Science-guided Machine Learning (Part 1):

Overview of Research Themes and Guiding Principles

Anuj Karpatne

Assistant Professor, Computer Science

Virginia Tech

karpatne@vt.edu

https://people.cs.vt.edu/karpatne/

SANGHANI CENTER FOR ARTIFICIAL INTELLIGENCE & DATA ANALYTICS VIRGINIA TECH

PGML Lab Physics-guided Machine Learning

Golden Age of Machine Learning / Artificial Intelligence

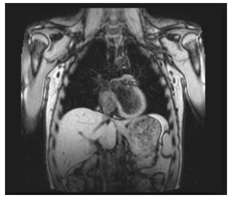


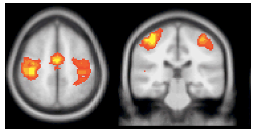



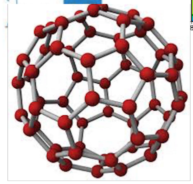
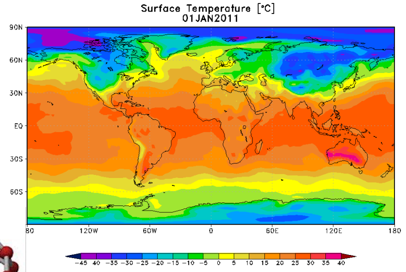

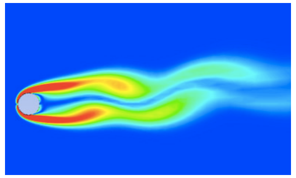


- Hugely successful in commercial applications:

IMAGENET

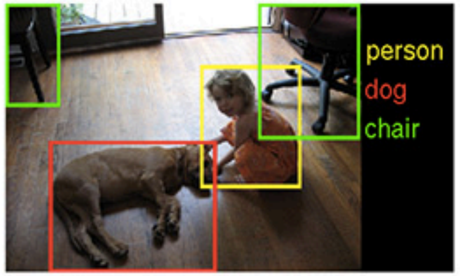

DeepMind

Google Ads

facebook

NETFLIX

Google AI algorithm masters ancient game of Go

Golden Age of Machine Learning / Artificial Intelligence

- Promise of Machine Learning (ML) in Accelerating Scientific Discovery

Will the rapidly growing area of **“black-box”** ML models make existing theory-based models obsolete?

- But disappointing results in scientific domains!
  - Require lots of labeled data
  - Unable to provide valuable physical insights

Science **The Parable of Google Flu: Traps in Big Data Analysis**

- Predicted flu using Google search queries
- Overestimated by twice in later years

Science-based vs. Data-based Models

- Scientific Rules and Equations

$$\frac{\partial \rho}{\partial t} = -\nabla.(\rho \mathbf{u})$$

$$\frac{\partial \rho \mathbf{u}}{\partial t} = -\nabla.\left(\frac{1}{\rho}(\rho\mathbf{u})\otimes(\rho\mathbf{u}) + p\mathbf{I}\right) + \rho\mathbf{g}$$

$$\frac{\partial E}{\partial t} = -\nabla.\left(\frac{1}{\rho}(E+p)(\rho\mathbf{u})\right) + \mathbf{u}.\rho\mathbf{g}$$

$$\mathbf{H}\Psi = E\Psi$$

Contain knowledge gaps in describing certain processes (turbulence, groundwater flow)

