∂³AWN
∂ata-∂riven Atmospheric & Water ∂yNamics

The Abdus Salam International Centre for Theoretical Physics

Towards Reliable Hybrid AI-Climate Modeling

COLLABORATIVE WORK

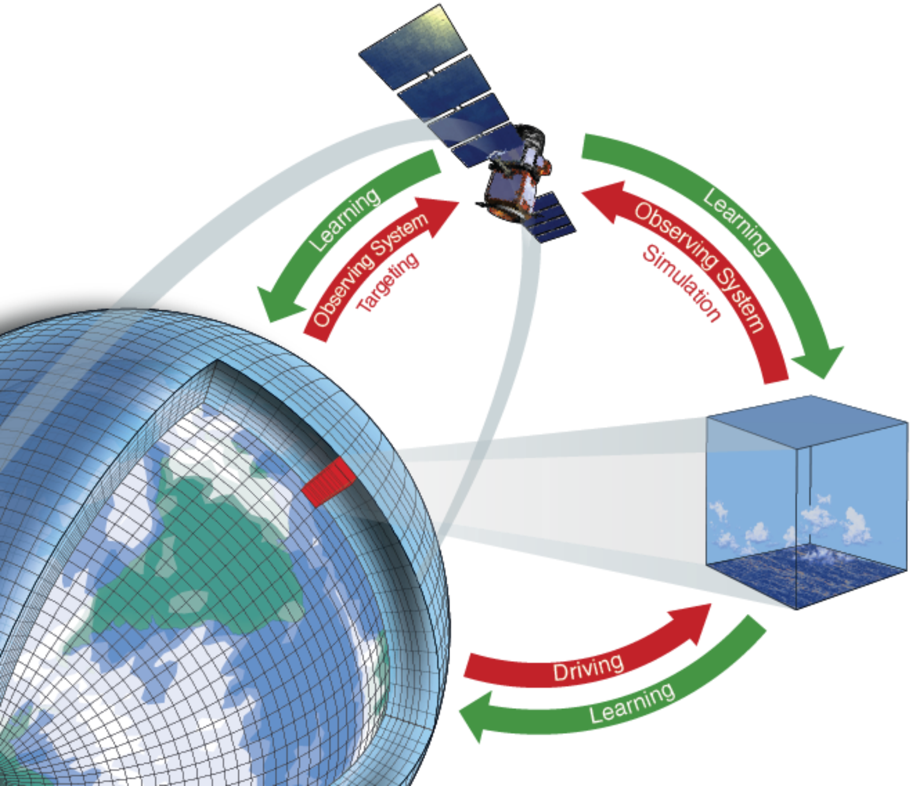

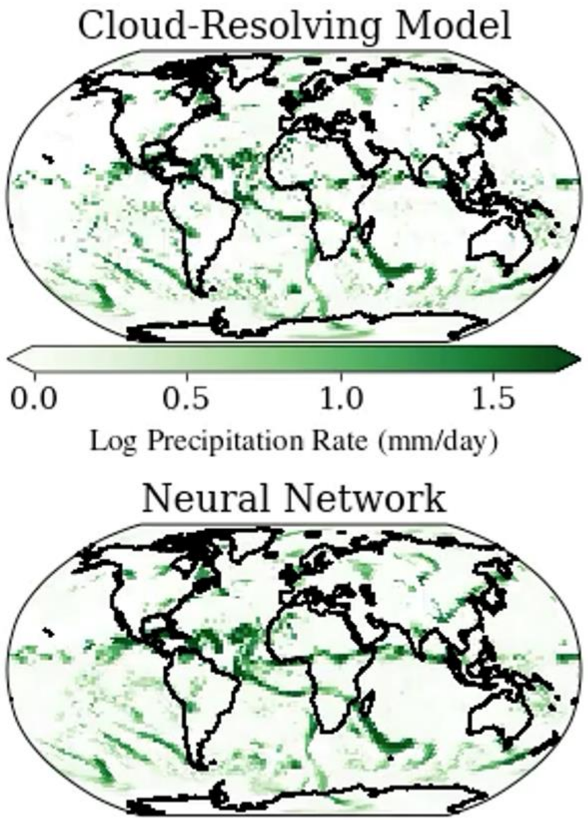

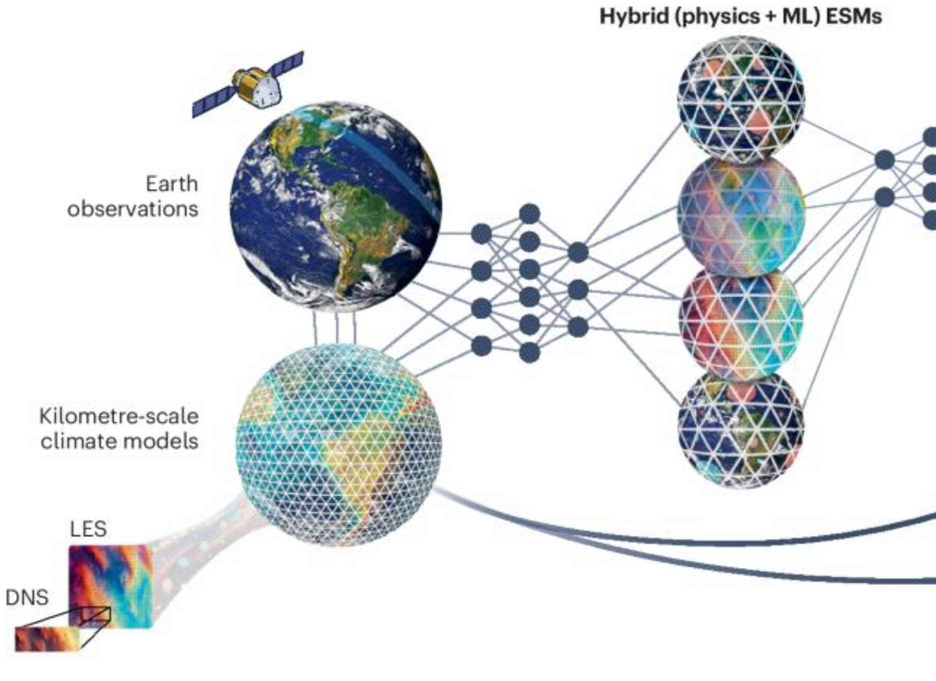

Image credits: *Climate Modeling Alliance*
Mooers et al. (2021); Eyring et al. (2024)

∂^3AWN

∂ata-∂riven
Atmospheric & Water
∂yNamics

UNIL | Université de Lausanne

Expertise Center for
Climate Extremes (ECCE)

Swiss National
Science Foundation

Project funded by

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Atmospheric Physics + AI

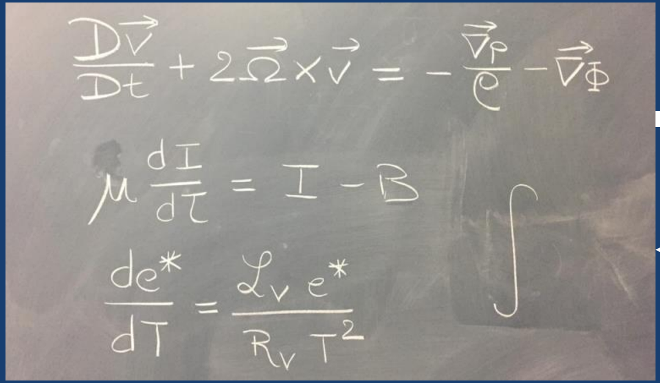

Physics-Guided ML

Data-Driven Discovery

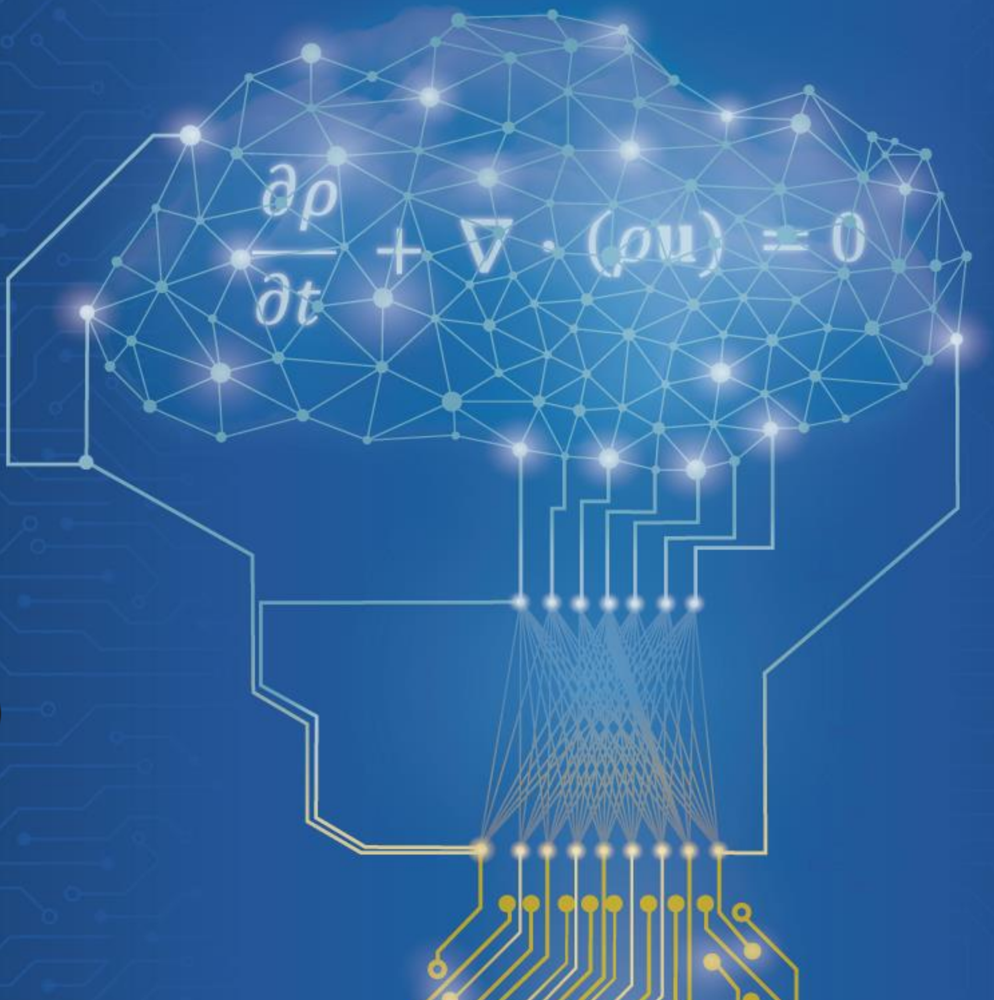

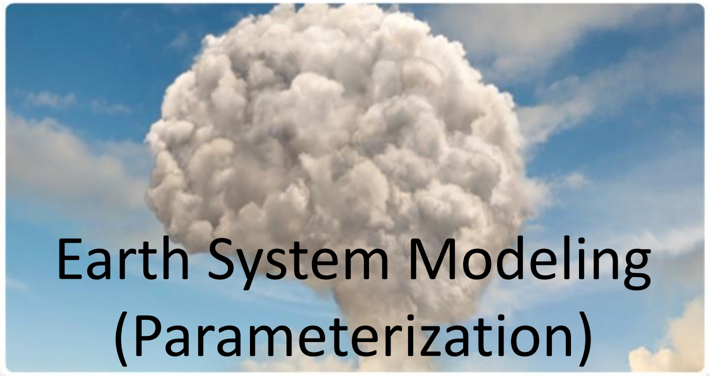

<u>Machine Learning</u> = Learning task from data without being explicitly programmed for task

<u>Neural Network</u> = Non-linear regression tool

High-dimensional Inputs **x**

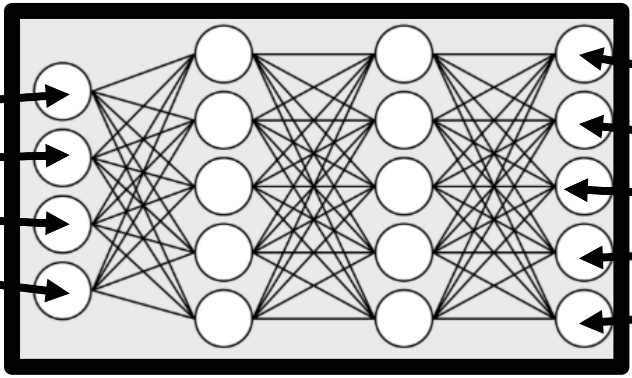

High-dimensional Outputs **y**

$$\boldsymbol{x} \mapsto \boldsymbol{y}$$

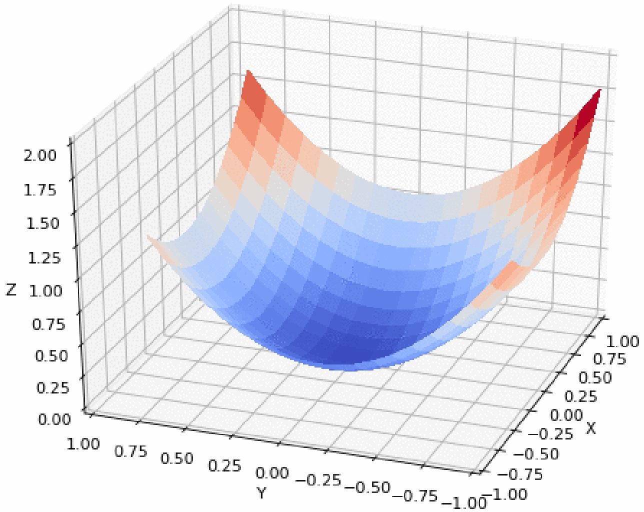

$$\min \ \text{Loss function}\left(\boldsymbol{y}_{\textbf{Predicted}}, \boldsymbol{y}_{\textbf{Truth}}\right)$$

⚠ Lots of data

<u>Image source</u>: Kathuria (Paperspace)

Largest uncertainties in long-term atmospheric projections linked to subgrid cloud processes

Processes that are not explicitly resolved by Earth System models must be **parameterized**

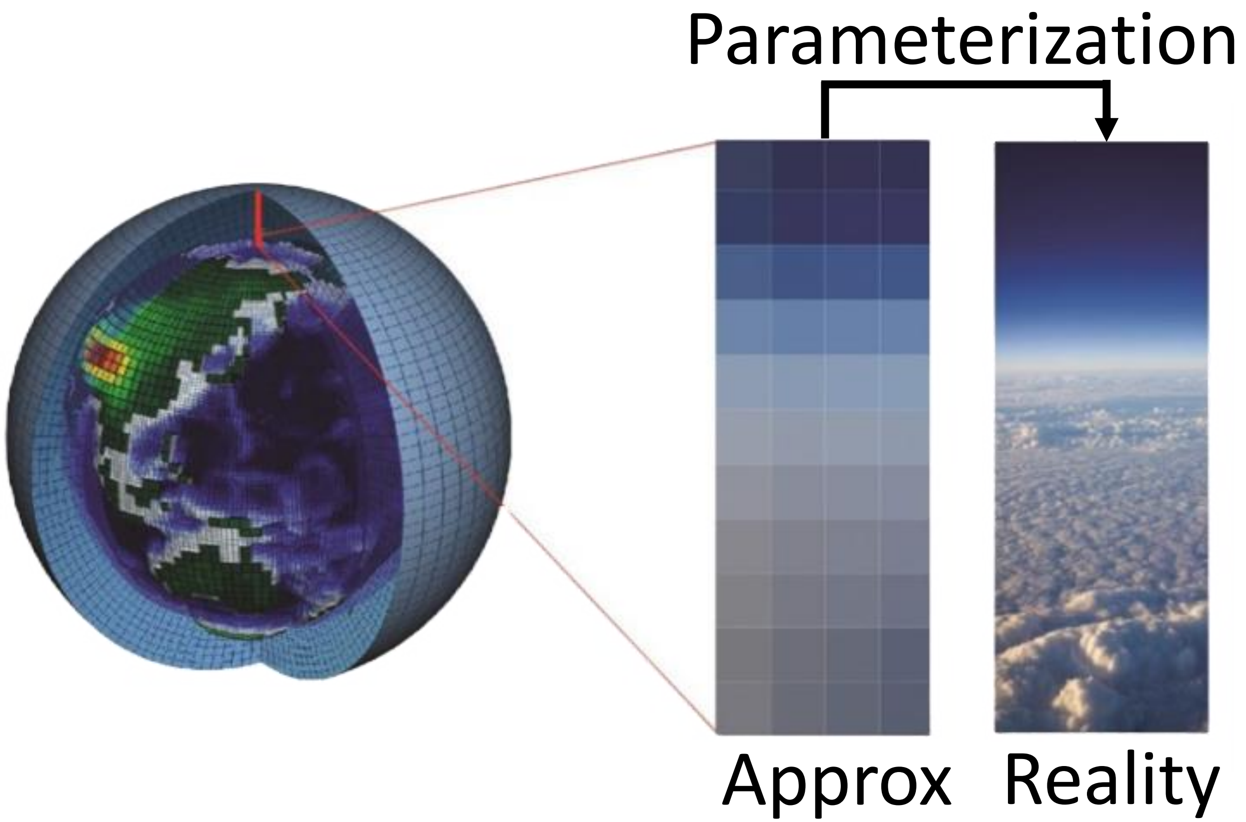

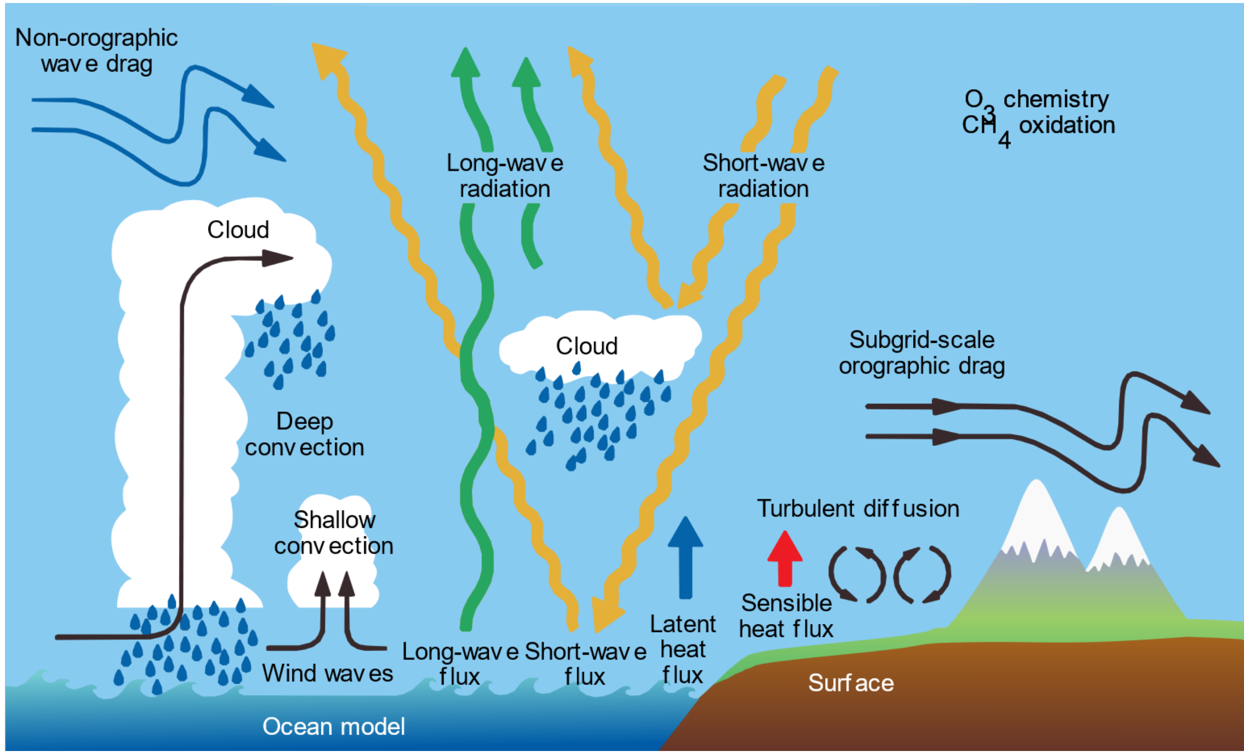

Schematic source: ECMWF

References: Humphrey et al. (2008), Mass et al. (2002), Zelinka et al. (2020), Boucher et al. (2014), Sherwood et al. (2014), Gentine et al. (2020)

We can machine learn a **parameterization** from data → **emulation** of observations and/or models

Grid-scale information

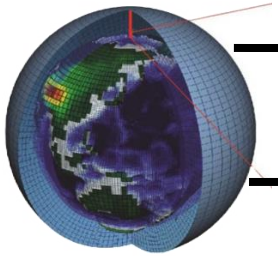

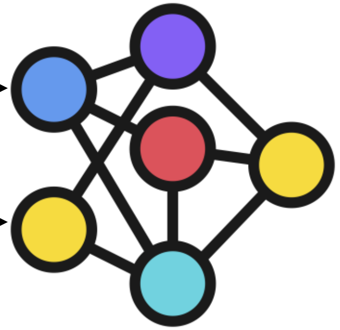

ML parameterization

Target process

Traditional parameterizations

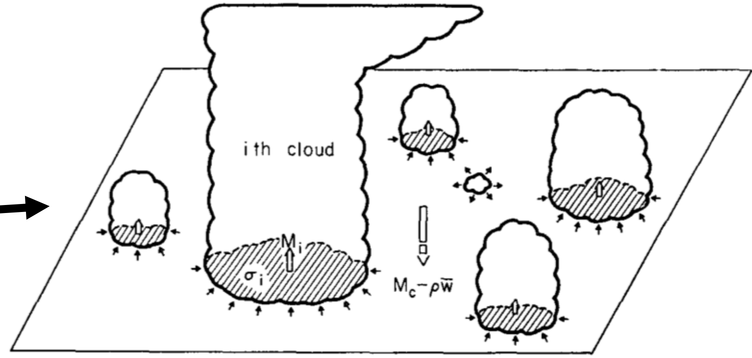

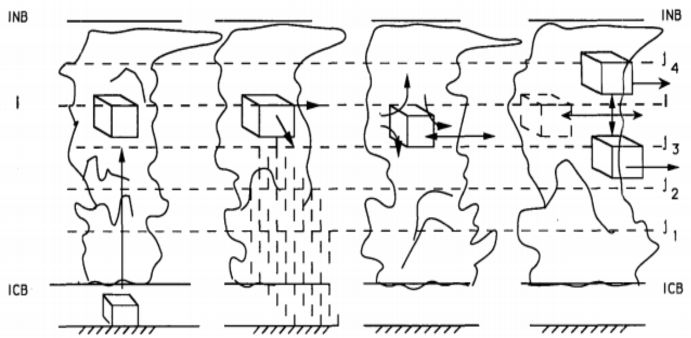

References: *Gentine et al. (2020), Bretherton (schematic), Arakawa & Schubert (1974), Emanuel (1999)*

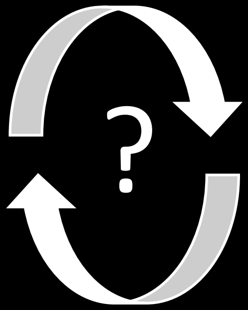

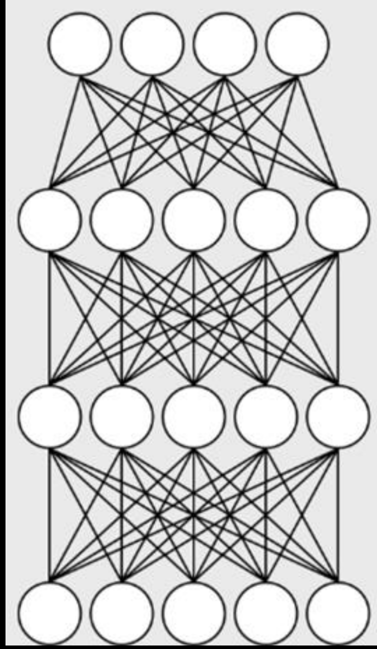

<u>Video source</u>: Global Storm-Resolving and Large-Domain Large-Eddy Simulations with ICON LEM. Deutsches Klimarechenzentrum
<u>Image source</u>: Pierre Gentine (LEAP)

Coarse graining enables learning the aggregate effects of subgrid clouds and turbulence, too costly to simulate in routine climate projections

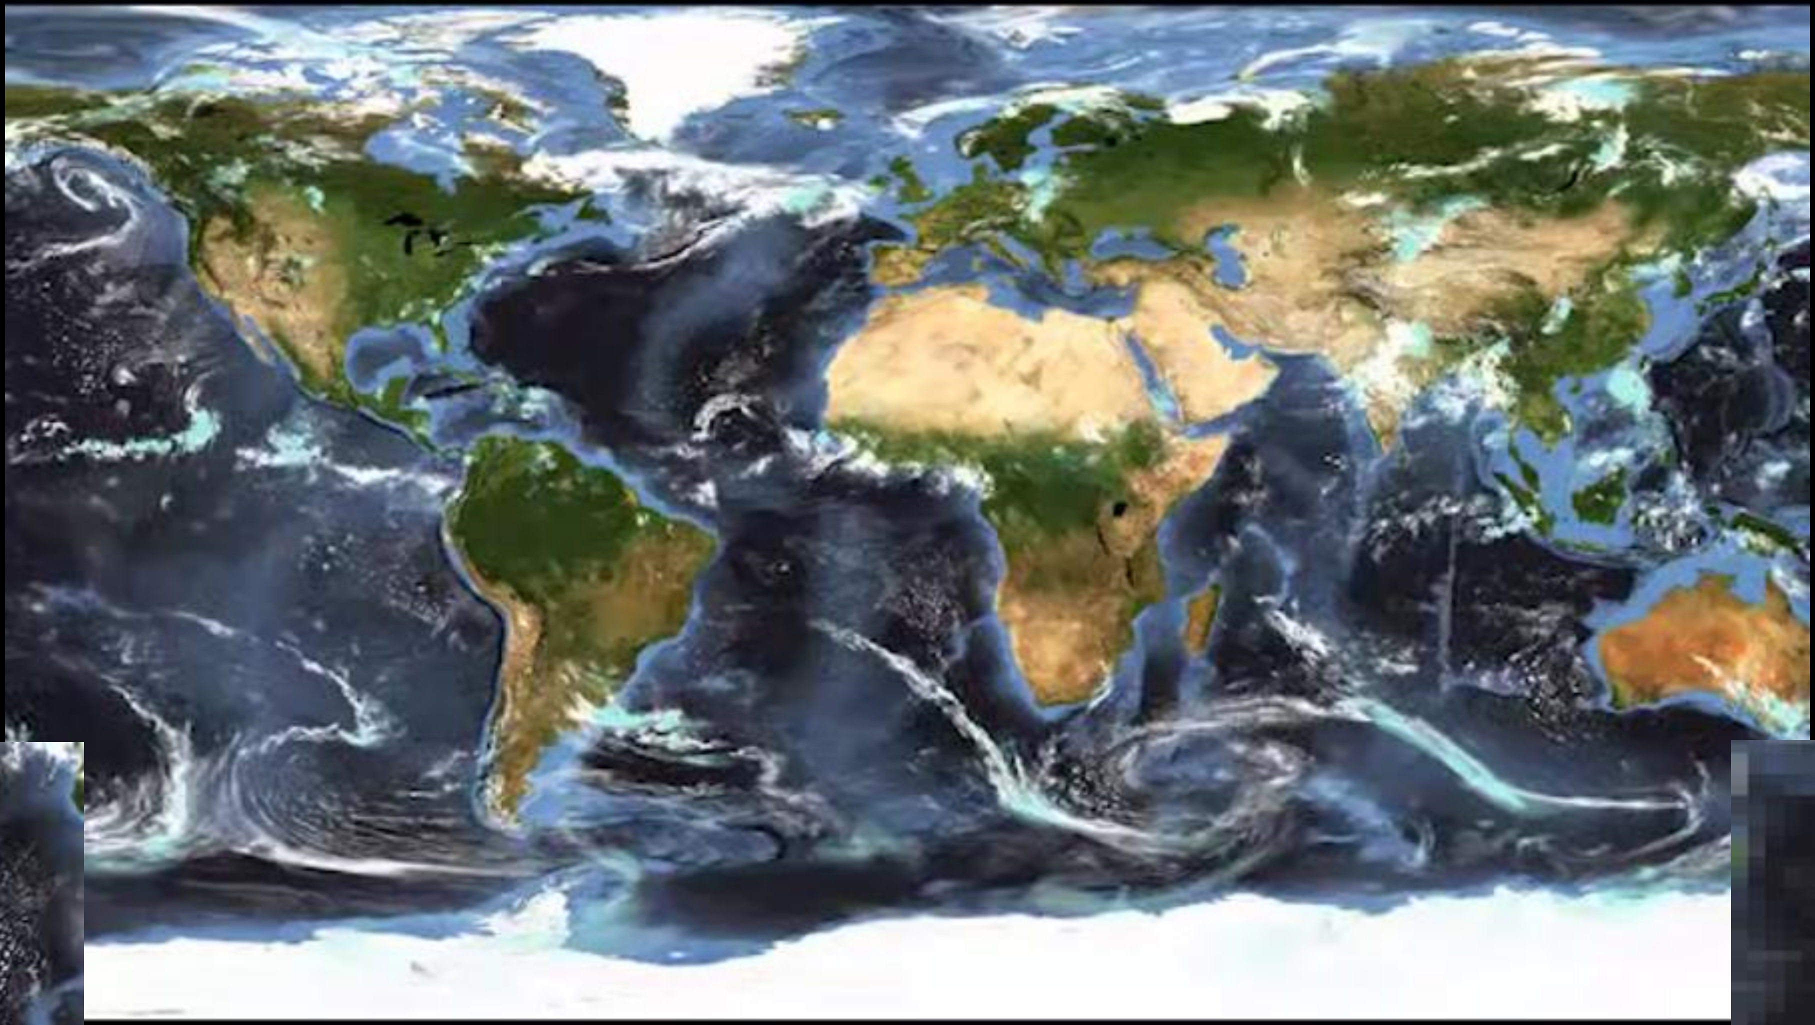

Coarse-Graining

<u>Video source</u>: P. Gentine (LEAP)

Neural Networks can emulate subgrid tendencies with high accuracy in realistic geography configurations

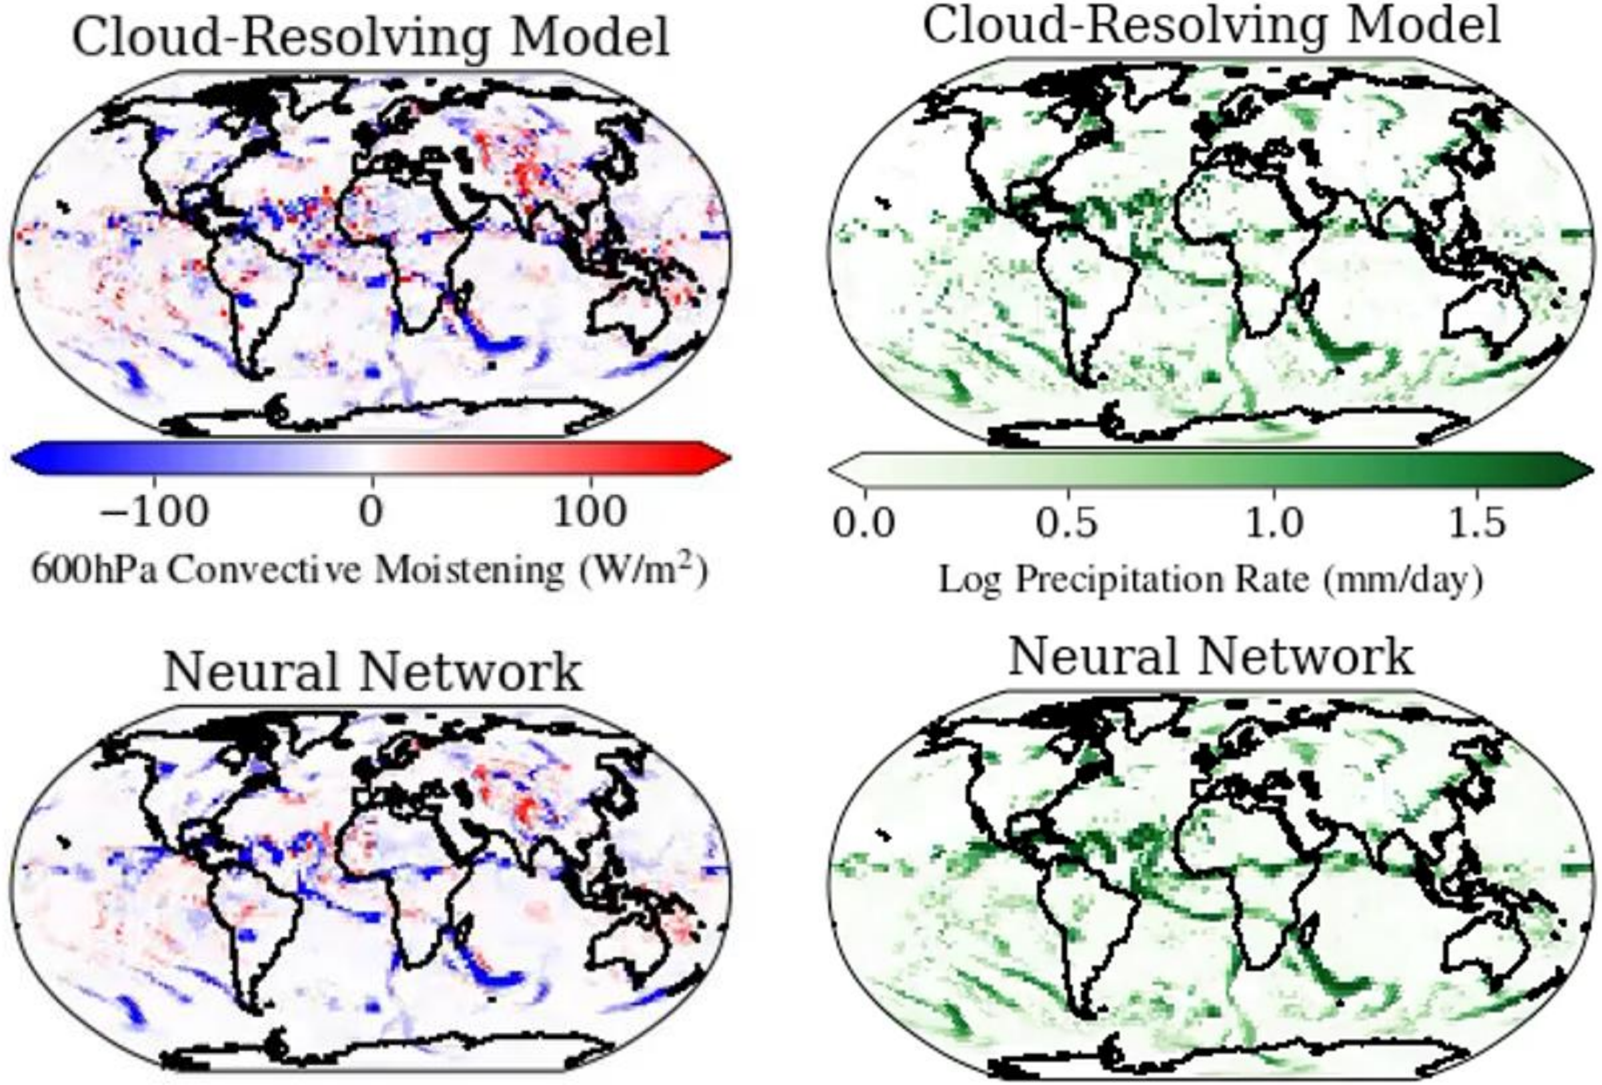

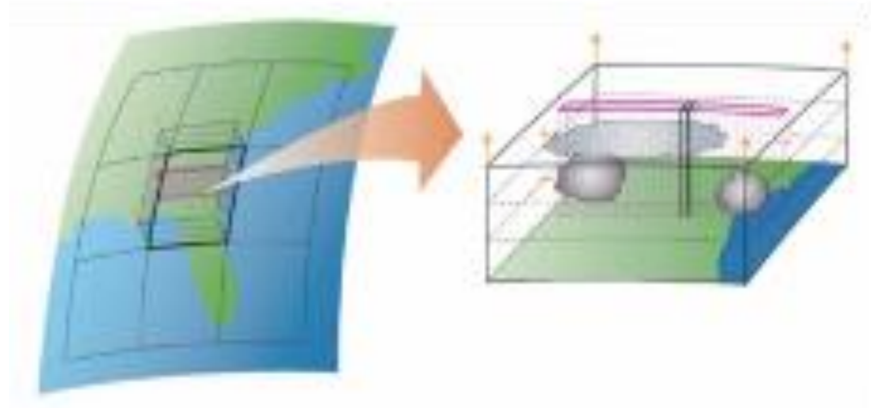

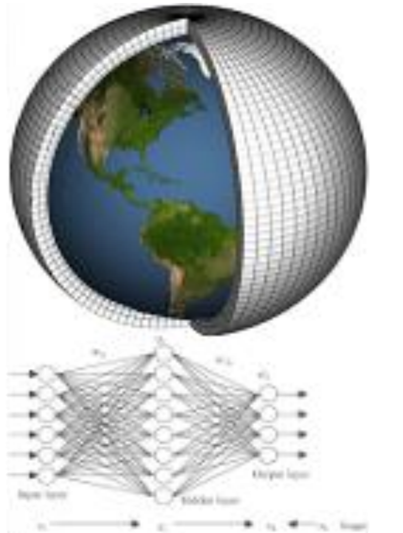

Video source:
Griffin Mooers (MIT)
Article:
Mooers et al. (2021)

Data-Driven Parameterizations are flourishing…

Online Function Learning with Ensemble Kalman Inversion

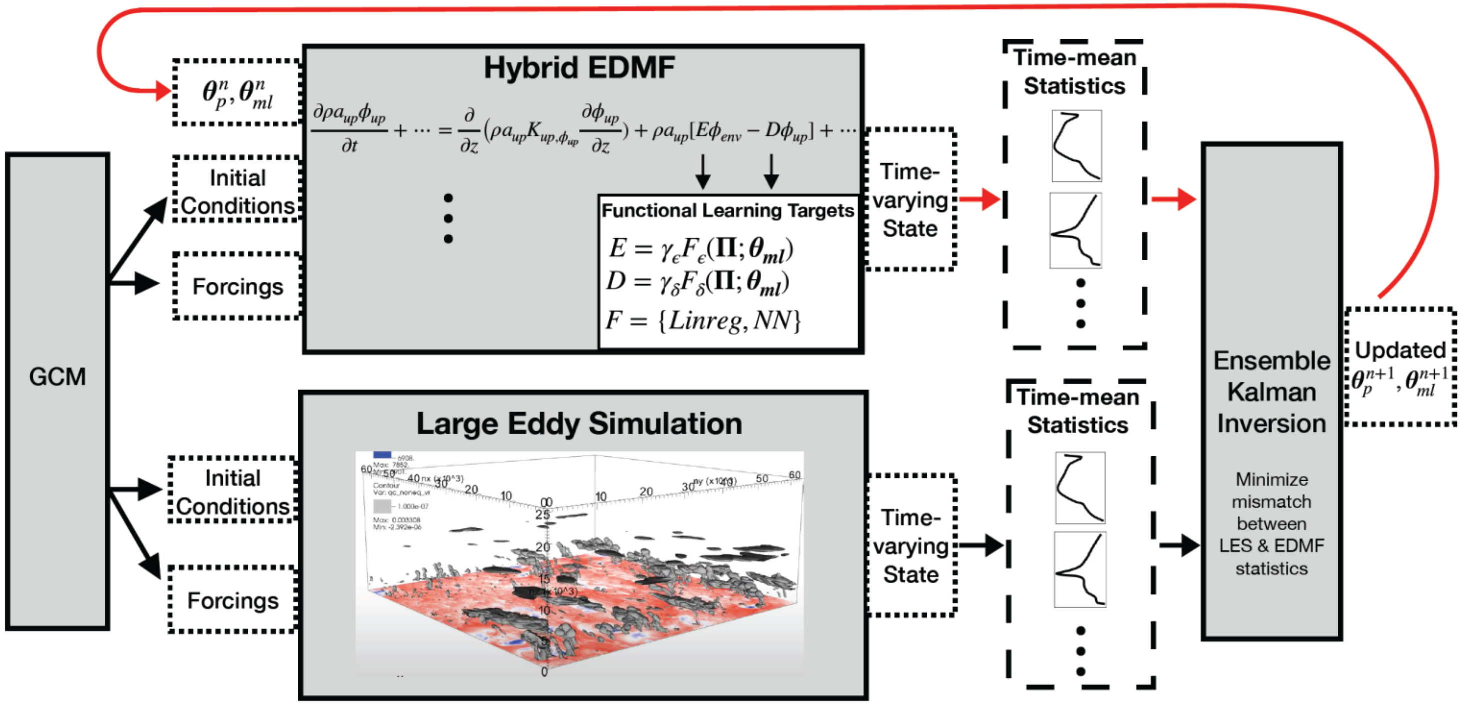

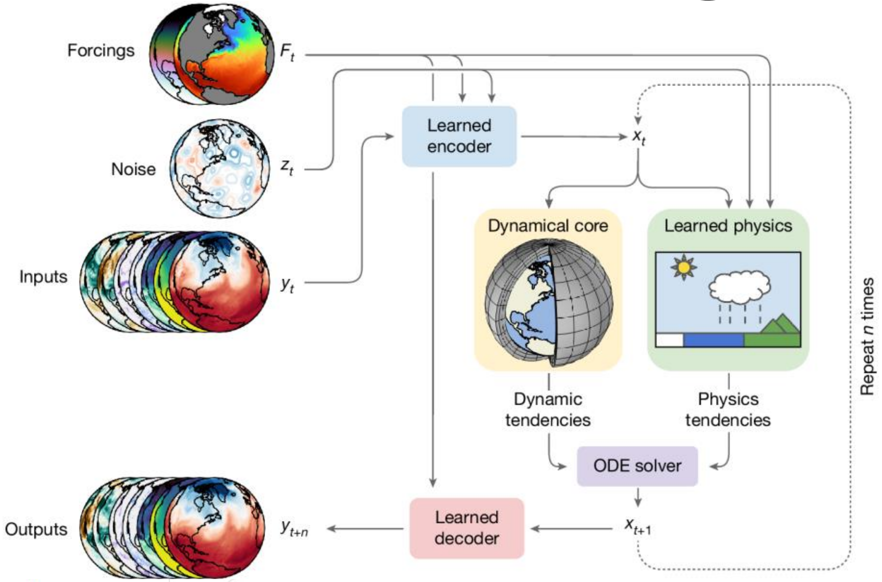

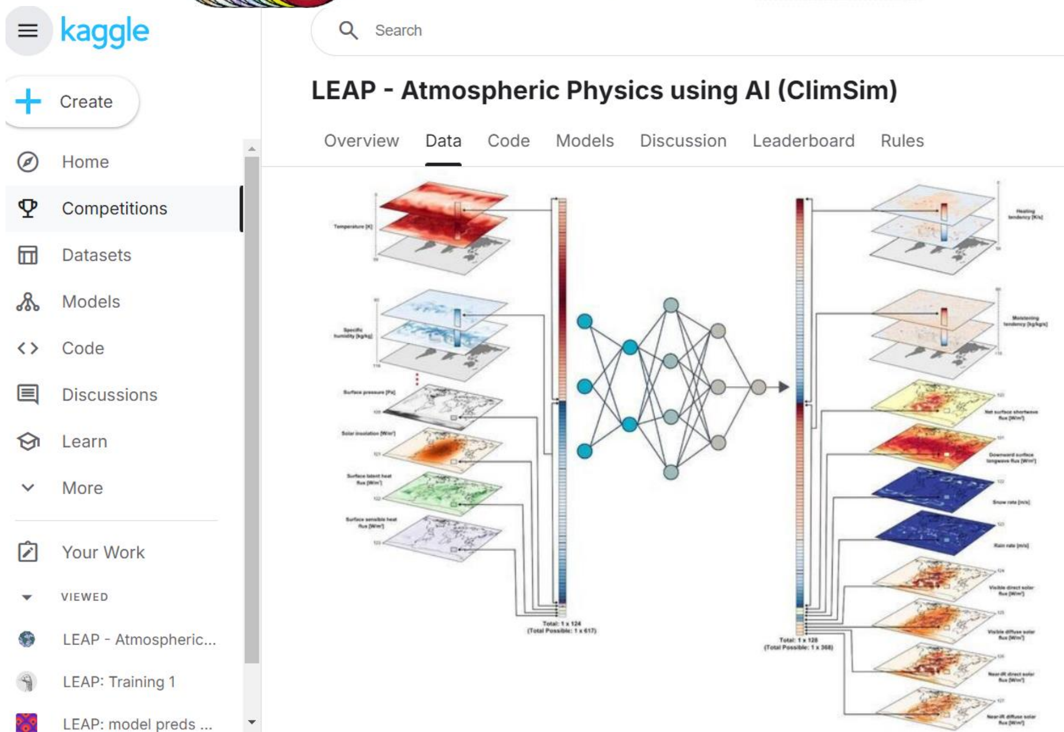

See: Christopoulos et al. (2024), Kochkov et al. (2024), Yu et al. (2023, 2024)

