural Resources, and from a web resource that collates federal and state agencies, academic moni- toring campaigns, and citizen data [18]. These temper- ature observations vary across depths and time, with some years and seasons being heavily sampled, while other time periods having little to no observations. The overall data comprised of 10,954 temperature observa- tions from 29 June 1990 to 3 Jan 2016. For each ob- servation, we used a set of 11 meteorological drivers as input variables, listed in Table 1. While many of these drivers were directly measured, we also used some domain-recommended ways of constructing derived fea- tures such as Growing Degree Days [16].

We used the General Lake Model (GLM) [7] as the physics-based approach for modeling lake temperature in our experimental studies. The GLM uses the drivers listed in Table 1 as input parameters and balances the energy and water budget of lakes or reservoirs on a daily or sub-daily timestep. It performs a 1D modeling (along depth) of a variety of lake variables (including water temperature) using a vertical Lagrangian layer scheme. We used the uncalibrated GLM model outputs as additional input attributes in our PGNN framework, along with the measured drivers.

Experimental Setup: We considered contigu- ous time windows for constructing training and test sets from the overall data, to ensure that the test set is indeed independent of the training set and the two data sets are not temporally auto-correlated. Since dif- ferent years have different number of observations, we roughly chose 40% of the overall data for testing (29 Oct 2005 to 9 Dec 2012), and used the remainder time periods (corresponding to 60% of the data) for training (29 June 1990 to 27 Oct 2005 and 11 Dec 2012 to 3 Jan 2016). A portion of the training set (corresponding to the year 2005) was held out as the validation set for find- ing the right choice of neural network hyper-parameters: λ_1, λ_2, λ_{Phy} . We used a network architecture with 2 hidden layers, with 50 and 30 hidden nodes in the first and second hidden layers, respectively. We used the stochastic gradient descent (SGD) algorithm with abatch size of 100 to run the backpropagation algorithm, and performed a time-varying decay of the learning rateusing an initial value of 0.01 and a decay parameter of 10^{-3}.

Evaluation: We consider the following baselinemethods to compare with our proposed framework:

PHY: We compare our performance with the ini- tial physics-based model, termed as PHY, that was used as an input in the PGNN framework. Ex- ploring the differences in the model outputs of PGNN and PHY can shed light on the deficien- cies of the generic physics-based model, and high- light the promise in using deep learning in conjunc- tion with physics-based models to improve model- ing performance.

pureDS: In order to understand the importance of combining physics with deep learning methods, we consider the baseline method of learning neu- ral network architectures in a purely data-driven fashion. This would correspond to only using the meteorological observations $\tilde{X}$ as input attributes in the network, and using a learning objective thatonly contains the training loss, $Loss_{Tr}$ and the L_1 and L_2 regularizers. This model is being termed as pureDS.

PGNN$_0$: In order to understand the contribu- tion of the physics-based learning objective used in PGNN, we consider an intermediate product of our framework, PGNN$_0$, as another baseline, which makes use of the physics-based model simulations, Y^{Phy}, as input attributes in the network architec- ture, but does not use the physics-based loss func- tion, $Loss_{Phy}$, in the learning objective. Hence, PGNN$_0$ differs from pureDS in its use of physics- based model simulations as input attributes, and differs from PGNN in its use of a purely data-drivenlearning objective.

To compare the performance of PGNN with differ- ent baseline schemes, we considered the following two evaluation measures:

RMSE: We use the root mean squared error (RMSE) of a model on the test set as an estimate of its generalization performance. The units of this metric is in ◦C.

Inconsistency: Apart from ensuring generaliz- ability, a key contribution of PGNN is to ensure the learning of physically consistent model predic- tions. Hence, apart from the RMSE, we use another critical evaluation measure to check the physical consistency of the results. In particular, we count the number of time-steps that have a positive pair- wise difference among consecutive depths, $\Delta_{i,t}$, andreport the percentage of such time-steps as the In- consistency measure.

the smallest RMSE than all other baseline methods. Compared with the physics-based model, PHY, that was used as a starting point for building PGNN, we areable to reduce the RMSE from 2.57◦C to 1.16◦C, whichis a substantial improvement in the field of limnology. To appreciate the significance of a drop in RMSE of 1.41◦C, note that a lake-specific calibration approach

that produced a median RMSE of 1.47◦C over 28 lakes is considered to be the state-of-the-art in the field [2]. For

this specific lake, we also used a lake-specific calibration of the physics-based model, termed as PHY*, which was prepared by running the GLM model using different combinations of model parameters and choosing the model parameter that showed the lowest RMSE on the overall data. This fine tuned model showed an RMSE of 1.26◦C on the overall data. Note that while this RMSE cannot be considered as an unbiased estimate of the performance of PHY◦ (since the same data was used for calibrating the model as well as for computing RMSE), note that the test RMSE of PGNN is still lower thanthat of PHY*. This shows the promise in augmenting simple physics-based models using data science methods for improved modeling performance, reducing the need for performing expensive model calibrations.