- Computational Models of Dynamical Systems

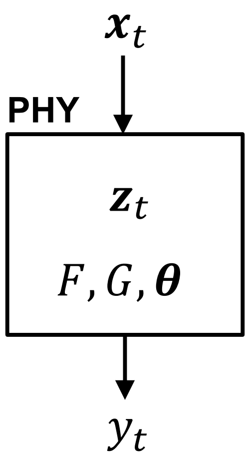

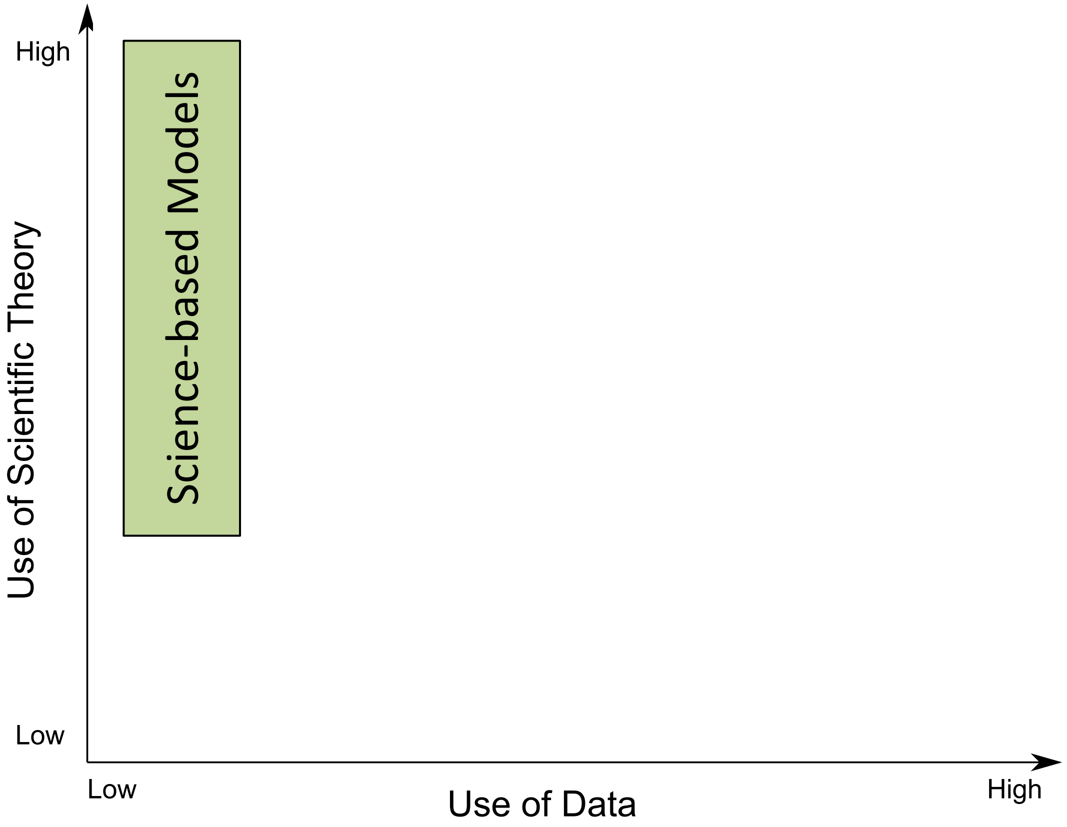

Limitations of Science-based Models

- Large number of parameters/states
- Incomplete or missing physics / process knowledge

Science-based vs. Data-based Models

Contain knowledge gaps in describing certain processes (turbulence, groundwater flow)

Take full advantage of data science methods without ignoring the treasure of accumulated knowledge in scientific "theories"

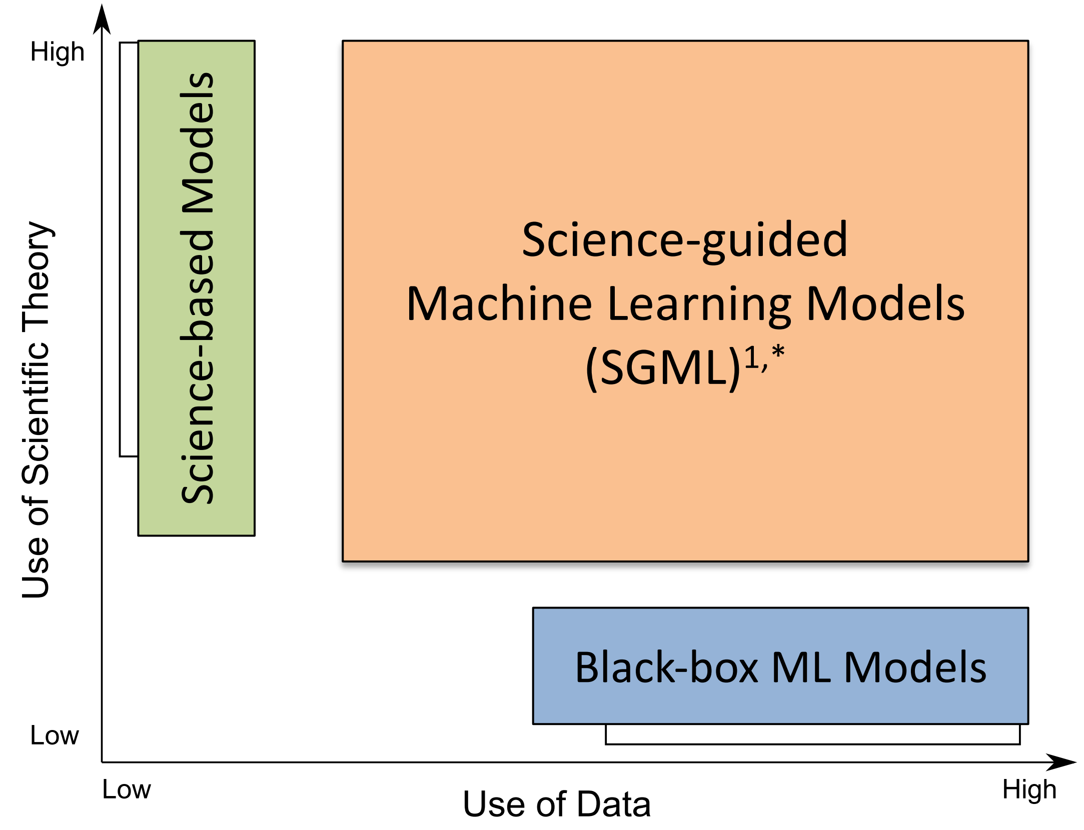

[1] Karpatne et al. "**Theory-guided data science: A new paradigm for scientific discovery**," TKDE 2017

Require large number of representative samples

Science-based vs. Data-based Models

*Work on this topic has been referred to by various names such as:

- Knowledge-guided ML
- Science-guided ML
- Physics-guided ML
- Physics-informed ML / Physics-informed NN
- Physics-aware AI
- Theory-guided Data Science

In these works, “**physics**” or "**physics-guided**" should be more generally interpreted as “**science”** or “**scientific knowledge”.**