Data-Driven Parameterizations are flourishing and **Hybrid AI-climate modeling** is within grasp…

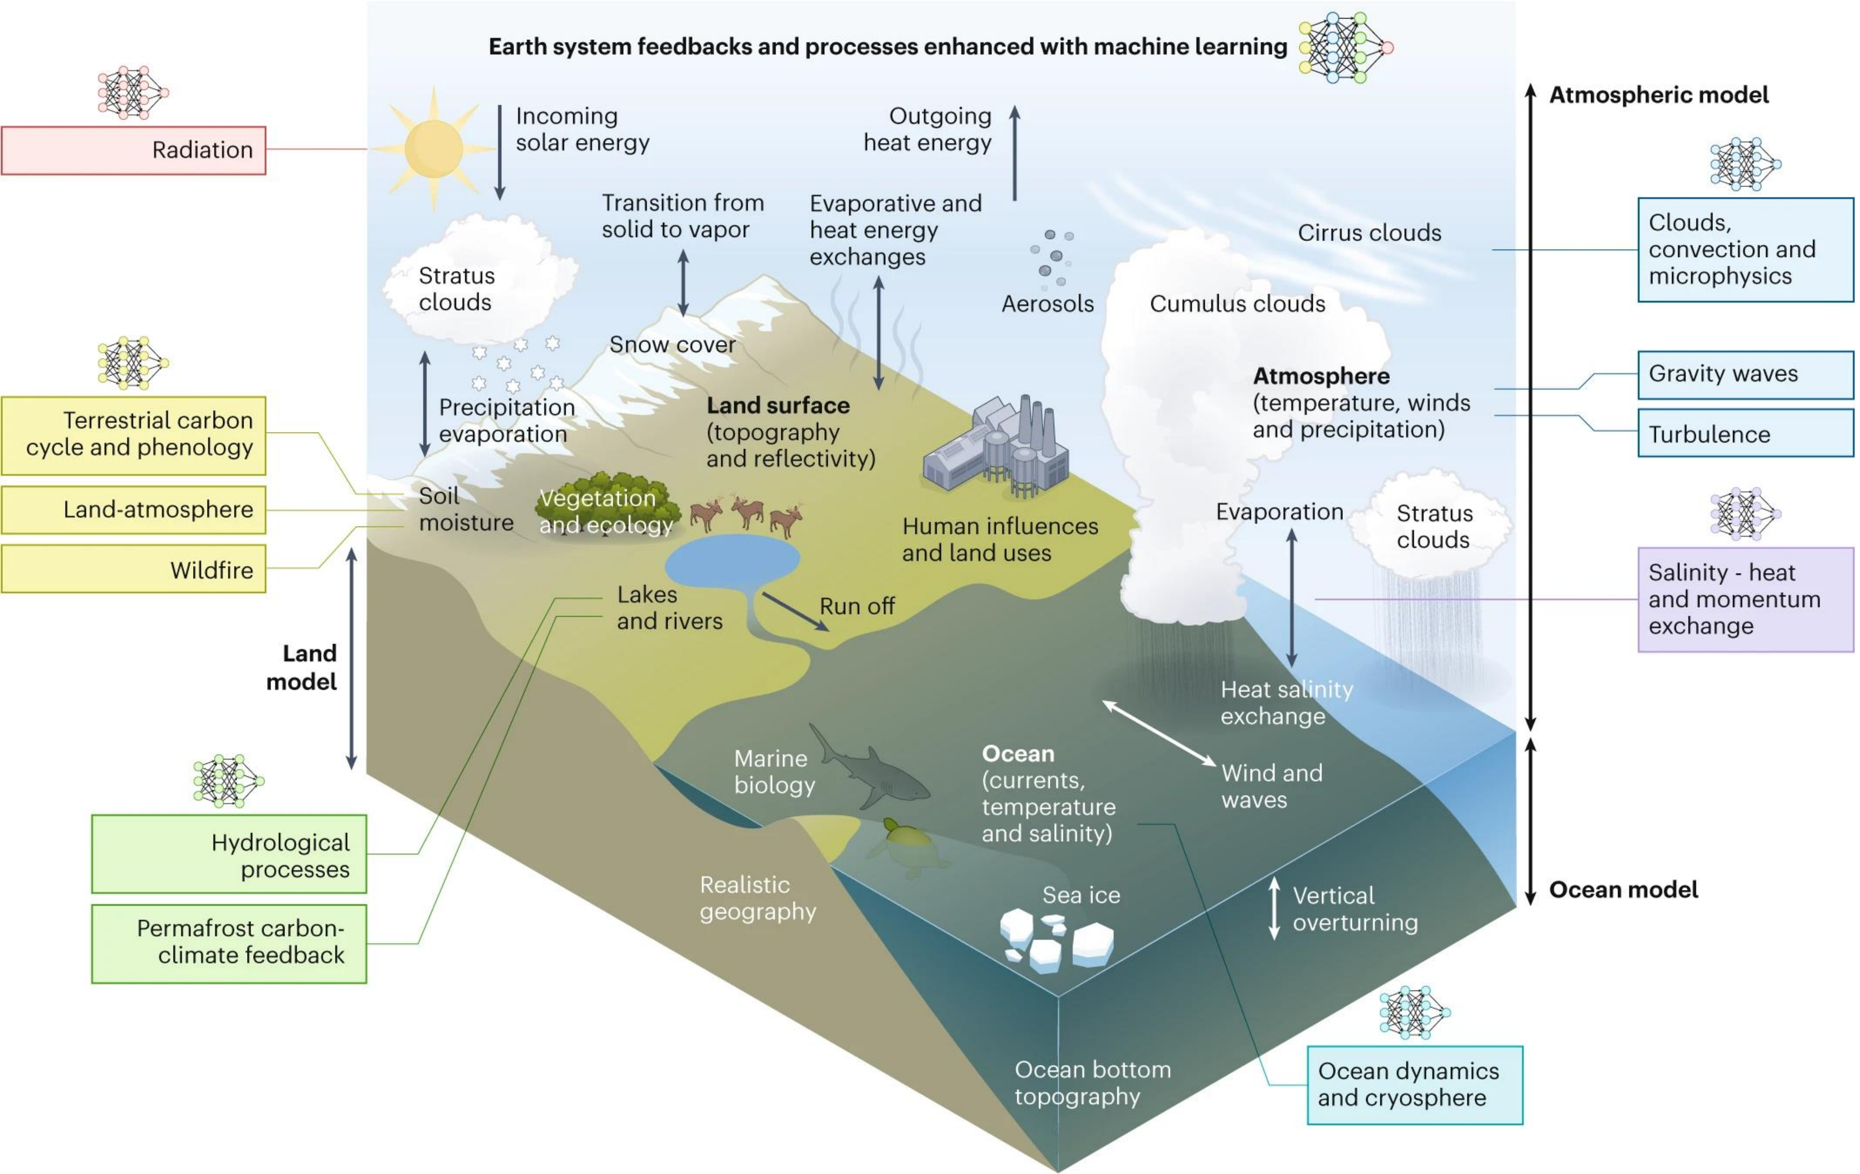

<u>Question</u>:

Why are hybrid AI-climate models not routine by now?

<u>*See*</u>*: Eyring et al. (2024)*

…but remain under-used in climate science

Cloud Fraction Parameterization (Grundner et al., 2024)

$$\mathcal{C}_{\text{frac}} = \text{NN}(p, q_v, q_l, q_i, T)$$

Subgrid-Scale Thermodynamics Parameterization (Beucler et al., 2024)

$$\begin{pmatrix} \dot{T} \\ \dot{q} \end{pmatrix} = \text{NN}(p, q_v, T, \text{LHF}, \text{SHF}, \mathcal{S}_0)$$

IS ML RELIABLE FOR CLIMATE PROJECTIONS?

HAVE WE LEARNED ANYTHING NEW USING ML?

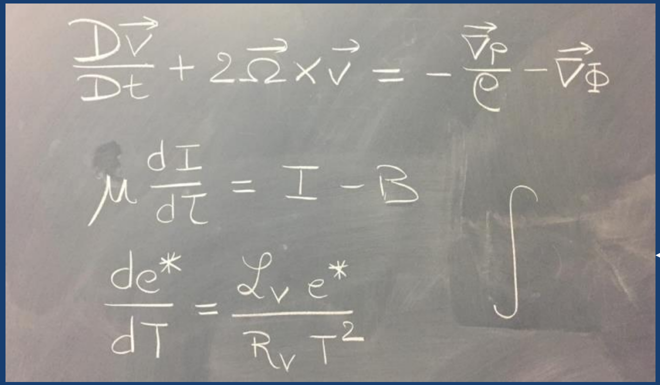

Data-Driven Discovery

Added value of ML measurable…

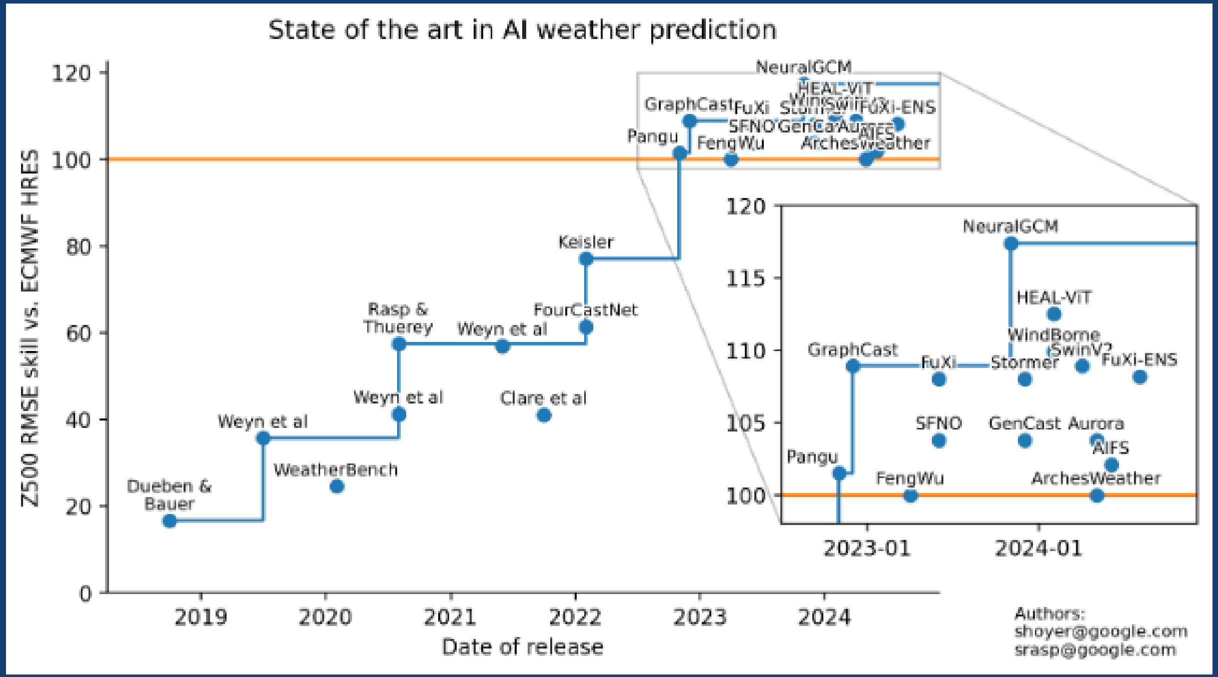

…but often challenging to explain

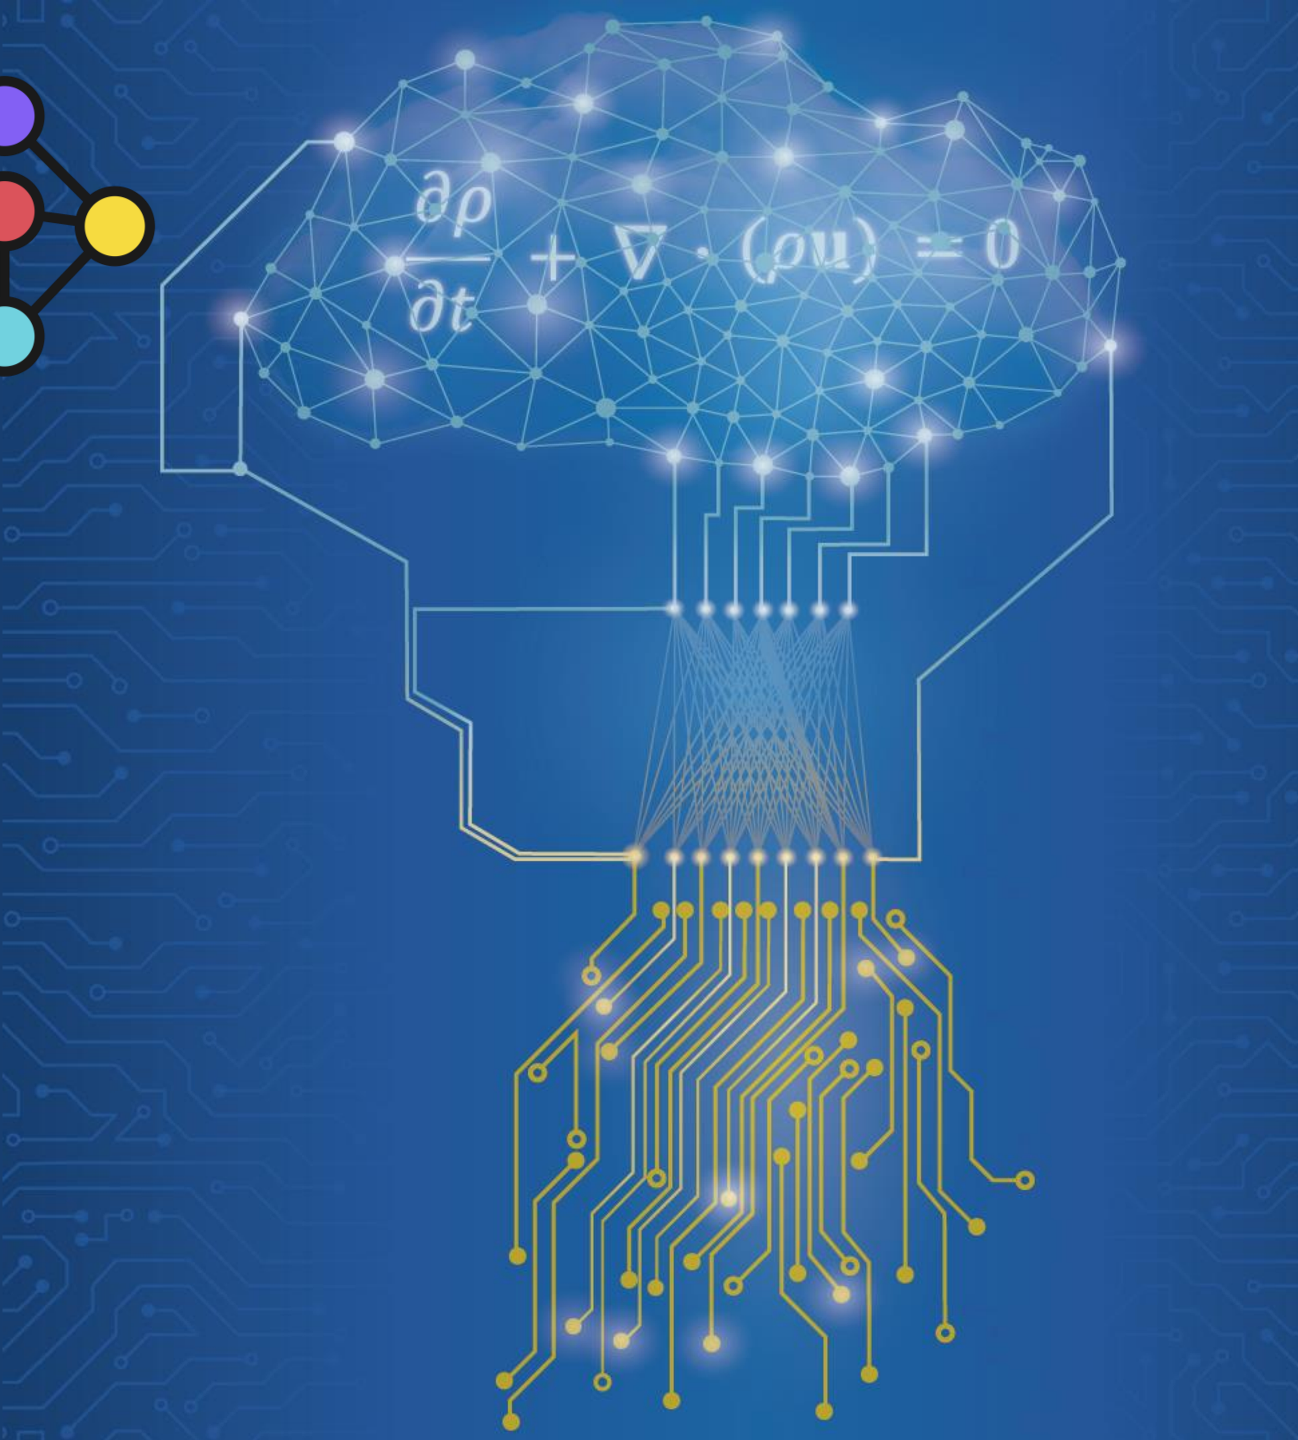

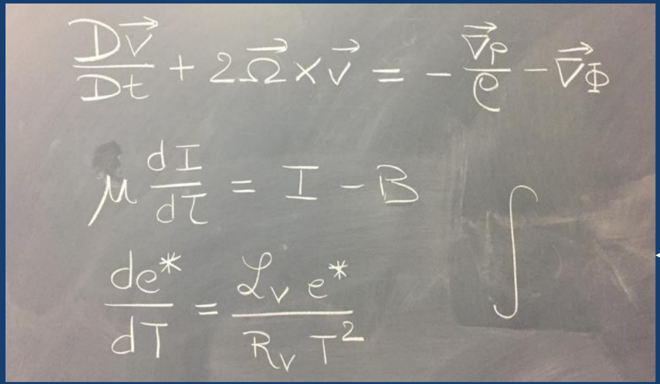

Data-Driven Discovery

Added value of ML measurable…

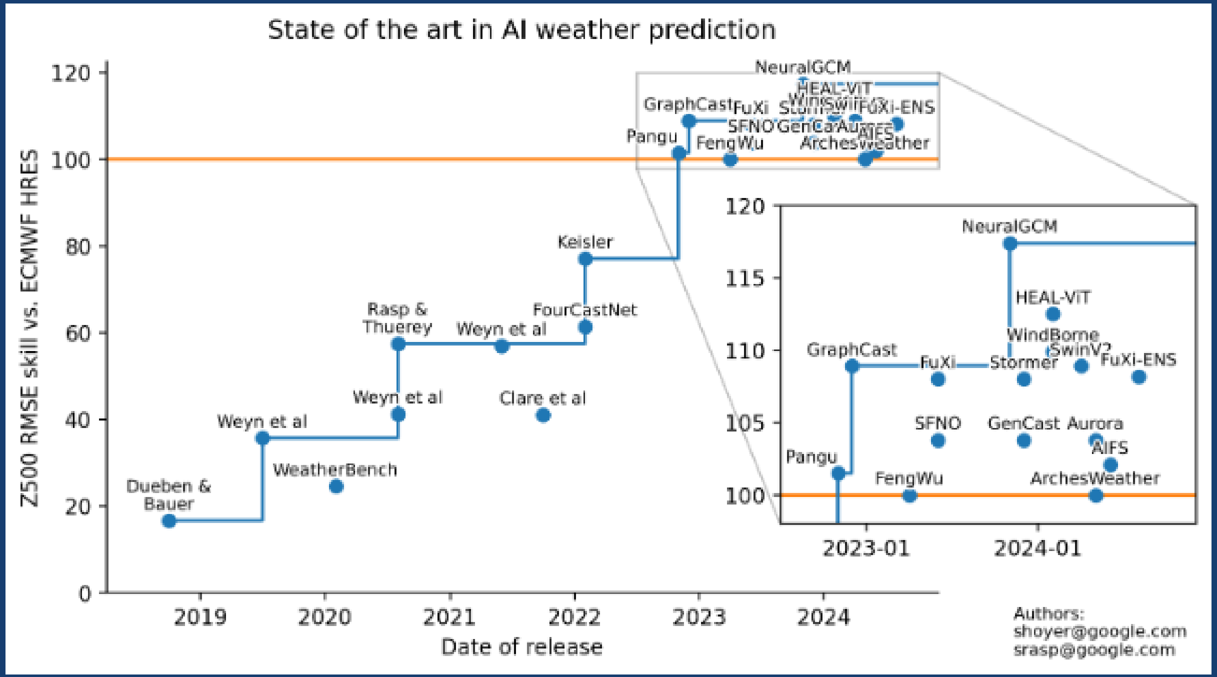

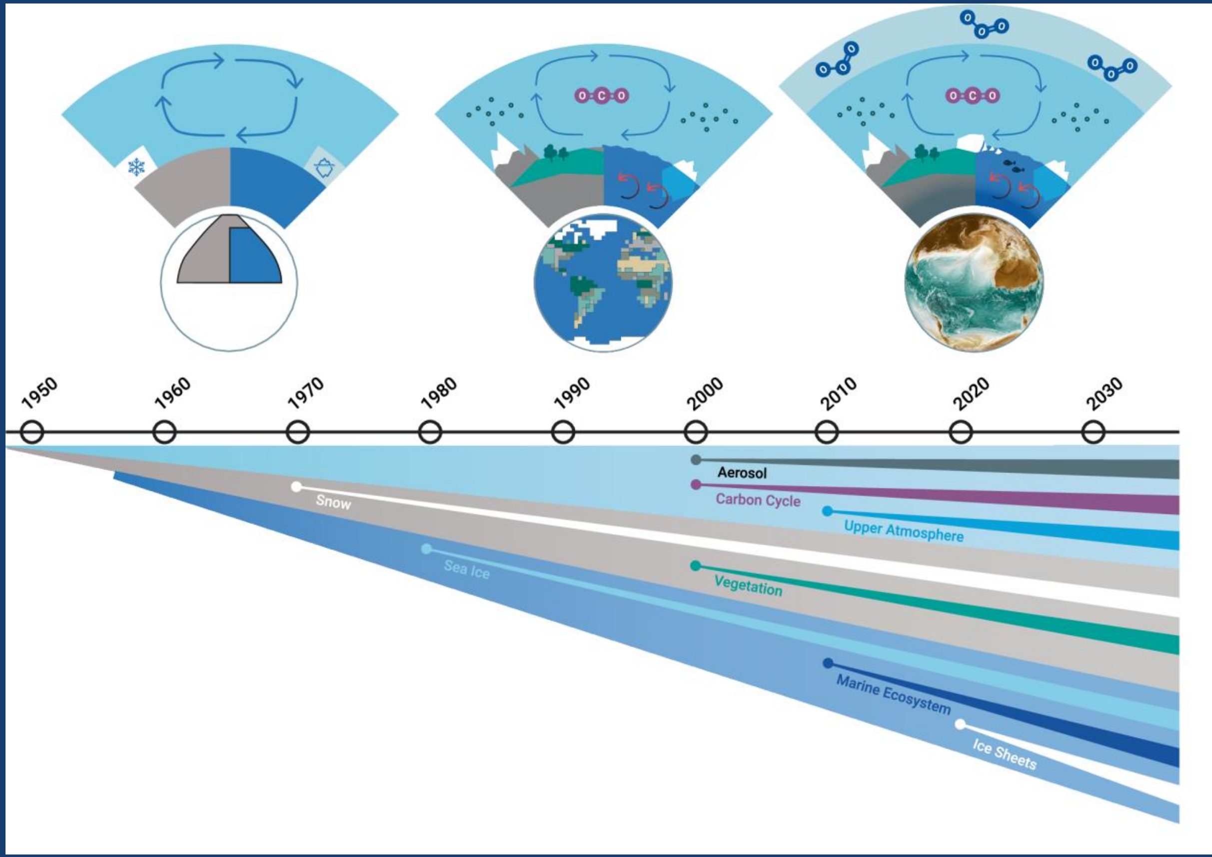

…but often challenging to explain

<u>Source</u>: *Bordoni et al. (2025), Bony et al. (2013);*
<u>See</u>: *Jeevanjee et al. (2017), Balaji (2022), ORNL (C. Jones, 2018)*

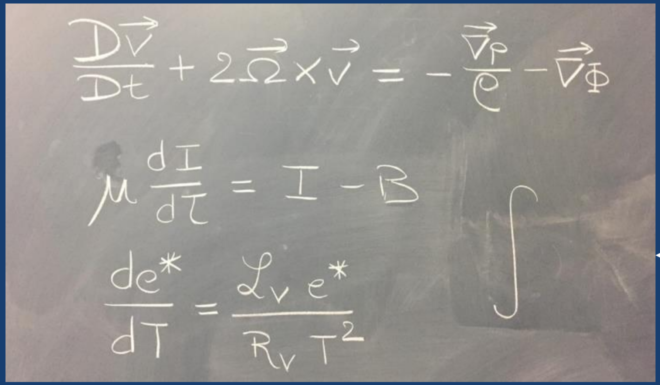

Data-Driven Discovery

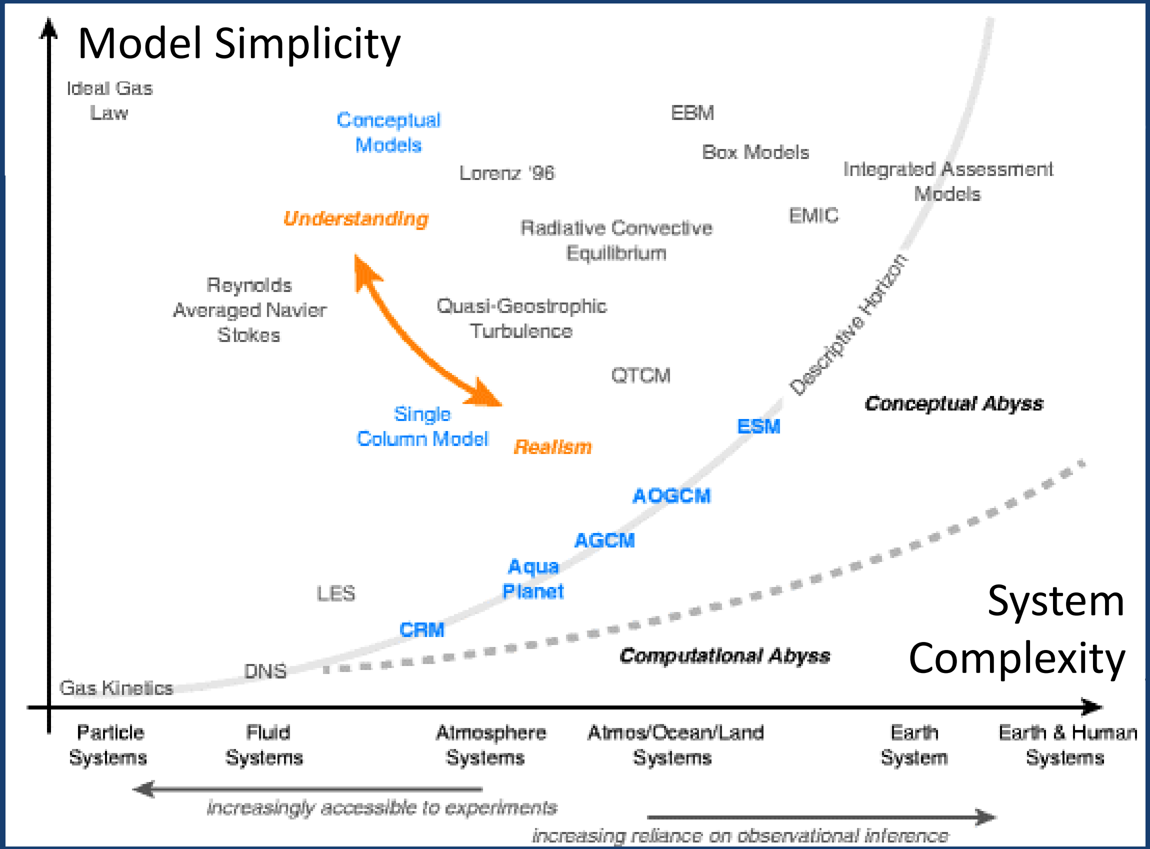

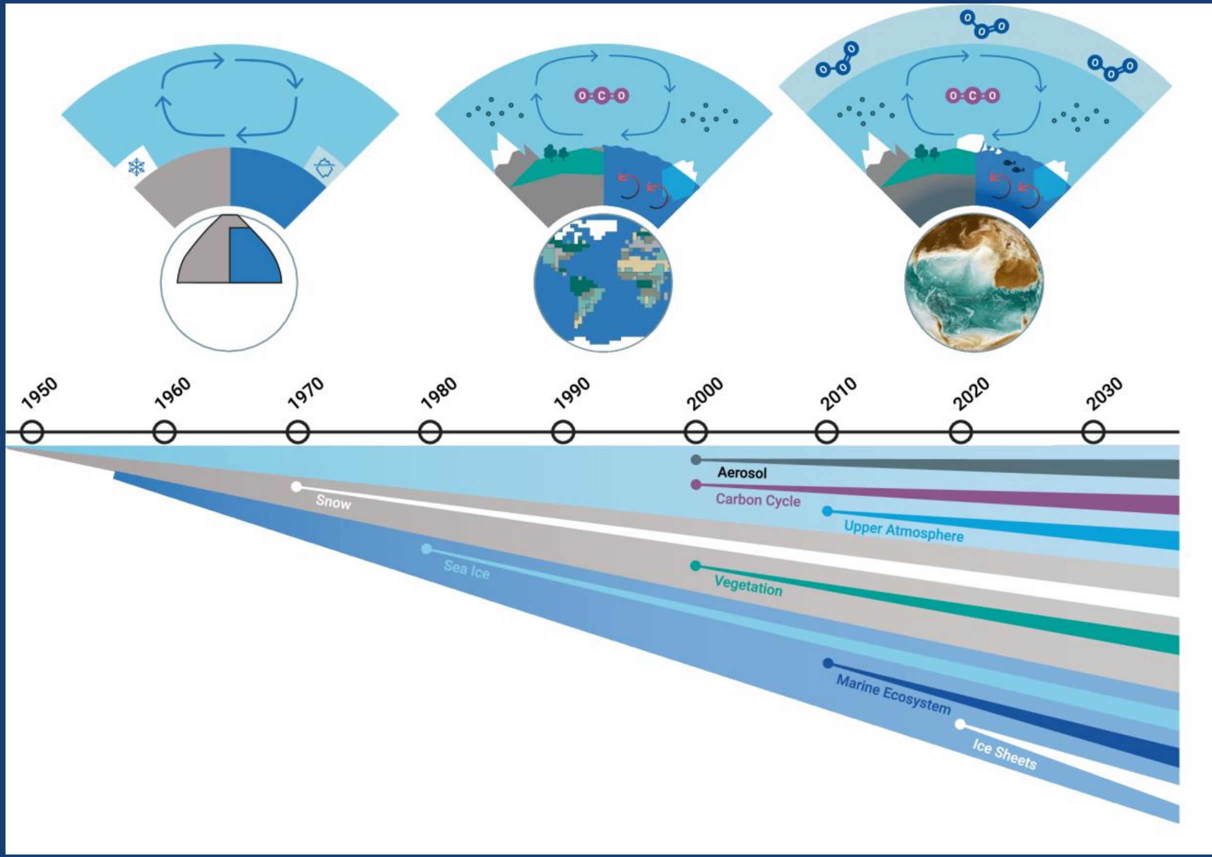

…but often challenging to explain

Source: Bordoni et al. (2025), Bony et al. (2013);
See: Jeevanjee et al. (2017), Balaji (2022), ORNL (C. Jones, 2018)

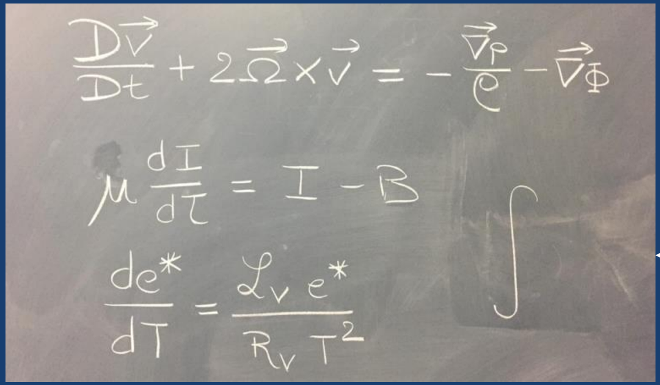

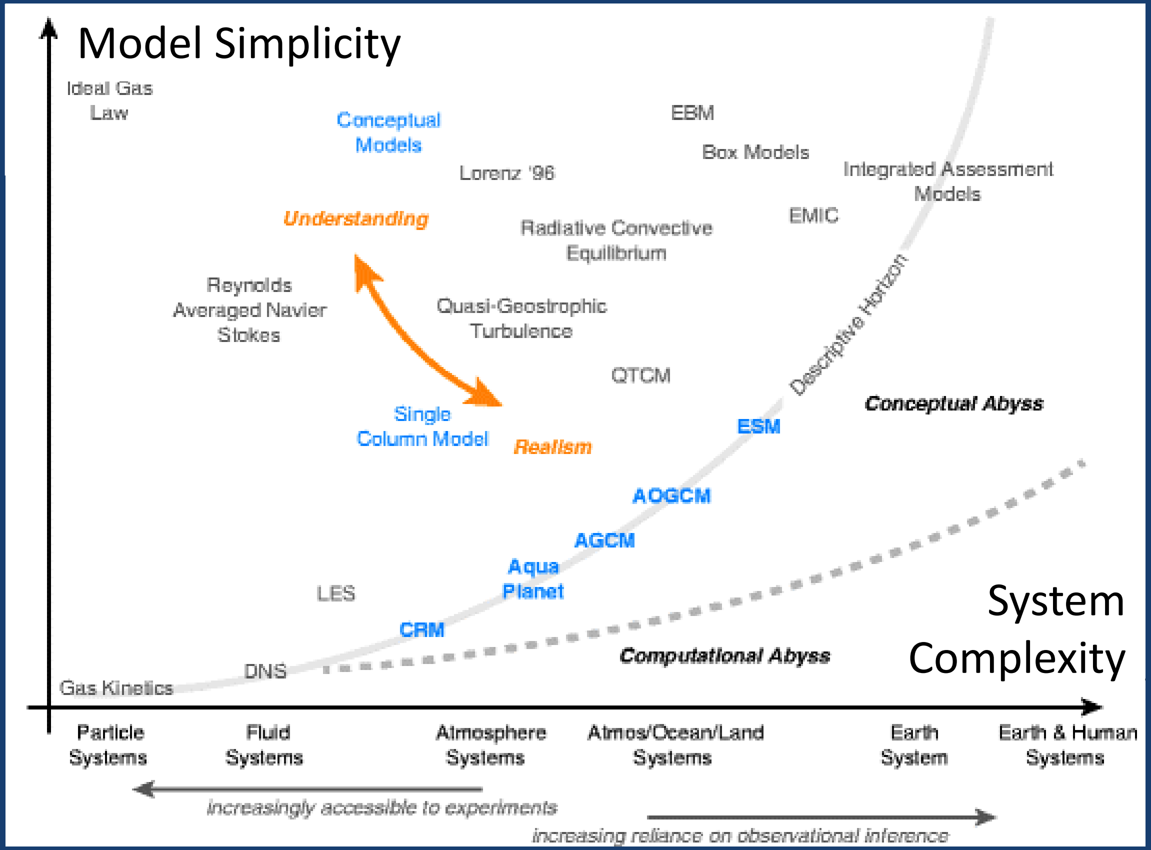

1. **Pareto-optimal model hierarchies**
2. Knowledge distillation
3. Challenges

Source: Bony et al. (2013); See: Jeevanjee et al. (2017), Balaji (2022), ORNL (C. Jones, 2018)

Idea: Work in a well-defined error-complexity plane

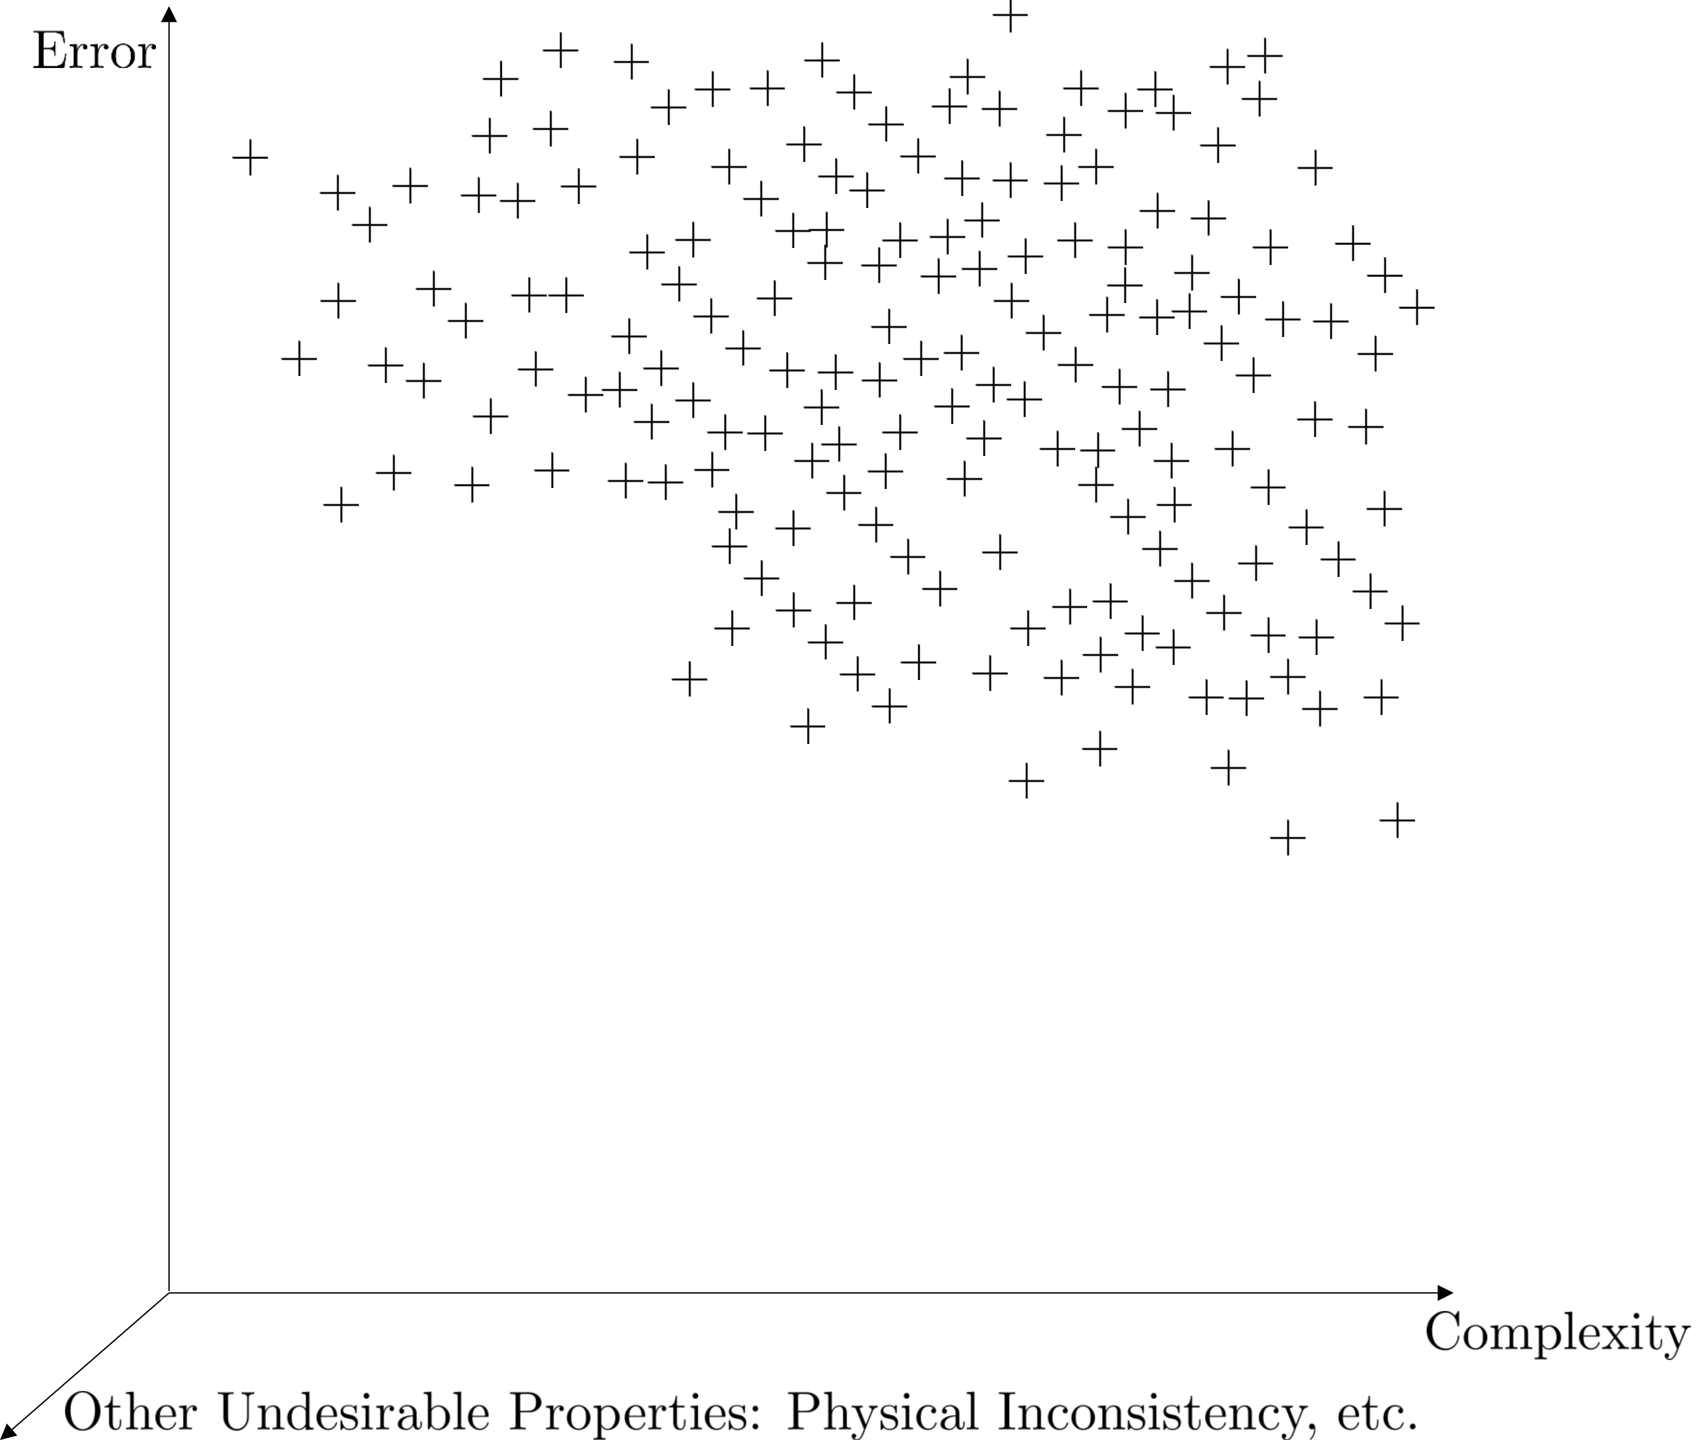

<u>Pareto Front</u>:

"When selecting a model from the Pareto front, switching to a different model means sacrificing the quality of at least one evaluation metric."

$$\mathrm{PF}_{\mathcal{E}} = \left\{ M_{\mathrm{opt}} \mid \nexists M \text{ s.t. } \begin{cases} \forall i\, \mathcal{E}_i(M) \le \mathcal{E}_i(M_{\mathrm{opt}}) \\ \exists j\, \mathcal{E}_j(M) < \mathcal{E}_j(M_{\mathrm{opt}}) \end{cases} \right\}$$

<u>See</u>: Censor (1977), Mietinnen (1999)

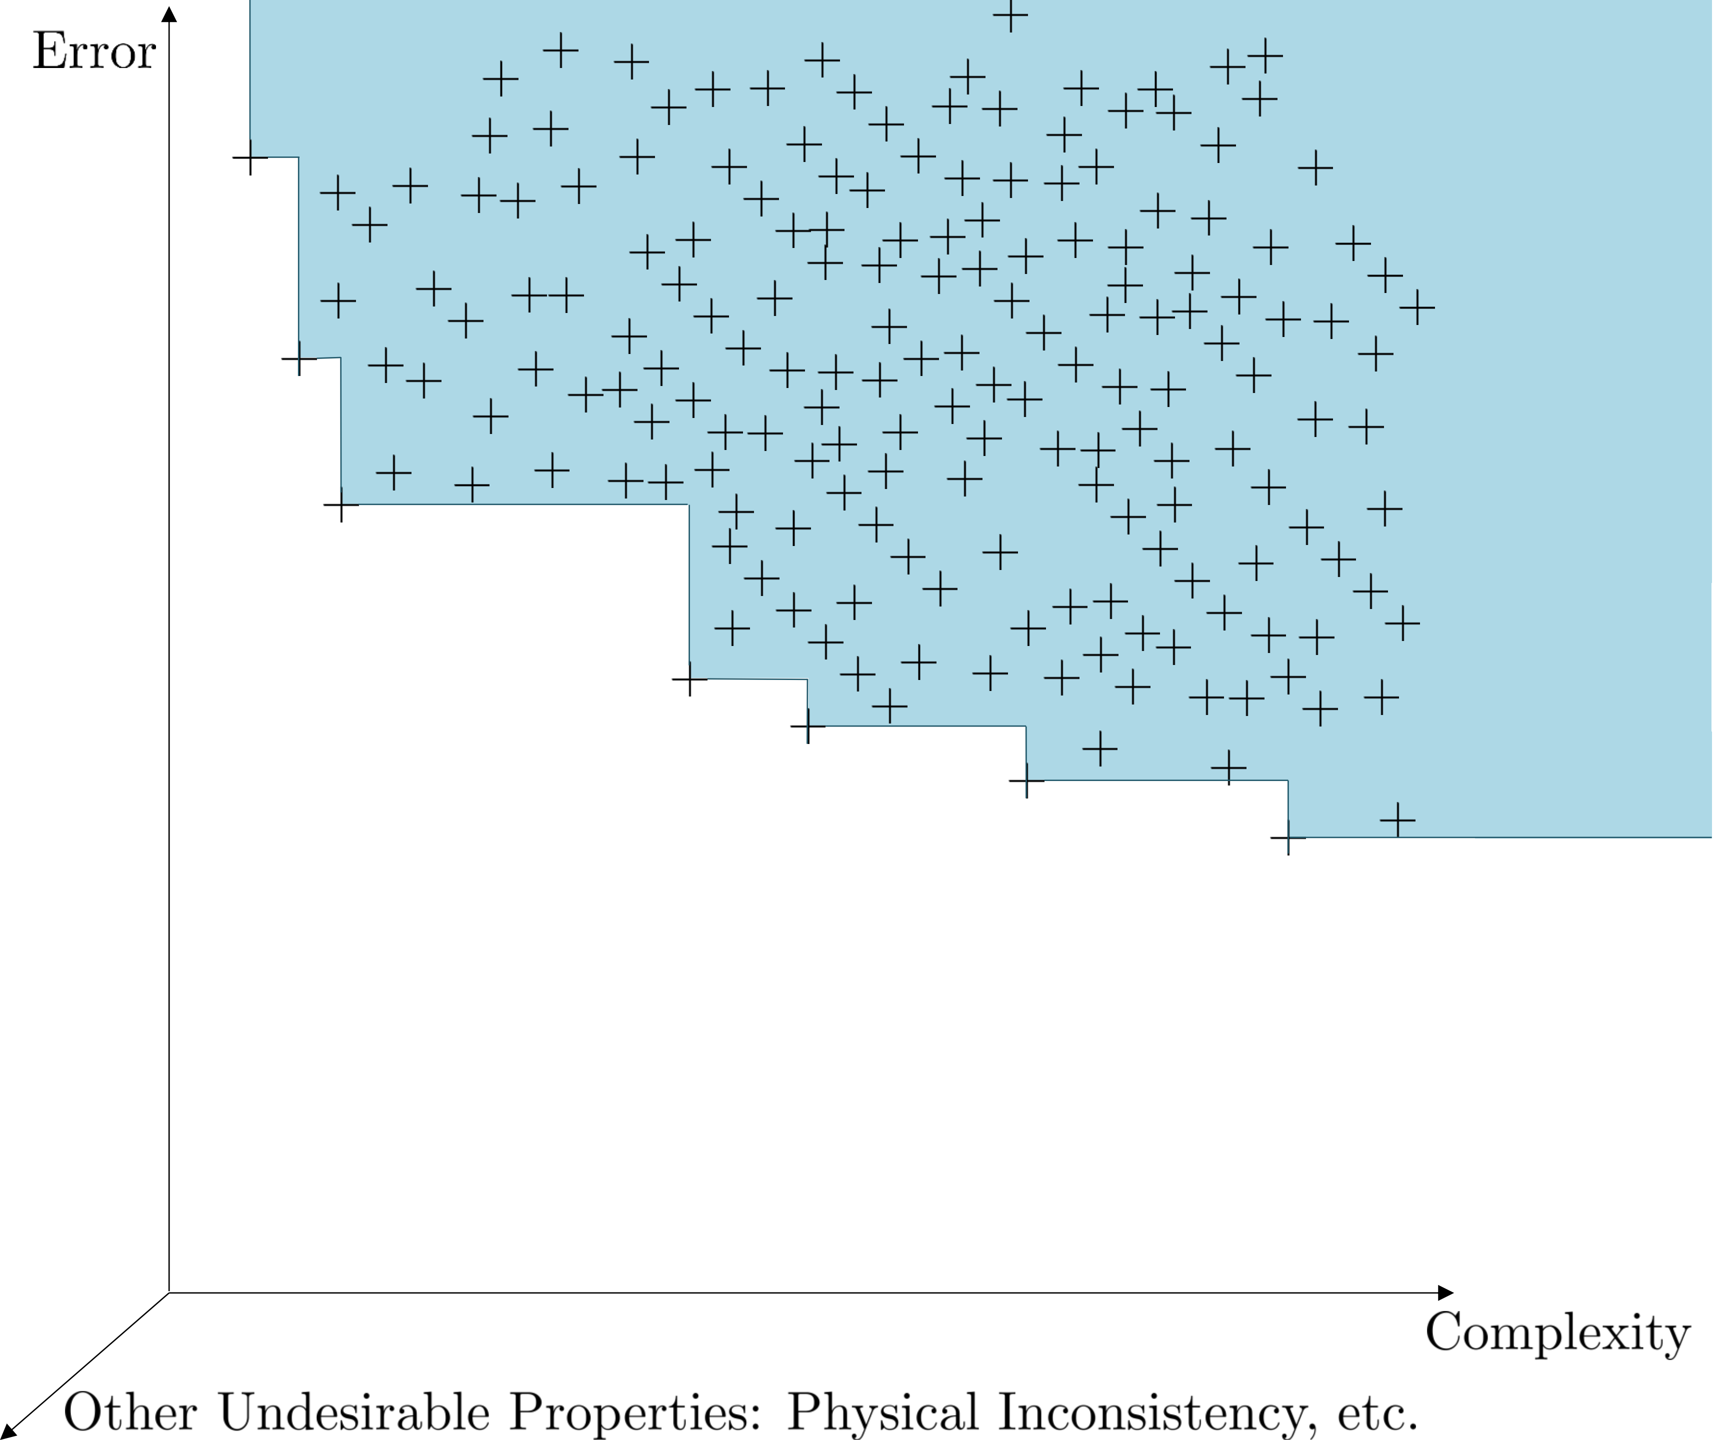

Pareto Front:

"When selecting a model from the Pareto front, switching to a different model means sacrificing the quality of at least one evaluation metric."

$$\mathrm{PF}_{\mathcal{E}} = \left\{ M_{\mathrm{opt}} \mid \nexists M \text{ s.t. } \begin{cases} \forall i\, \mathcal{E}_i(M) \leq \mathcal{E}_i(M_{\mathrm{opt}}) \\ \exists j\, \mathcal{E}_j(M) < \mathcal{E}_j(M_{\mathrm{opt}}) \end{cases} \right\}$$

See: Censor (1977), Mietinnen (1999)

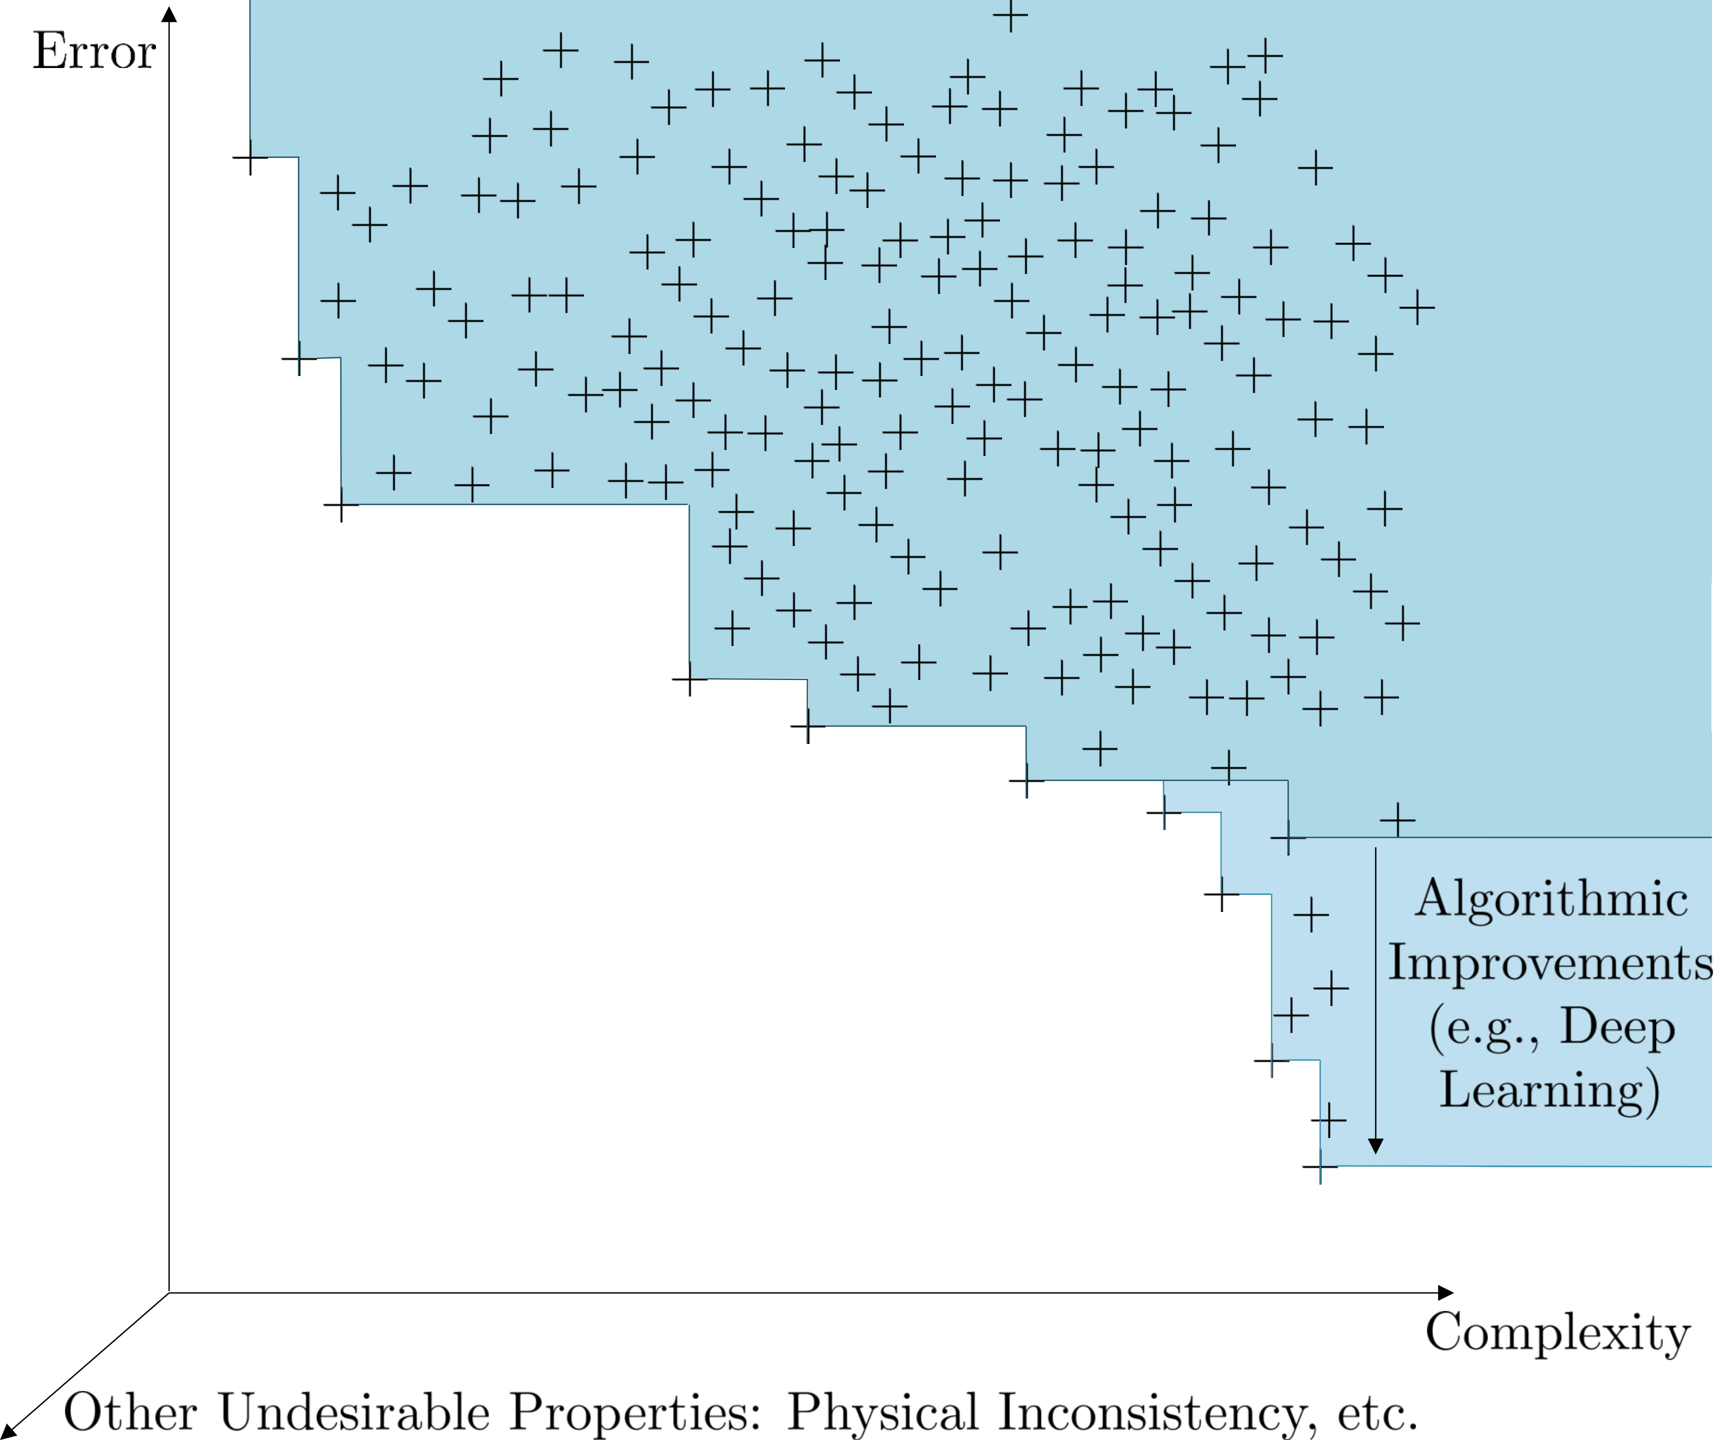

<u>Pareto Front</u>:

“When selecting a model from the Pareto front, switching to a different model means sacrificing the quality of at least one evaluation metric.”

$$\mathrm{PF}_{\mathcal{E}} = \left\{ M_{\mathrm{opt}} \mid \nexists M \text{ s.t. } \begin{cases} \forall i\, \mathcal{E}_i(M) \leq \mathcal{E}_i(M_{\mathrm{opt}}) \\ \exists j\, \mathcal{E}_j(M) < \mathcal{E}_j(M_{\mathrm{opt}}) \end{cases} \right\}$$

<u>See</u>: Censor (1977), Mietinnen (1999)

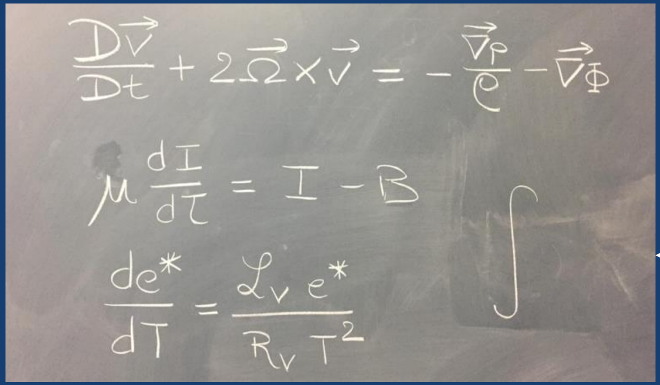

Data-Driven Discovery

$$Y = M\left[X_{\vec{x},t}\right]$$

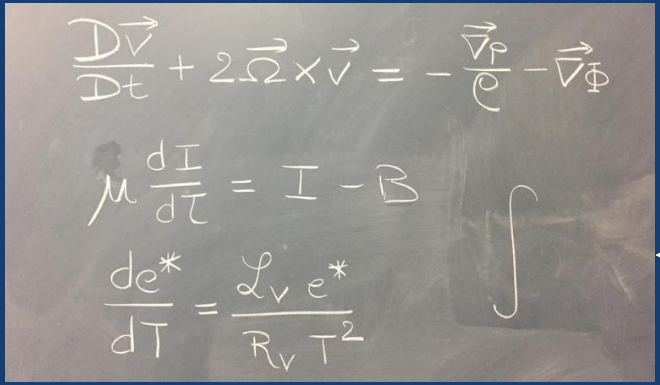

Data-Driven Discovery

Distillable value:

1) Functional representation

$$Y = M\left[X_{\vec{x},t}\right]$$

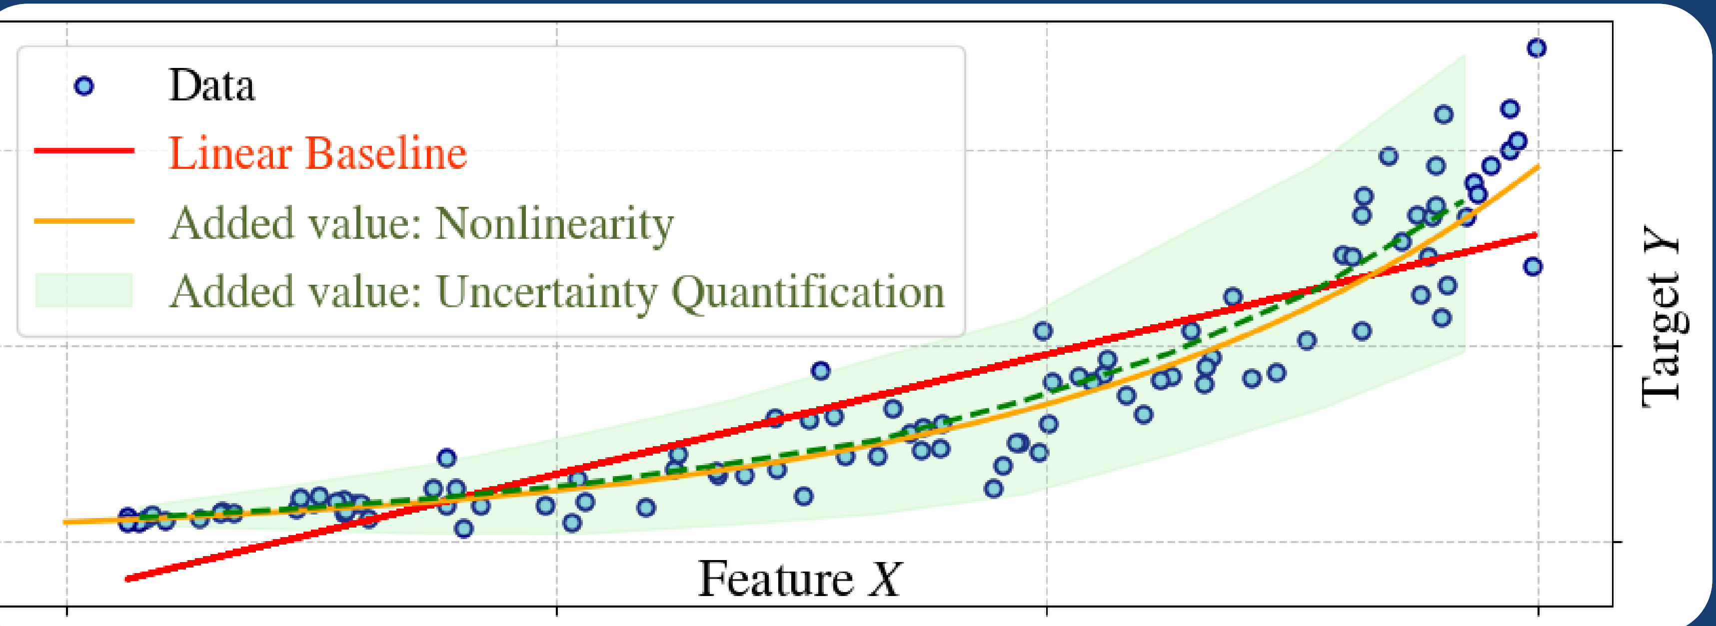

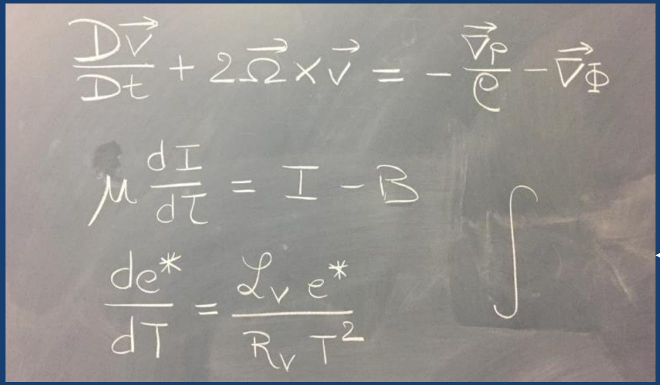

Data-Driven Discovery

Distillable value:

1) Functional representation

2) Feature assimilation

$$Y = M\left[X_{\vec{x}}, t\right]$$

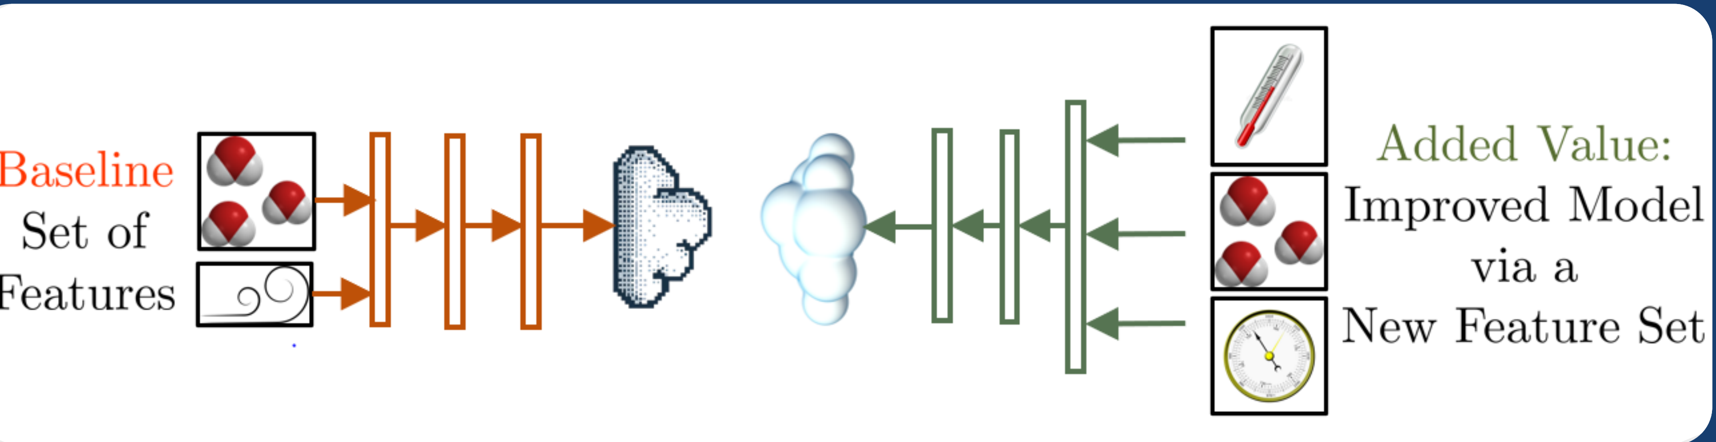

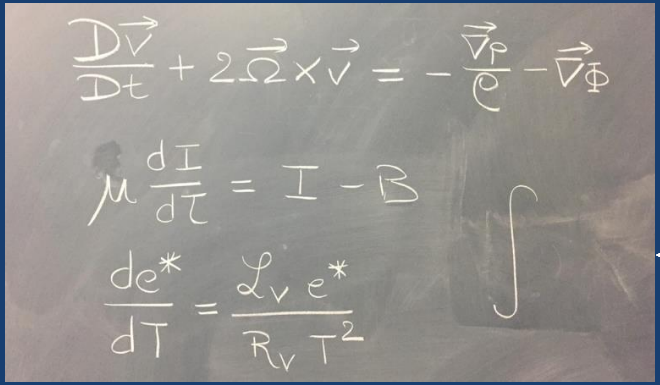

Data-Driven Discovery

$$Y = M\left[X_{\vec{x}}, t\right]$$

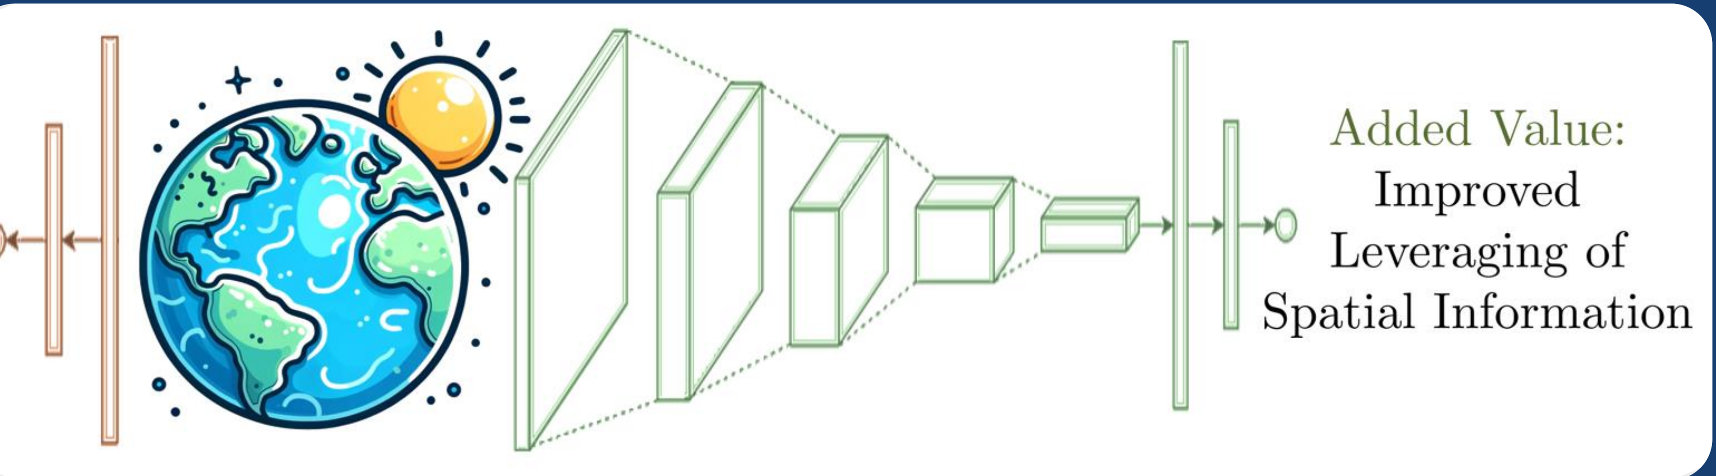

Distillable value:

1) Functional representation

2) Feature assimilation

3) Spatial connectivity

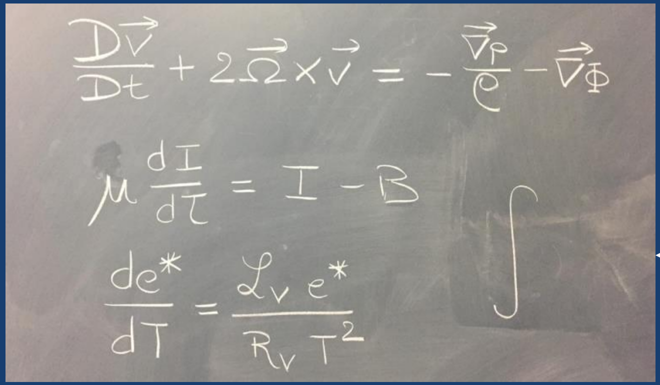

Data-Driven Discovery

$$Y = M\left[X_{\vec{x}}, t\right]$$

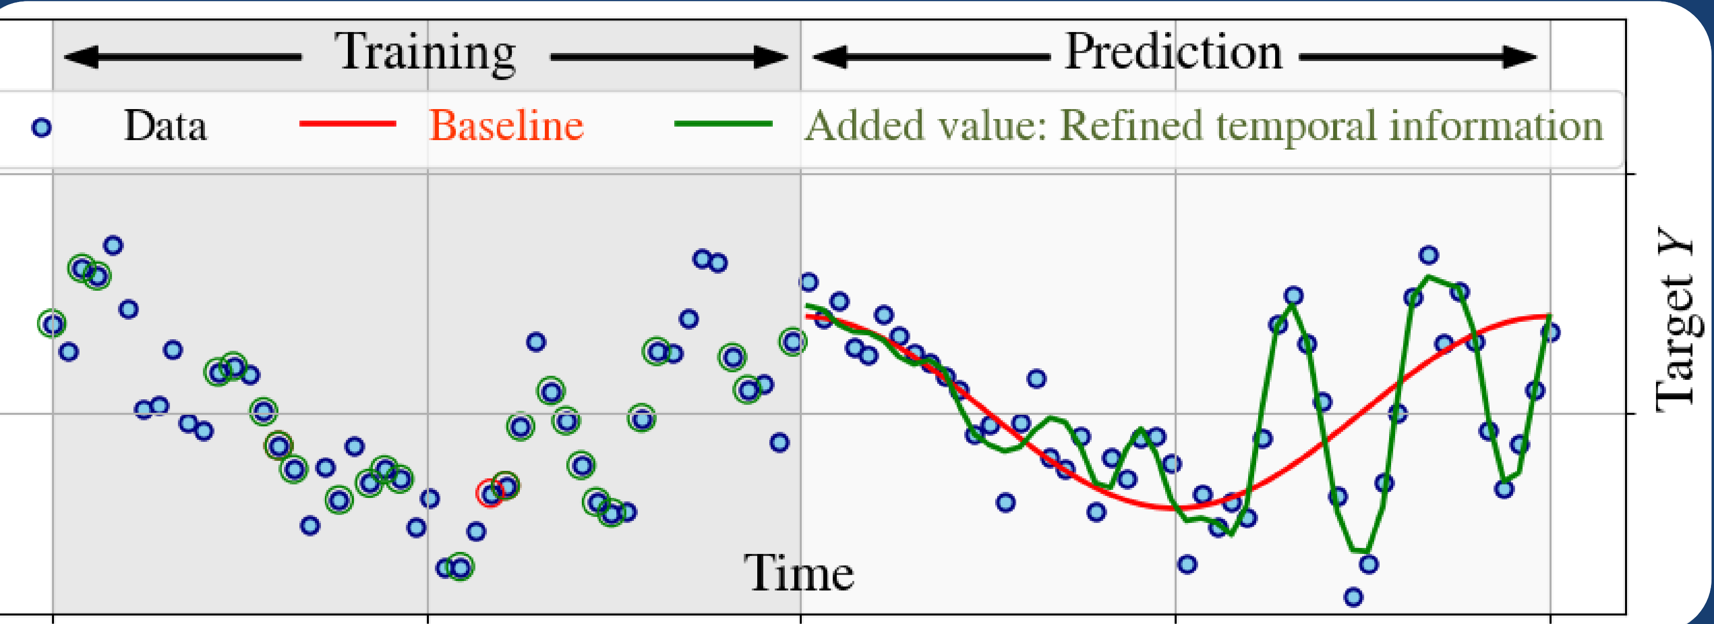

Distillable value:

1) Functional representation
2) Feature assimilation
3) Spatial connectivity
4) Temporal connectivity