Another highlight of the results in Table 2 is that PGNN not only achieves the lowest test RMSE, it is able to do so while incurring the lowest Inconsistency among all data science methods, thus representing a generaliz- able as well as physically consistent solution. Note that if we apply the black-box data science model, pureDS, we indeed are able to achieve a lower RMSE than the physics-based model, PHY. However, this improvement in RMSE is achieved at the cost of a large value of Incon-sistency in the model predictions of pureDS (almost 50%of the time-steps have inconsistent depth-density rela- tionships in its predictions). This makes the pureDS unfit for use in the process of scientific discovery, be- cause although it is able to somewhat improve the pre- dictions of the target variable (i.e. temperature), it is incurring large errors in capturing the physical relation- ships of temperature with other variables, leading to non-meaningful results.

When the neural network is fed with the output from the physics-based model, PHY, we can see that the performance of the resultant model, $PGNN_0$ im- proves in comparison with pureDS both with respect to RMSE as well as Inconsistency. This is because the output of PHY (although with a high RMSE) contains vital physical information about the dynamics of lake temperature, which when coupled with powerful data science frameworks such as deep learning, can result in major improvements in RMSE. Since the output of PHYis inherently designed to be physically consistent, using

PHY also helps in achieving a lower value of Inconsis- tency in $PGNN_0$. However, this value is still close to 20%, which is considerably high from an operational perspective. It is only by the use of physics-based loss functions that we can achieve not only a lower RMSE than $PGNN_0$, but a substantially lower Inconsistency. Hence, the framework of PGNN shows promise in im- proving both the physical consistency as well as gener- alizability of the model predictions.

Conclusions and Future Work

This paper presented a novel framework for learning physics-guided neural networks (PGNN), by using the outputs of physics-based model simulations as well asby leveraging physical relationships to enforce scientific consistency on the neural network predictions. By an- choring deep learning methods with scientific knowl- edge, we are able to show that the proposed frame- work not only generates physically meaningful results, but also helps in achieving better generalizability than black-box data science methods.

We anticipate this paper to be the first stepping stone in the broader theme of research on using physics- based learning objectives in the training of data science models. While the specific formulation of PGNN ex- plored in this paper was developed for the example prob- lem of modeling lake temperature, similar developmentscould be explored in a number of other scientific and engineering disciplines where known forms of physical relationships can be used to guide the learning of data science models to physically consistent solutions. This paper paves the way towards learning neural networks by not only improving their ability to solve a given task, but also being cognizant of the physical relationships of the model outputs with other tasks, thus producing a more holistic view of the physical problem.

There are a number of directions of future research that can be explored as a continuation of this work.First, for the specific problem of lake temperature mod- eling, given the spatial and temporal nature of the prob- lem domain, a natural extension would be to exploit the spatial and temporal dependencies in the test instances,e.g., by using recurrent neural network based architec- tures. Second, the analysis of the physically consistent model predictions produced by PGNN could be usedto investigate the modeling deficiencies of the baseline physics-based model in detail. Finally, while this pa-per explored the use of physical relationships between temperature, density, and depth of water in the learn- ing of multi-layer perceptrons, other forms of physical relationships in different neural network models can be explored as future work. Of particular value would beto develop generative models that are trained to not only capture the structure in the unlabeled data, but are also guided by physics-based models to discover andemulate the known laws of physics. The paradigm of PGNN, if effectively utilized, could help in combining the strengths of physics-based and data science models,and opening a novel era of scientific discovery based on both physics and data.

References

C. Anderson. The end of theory. *Wired magazine*,16(7):16–07, 2008.

X. Fang, S. R. Alam, H. G. Stefan, L. Jiang,

P. C. Jacobson, and D. L. Pereira. Simulations of water quality and oxythermal cisco habitat inminnesota lakes under past and future climate scenarios. *Water Quality Research Journal*, 47(3- 4):375–388, 2012.

J. Friedman, T. Hastie, and R. Tibshirani. *The elements of statistical learning*, volume 1. Springer series in statistics Springer, Berlin, 2001.

D. Graham-Rowe, D. Goldston, C. Doctorow,

M. Waldrop, C. Lynch, F. Frankel, R. Reid, S. Nel- son, D. Howe, S. Rhee, et al. Big data: science in the petabyte era. *Nature*, 455(7209):8–9, 2008.