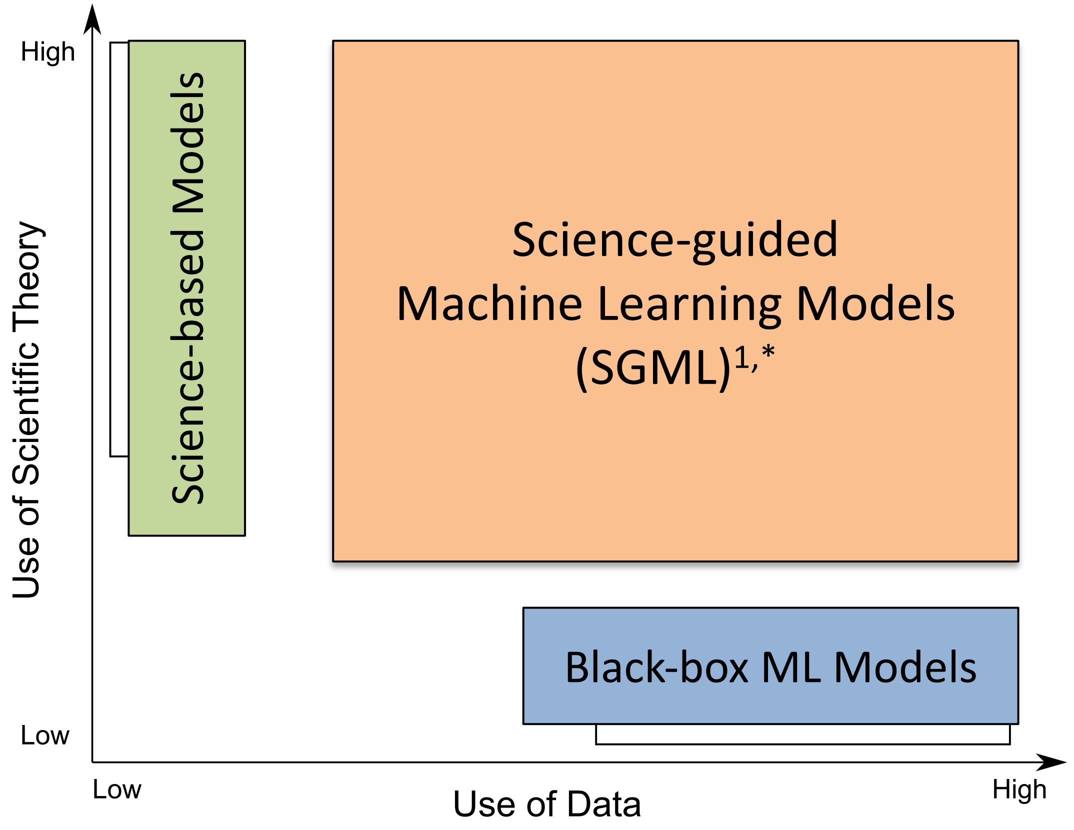

[1] Karpatne et al. “**Theory-guided data science: A new paradigm for scientific discovery**,” TKDE 2017

Require large number of representative samples

Recent Developments in SGML

Defense Advanced Research Projects Agency › Program Information

Physics of Artificial Intelligence (PAI)

DARPA

The Physics of Artificial Intelligence (PAI) program is part of a broad DAPRA initiative t
and adversarial spoofing, and that incorporate domain-relevant knowledge through gen

It is anticipated that AI will play an ever larger role in future Department of Defense (Do
processing, to control and coordination of composable systems. However, despite rapi
subfield of machine learning – AI's successful integration into numerous DoD applicatio
development of causal, predictive models and dealing with incomplete, sparse, and noi

To facilitate better incorporation of AI into DoD systems, the PAI program is exploring n
physics, mathematics, and prior knowledge relevant to DoD application domains. PAI a
will help to overcome the challenges of sparse data and will facilitate the development

CCC Computing Community Consortium Catalyst

Catalyzing the computing research community and enabling the pursuit of innovative, high-impact research.

ABOUT VISIONING LEADERSHIP DEVELOPMENT TASK FORCES RESOURCES EVENTS BLOG CCC BY

Visioning Activity

Artificial Intelligence Roadmap

In fall 2018, the Computing Community Consortium (CCC) initiated an effort to create a 20-Year Roadmap for Artificial Intelligence, led by Yolanda Gil (University of Southern California and President of AAAI) and Bart Selman (Cornell University and President-Elect of AAAI). The goal of the initiative was to identify challenges, opportunities, and pitfalls in the AI landscape, and to create a compelling report to inform future decisions, policies, and investments in this area.

The Roadmap was based on broad community input gathered via a number of forums and communication channels: three topical workshops during the fall and winter of 2018/2019, a Town Hall at the annual meeting of the AAAI, and feedback from other groups of stakeholders in industry, government, academia,

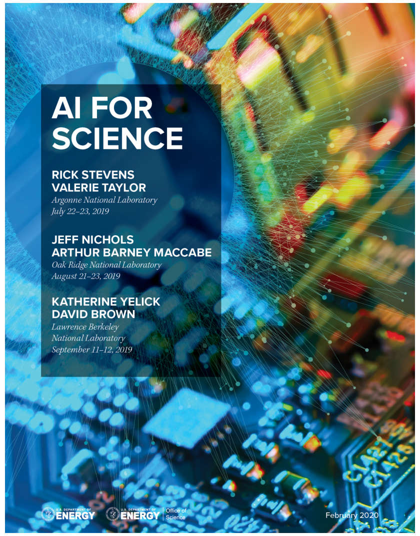

Report on DOE Town halls on "AI for Science"

Many conferences/workshops

- 2020 AAAI Fall Symposium on Physics-guided AI
- 2020 and 2021 AAAI Spring Symposium on ML in Physical Sciences

Physics-Informed Learning Machines for Multiscale and Multiphysics Problems

Pacific Northwest NATIONAL LABORATORY

Machine Learning for Physics and the Physics of Learning

ipam

PHYSICS INFORMED MACHINE LEARNING

Workshop by Los Alamos National Laboratory, 2016, 2018, 2020

Integrating Physics-Based Modeling With Machine Learning: A Survey

JARED WILLARD* and XIAOWEI JIA*, University of Minnesota
SHAOMING XU, University of Minnesota
MICHAEL STEINBACH, University of Minnesota
VIPIN KUMAR, University of Minnesota

Surveys more than 300 papers

https://arxiv.org/pdf/2003.04919.pdf

Guiding Principles of SGML

- **How can Science help ML?**

- **How can ML advance Science?**

Guiding Principles of SGML

- **How can Science help ML?**
  - Guide the learning of AI models to *scientifically consistent* solutions
  - Ensure *generalizability* even when training data is limited

$$\text{Generalization Error} \propto \text{Training Error} + \text{Complexity} + \textbf{Scientific Inconsistency}$$