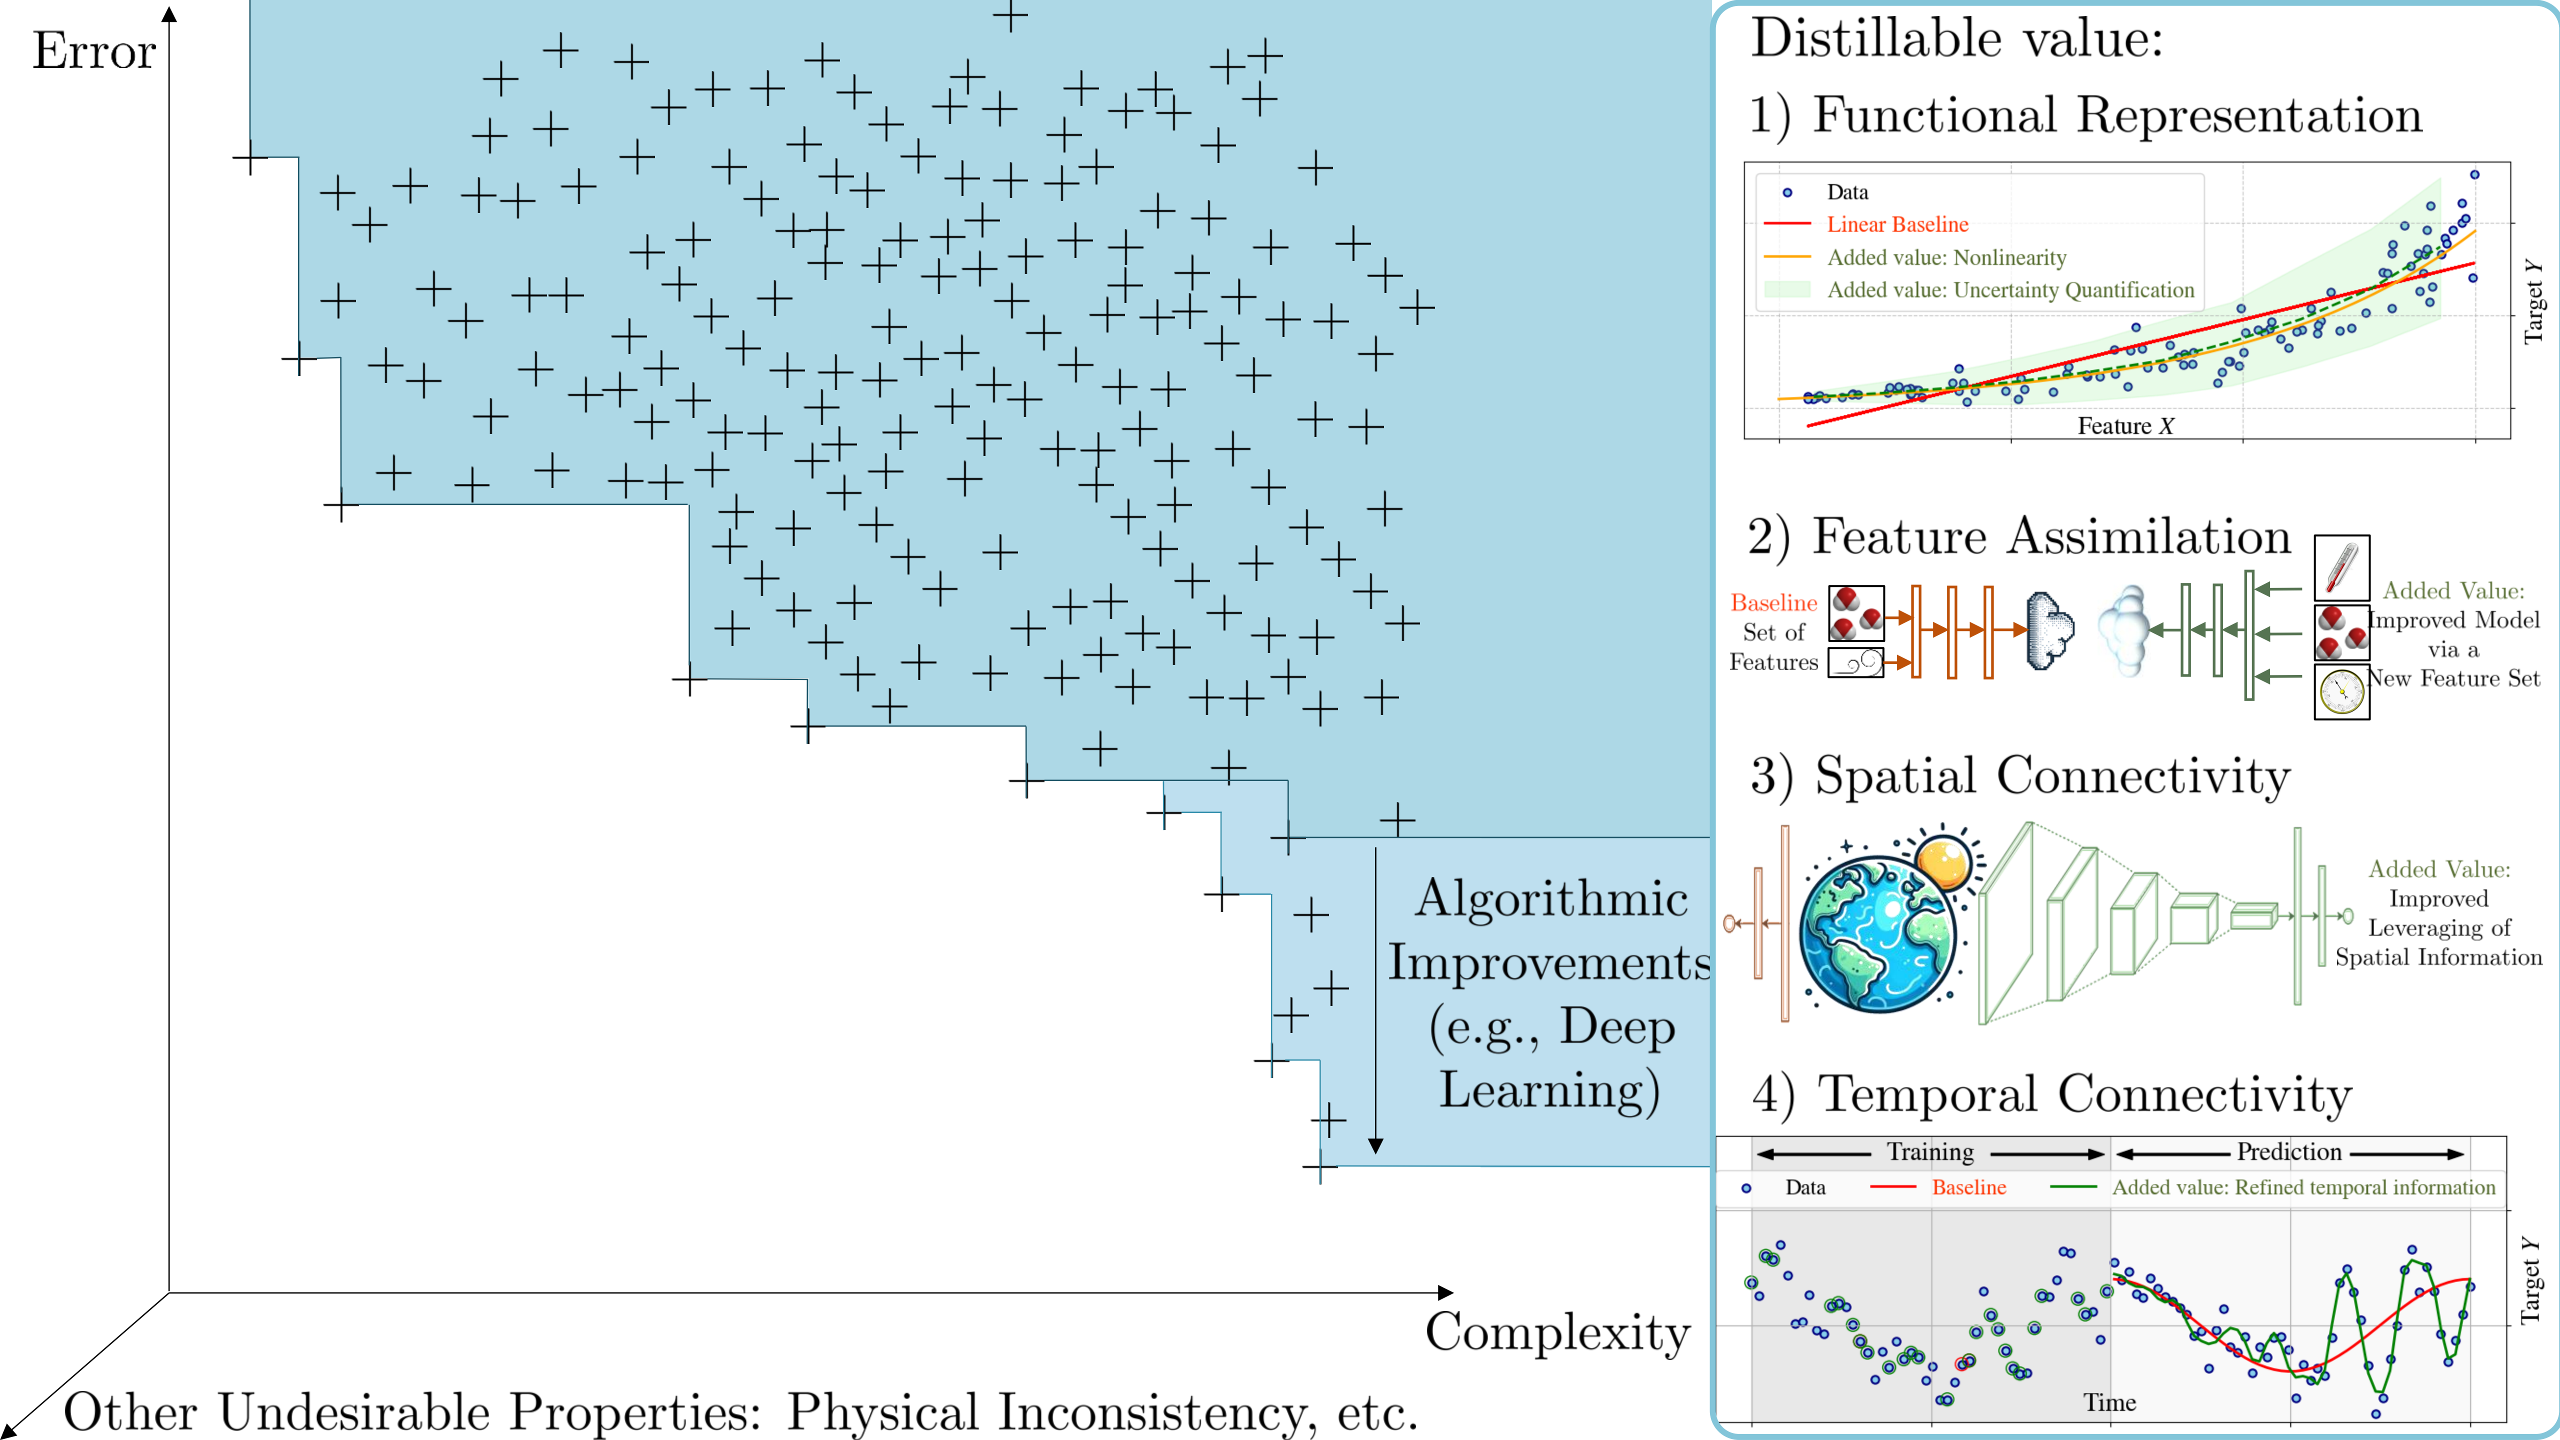

Error
Complexity
Other Undesirable Properties: Physical Inconsistency, etc.
Algorithmic Improvements (e.g., Deep Learning)
Distillable value:
1) Functional Representation
Data
Linear Baseline
Added value: Nonlinearity
Added value: Uncertainty Quantification
Feature X
Target Y
2) Feature Assimilation
Baseline Set of Features
Added Value: Improved Model via a New Feature Set
3) Spatial Connectivity
Added Value: Improved Leveraging of Spatial Information
4) Temporal Connectivity
Training
Prediction
Data
Baseline
Added value: Refined temporal information
Time
Target Y

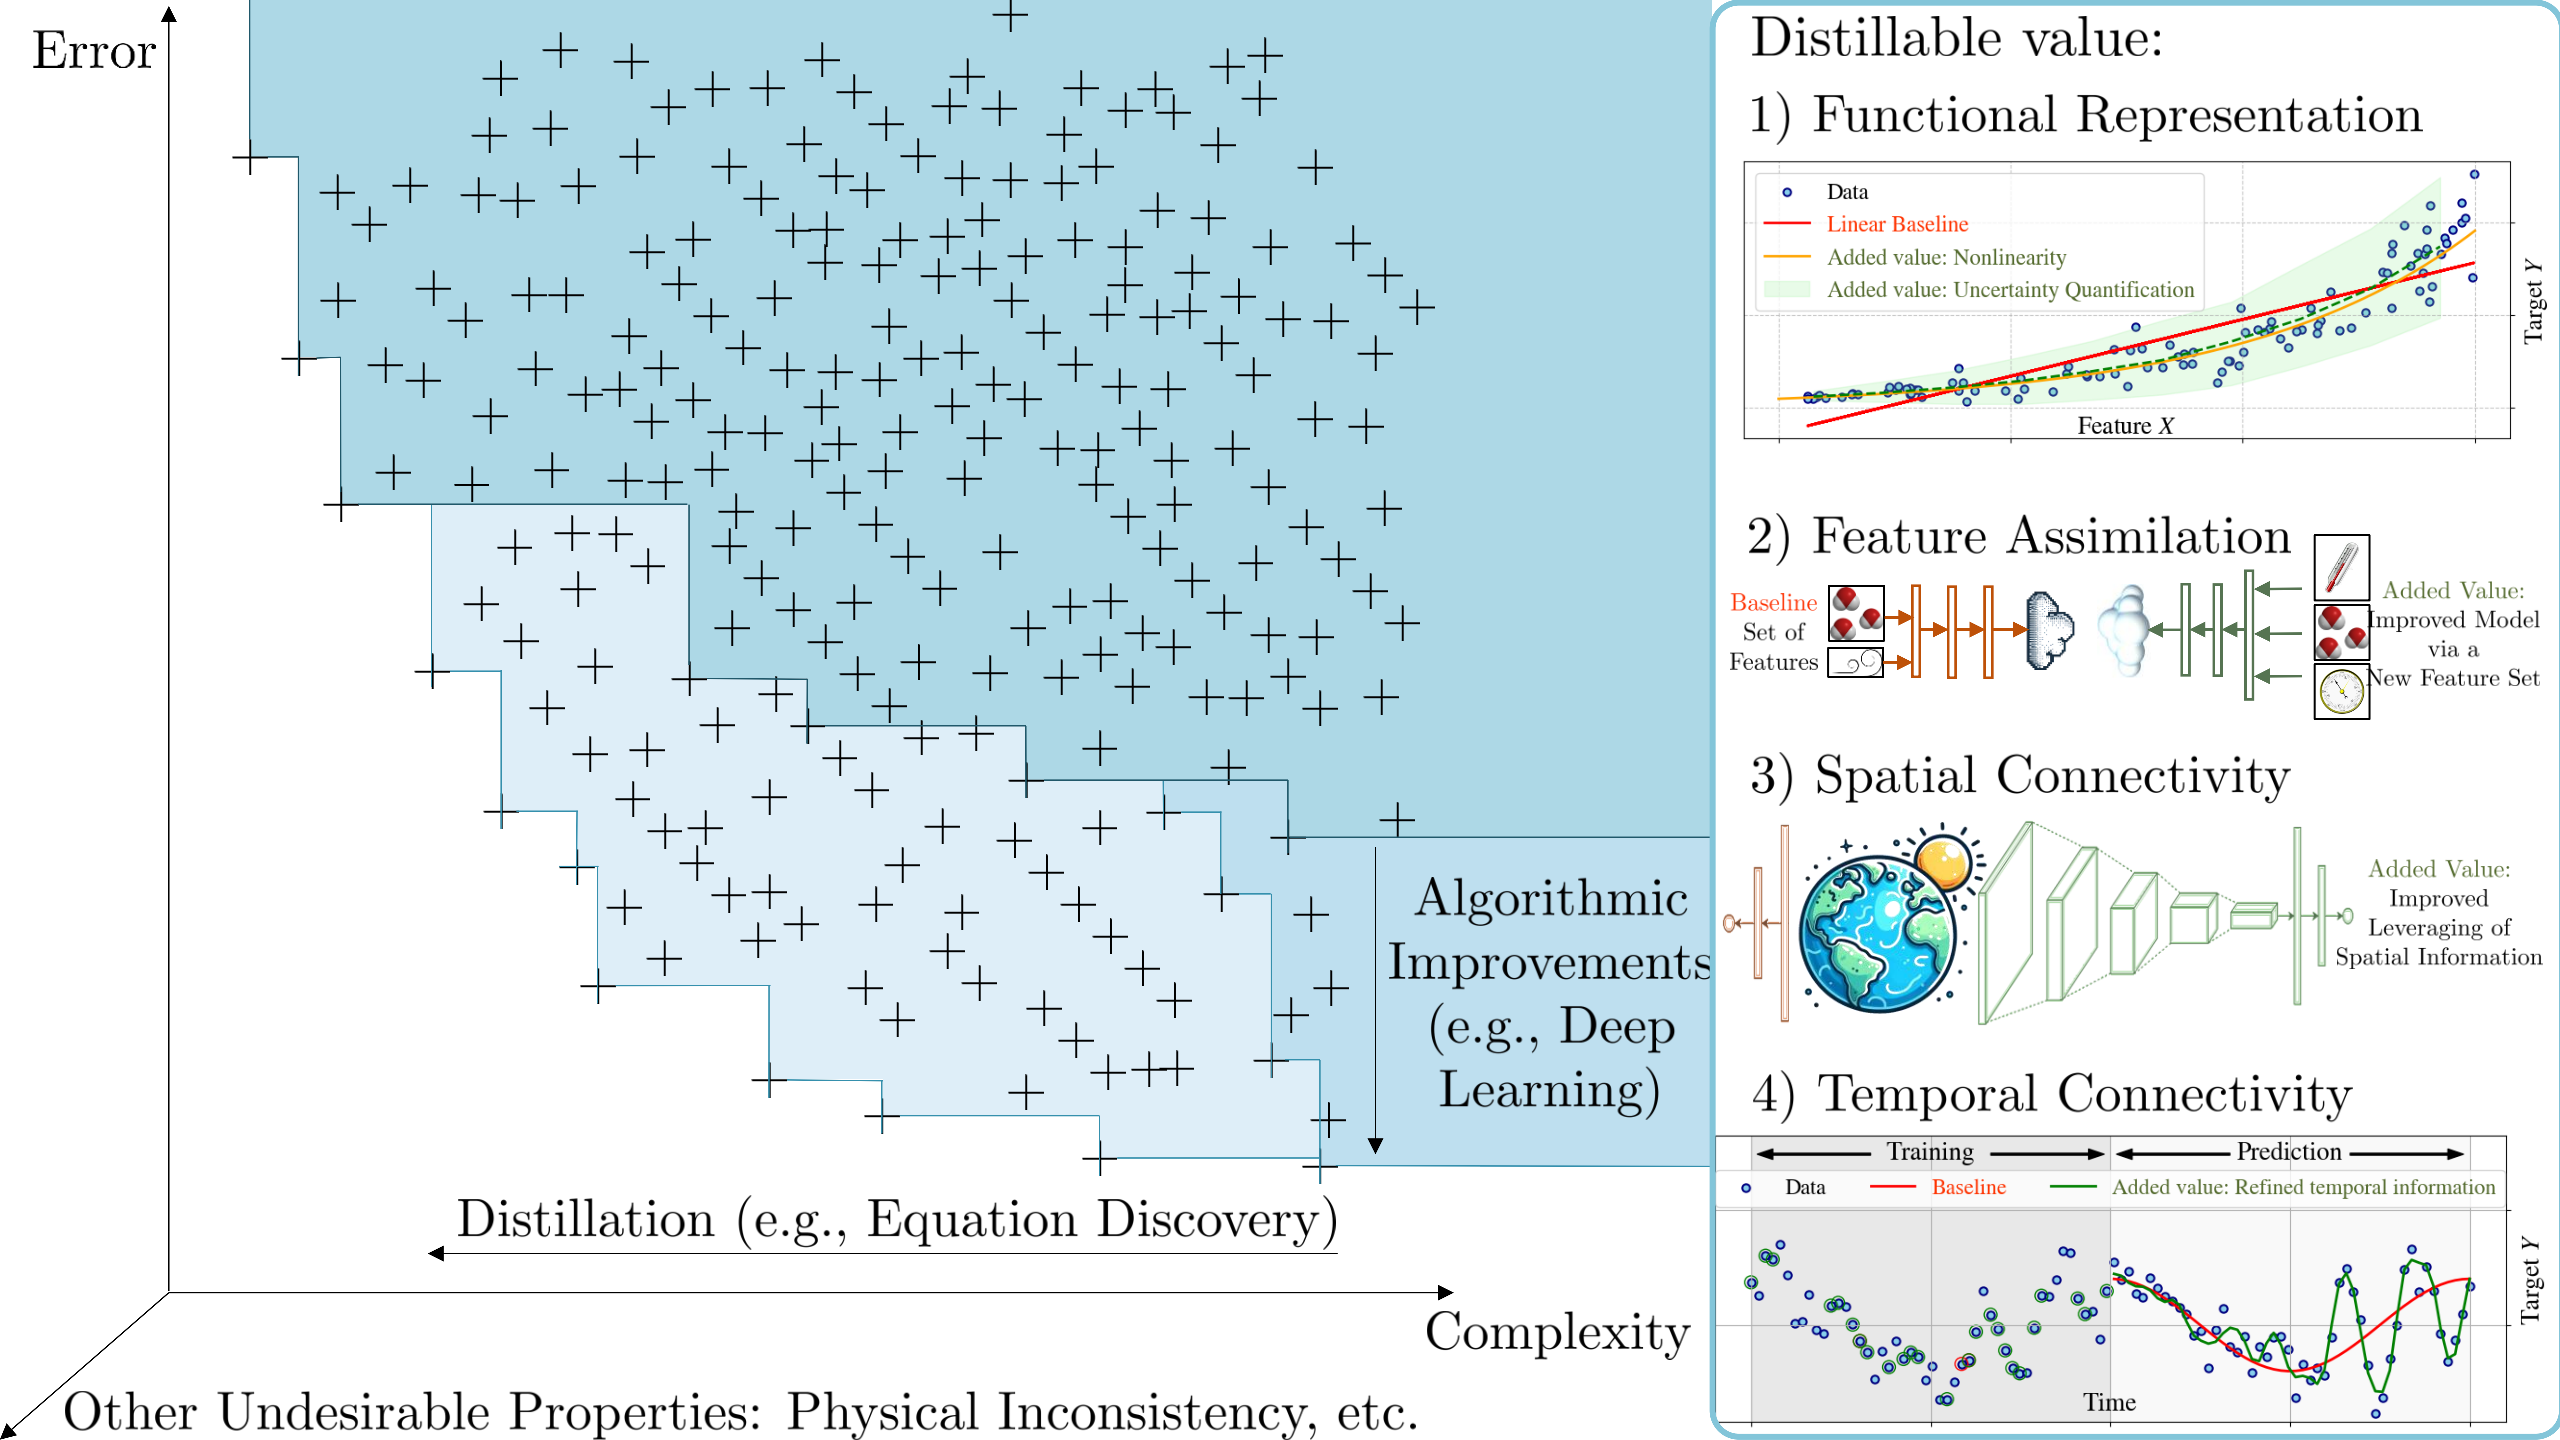

Error
Complexity
Other Undesirable Properties: Physical Inconsistency, etc.
Distillation (e.g., Equation Discovery)
Algorithmic Improvements (e.g., Deep Learning)
Distillable value:
1) Functional Representation
Data
Linear Baseline
Added value: Nonlinearity
Added value: Uncertainty Quantification
Feature X
Target Y
2) Feature Assimilation
Baseline Set of Features
Added Value: Improved Model via a New Feature Set
3) Spatial Connectivity
Added Value: Improved Leveraging of Spatial Information
4) Temporal Connectivity
Training
Prediction
Data
Baseline
Added value: Refined temporal information
Time
Target Y

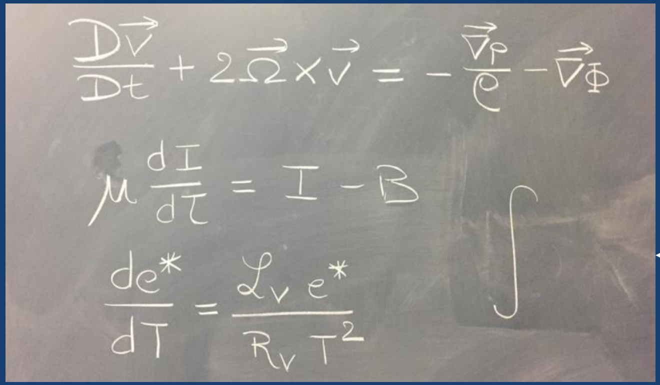

1. Pareto-optimal model hierarchies

2. **Knowledge distillation**

3. Challenges

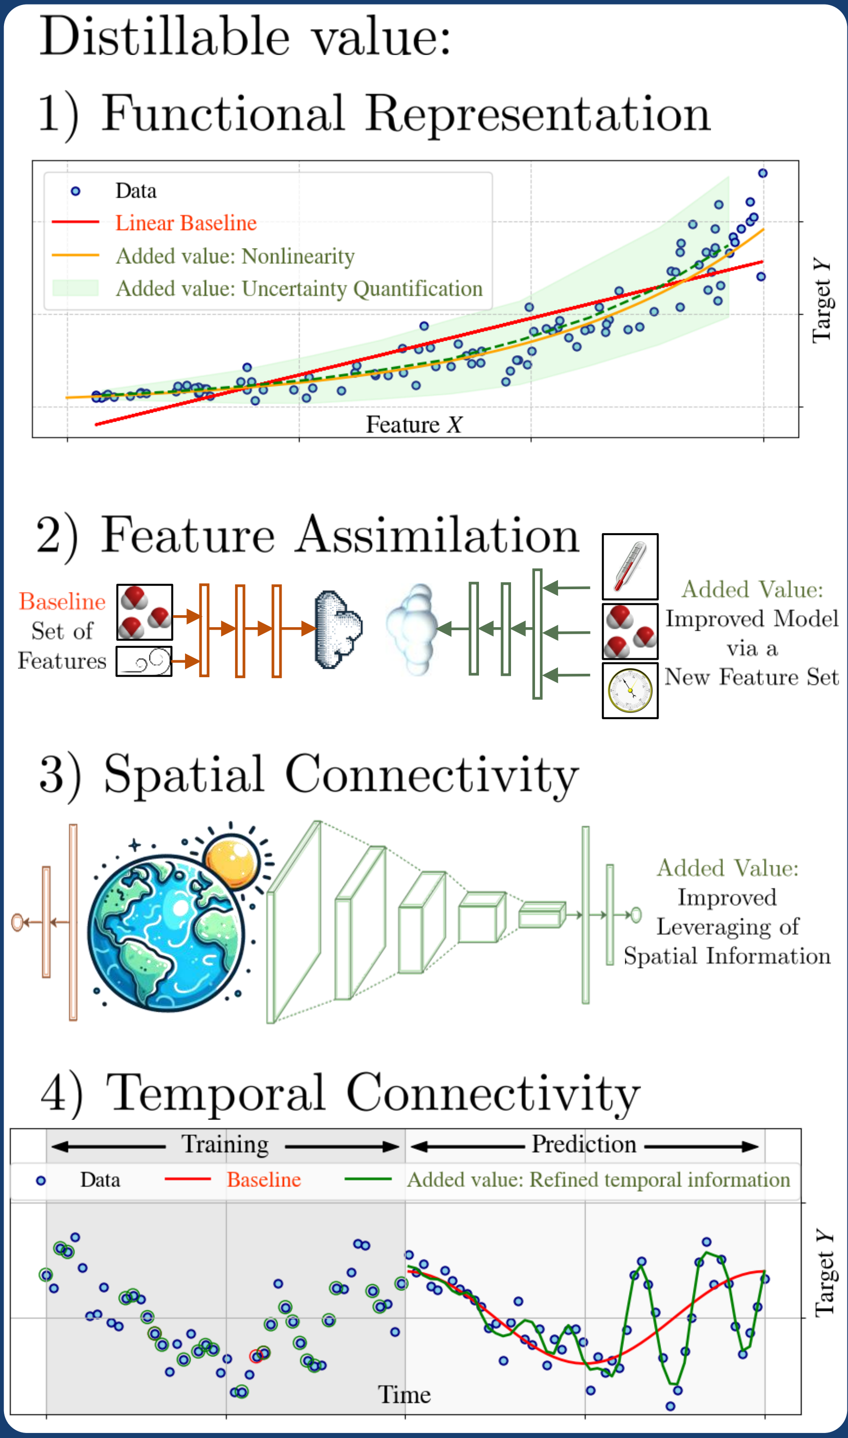

Cloud cover parameterization maps the grid-scale environment to the fraction occupied by clouds

<u>Motivation</u>: Reducing cloud-related biases via storm-resolving simulations

<u>Data</u>: 2.5km-res, 59-layer, global storm-resolving ICON runs (DYAMOND)

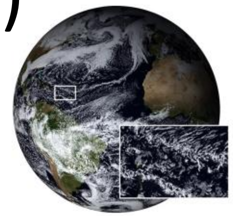

<u>Original Cloud Cover</u>

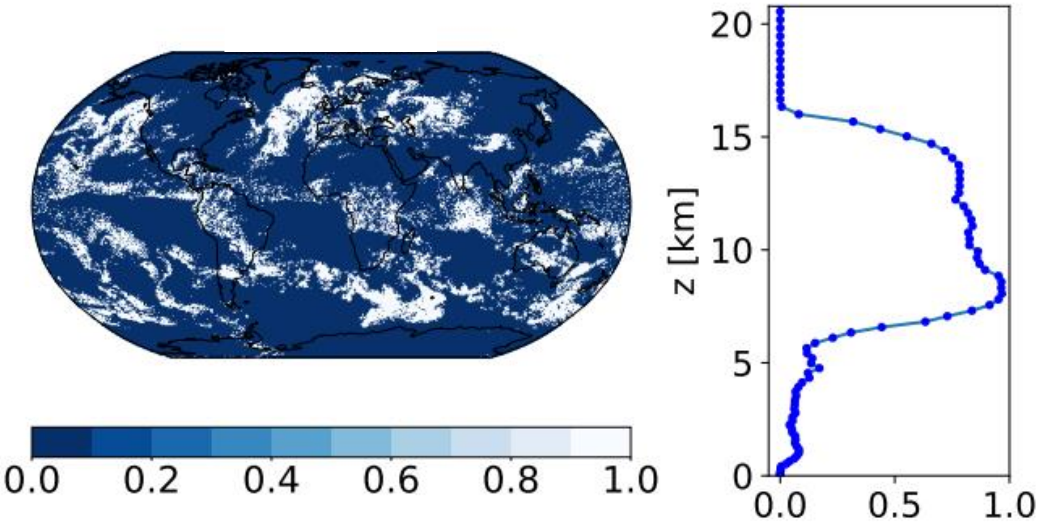

Coarse Graining

<u>80km-res "High-fidelity" Cl. Cov.</u>

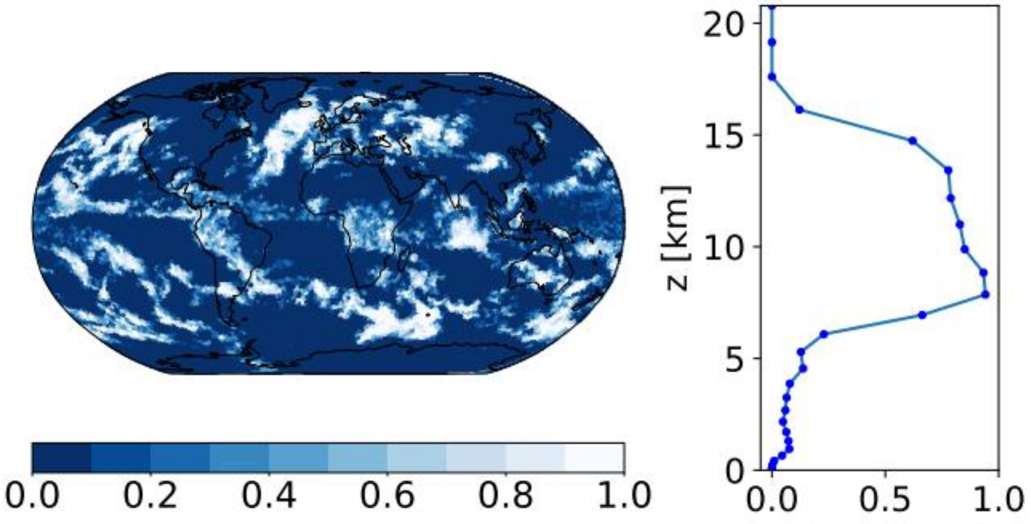

<u>Source</u>: Grundner, Beucler et al. (2022), Giorgetta et al. (2022), Stevens et al. (2019)

Cloud cover parameterization maps the grid-scale environment to the fraction occupied by clouds

<u>80km-res Coarse Environment</u>

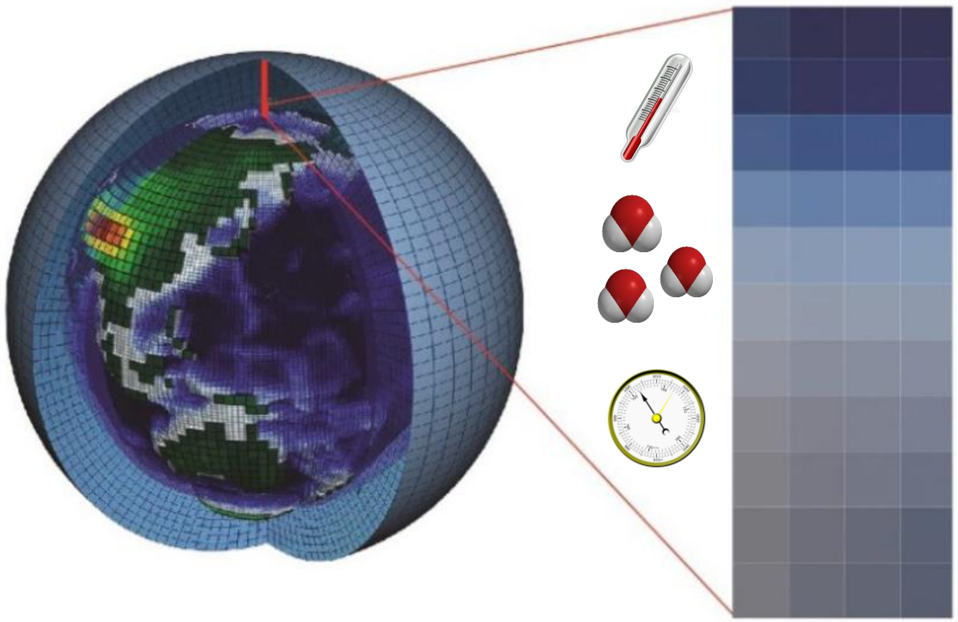

Goal

<u>80km-res “High-fidelity” Cl. Cov.</u>

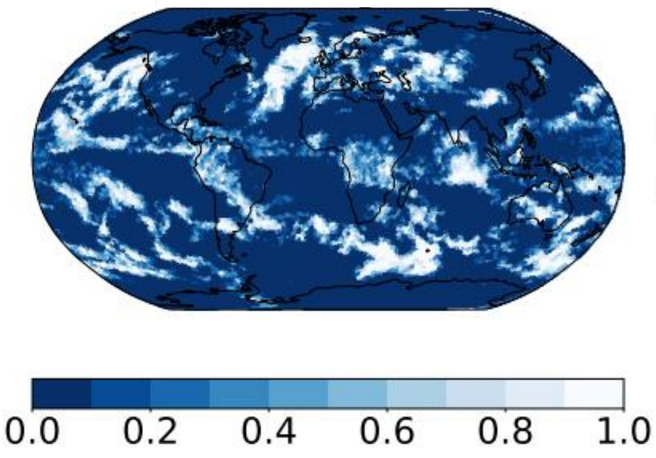

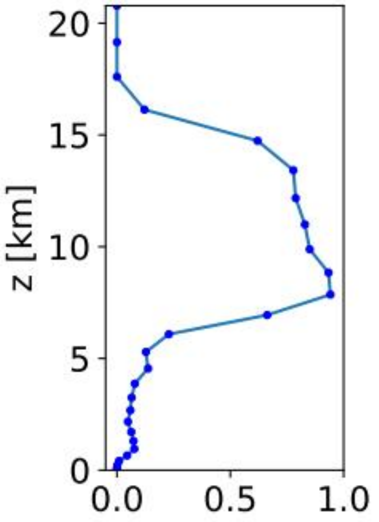

<u>Movie from</u>: Monsoon IV (Olbinski, 2017)

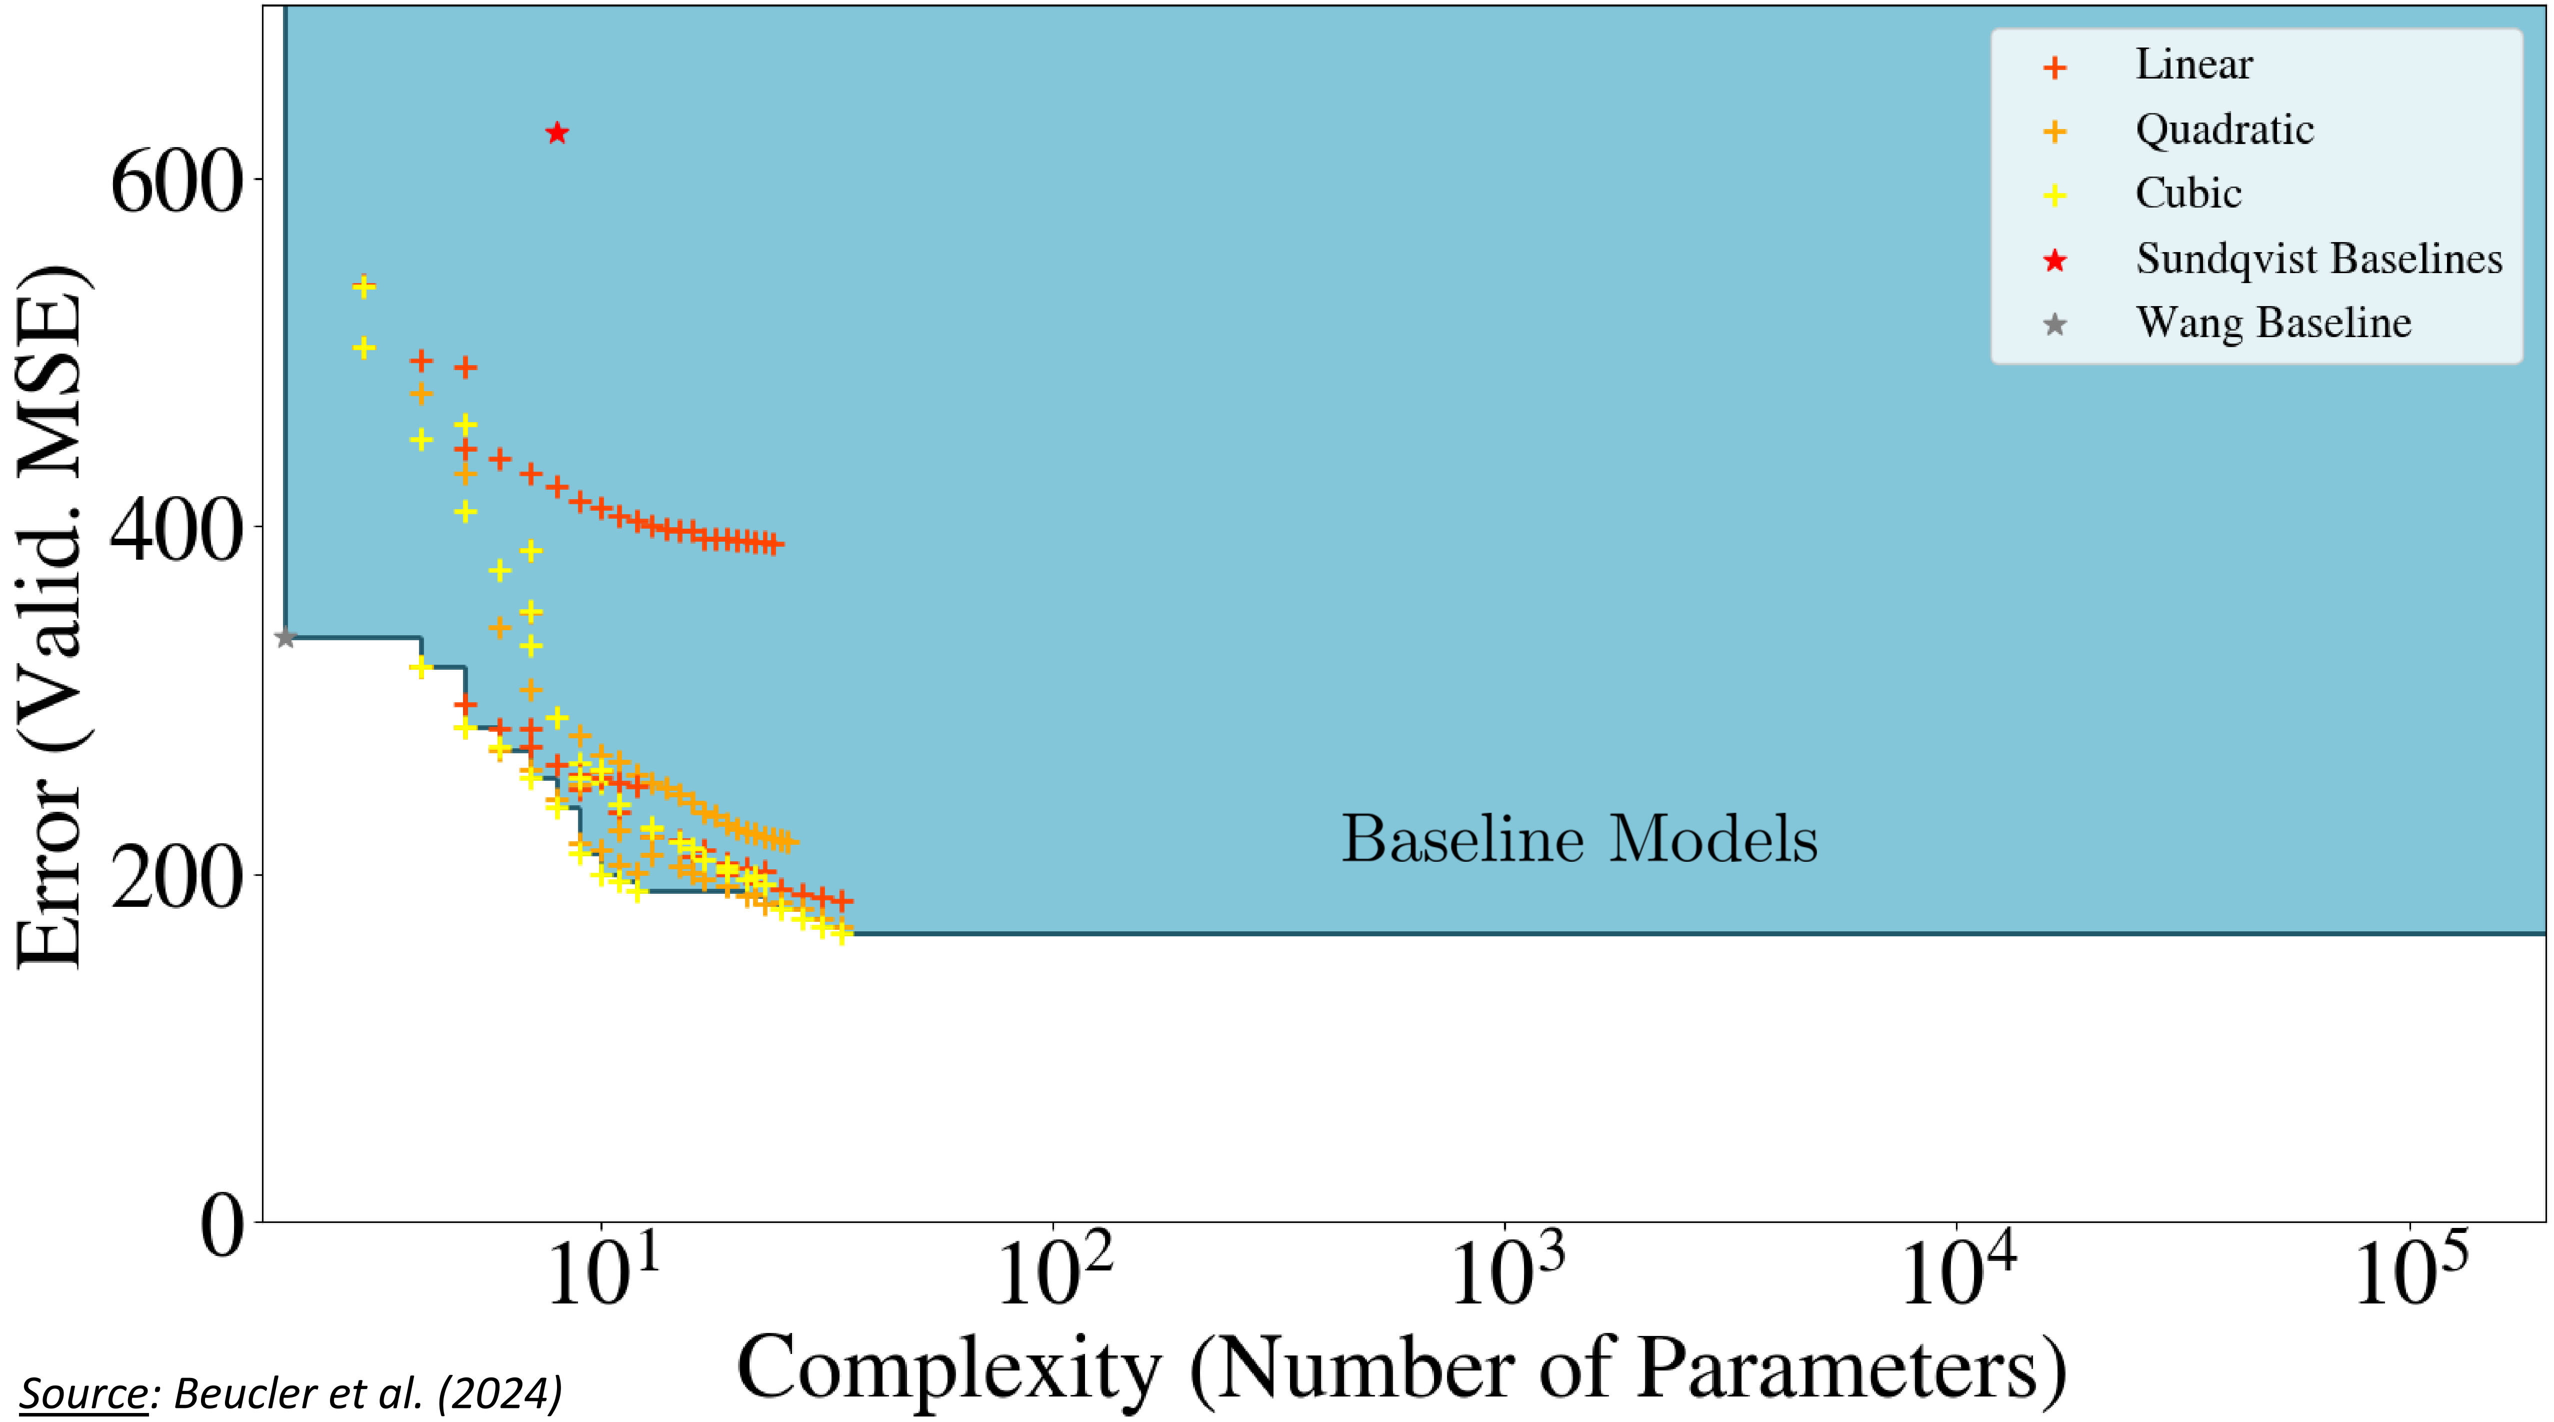

<u>Source</u>: *Beucler et al. (2024)*

Neural Nets achieve root-mean squared errors < 7%

Neural Network
Estimate

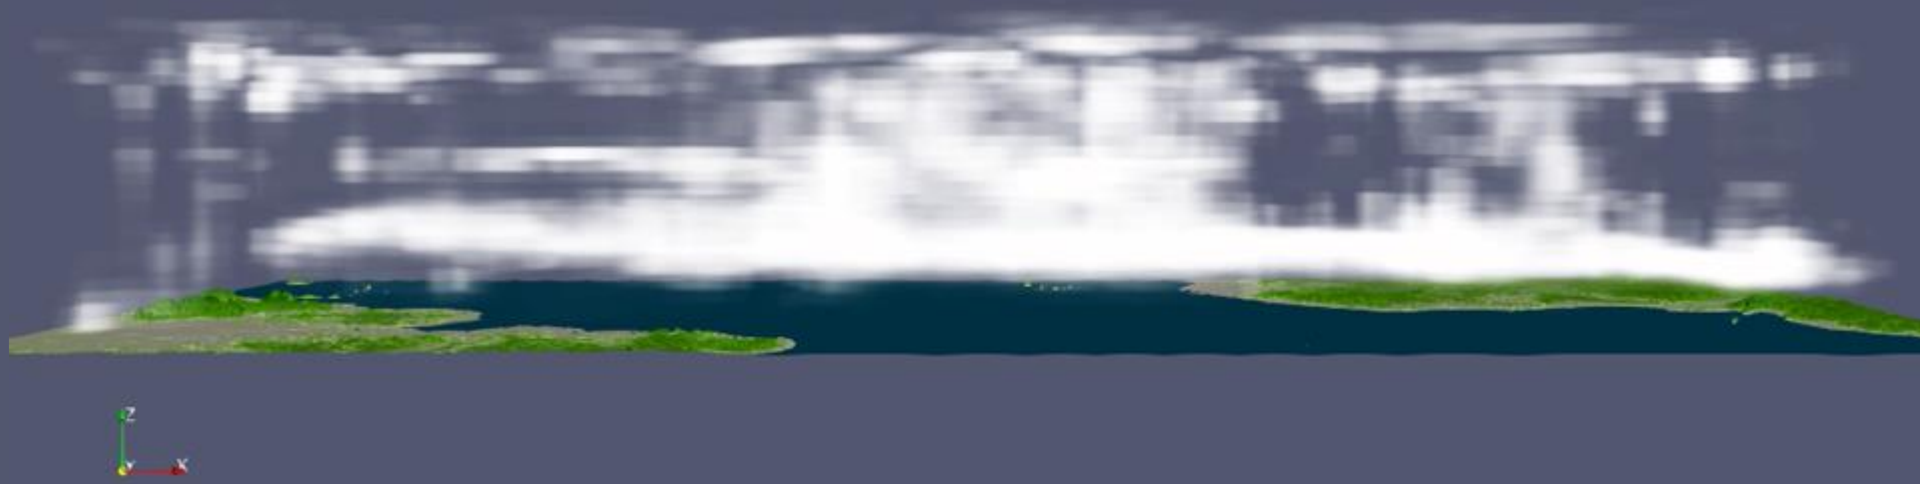

Reference
(Coarse-Grained
High-resolution
simulation)