H. V. Gupta and G. S. Nearing. Debatesthe fu- ture of hydrological sciences: A (common) pathforward? using models and data to learn: A sys- tems theoretic perspective on the future of hydro- logical science. *Water Resources Research*, 50(6): 5351–5359, 2014.

T. D. Harris and J. L. Graham. Predict- ing cyanobacterial abundance, microcystin, and geosmin in a eutrophic drinking-water reservoir us-ing a 14-year dataset. *Lake and Reservoir Manage- ment*, 33(1):32–48, 2017.

M. Hipsey, L. Bruce, and D. Hamilton. Glmgeneral lake model: Model overview and user information. *Perth (Australia): University of Western AustraliaTechnical Manual*, 2014.

T. Jonathan, A. Gerald, et al. Special issue: dealing with data. *Science*, 331(6018):639–806, 2011.

A. Karpatne, G. Atluri, J. H. Faghmous, M. Stein-bach, A. Banerjee, A. Ganguly, S. Shekhar,

N. Samatova, and V. Kumar. Theory-guided data science: A new paradigm for scientific discovery from data. *IEEE Transactions on Knowledge and Data Engineering*, 29(10):2318–2331, 2017.

U. Lall. Debatesthe future of hydrological sciences: A (common) path forward? one water. one world. many climes. many souls. *Water Resources Re-search*, 50(6):5335–5341, 2014.

D. Lazer, R. Kennedy, G. King, and A. Vespignani.

The Parable of Google Flu: Traps in Big DataAnalysis. *Science (New York, N.Y.)*, 343(6176): 1203–5, Mar. 2014. ISSN 1095-9203. doi: 10.

1126/science.1248506. URL http://www.ncbi.nlm.nih.gov/pubmed/24626916.

J. J. Magnuson, L. B. Crowder, and P. A. Medvick.Temperature as an ecological resource. *American Zoologist*, 19(1):331–343, 1979.

J. L. Martin and S. C. McCutcheon. *Hydrodynam- ics and transport for water quality modeling*. CRCPress, 1998.

J. J. McDonnell and K. Beven. Debatesthe fu- ture of hydrological sciences: A (common) path for-ward? a call to action aimed at understanding ve- locities, celerities and residence time distributions of the headwater hydrograph. *Water Resources Re- search*, 50(6):5342–5350, 2014.

H. W. Paerl and J. Huisman. Blooms like it hot.

Science, 320(5872):57–58, 2008.

I. C. Prentice, W. Cramer, S. P. Harrison, R. Lee- mans, R. A. Monserud, and A. M. Solomon. Special paper: a global biome model based on plant phys- iology and dominance, soil properties and climate. *Journal of biogeography*, pages 117–134, 1992.

F. J. Rahel and J. D. Olden. Assessing the effects of climate change on aquatic invasive species. *Con- servation biology*, 22(3):521–533, 2008.

[1] E. K. Read, L. Carr, L. De Cicco, H. A. Dugan,

P. C. Hanson, J. A. Hart, J. Kreft, J. S. Read, and L. A. Winslow. Water quality data for national- scale aquatic research: The water quality portal. *Water Resources Research*, 53(2):1735–1745, 2017.

J. S. Read, L. A. Winslow, G. J. Hansen, J. Van Den Hoek, P. C. Hanson, L. C. Bruce, and C. D. Markfort. Simulating 2368 temperate lakes reveals weak coherence in stratification phenology. *Ecolog-ical modelling*, 291:142–150, 2014.

J. J. Roberts, K. D. Fausch, D. P. Peterson, and M. B. Hooten. Fragmentation and thermal risks from climate change interact to affect persistence of native trout in the colorado river basin. *GlobalChange Biology*, 19(5):1383–1398, 2013.

J. J. Roberts, K. D. Fausch, M. B. Hooten, and D. P. Peterson. Nonnative trout invasions com- bined with climate change threaten persistence of isolated cutthroat trout populations in the south- ern rocky mountains. *North American Journal of Fisheries Management*, 37(2):314–325, 2017.

T. J. Sejnowski, P. S. Churchland, and J. A. Movshon. Putting big data to good use in neu- roscience. *Nature neuroscience*, 17(11):1440–1441,2014.

P.-N. Tan, M. Steinbach, and V. Kumar. *Intorduc- tion to Data Mining*. Addison-Wesley, 2005.

V. N. Vapnik and V. Vapnik. *Statistical learningtheory*, volume 1. Wiley New York, 1998.