- **How can ML advance Science?**

Guiding Principles of SGML

- **How can Science help ML?**
  - Guide the learning of AI models to *scientifically consistent* solutions
  - Ensure *generalizability* even when training data is limited

$$\text{Generalization Error} \propto \text{Training Error} + \text{Complexity} + \textbf{Scientific Inconsistency}$$

- **How can ML advance Science?**
  - Discover new scientific laws, model parameters from data
  - Augment or replace components of science-based models

Research Themes in SGML

Science helps ML

ML advances Science

1. Science-guided Design
    - Choice of Response Function
    - Design of Model Architecture
    - …

2. Science-guided Learning
    - Using Loss Functions, Constraints, Priors, Training Labels
    - …

3. Science-guided Refinement
    - Post-processing
    - Pruning
    - …

4. Discovery of Scientific Laws from Data
    - Symbolic Regression, Autoencoders, …

5. Inferring Parameters in Science-based Models
    - Model Calibration, Inverse Modeling, Data Assimilation, …

6. Hybrid-Science-ML Modeling
    - Residual Modeling, Augmenting system components using ML, Pretraining, …

Research Theme 1:

Science-guided Model Design

- Choice of Model Architecture Governed by Scientific Knowledge

Encoding invariances/symmetries in ANN architecture

- Rotational/translational invariance in ANN structure
- Symmetries linked to conservation of physical quantities

Kondor et al. arXiv 2018; Ling et al. *JFM, 2016*

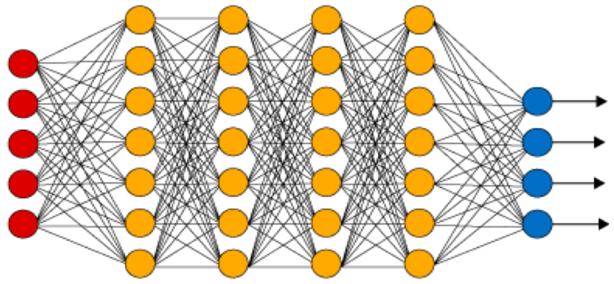

Hard-wiring physics in LSTM connections

- Application in Lake Temperature Modeling
- Provides better uncertainty quantification when used with MC Dropout, while ensuring physical consistency

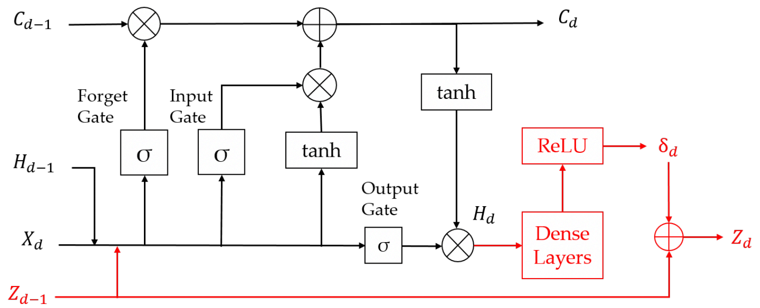

Daw et al. SDM 2020

Using Physical Intermediates at Hidden Layers of NN

- Application in Fluid Dynamics: Drag Force Prediction

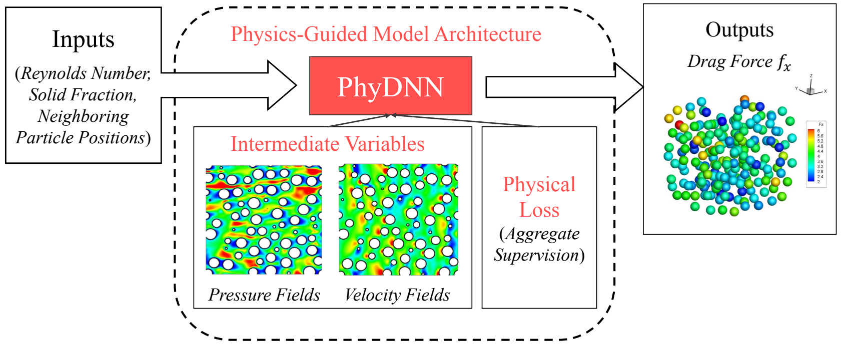

Muralidhar et al. SDM 2020

Research Theme 1:

Science-guided Model Design Summary

- Requires knowledge of:
  - what intermediate features should be expressed at hidden nodes/layers
  - high-level properties of ANN architecture such as invariances
- Related to the field of “explainable/trustworthy/interpretable” AI
- Questions:
  - Are there known physical pathways from inputs to outputs in your problem where the intermediates can be modeled in the hidden layers?
  - What kind of symmetries/structures do you have in your application domain that can be encoded?

Research Theme 2:

Science-guided Model Learning

- Traditionally, “simpler” models are preferred for generalizability
  - Basis of several statistical principles such as bias-variance trade-off

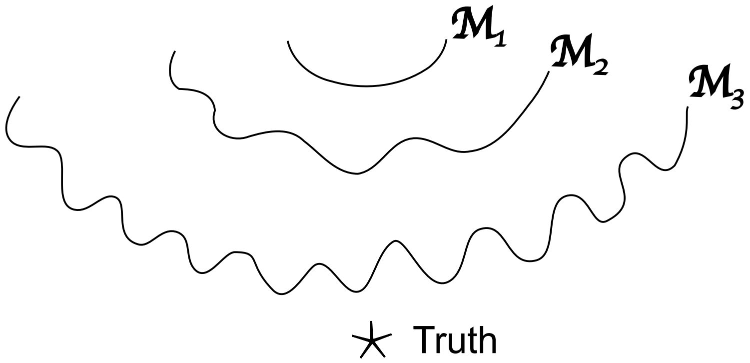

$\mathcal{M}1$ (less complex model):
High bias—Low variance

$\mathcal{M}3$ (more complex model):
Low bias—High variance

Generalization Performance $\propto$ Accuracy + Simplicity

Research Theme 2:

Science-guided Model Learning

- Traditionally, "simpler" models are preferred for generalizability
  - Basis of several statistical principles such as bias-variance trade-off

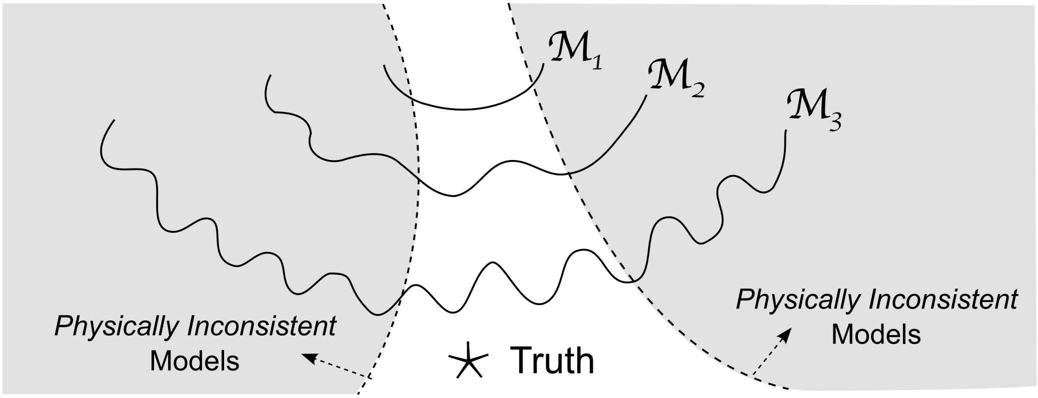

$\mathcal{M}_1$ (less complex model): High bias—Low variance

$\mathcal{M}_3$ (more complex model): Low bias—High variance

- In scientific problems, **"scientific consistency"** can be used as another measure of generalizability
  - Can help in pruning large spaces of inconsistent solutions
  - Result in generalizable *and* scientific meaningful results

Generalization Performance $\propto$ Accuracy + Simplicity + ***Consistency***

Research Theme 2:

Science-guided Learning

- Learning algorithm ensures the selection of *scientifically consistent* models

Generalization Performance $\propto$ Accuracy + Simplicity + ***Consistency***

- Methods:
  - **Physics-guided Loss Functions / Regularizers**
  - Physics-guided Priors
  - Physics-guided Constraints
  - Physics-guided Initialization
  - …

Research Theme 2:

Science-guided Learning Examples

Physics-guided Loss Functions in ANN Learning

$$\text{Prediction Loss }(y, \hat{y}) \; + \; \lambda\, \mathrm{R}(\theta) + \; \lambda_{\mathrm{PHY}}\ \textbf{Physics-guided Loss }(\hat{y})$$

D_{Tr} $\qquad$ D_U

Advantages:

- PG-Loss can be applied even on unlabeled data
- ANN results are scientifically consistent and generalizable even in paucity of labels

Physics-guided NN for lake temperature modeling

- Incorporate density-depth physics and energy conservation as PG loss (along with pred-loss)

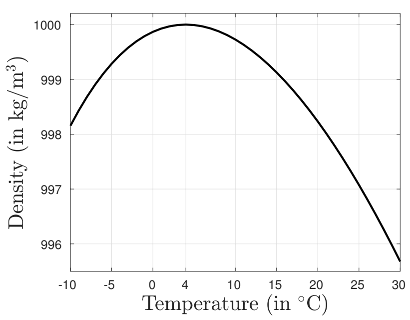

Karpatne et al. *arXiv* 2017

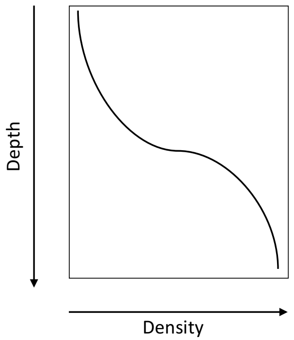

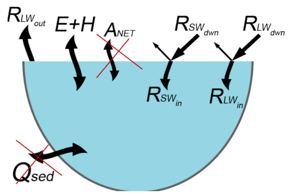

Jia et al. *SDM* 2019.

Physics-informed NN for Solving PDEs

- Label-free learning only using PG loss
- PDE loss computed using Automatic differentiation

Raissi et al. *arxiv* 2019

Physics-guided learning in quantum science with competing physics objectives

- Adaptive tuning of physics loss in objective function (over training epochs) leads to better generalizability

Bu et al. *Arxiv* 2020.

Research Theme 2:

Science-guided Learning Summary

$$\text{Prediction Loss } (y, \hat{y}) + \lambda \, \mathrm{R}(\theta) + \lambda_{\mathrm{PHY}} \, \textbf{Physics-guided Loss } (\hat{y})$$

D_{Tr} D_U

- Questions
  - What kinds of scientific knowledge can be used as loss functions (or constraints, priors, …) in your work?
  - How can we compute or evaluate PG-loss during model training?
  - What is the trade-off between minimizing prediction loss and PG-loss?
  - How many labeled examples do we need in D_{Tr} and how representative should D_U be?
  - Can we work in multi-physics problems with competing PG loss functions?