Source: Grundner, Beucler et al. (2022)

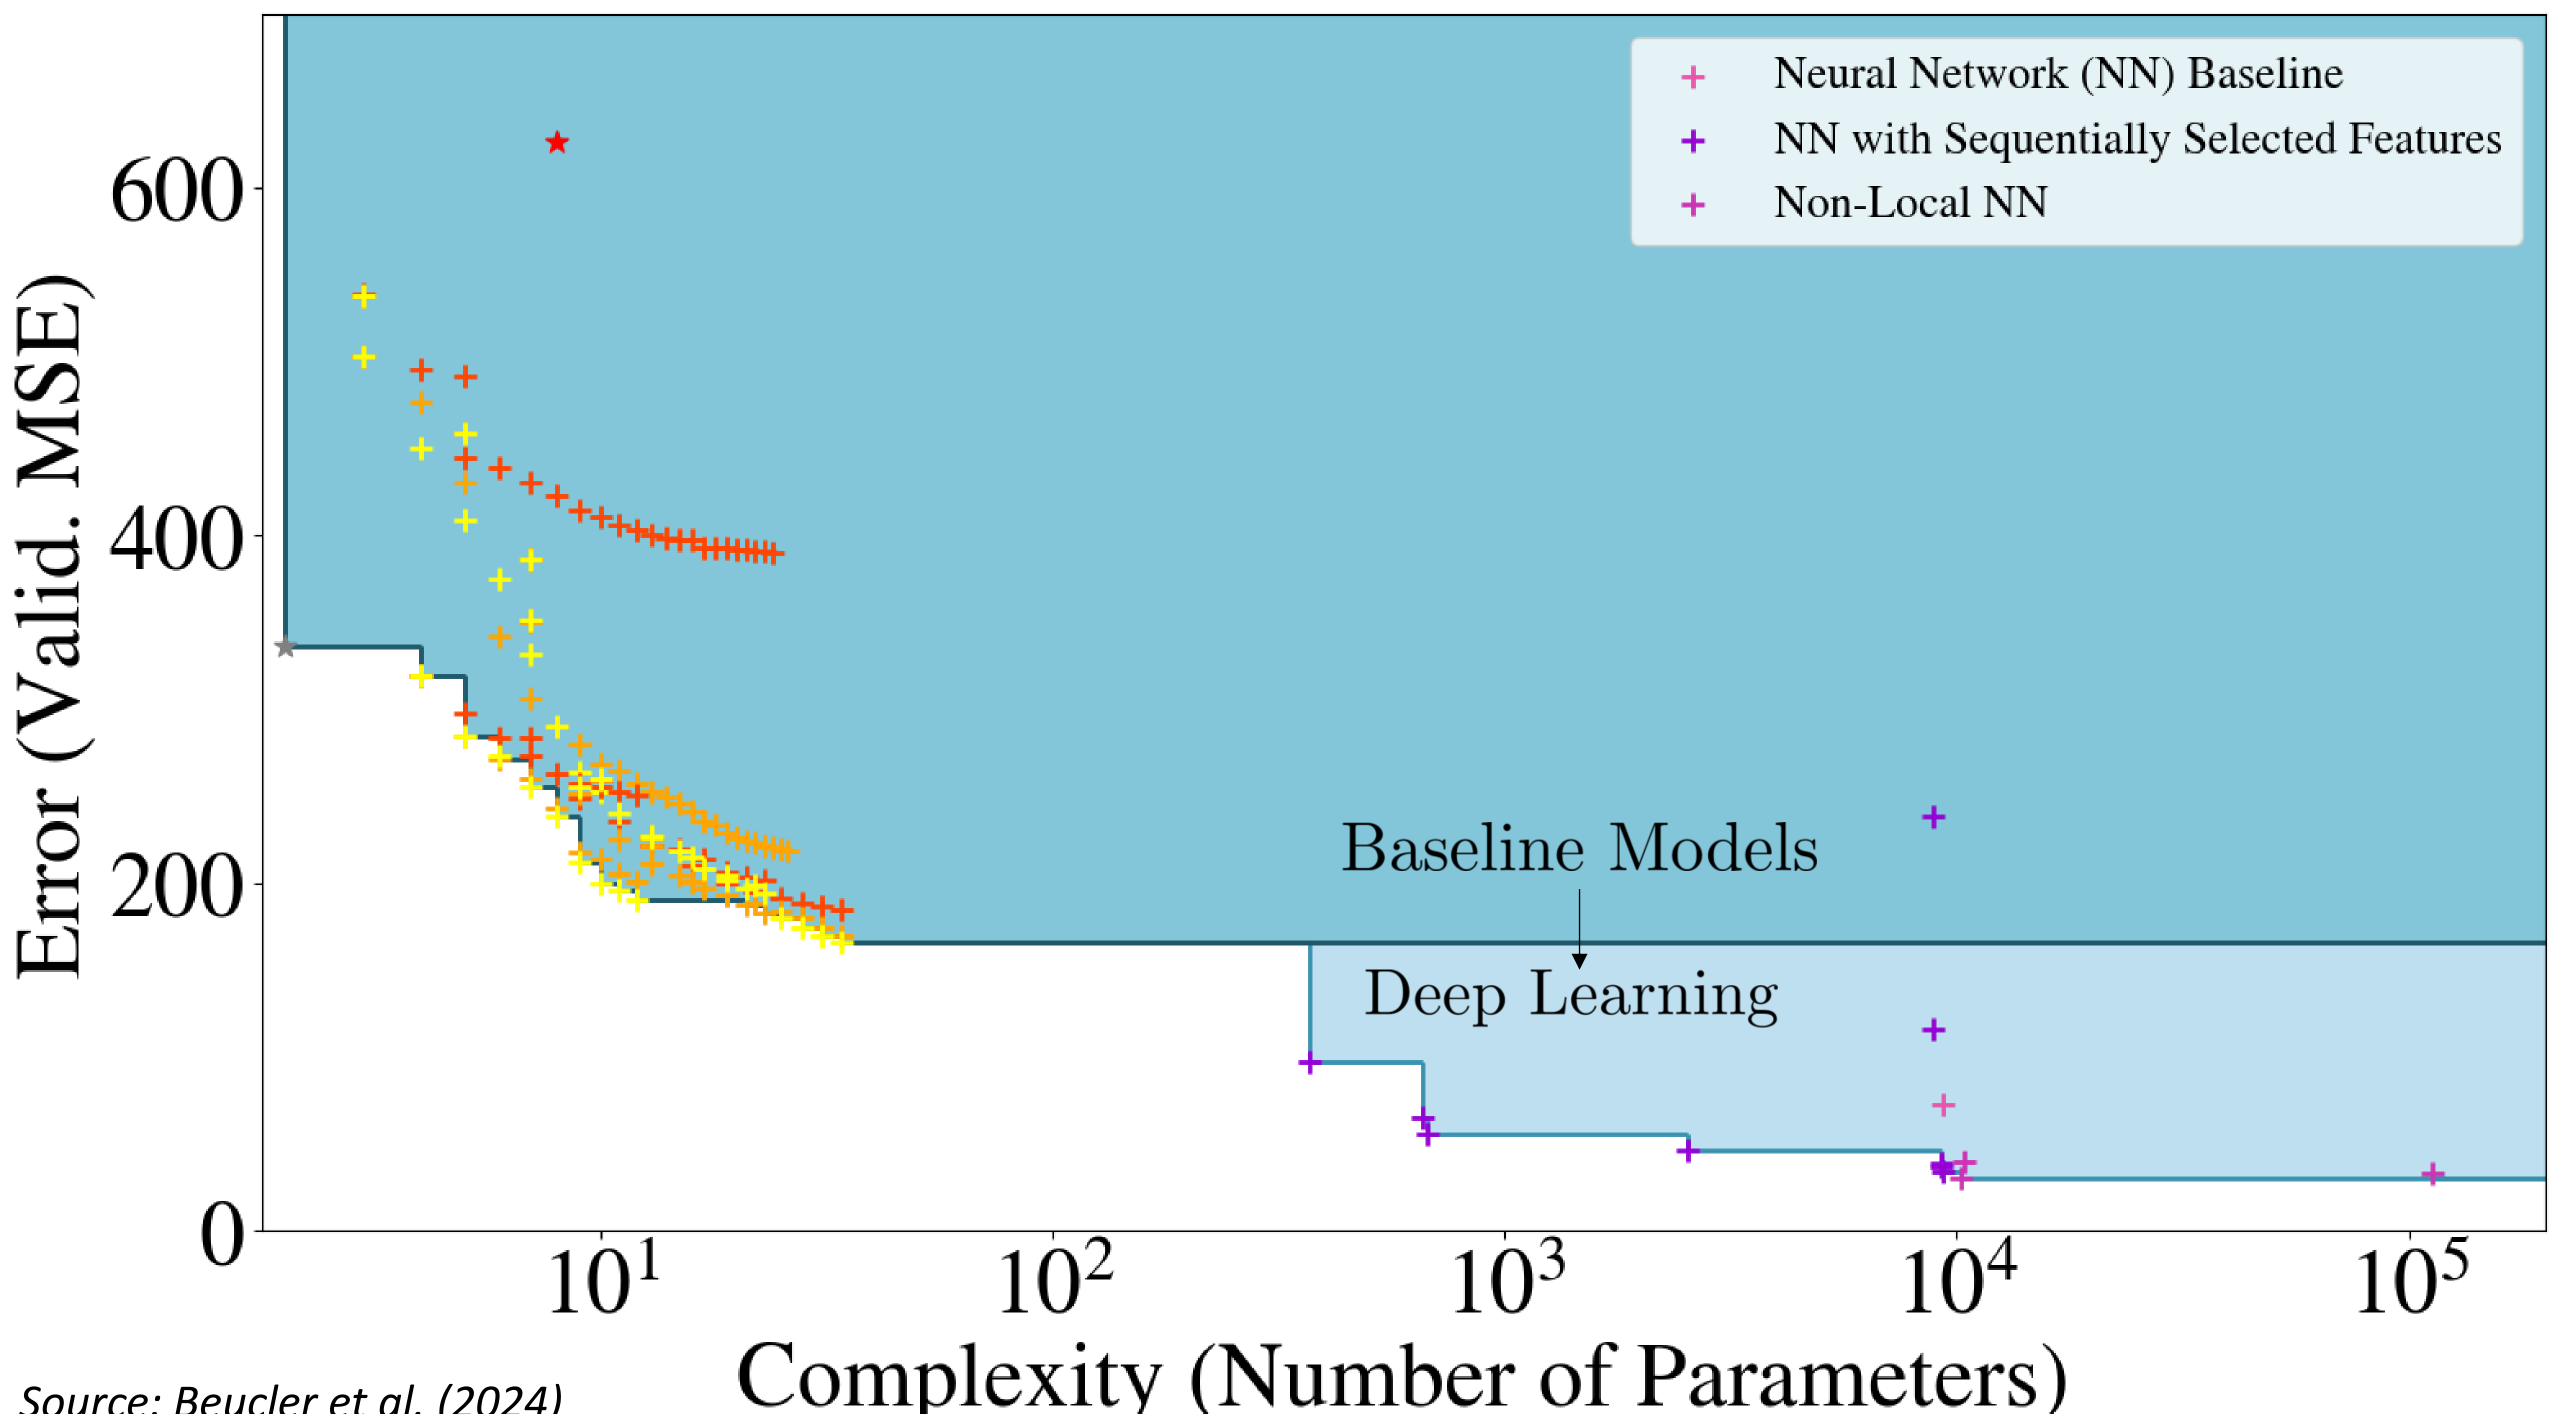

Source: Beucler et al. (2024)

And guide the discovery of new equations for cloud cover
Example of **transparent machine learning**…

$$C_{\text{PySR}} = \underbrace{\mathcal{I}_1(\text{RH}, T)}_{\text{Humidity/Temperature}} + \underbrace{\mathcal{I}_2\left(\frac{d\text{RH}}{dz}\right)}_{\text{Inversion}} + \underbrace{\mathcal{I}_3(q_\ell, q_i)}_{\text{Condensates}}$$

$$\begin{aligned}\mathcal{I}_1 \quad &= \overline{C} + \left(\frac{\partial C}{\partial \text{RH}}\right)_{\overline{\text{RH}},\overline{T}}\left(\text{RH} - \overline{\text{RH}}\right) - \left(\frac{\partial C}{\partial T}\right)_{\overline{\text{RH}},\overline{T}}\left(T - \overline{T}\right) \\ &+ \frac{1}{2}\left(\frac{\partial^2 C}{\partial \text{RH}^2}\right)_{\overline{\text{RH}},\overline{T}}\left(\text{RH} - \overline{\text{RH}}\right)^2 \\ &+ \frac{1}{2}\left(\frac{\partial C}{\partial \text{RH}\partial T^2}\right)_{\overline{\text{RH}},\overline{T}}\left(T - \overline{T}\right)^2\left(\text{RH} - \overline{\text{RH}}\right).\end{aligned}$$

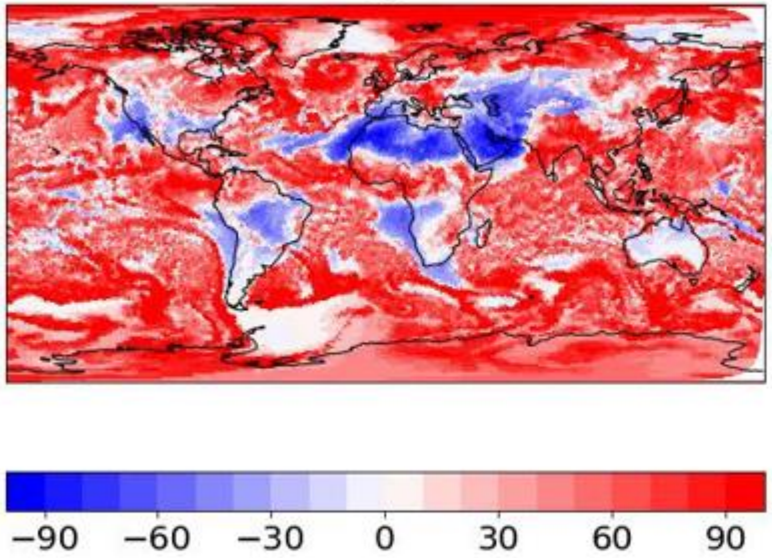

$$\mathcal{I}_2 = H^3_{\text{PySR}}\left[\frac{d\text{RH}}{dz} + \frac{3}{2}\left(\frac{d\text{RH}}{dz}\right)_{\max C}\right]\left(\frac{d\text{RH}}{dz}\right)^2$$

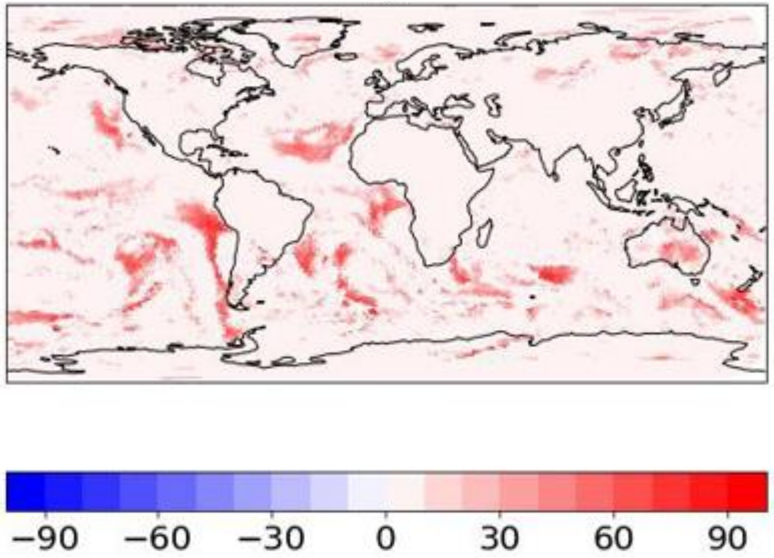

$$\mathcal{I}_3 = -\frac{1}{\epsilon_{\text{PySR}}} \times \frac{1}{1 + 2\epsilon_{\text{PySR}}\left(\lambda_\ell q_\ell + \lambda_i q_i\right)}$$

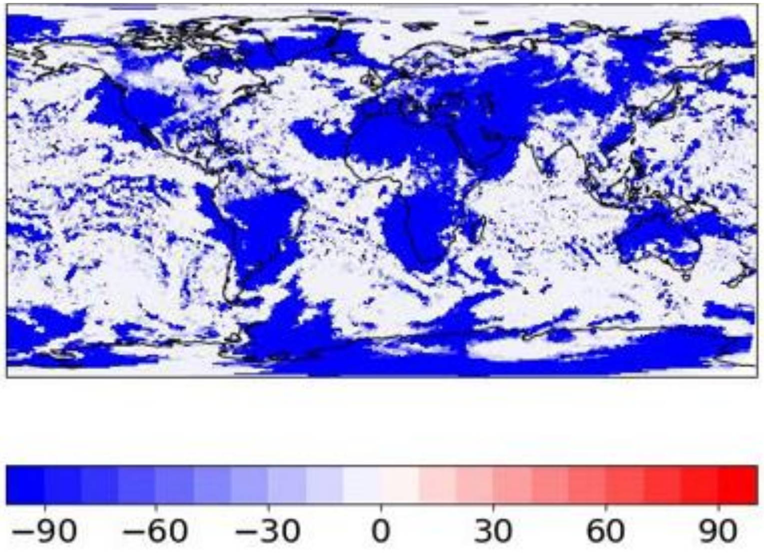

Source: *Grundner, Beucler et al. (2024)*; *Video source*: *PySR (2025)*

Unexpected discovery: The faster the subgrid distribution tends → 0, the more sensitive cloud cover is to cloud water concentration

$$\boxed{I_3(q_c, q_i) \stackrel{\text{def}}{=} \frac{-1}{q_c/a_8 + q_i/a_9 + \epsilon}}$$

$\neq$

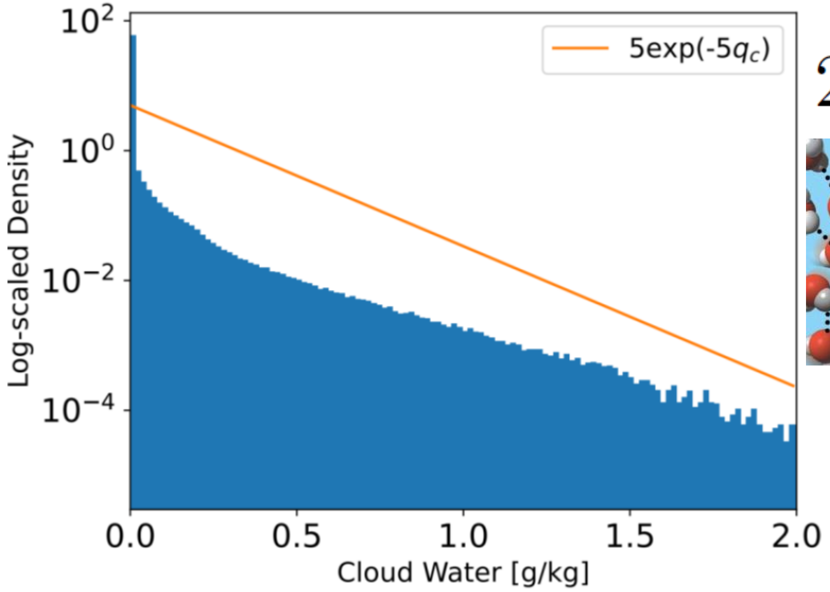

$$2\lambda_l \approx 1/(a_8\epsilon^2)$$

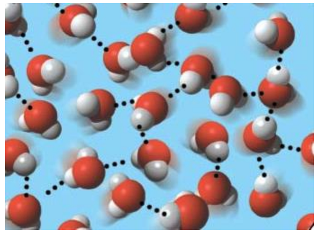

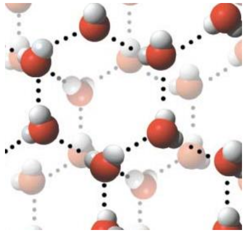

$$2\lambda_i \approx 1/(a_9\epsilon^2)$$

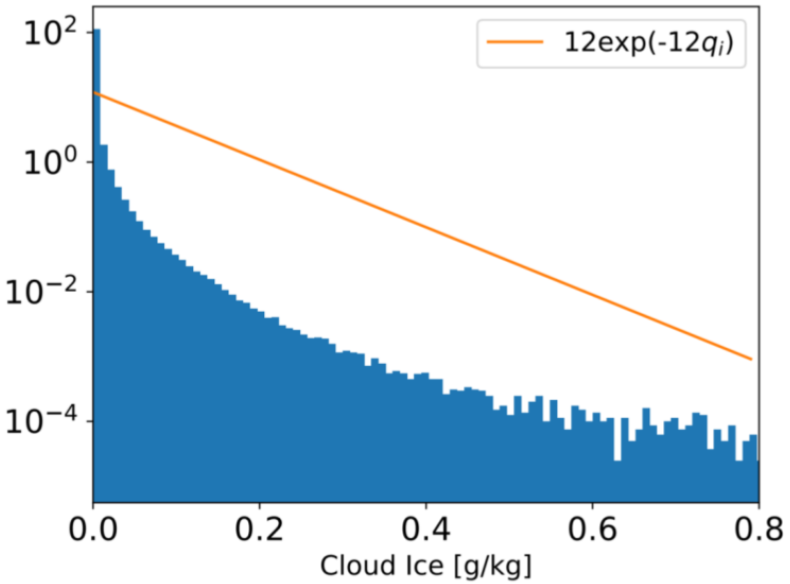

Source: Grundner et al. (2024), 2013 Pearson Ed.

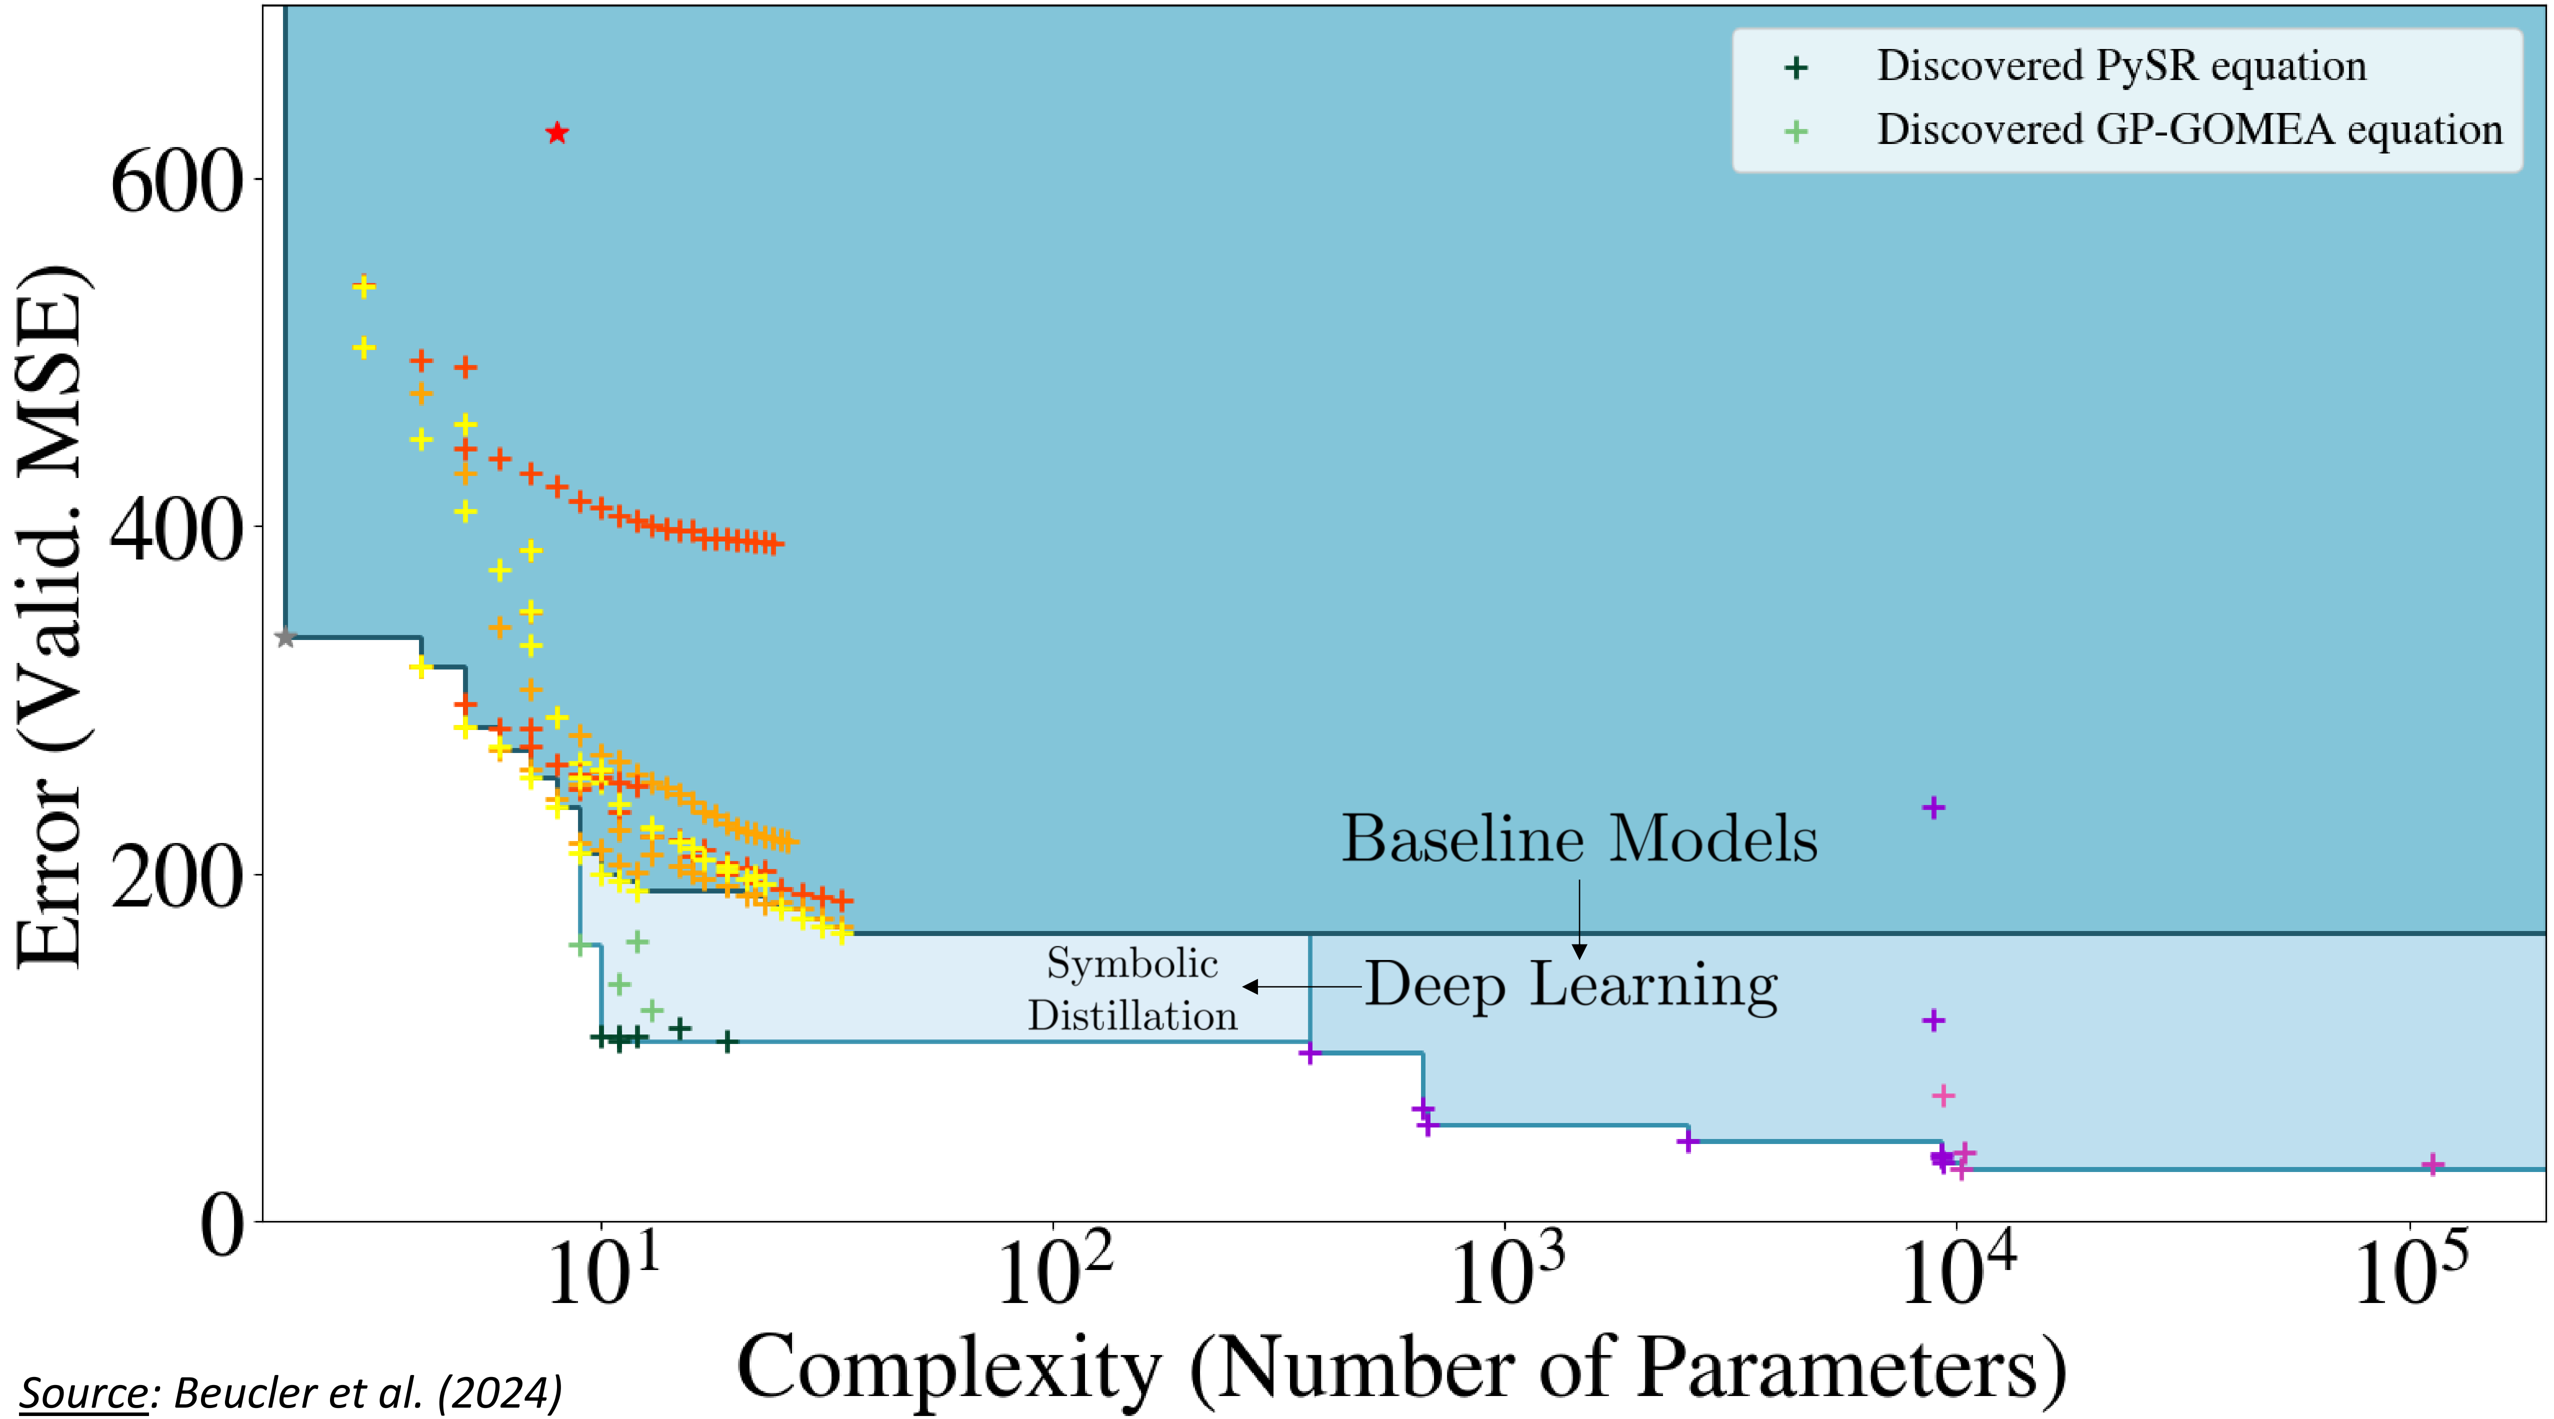

<u>Source</u>: *Beucler et al. (2024)*

Example of **transparent ML** that reduces biases of the resulting atmospheric simulation (below: outgoing TOA LW fluxes)

$$C_{\text{PySR}} = \underbrace{\mathcal{I}_1(\text{RH}, T)}_{\text{Humidity/Temperature}} + \underbrace{\mathcal{I}_2\left(\frac{d\text{RH}}{dz}\right)}_{\text{Inversion}} + \underbrace{\mathcal{I}_3(q_\ell, q_i)}_{\text{Condensates}}$$

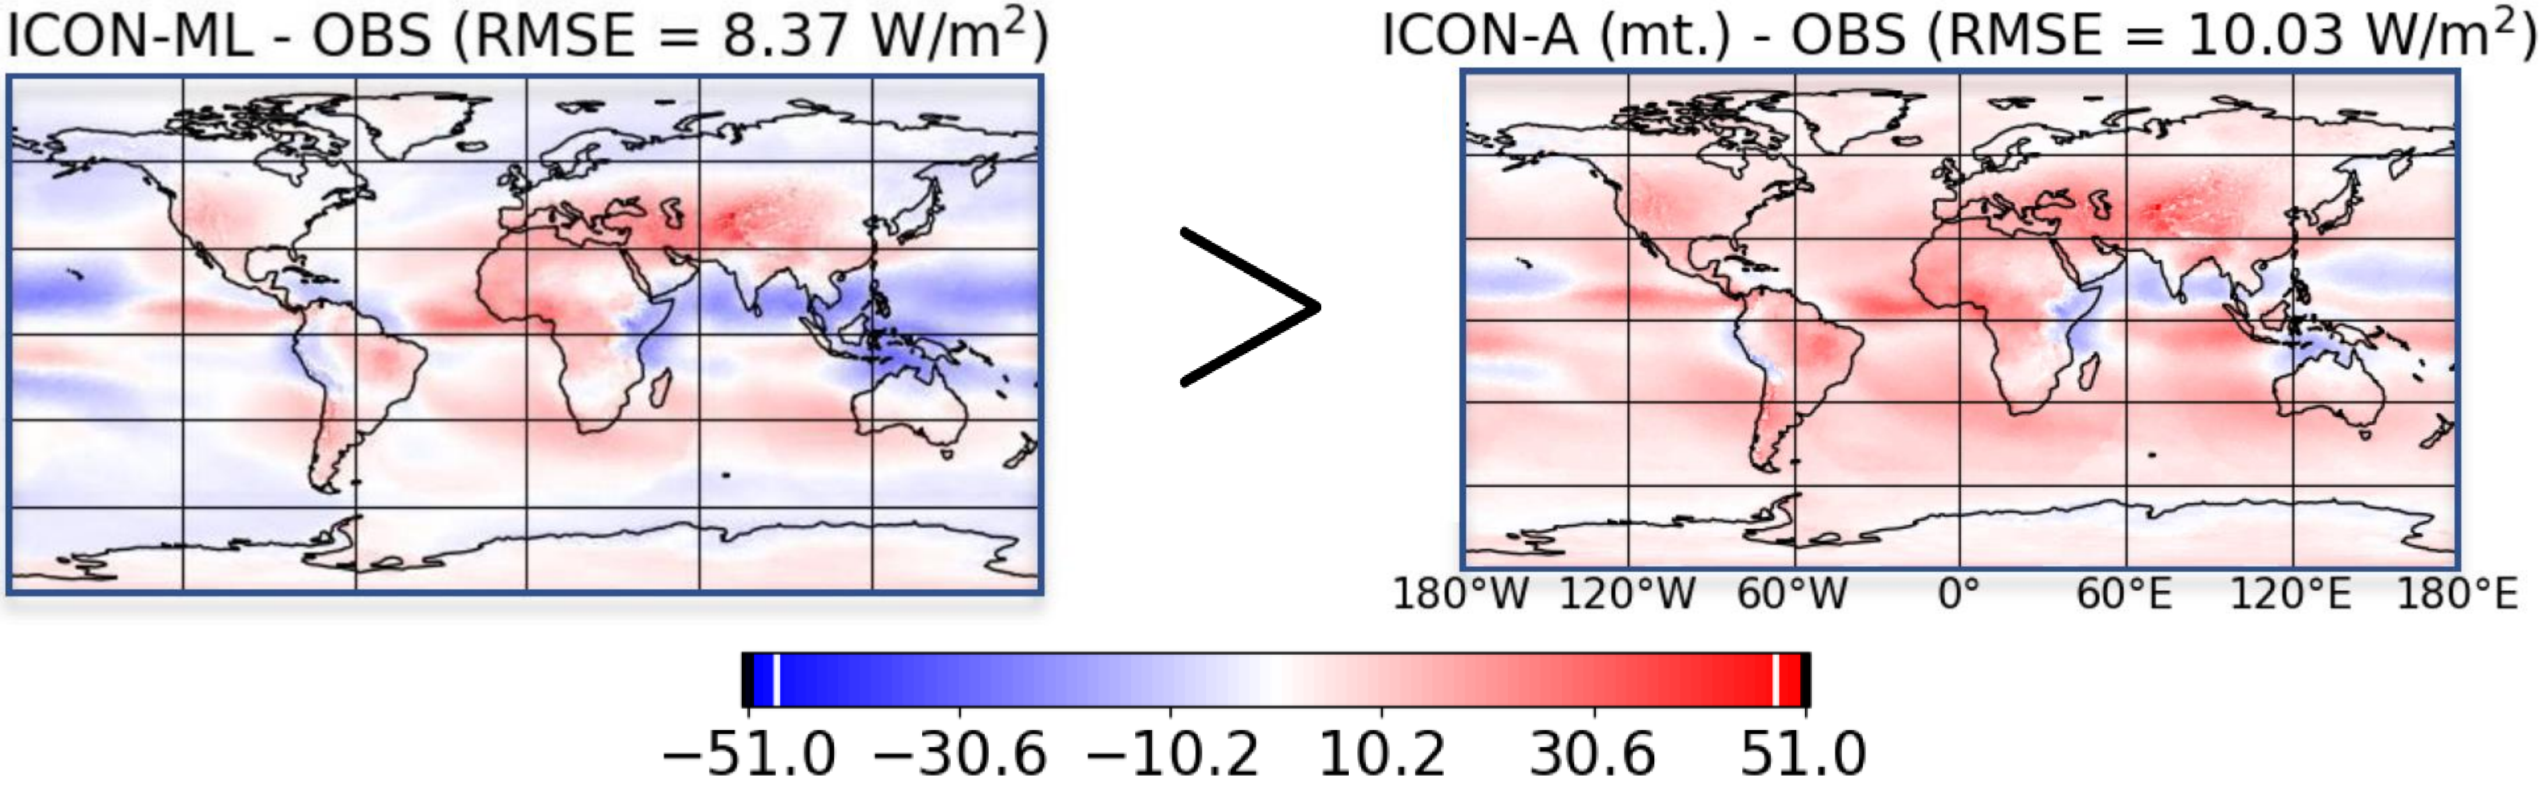

Source: Grundner, Beucler, … & Eyring (Submitted, Preprint coming soon)