Research Theme 3:

Science-guided Refinement

- ML results on test set can be post-processed (pruned/improved) using domain theories

Example in Material Science: Discovery of novel materials

Hautier et al. 2010, Fischer et al. 2006, Curtarolo et al. 2013

- Computationally expensive (DFT) methods used for checking material properties
- Prune discovered materials using theory-based models
- **Discovery of hundred new ternary oxides**

Post-processing using science-based models

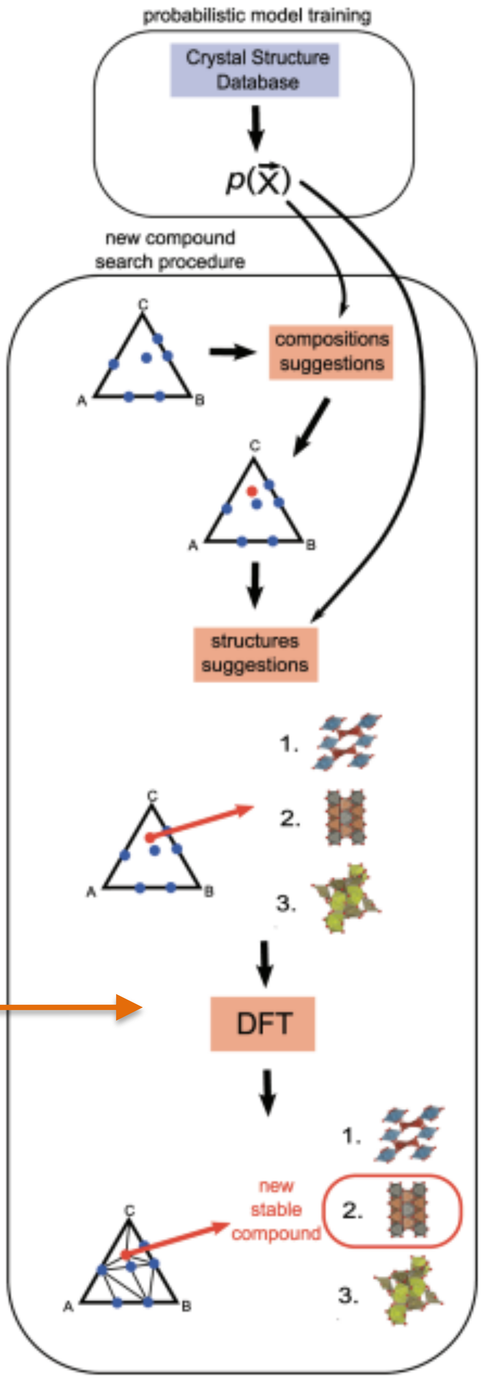

Hautier et al. 2010

Research Theme 3:

Science-guided Refinement Summary

- Questions:
  - Which ML areas can we apply science-guided refinement (e.g., in predictive learning, generative modeling, …)?
  - How can you reduce the computational costs of post-processing by ensuring ML results are scientifically consistent in the first place?

Research Theme 4:

Discovery of Scientific Laws from Data

- Learn an ML model to discover the governing equations of a scientific process from data using a **parsimonious** representation
  - Use known library of operators common in scientific community

- Common techniques: sparse regression, symbolic regression, autoencoders, ….

Research Theme 4:

Discovery of Scientific Laws from Data Examples

Sparse Identification of Nonlinear Dynamics (SINDy) Brunton et al. 2015

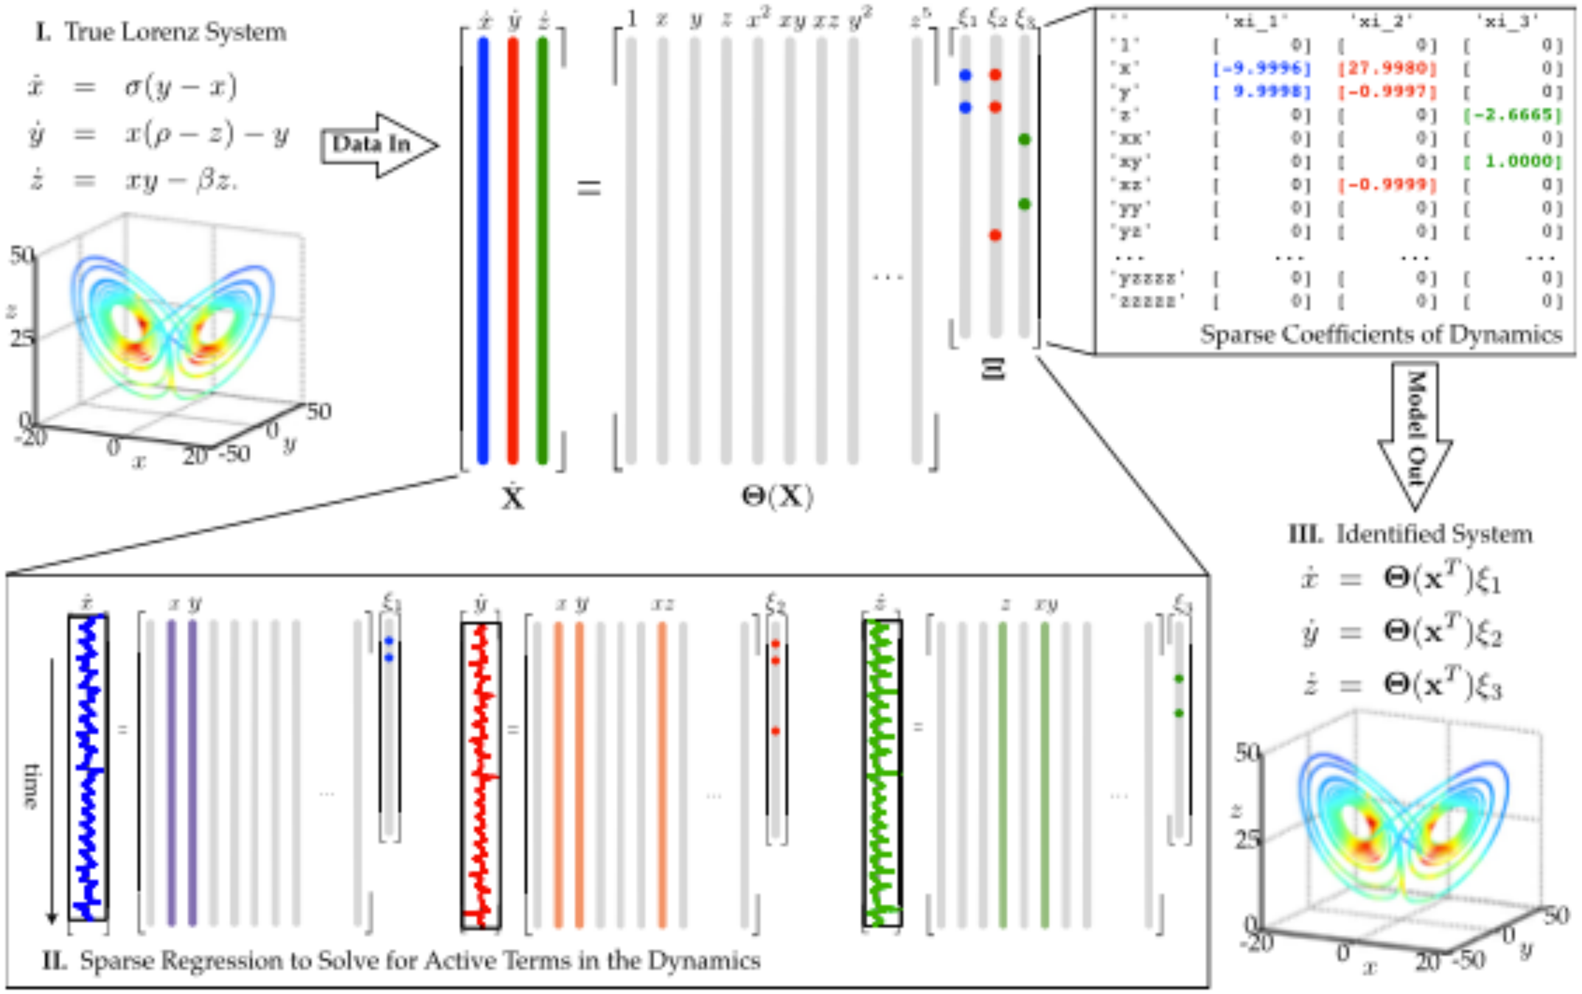

Research Theme 4:

Discovery of Scientific Laws from Data Summary

- Questions:
  - Given two parsimonious representations of a data, how do we favor one over other?
  - How can we find the right library/basis set of operators to use in a problem?
  - How can these methods handle noisy data in real-world problems?

Research Theme 5:
Inferring Parameters in Science-based Models

- Use ML methods to learn latent parameters (static or time-dependent) in numerical models of complex systems
- Techniques in this area include inverse modeling, data assimilation, system identification, …

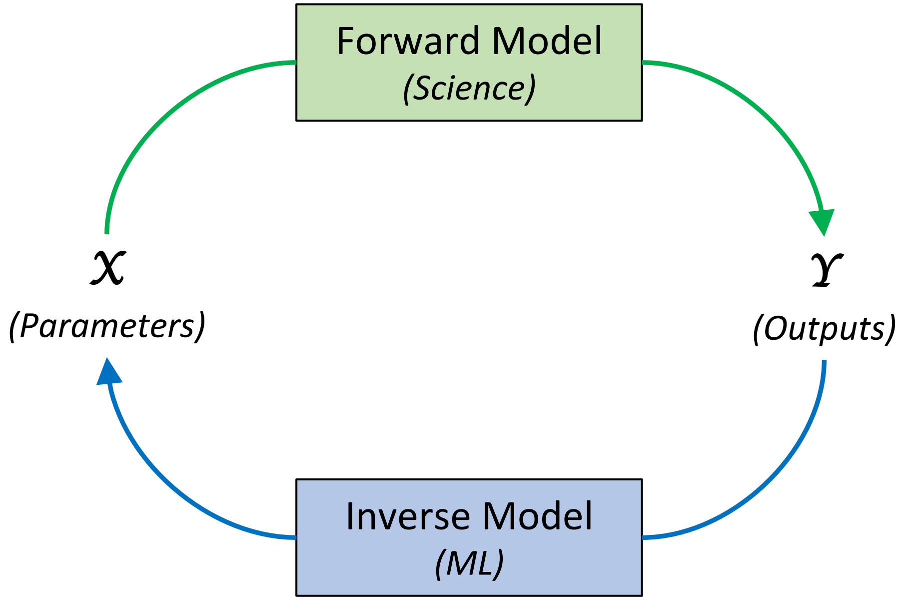

Illustrative Applications:

- Geosciences
- Seismology
- Material Science
- Imaging
- …

Research Theme 5:
Inferring Parameters in Science-based Models Summary

- Questions:
  - How can we inform the inverse ML model with the knowledge encoded in the forward science model?
  - How can inverse ML models dynamically update themselves as the distribution of forward model outputs change during deployment?

Research Theme 6:
Hybrid-Science-ML Modeling

- Components of science-based models can be augmented or replaced by ML models

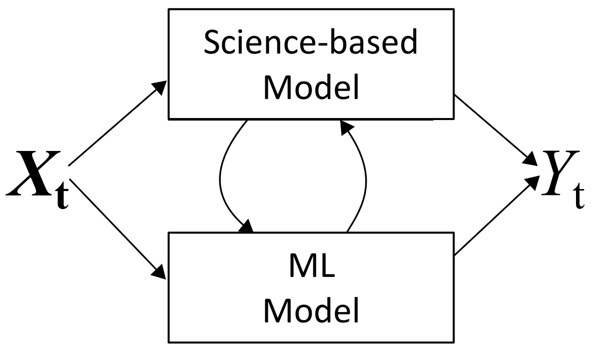

- ML methods can be used to predict:
  - Patterns of residuals (errors) in model outputs
  - Intermediate modeling quantities that are not fully known

Applications:
- Geosciences
- Climate Science
- Turbulence Modeling

Research Theme 6:

Hybrid-Science-ML Modeling Summary

- Questions:
  - How can we identify missing/incomplete components/quantities in science-based models that can be replaced by ML models?
  - How can we ascertain the source of modeling errors (inaccurate parameters or inherent biases)?