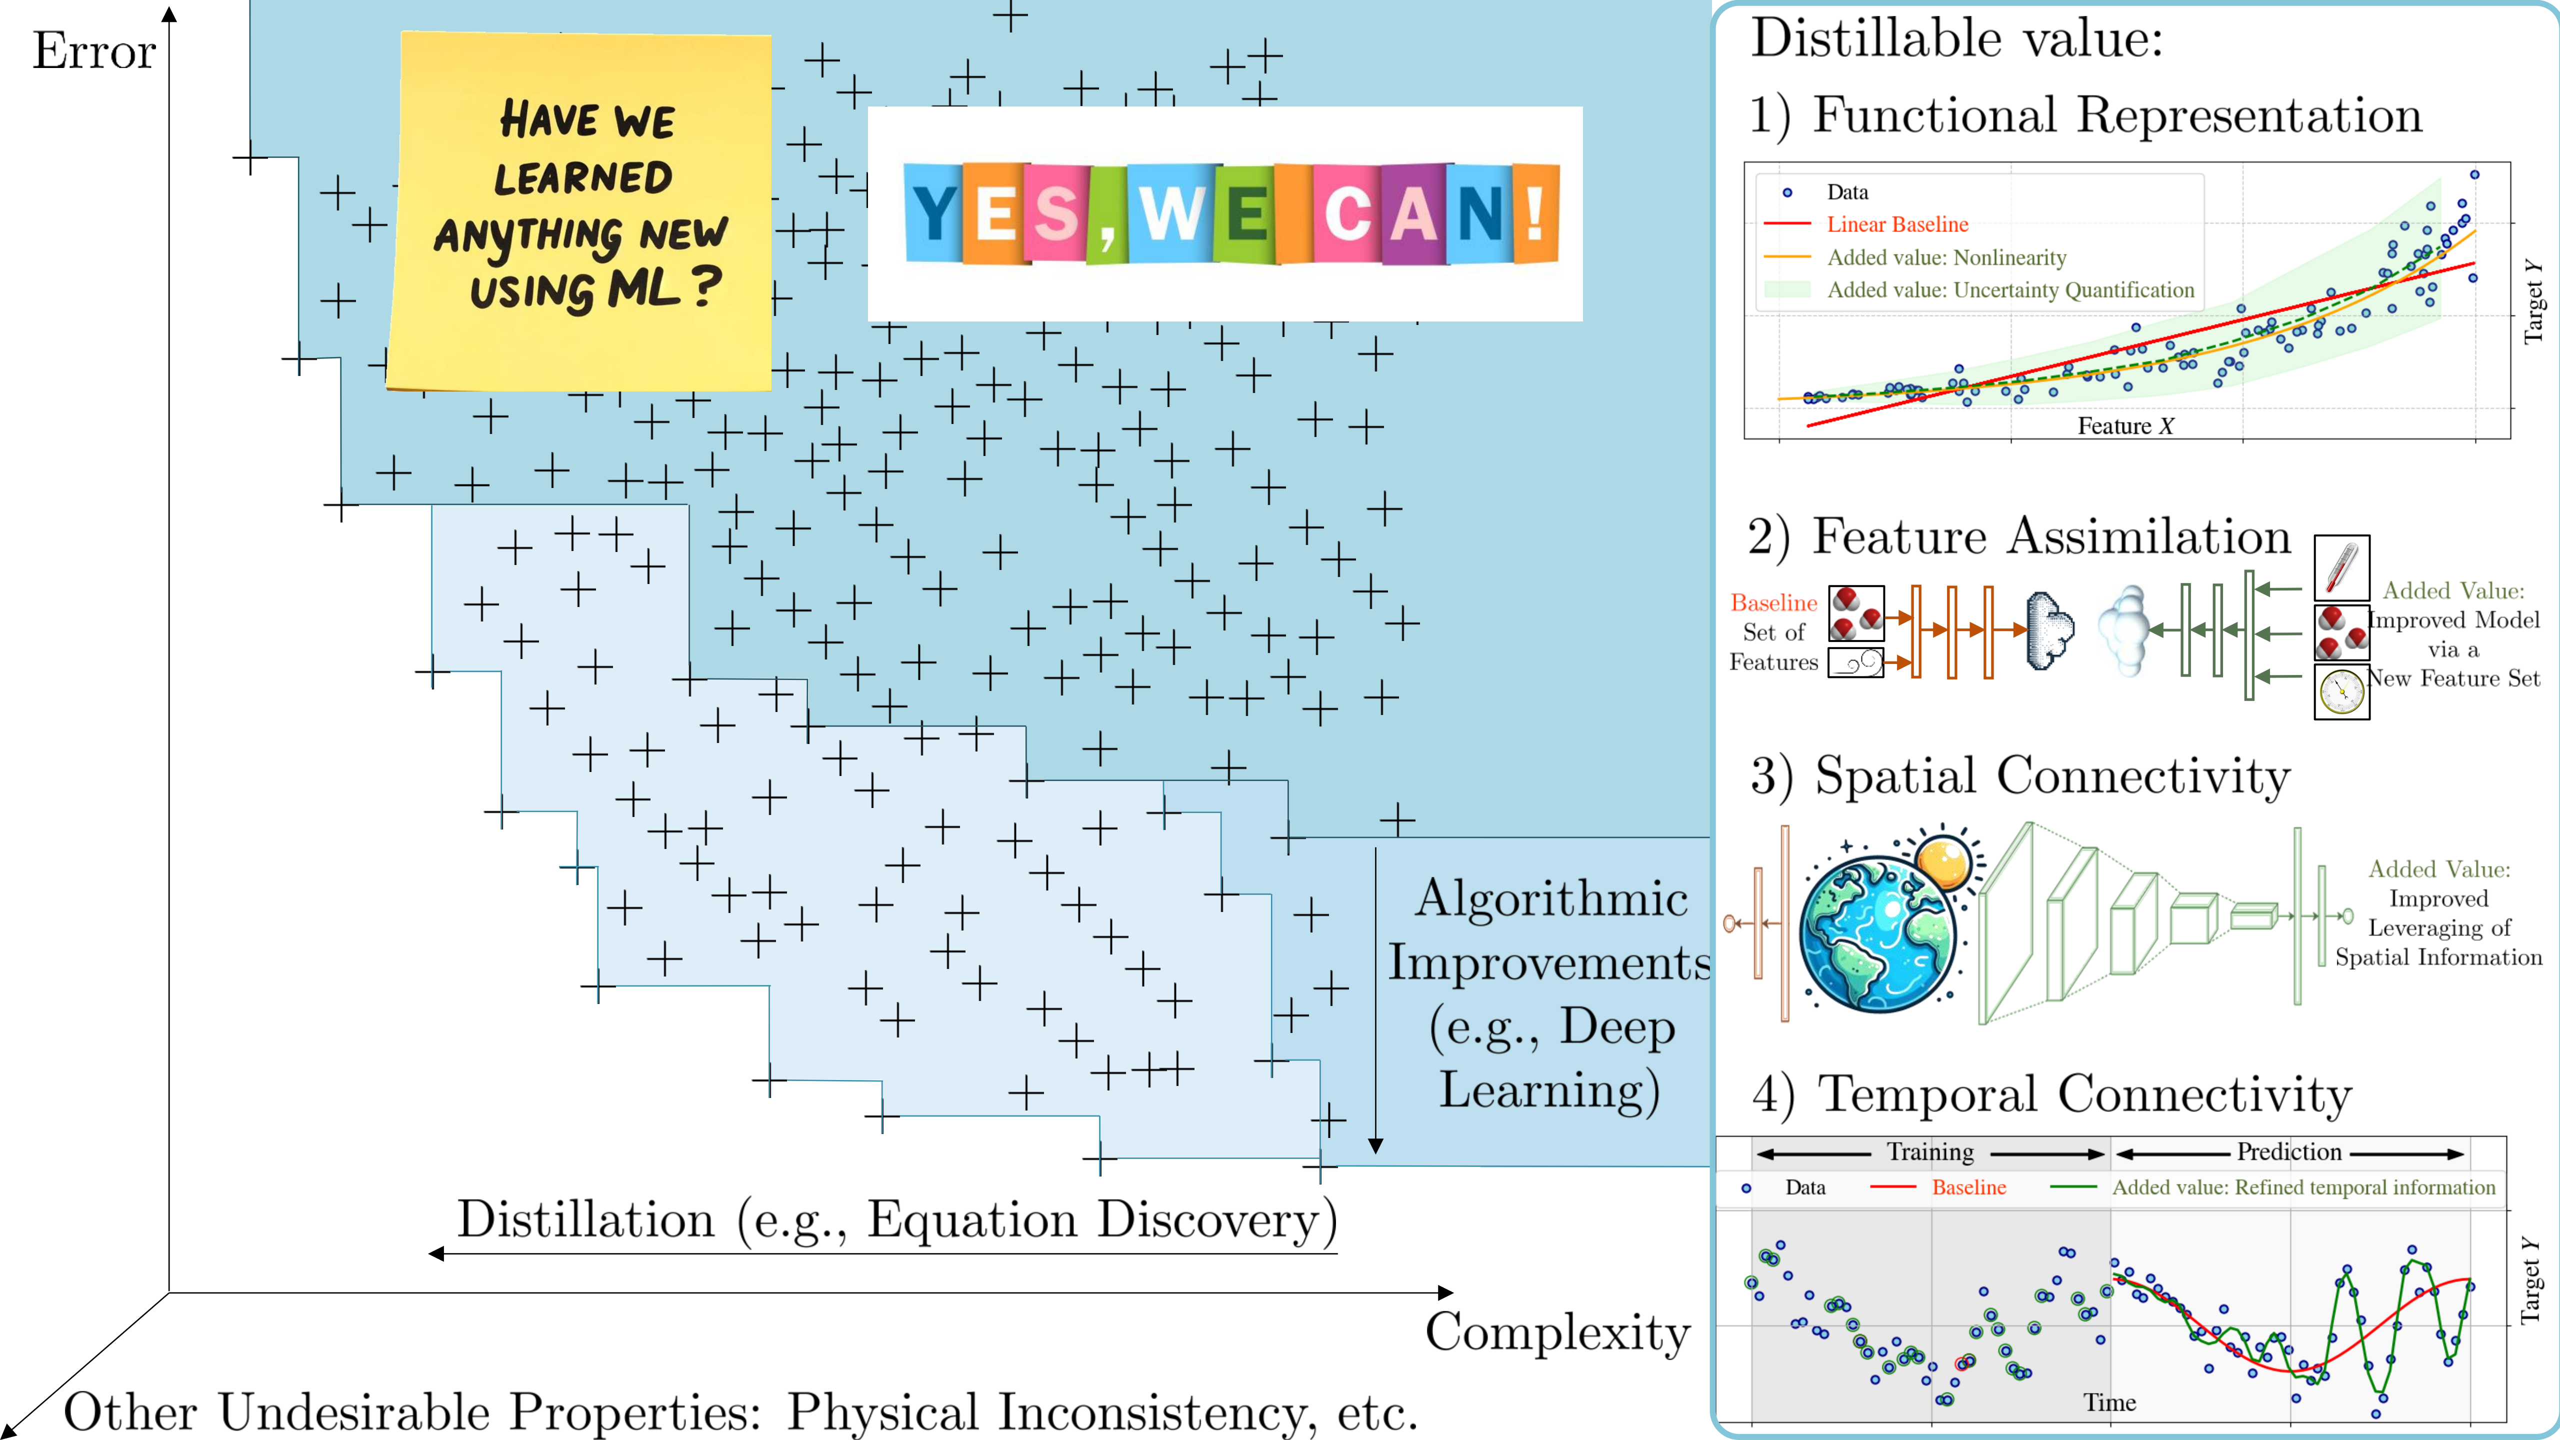
Error
HAVE WE LEARNED ANYTHING NEW USING ML?
YES, WE CAN!
Algorithmic Improvements (e.g., Deep Learning)
Distillation (e.g., Equation Discovery)
Complexity
Other Undesirable Properties: Physical Inconsistency, etc.
Distillable value:
1) Functional Representation
Data
Linear Baseline
Added value: Nonlinearity
Added value: Uncertainty Quantification
Feature X
Target Y
2) Feature Assimilation
Baseline Set of Features
Added Value: Improved Model via a New Feature Set
3) Spatial Connectivity
Added Value: Improved Leveraging of Spatial Information
4) Temporal Connectivity
Training
Prediction
Data
Baseline
Added value: Refined temporal information
Time
Target Y

Group activity:

1) Form groups of approximately 5 people.
2) For 15 minutes, discuss and propose equations or conceptual solutions to:

- How can we make hybrid climate–AI models more reliable for climate change projections? Think of the key outcomes we ultimately care about?
- How can AI help advance our understanding of the climate system?

Next 15 min: Each group presents a 1-minute summary of their ideas on the board.

…but remain under-used in climate science

$$c_{\text{frac}} = (p, q_v, q_l, q_i, T)$$

Cloud Fraction Parameterization (Grundner et al., 2024)

Subgrid-Scale Thermodynamics Parameterization (Beucler et al., 2024)

$$\begin{pmatrix} \dot{T} \\ \dot{q} \end{pmatrix} = (p, q_v, T, \mathrm{LHF}, \mathrm{SHF}, S_0)$$

IS ML RELIABLE FOR CLIMATE PROJECTIONS?

HAVE WE LEARNED ANYTHING NEW USING ML?

Improving the representation of subgrid-scale thermodynamics in CAM

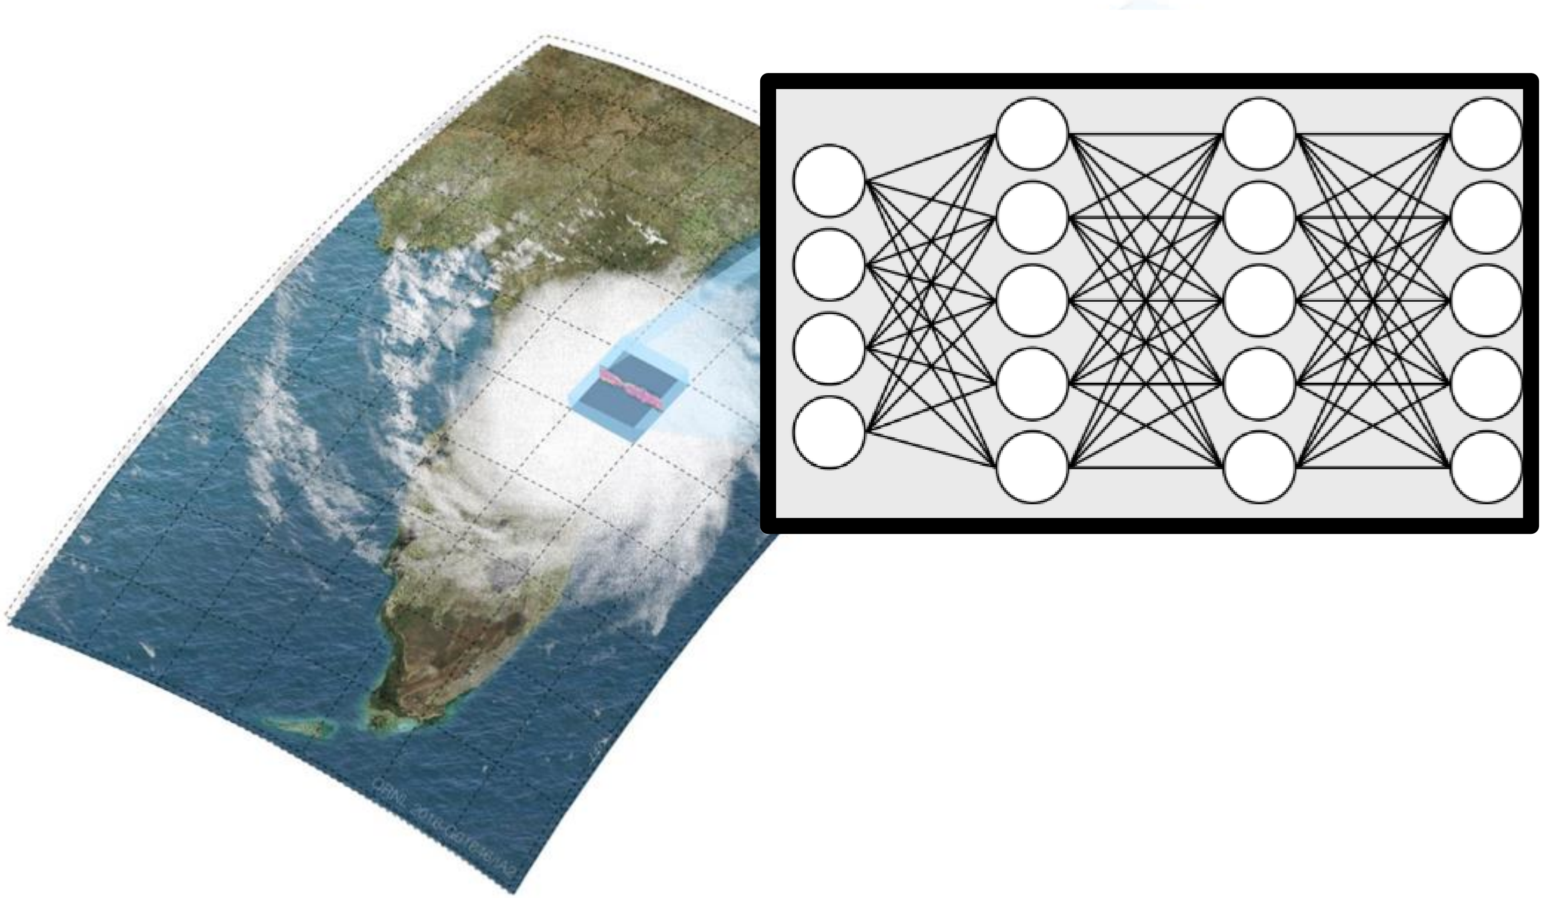

Once trained, neural networks accelerate the simulation 20x

Setup: Super-Parameterized Community Atmosphere Model v3.0

Image source: e3sm.org, Model source: Khairoutdinov et al. (2004)

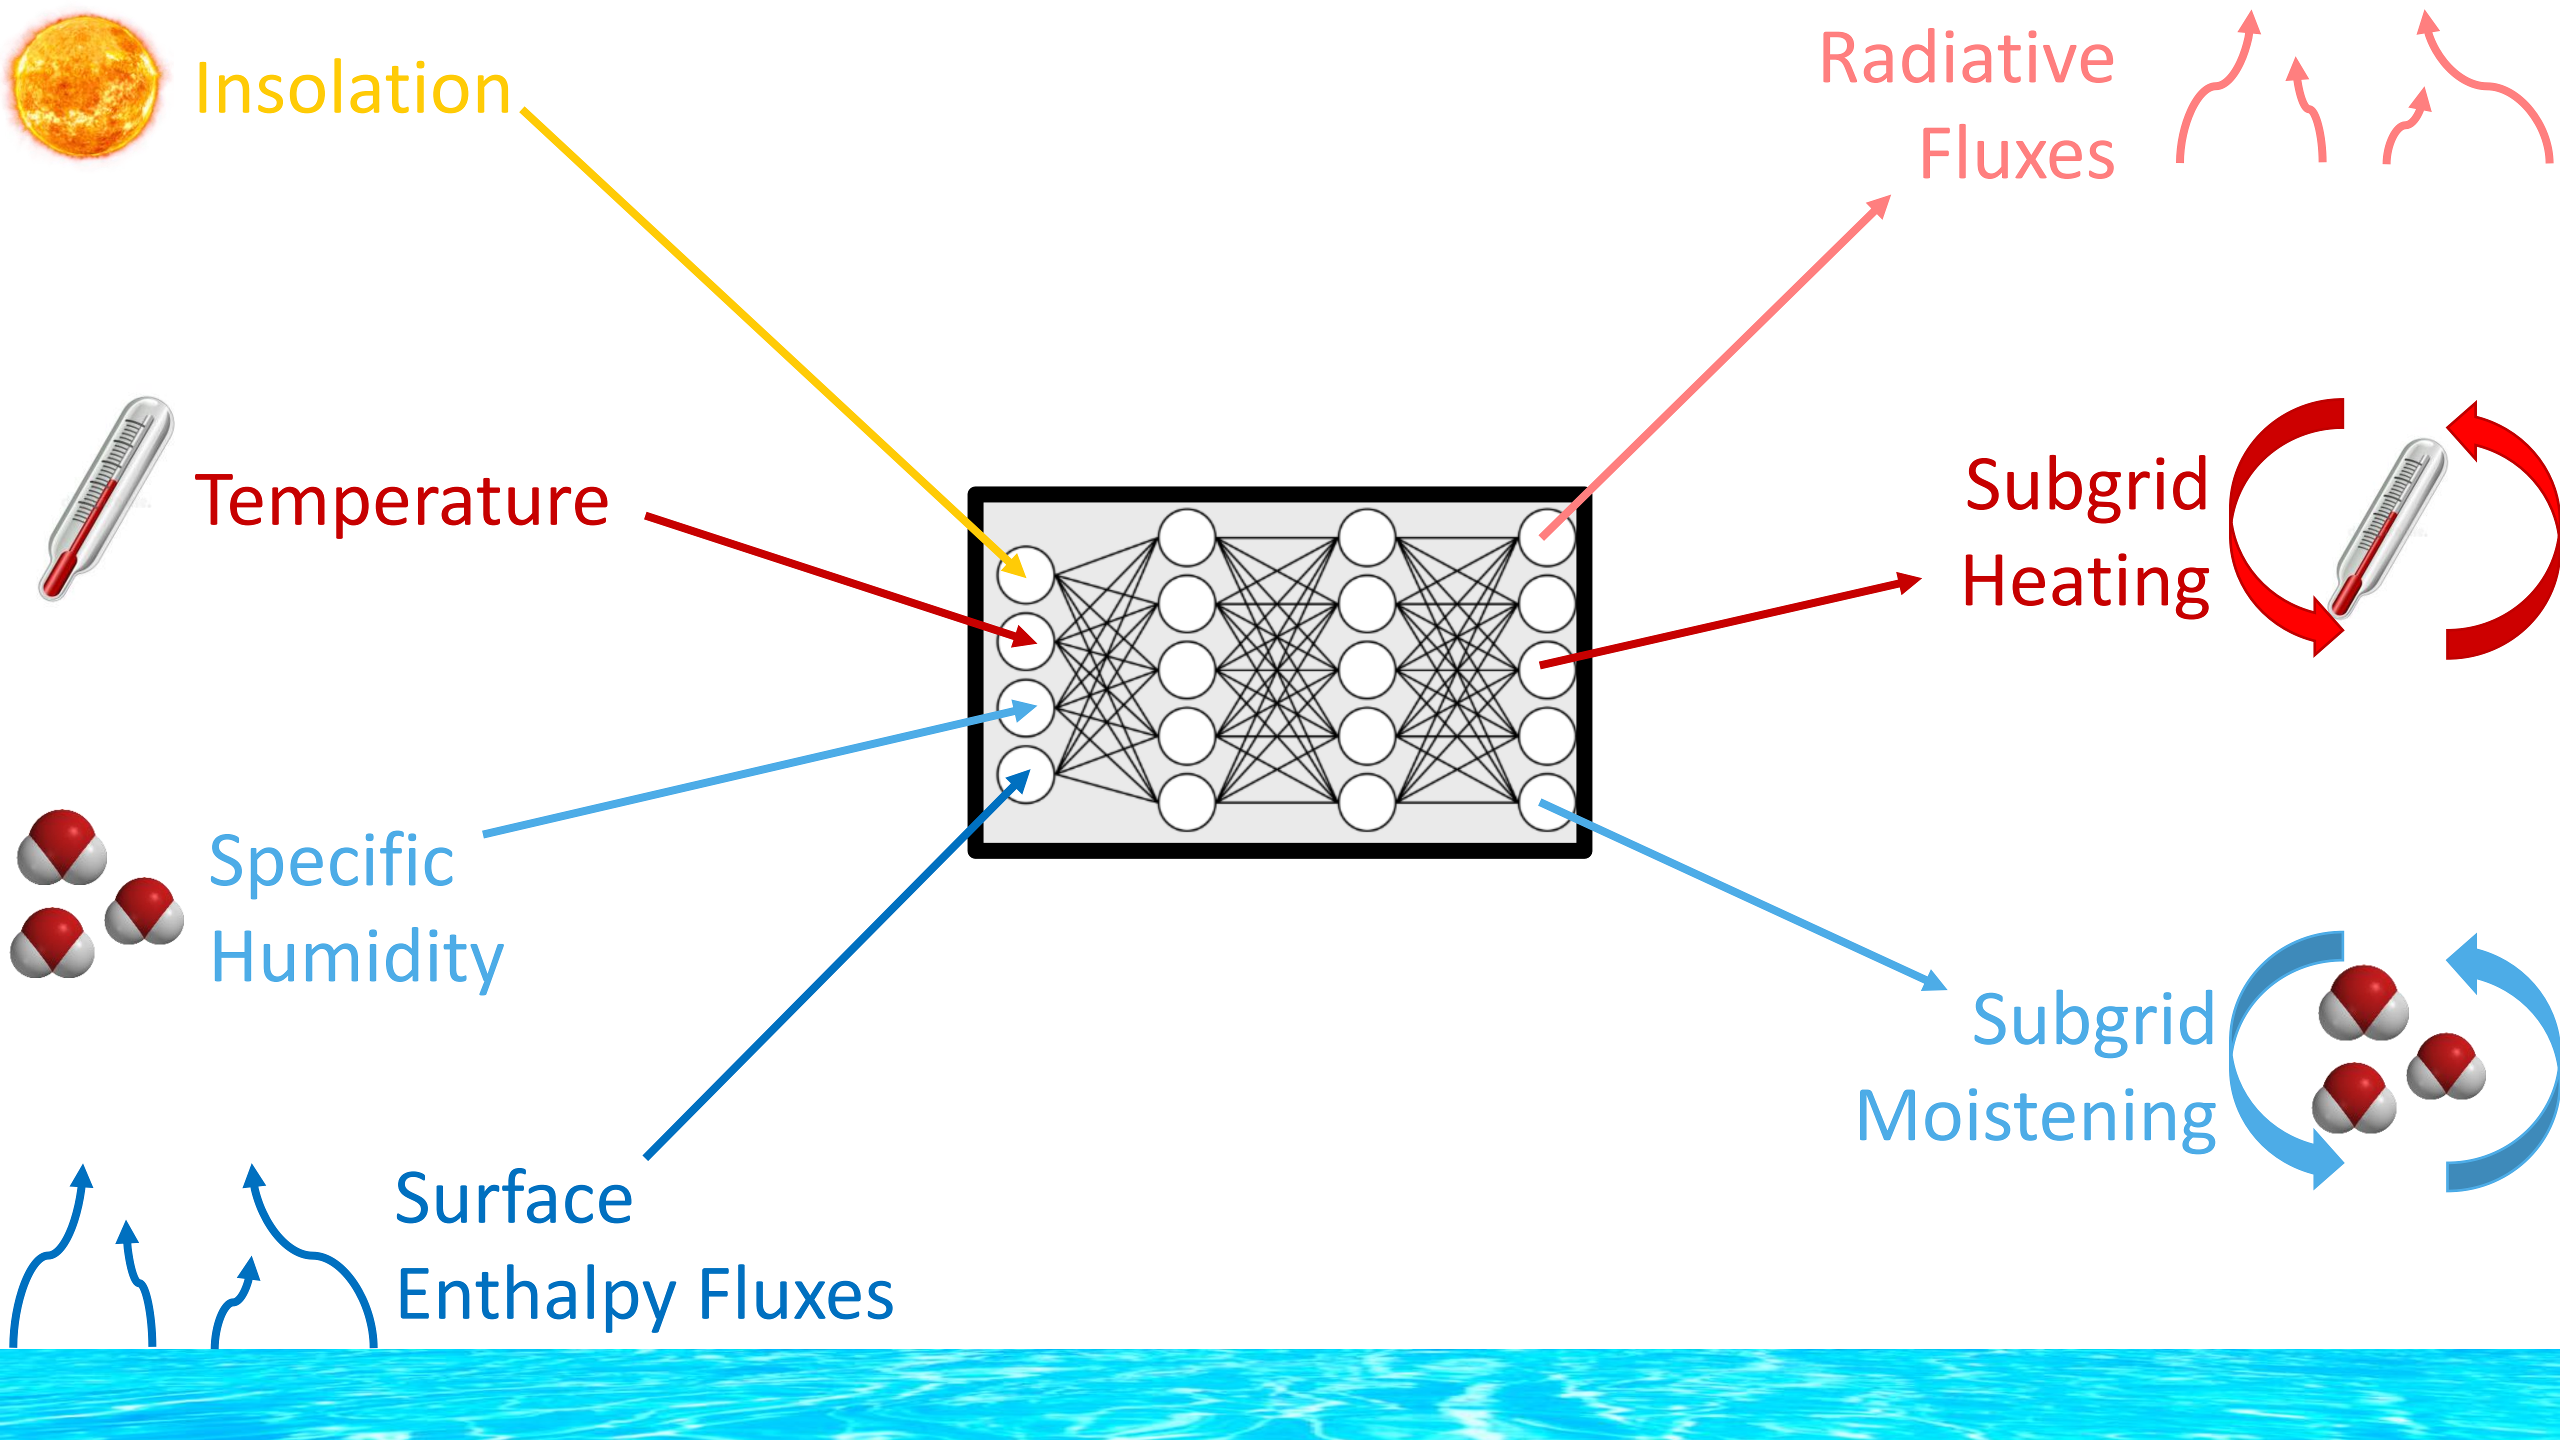
Insolation
Temperature
Specific
Humidity
Surface
Enthalpy Fluxes
Radiative
Fluxes
Subgrid
Heating
Subgrid
Moistening

Problem: Neural Nets often fail to generalize out-of-distribution

Daily-mean Tropical prediction in reference climate

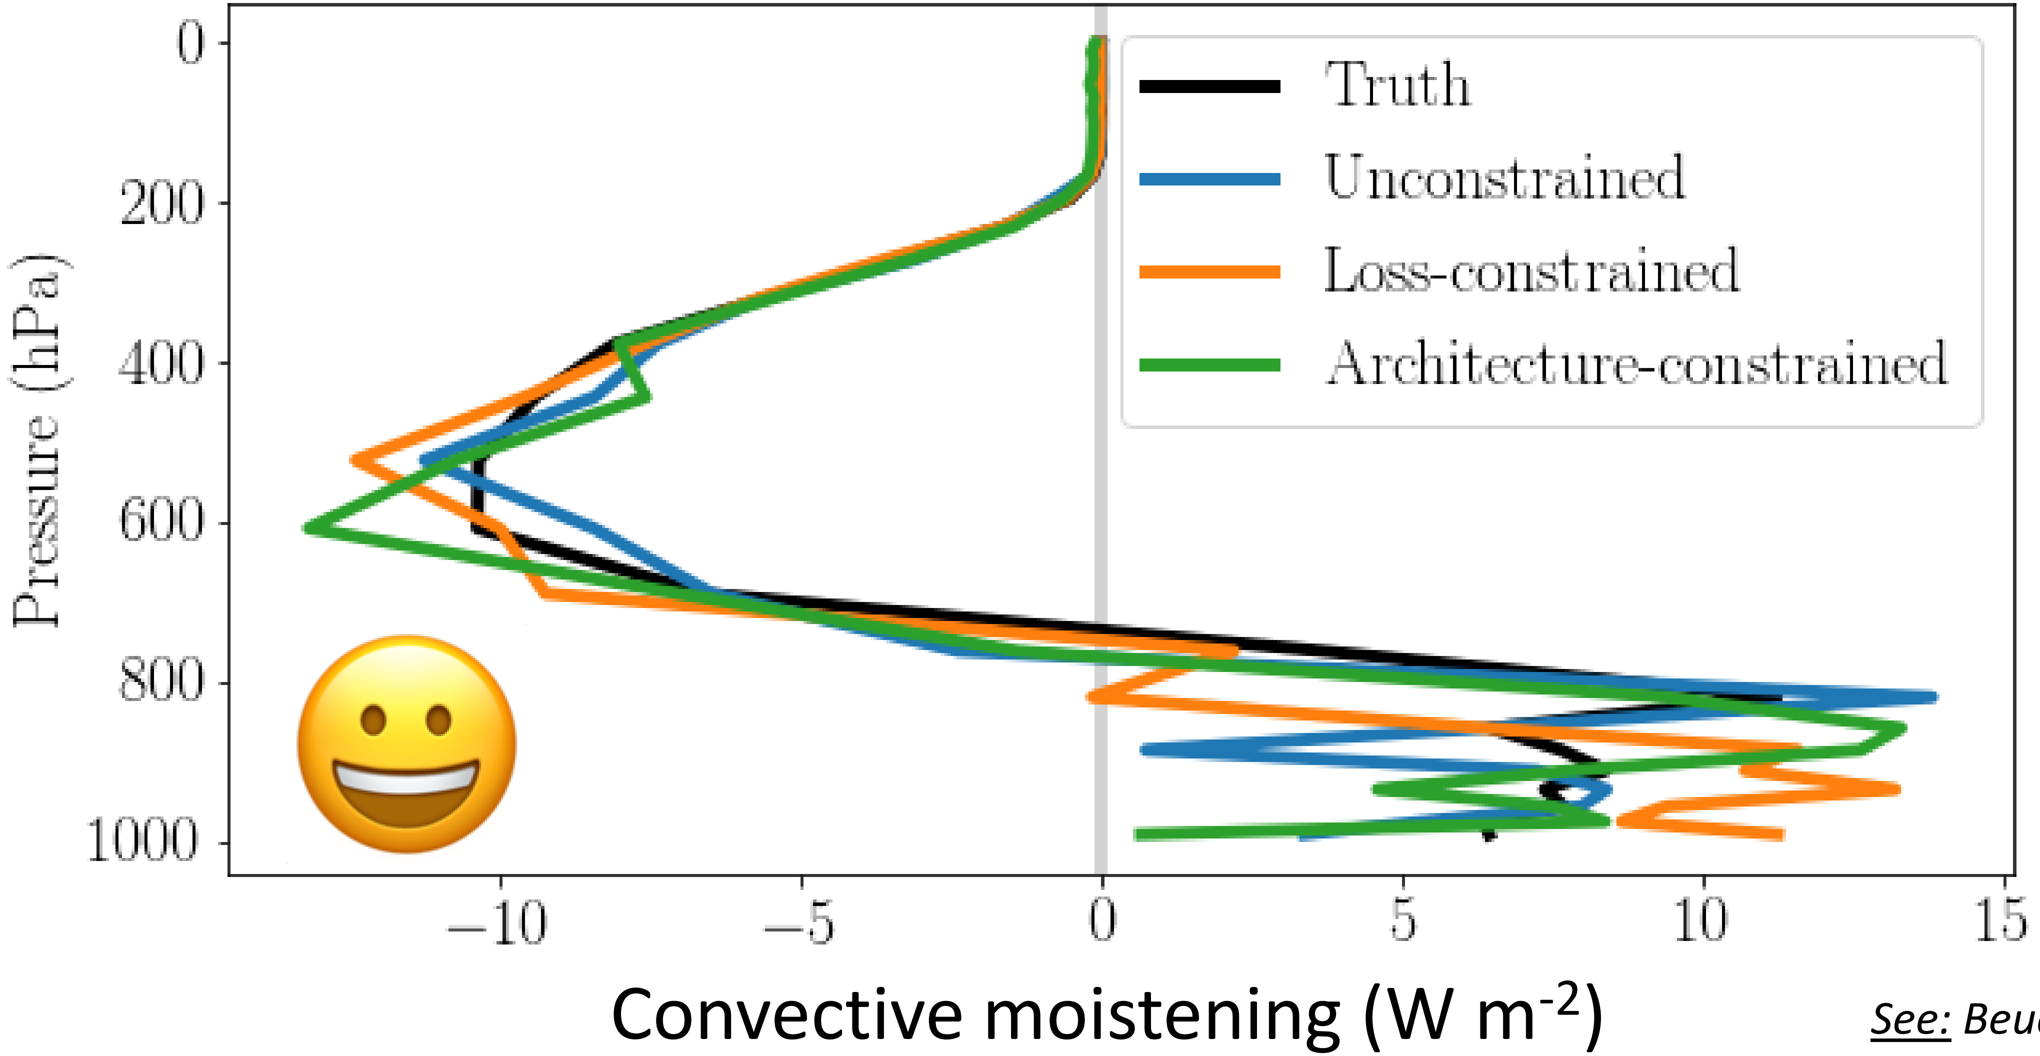

See: Beucler et al. (2020)

Problem: Neural Nets often fail to generalize out-of-distribution

Daily-mean Tropical prediction in (+4K) warming experiment

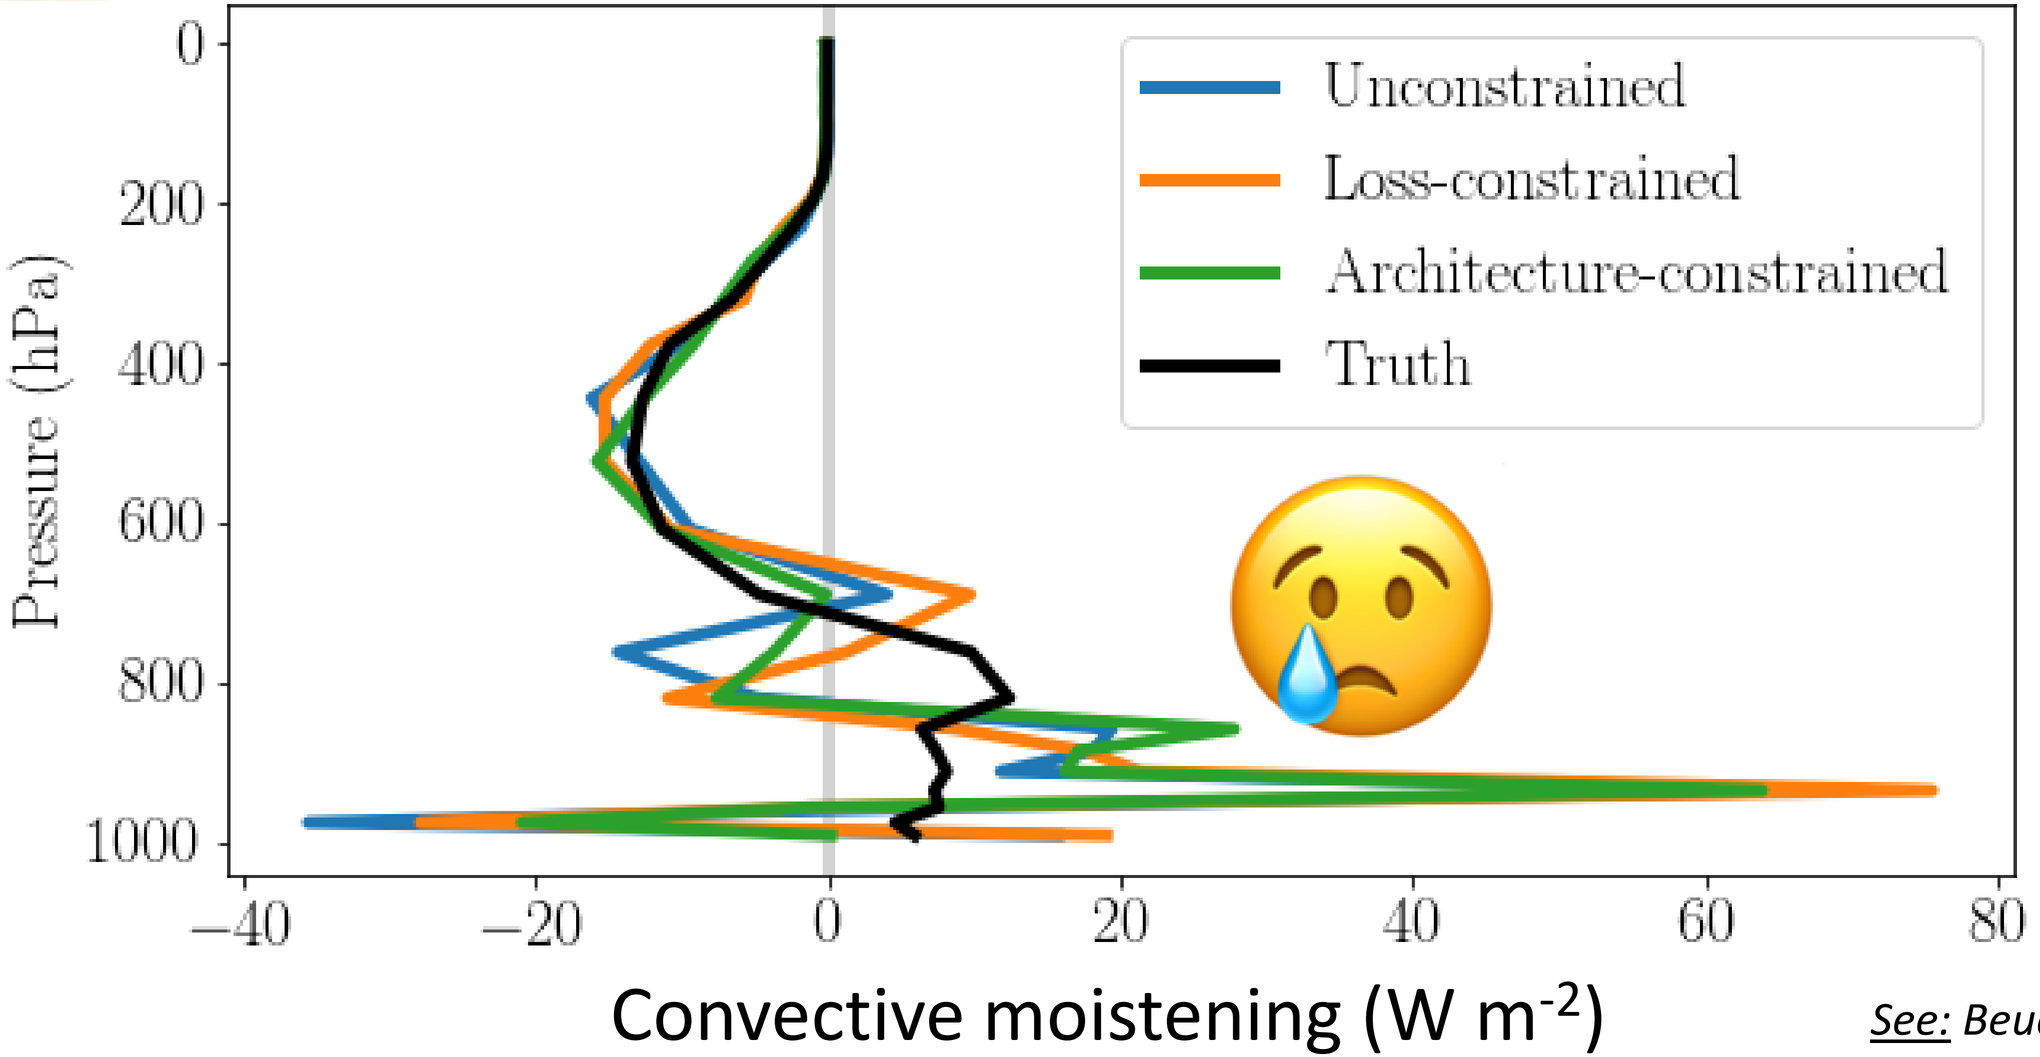

See: Beucler et al. (2020)

Idea: Break the model even more!