Alternate Categorization of SGML

Courtesy of Willard, Jared, Xiaowei Jia, Shaoming Xu, Michael Steinbach, and Vipin Kumar. "Integrating physics-based modeling with machine learning: A survey." arXiv preprint arXiv:2003.04919 (2020).

SGML Methods (ways of integrating science + ML)

SGML Objectives

	Basic ML	Physics-guided Loss Function	Physics-guided Initialization	Physics-guided Architecture	Residual Model	Hybrid Model
Improving predictions beyond physical model						
Parameterization						
Reduced Order Models						
Downscaling						
Uncertainty Quantification						
Inverse Modeling						
Discovering Governing Equations						
Solve PDEs						
Data Generation						

Alternate Categorization of SGML

Courtesy of Willard, Jared, Xiaowei Jia, Shaoming Xu, Michael Steinbach, and Vipin Kumar. "Integrating physics-based modeling with machine learning: A survey." arXiv preprint arXiv:2003.04919 (2020).

SGML Methods (ways of integrating science + ML)

SGML Objectives

	Basic ML	Physics-guided Loss Function	Physics-guided Initialization	Physics-guided Architecture	Residual Model	Hybrid Model
Improving predictions beyond physical model						
Parameterization						
Reduced Order Models						
Downscaling						
Uncertainty Quantification						
Inverse Modeling						
Discovering Governing Equations						
Solve PDEs						
Data Generation						

SGML Method: Physics-guided Initialization

- Initializing neural network weights (pre-training) using physics
  - Pre-train ANN model using physical simulations (approximate but cheap)
  - Fine-tune pretrained model on ground-truth (accurate but costly)

- Question: How many physical simulations are sufficient to seed the pre-training of ANN models to generalizable solutions?

SGML Objective:
Improving predictions beyond physical model

- Given a physical model mappings inputs to outputs that is either:
  - *Imperfect* due to incomplete/approximate physics
  - *Expensive* to run at required operational scales
- Goal:
  - Build a surrogate ML model to augment or emulate the physical model

SGML Objective: Solve PDEs

- Given the governing equations of a system in terms of PDEs where:
  - Initial/boundary conditions and equation parameters are known, but
  - Solving PDEs directly using Finite Difference Methods is computationally expensive

- Goal:
  - Build a surrogate ML model to predict PDE solutions as outputs, using initial/boundary conditions and parameters as inputs
  - Optionally use ground-truth solutions of exact PDE solvers as labeled supervision

SGML Objectives:
Data Generation and Uncertainty Quantification

- Given input-output data simulated by a physics-based model:
  - Involving high-dimensional spaces and stochastic processes
- Goal:
  - Learn a generative ML model to produce a similar distribution of data as those simulated by the physics-based model
  - Produce uncertainty estimates of outputs given inputs using ML models

SGML Objectives: Parameterization, Downscaling, and Reduced Order Modeling

- Parameterization:
  - Replacing some components of complex physical models (that are inexact or too expensive to run) using ML-based parameterized approximations
  - Need to ensure parameterizations are not overly complex (otherwise, can loose interpretability)

- Downscaling:
  - Using ML models to predict fine-scale variables as outputs, using coarse-scale physical simulations as inputs

- Reduced Order Modeling (ROM):
  - Learn approximate but fast representations of complex models using dimensionality reduction techniques

	Basic ML	Physics-Guided Loss Function	Physics-Guided Initialization	Physics-Guided Architecture	Residual Model	Hybrid Model
Improve prediction beyond physical model	[287] [35] [102] [104] [271]	[247] [4] [199] [139][128] [160] [183] [66] [74] [119] [153] [182] [213] [295] [109] [113] [165] [79] [285]	[158] [276] [162] [248] [31] [230] [239] [114] [119] [213] [276]	[152] [56] [17] [238] [6] [11] [183] [182] [193] [196] [221] [113] [260] [154] [293] [232] [292] [289] [291] [246] [228]	[247] [275] [258] [223] [257] [271] [153]	[95] [99] [222] [105] [128] [236] [50] [156] [191] [284] [294] [255] [280] [97] [165]
Parameterization	[198] [157] [203] [135] [169] [95] [43] [90] [186] [212] [29] [33] [189]	[292] [23] [24] [285]		[21] [23]		[294]
Reduced Order Models	[244] [179] [101] [273] [225] [286] [173]	[188] [10] [148]		[125] [262] [171] [74] [188] [190]	[125] [223] [224] [257] [101] [190]	[57]
Downscaling	[237] [254] [81] [232]					
Uncertainty quantification	[275] [142] [250] [283]	[270] [281] [266] [282] [301] [89]	[276] [162]	[58] [277] [266]		[67]
Inverse modeling	[59] [47] [161] [197]	[210] [122]		[27] [77] [240] [132]	[111]	[192] [105] [121] [241] [39] [38] [68] [229]
Discover Governing Equations	[30] [227] [37] [167] [168] [216] [123] [200][205] [52] [141] [290]	[207] [155]		[251]		
Solve PDEs	[147] [65] [140] [9] [215] [42] [264] [106] [235] [107] [131]	[209] [234] [279] [281] [60] [177] [208] [231] [301] [71] [89] [129] [268] [194] [176] [16]		[48] [218] [44] [60] [172] [235] [19] [51] [131] [76] [278] [194] [176] [206]		[166]
Data Generation	[187] [78]	[62] [40] [28] [133] [268] [288] [298]		[49] [274] [231]		

What is Coming Up Next?

- Next Class:
  - Case Studies, Recent Progress, and Future Prospects