Image source: IT Biz Advisor

Generalization Experiment: +8K surface warming

Training and Validation on
cold simulations

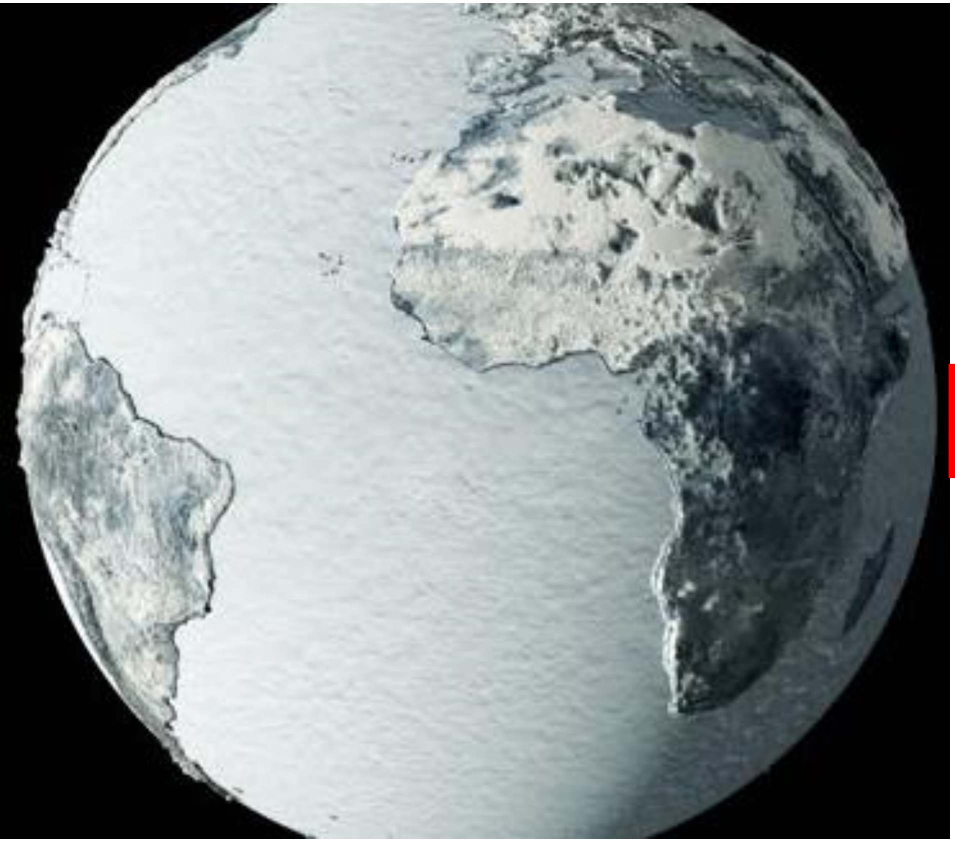

+8K

Test on warm simulations

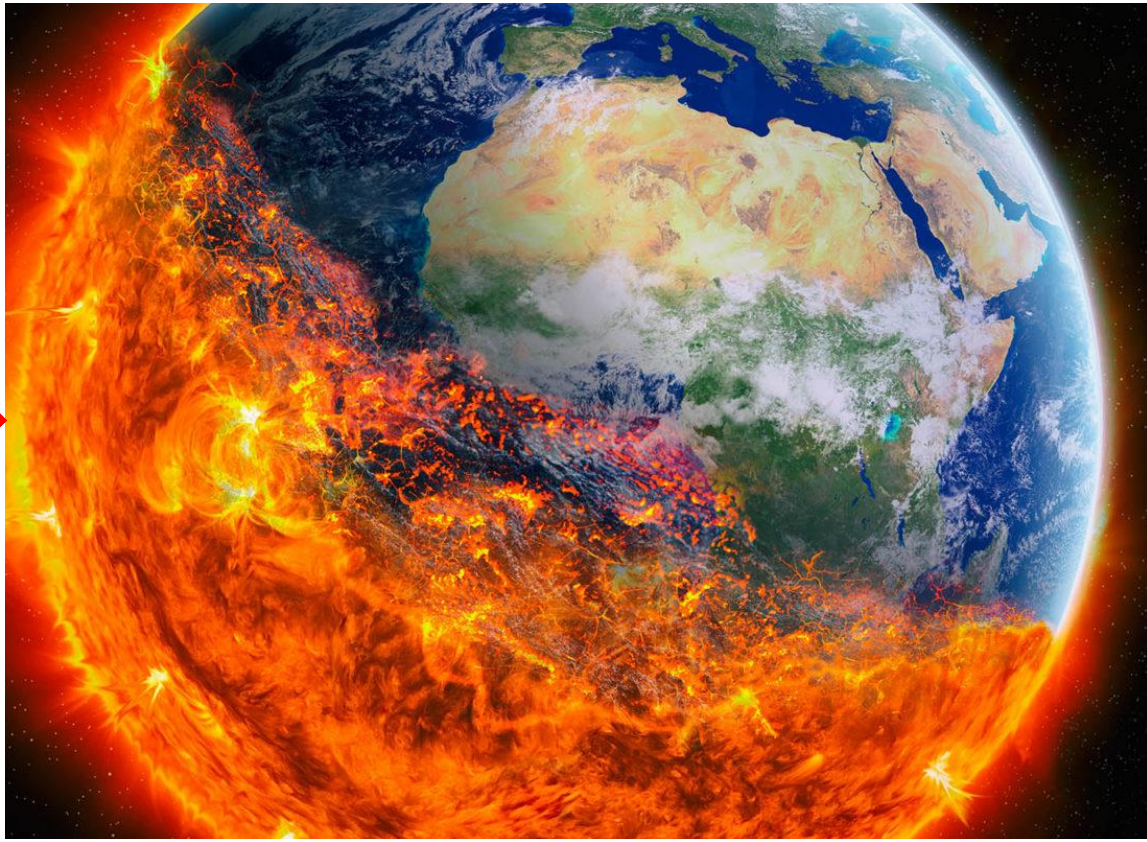

<u>Images</u>: Rashevskyi Viacheslav, Sebastien Decoret

IS ML RELIABLE FOR CLIMATE PROJECTIONS?

Problem: NNs fail to generalize to unseen climates

Daily-mean Tropical prediction in cold climate

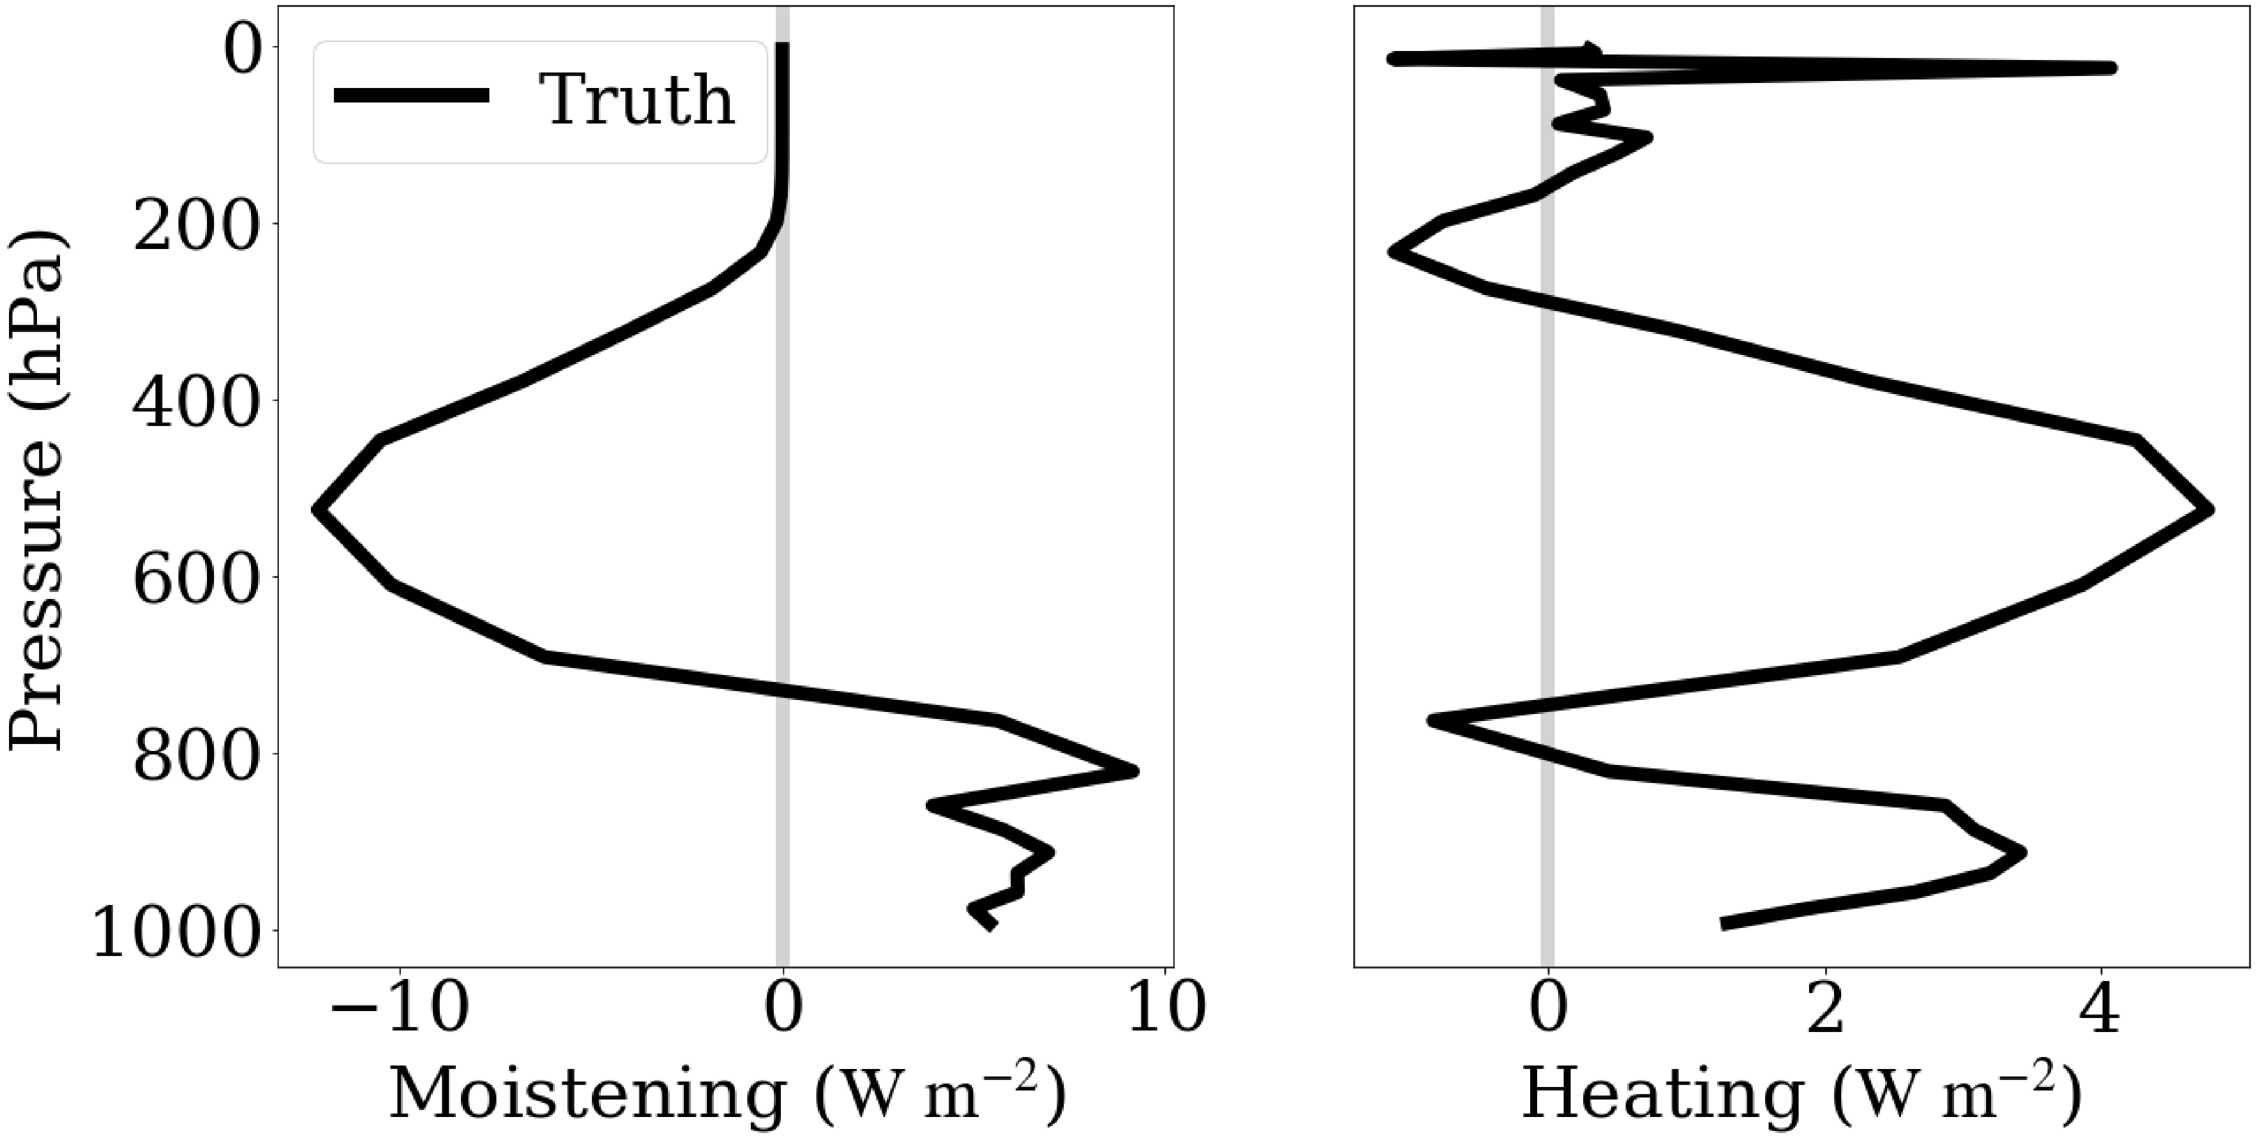

Problem: NNs fail to generalize to unseen climates

Daily-mean Tropical prediction in cold climate

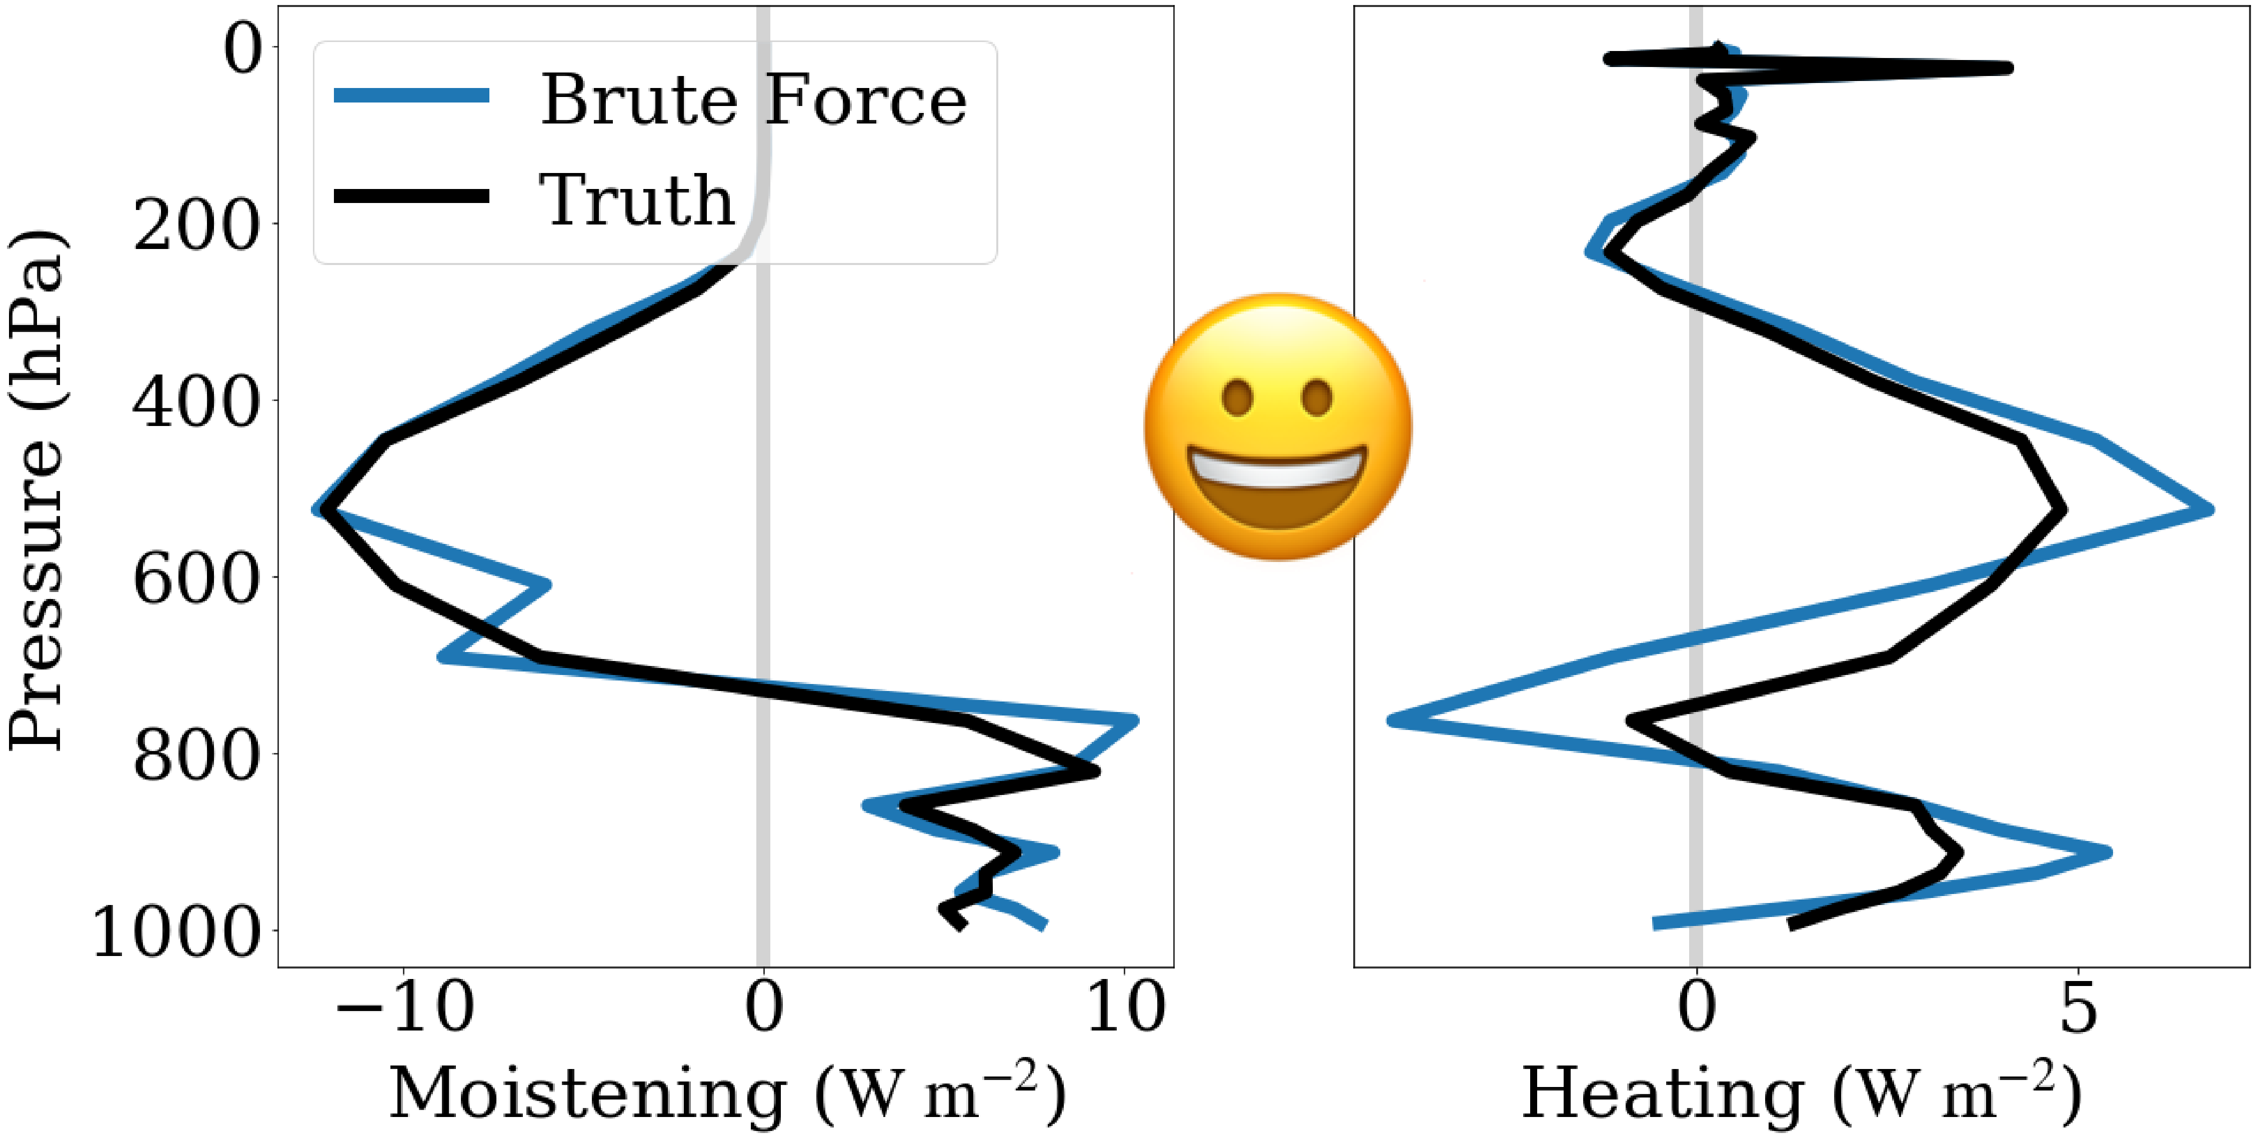

Daily-mean Tropical prediction in warm climate

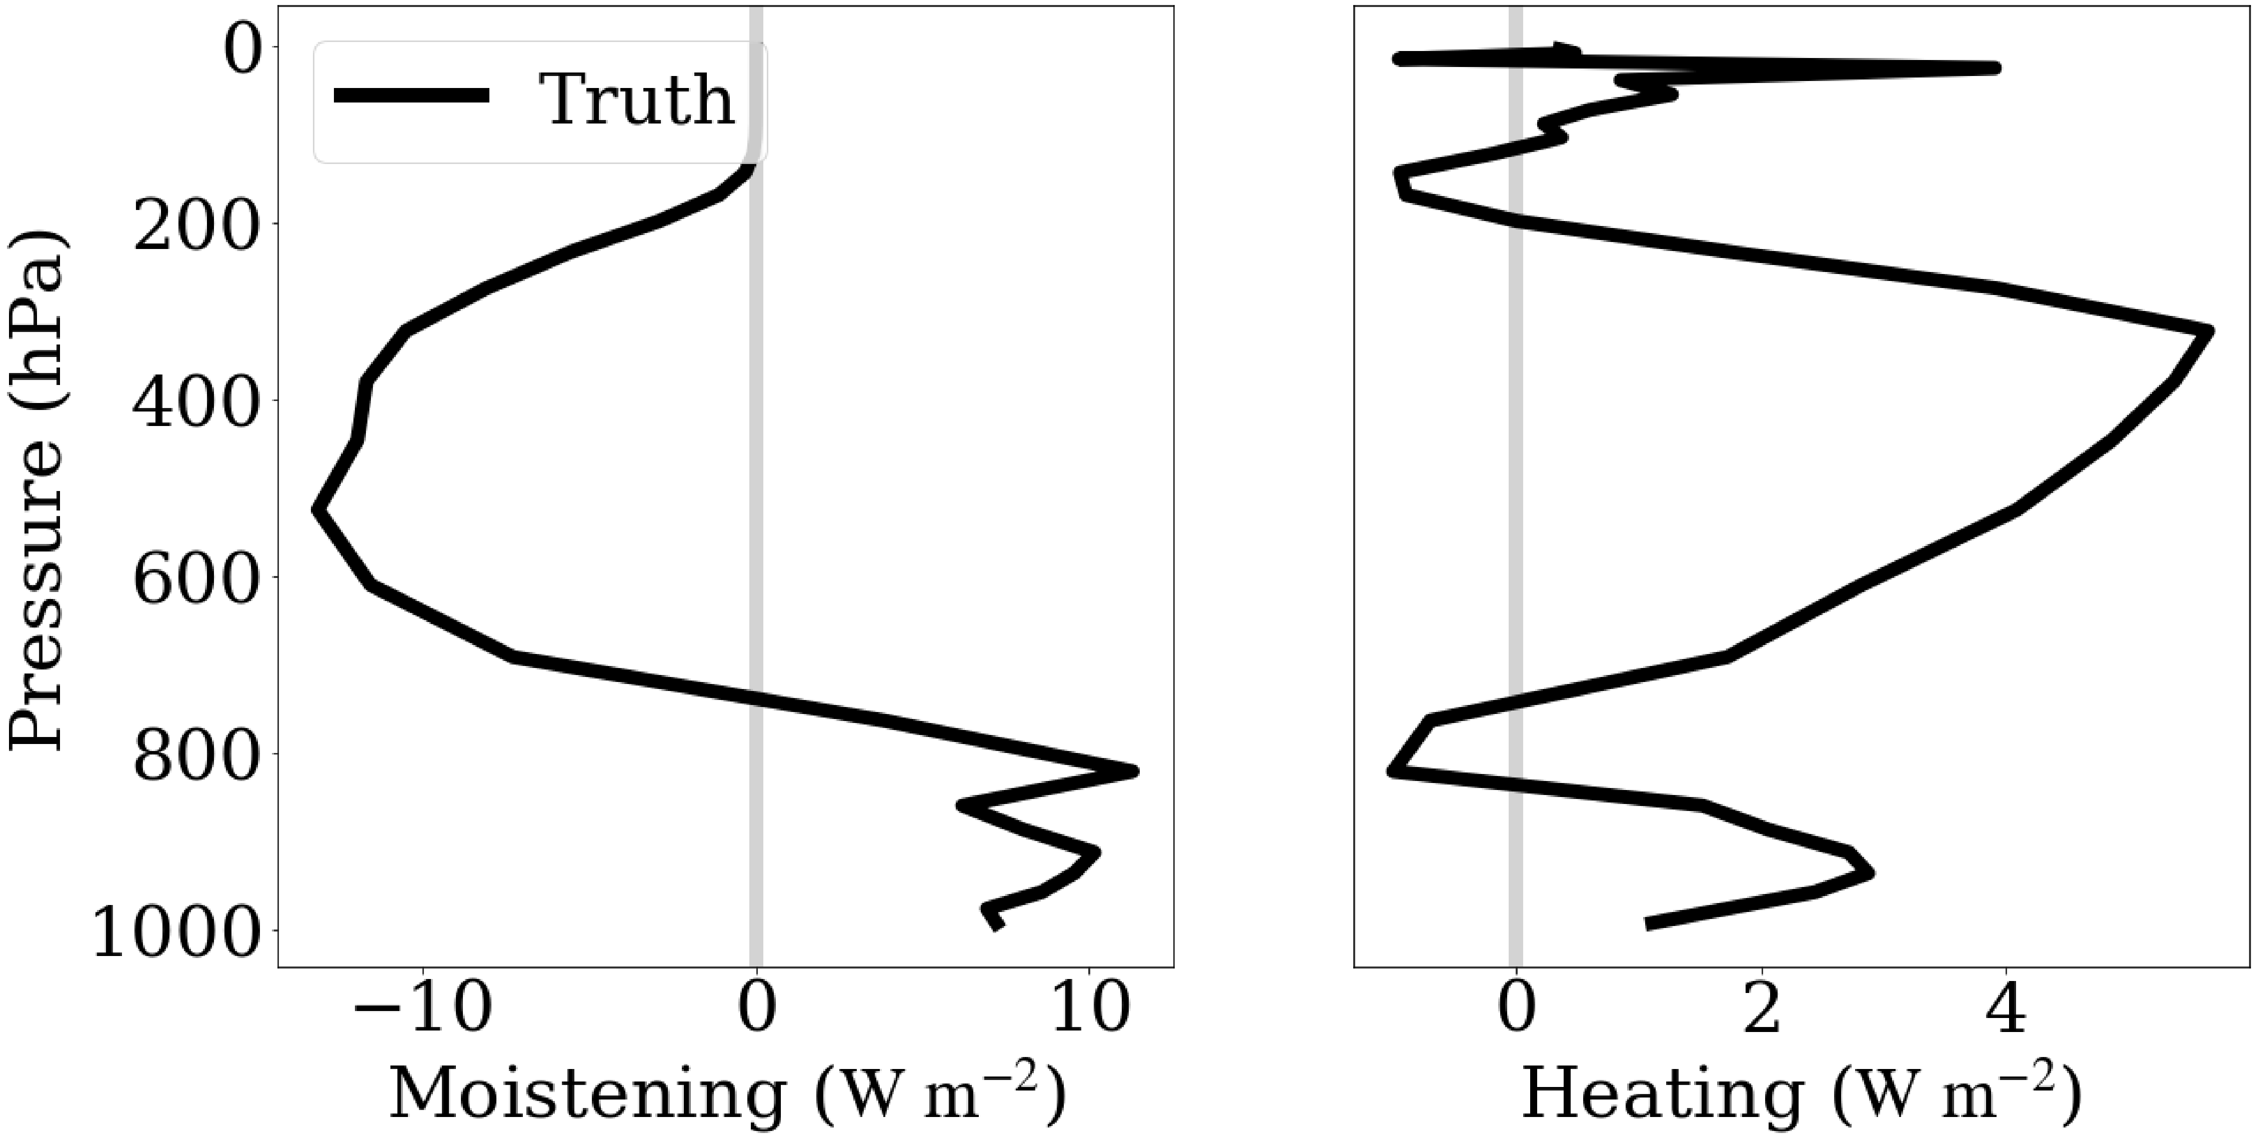

Daily-mean Tropical prediction in warm climate

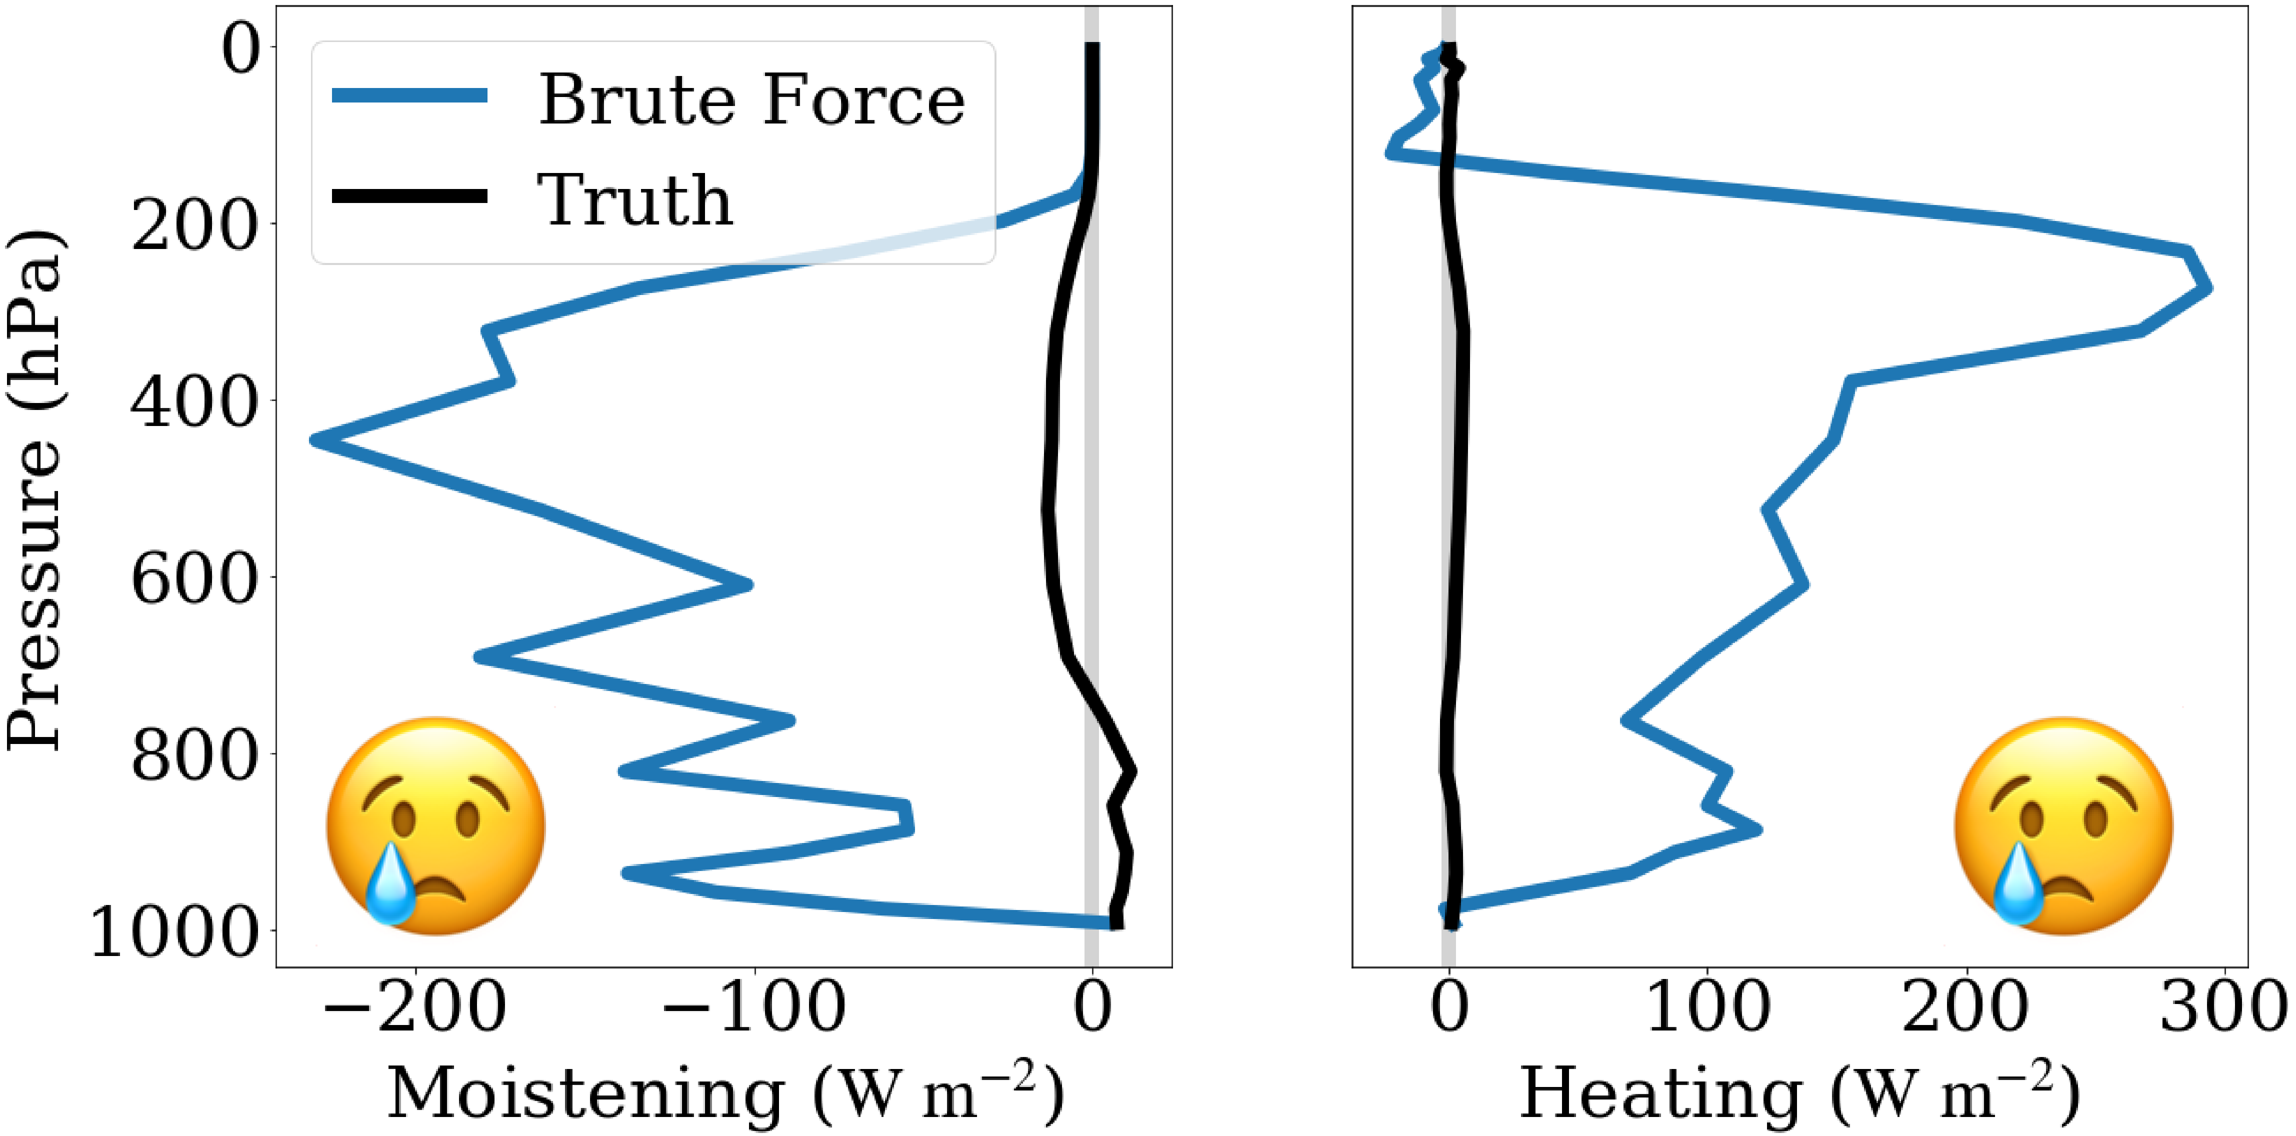

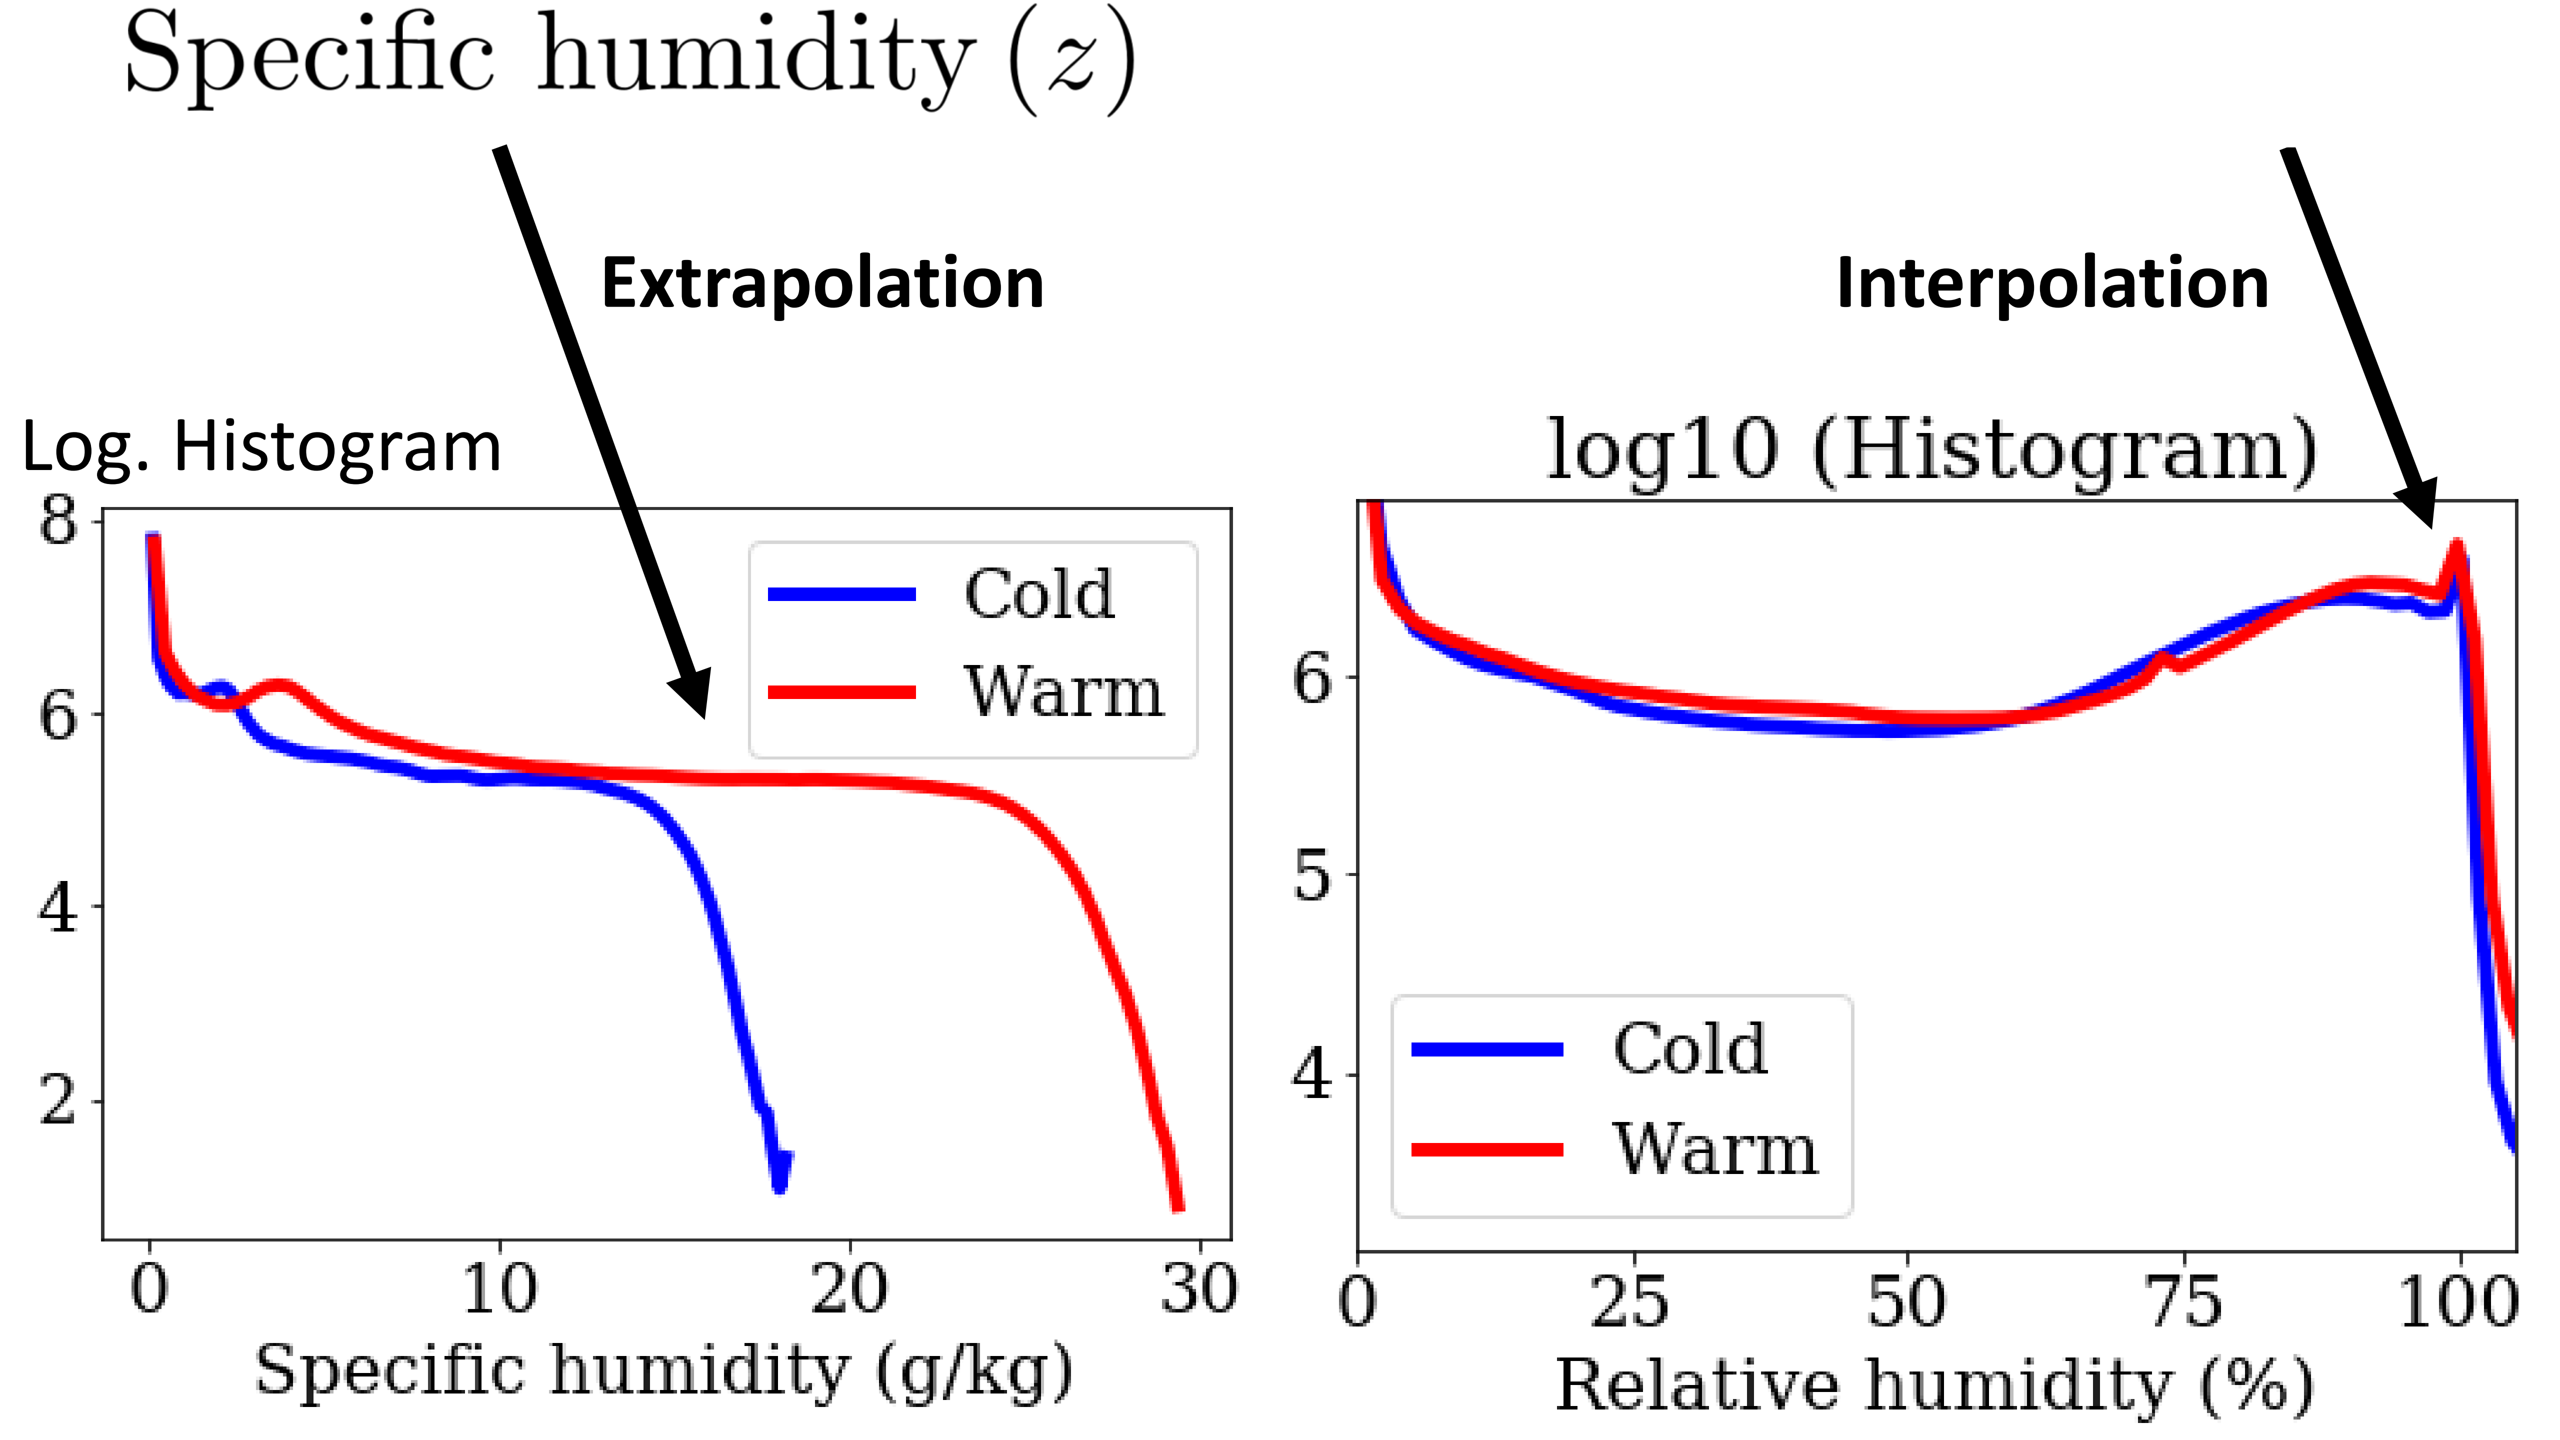
Specific humidity (z)
Extrapolation
Interpolation
Log. Histogram
8
6
4
2
Cold
Warm
0
10
20
30
Specific humidity (g/kg)
log10 (Histogram)
6
5
4
Cold
Warm
0
25
50
75
100
Relative humidity (%)

Specific humidity (z) $\rightarrow$ Relative humidity (z)

Generalization improves dramatically!

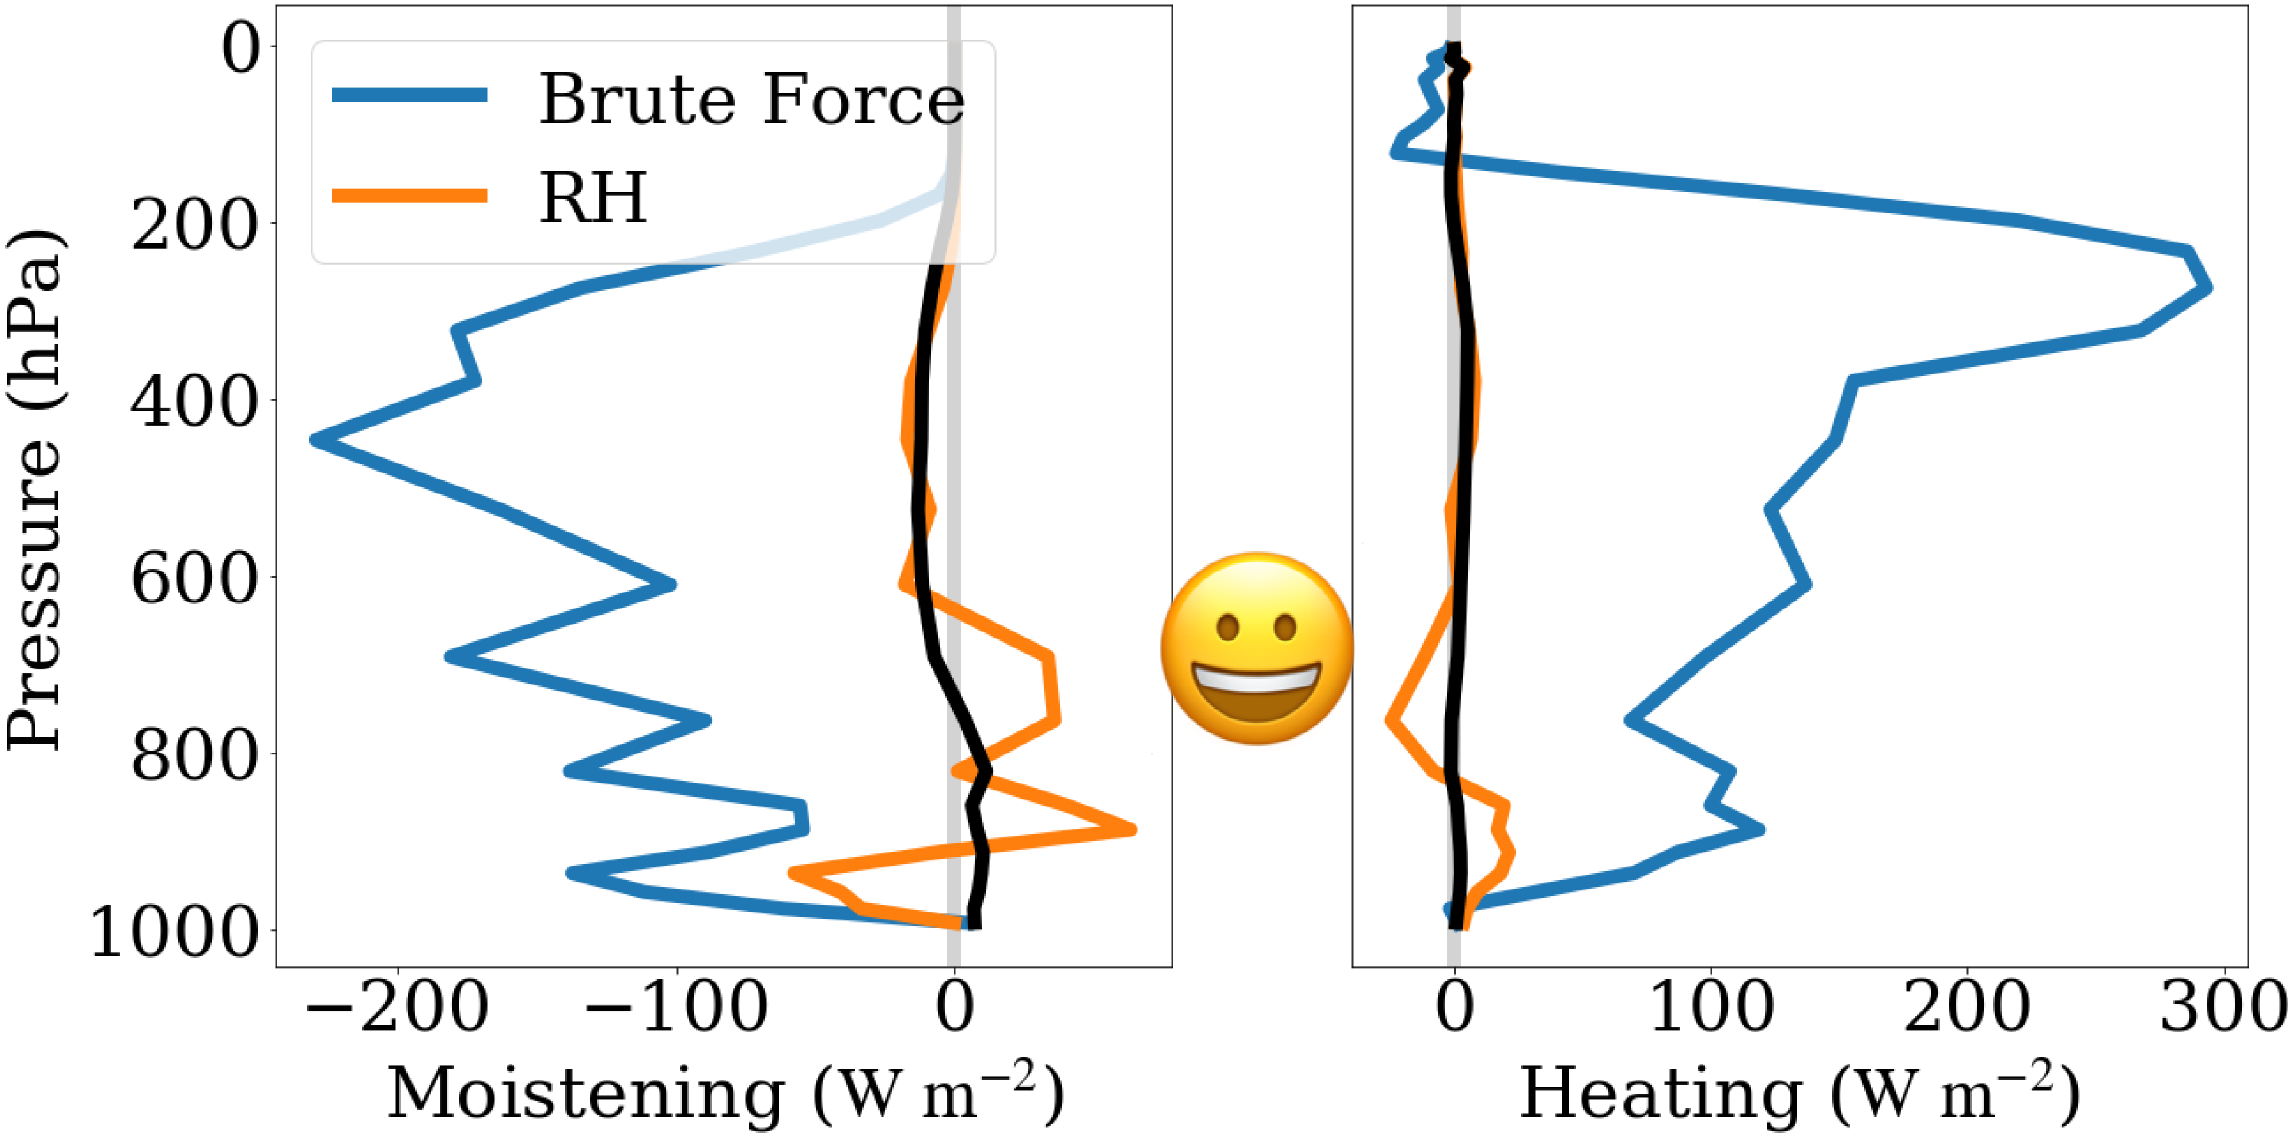

Physically transform the data to convert extrapolation into interpolation

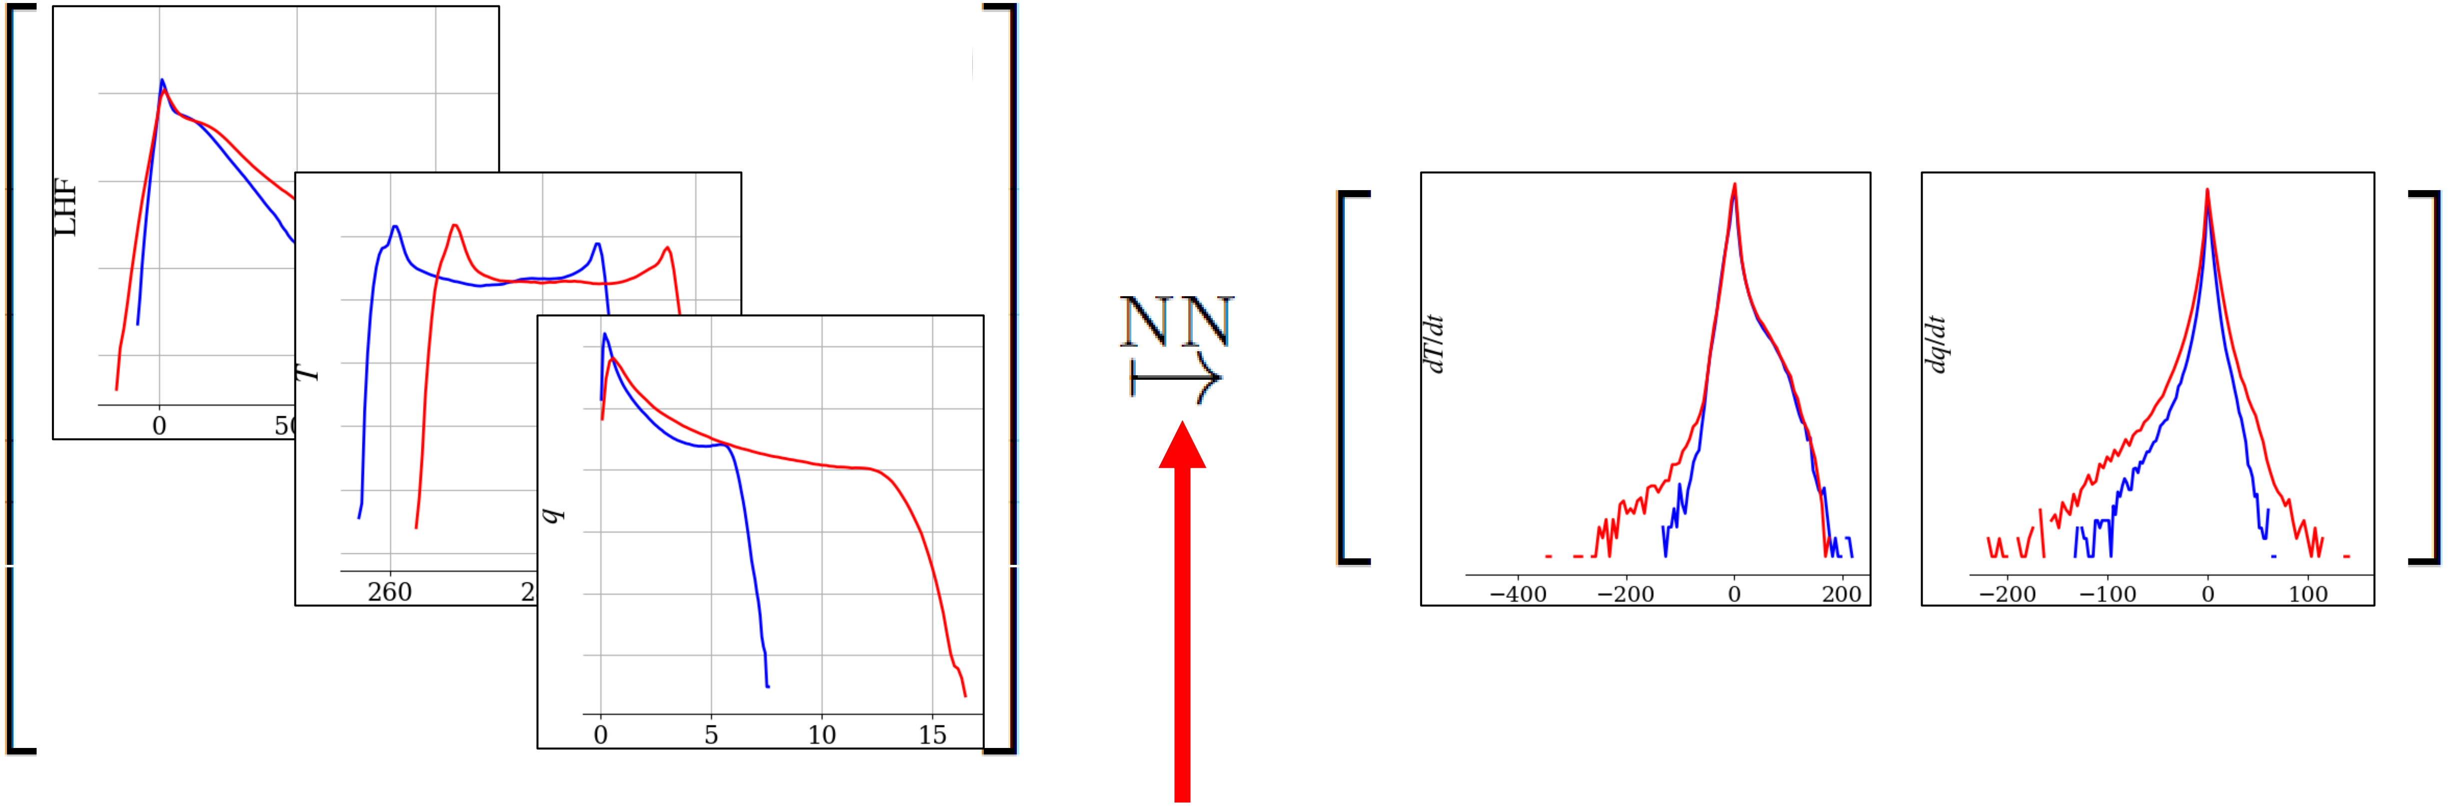

Raw Data: Not Climate-Invariant

Physically transform the data
to convert extrapolation into interpolation

Idea: Uncover **climate-invariant** mapping from climate to convection

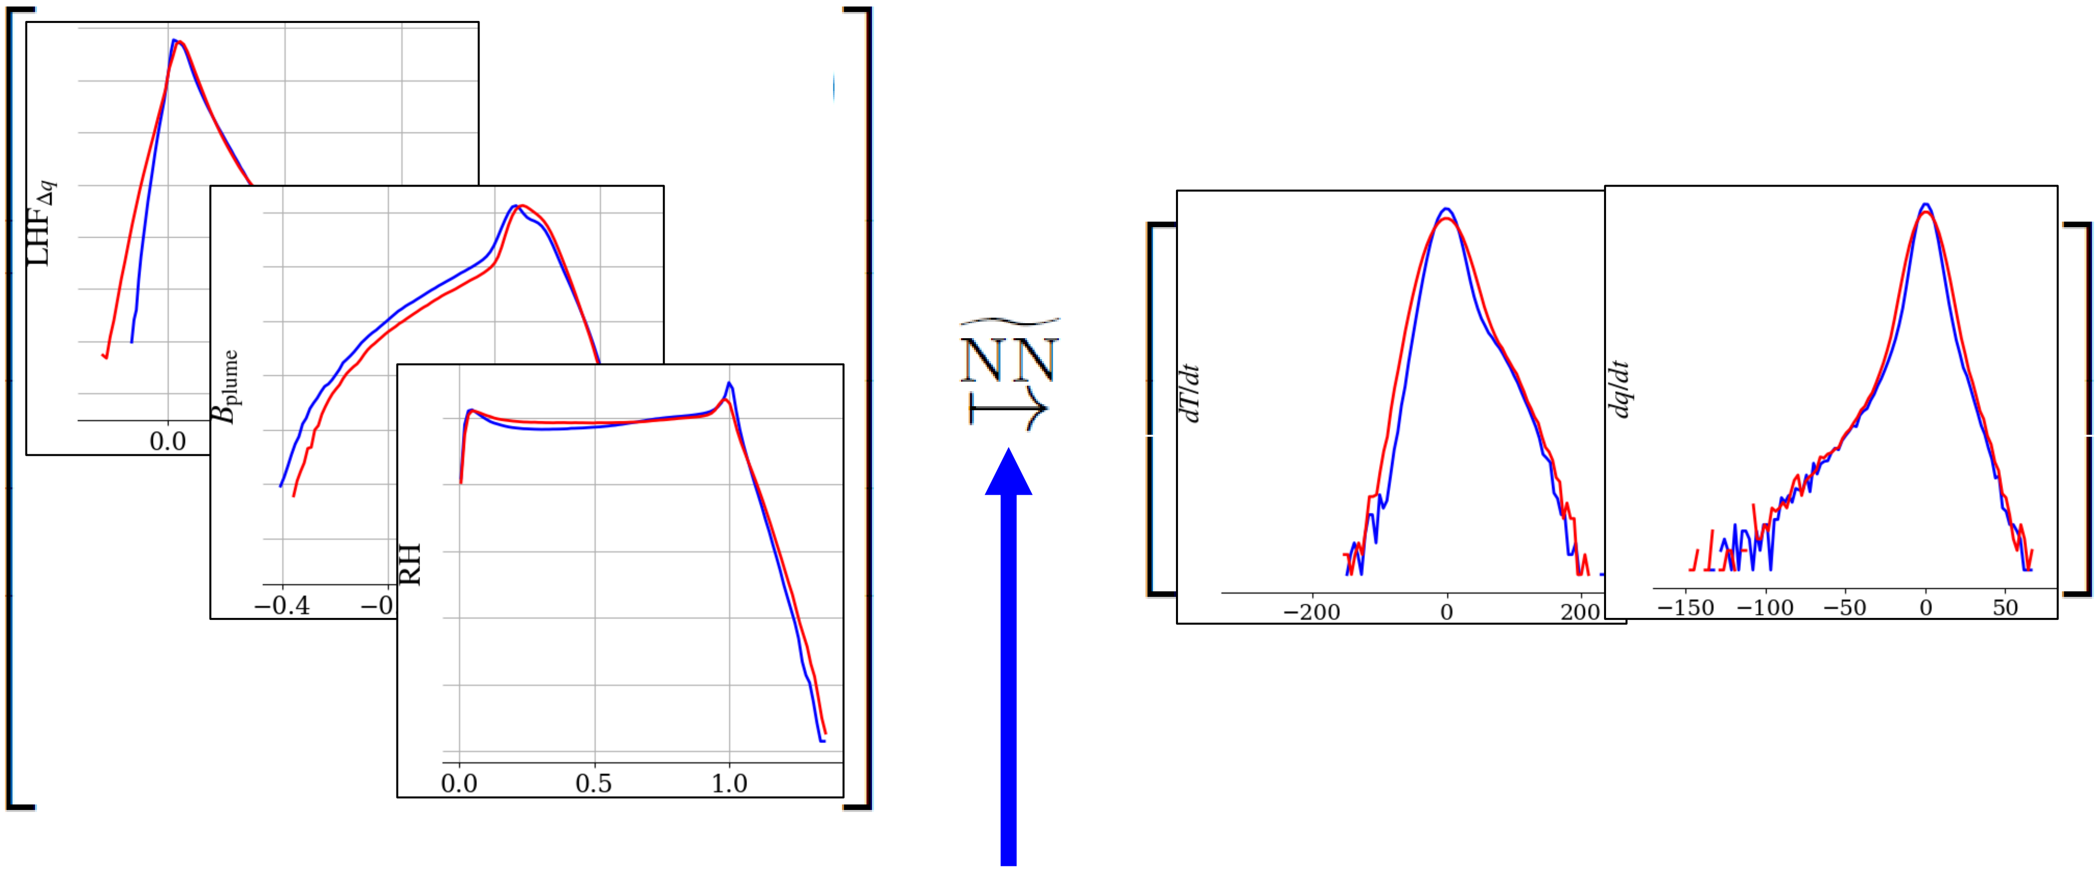

Physically-transformed data: Climate-Invariant

Climate-Invariant NNs generalization error close to NN trained in warm climate

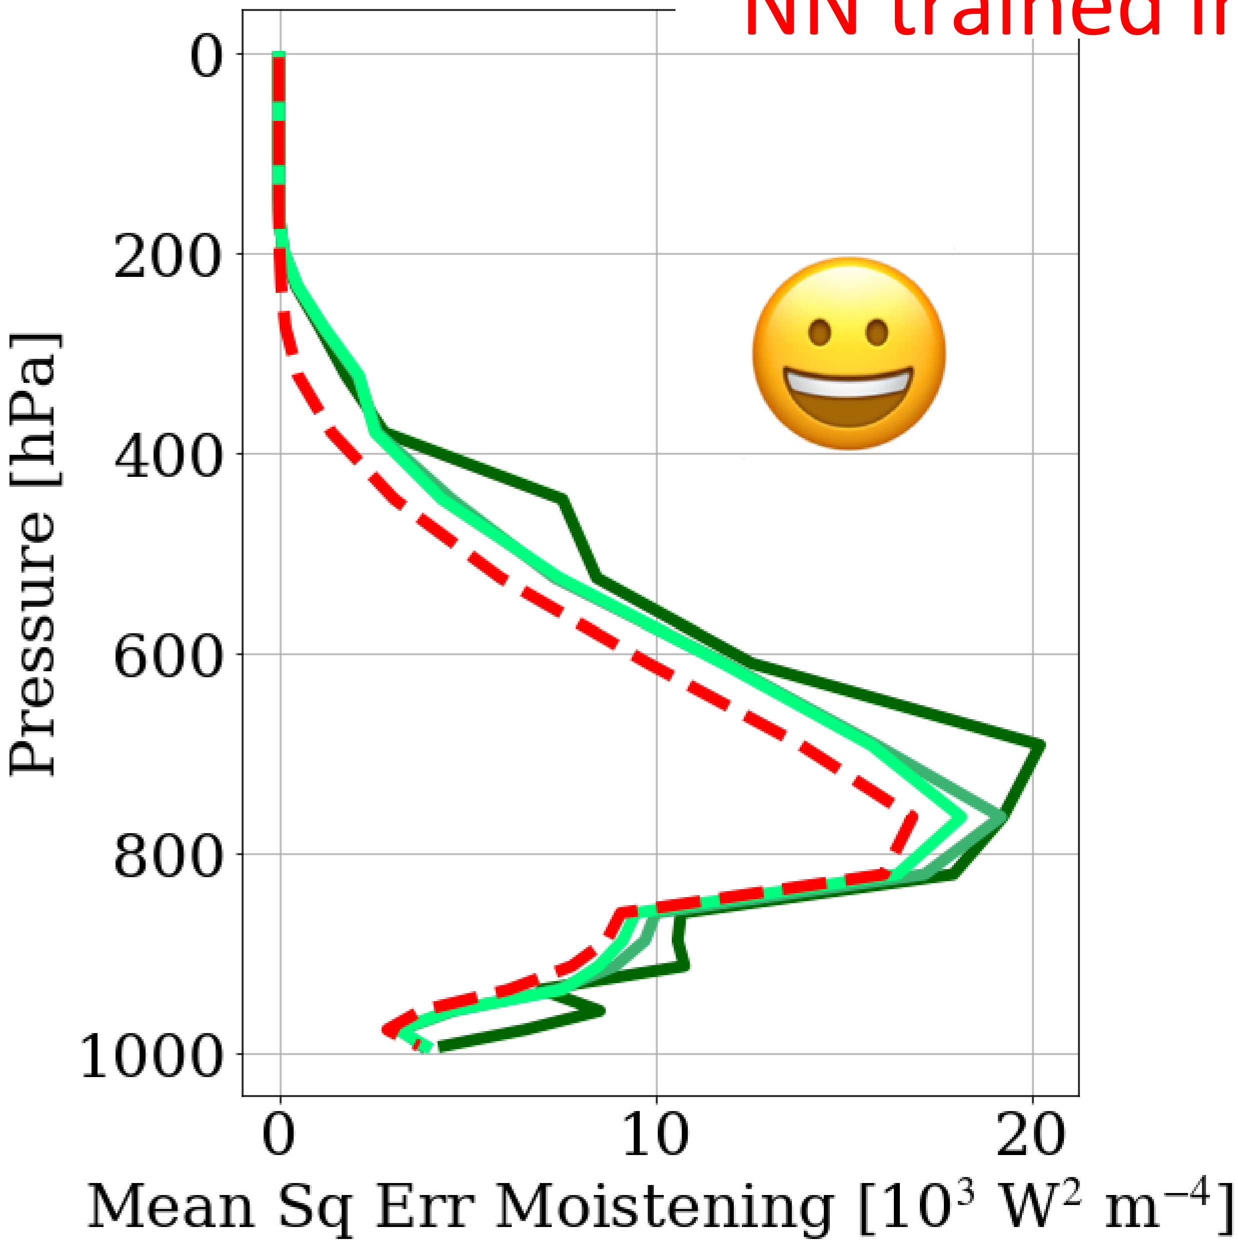

IS ML RELIABLE FOR CLIMATE PROJECTIONS?

Problem: Physically Transforming Inputs allows NNs to generalize from cold to warm climate

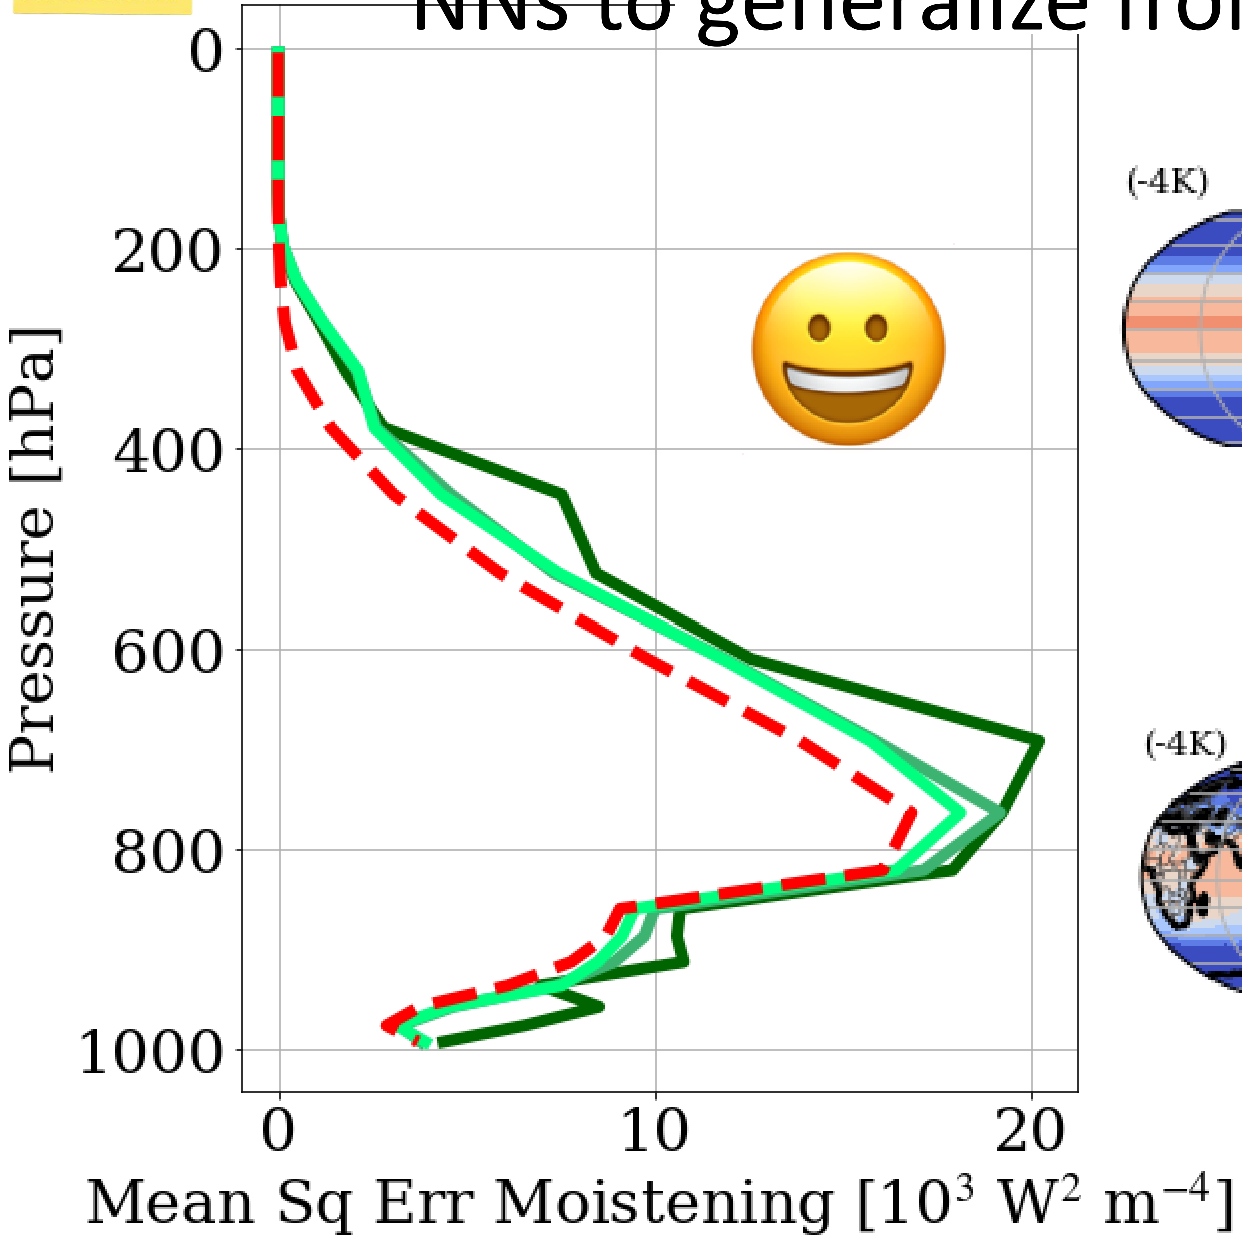

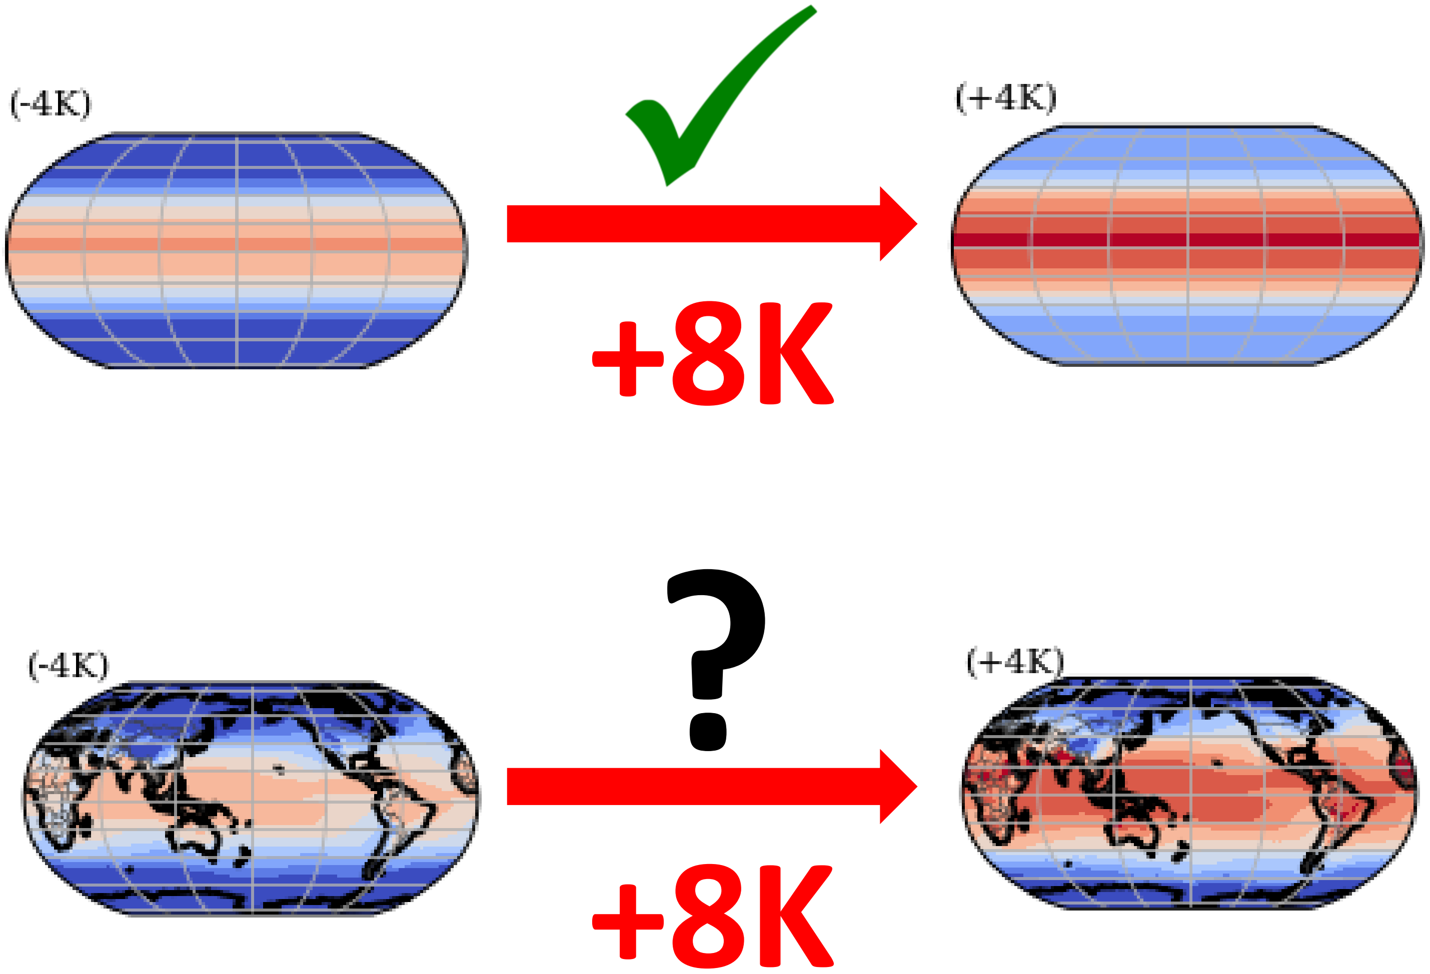

See: Beucler et al. (2024)

Physically-Informed Neural Networks Generalize Better Across Climates **in Earth-like configurations**

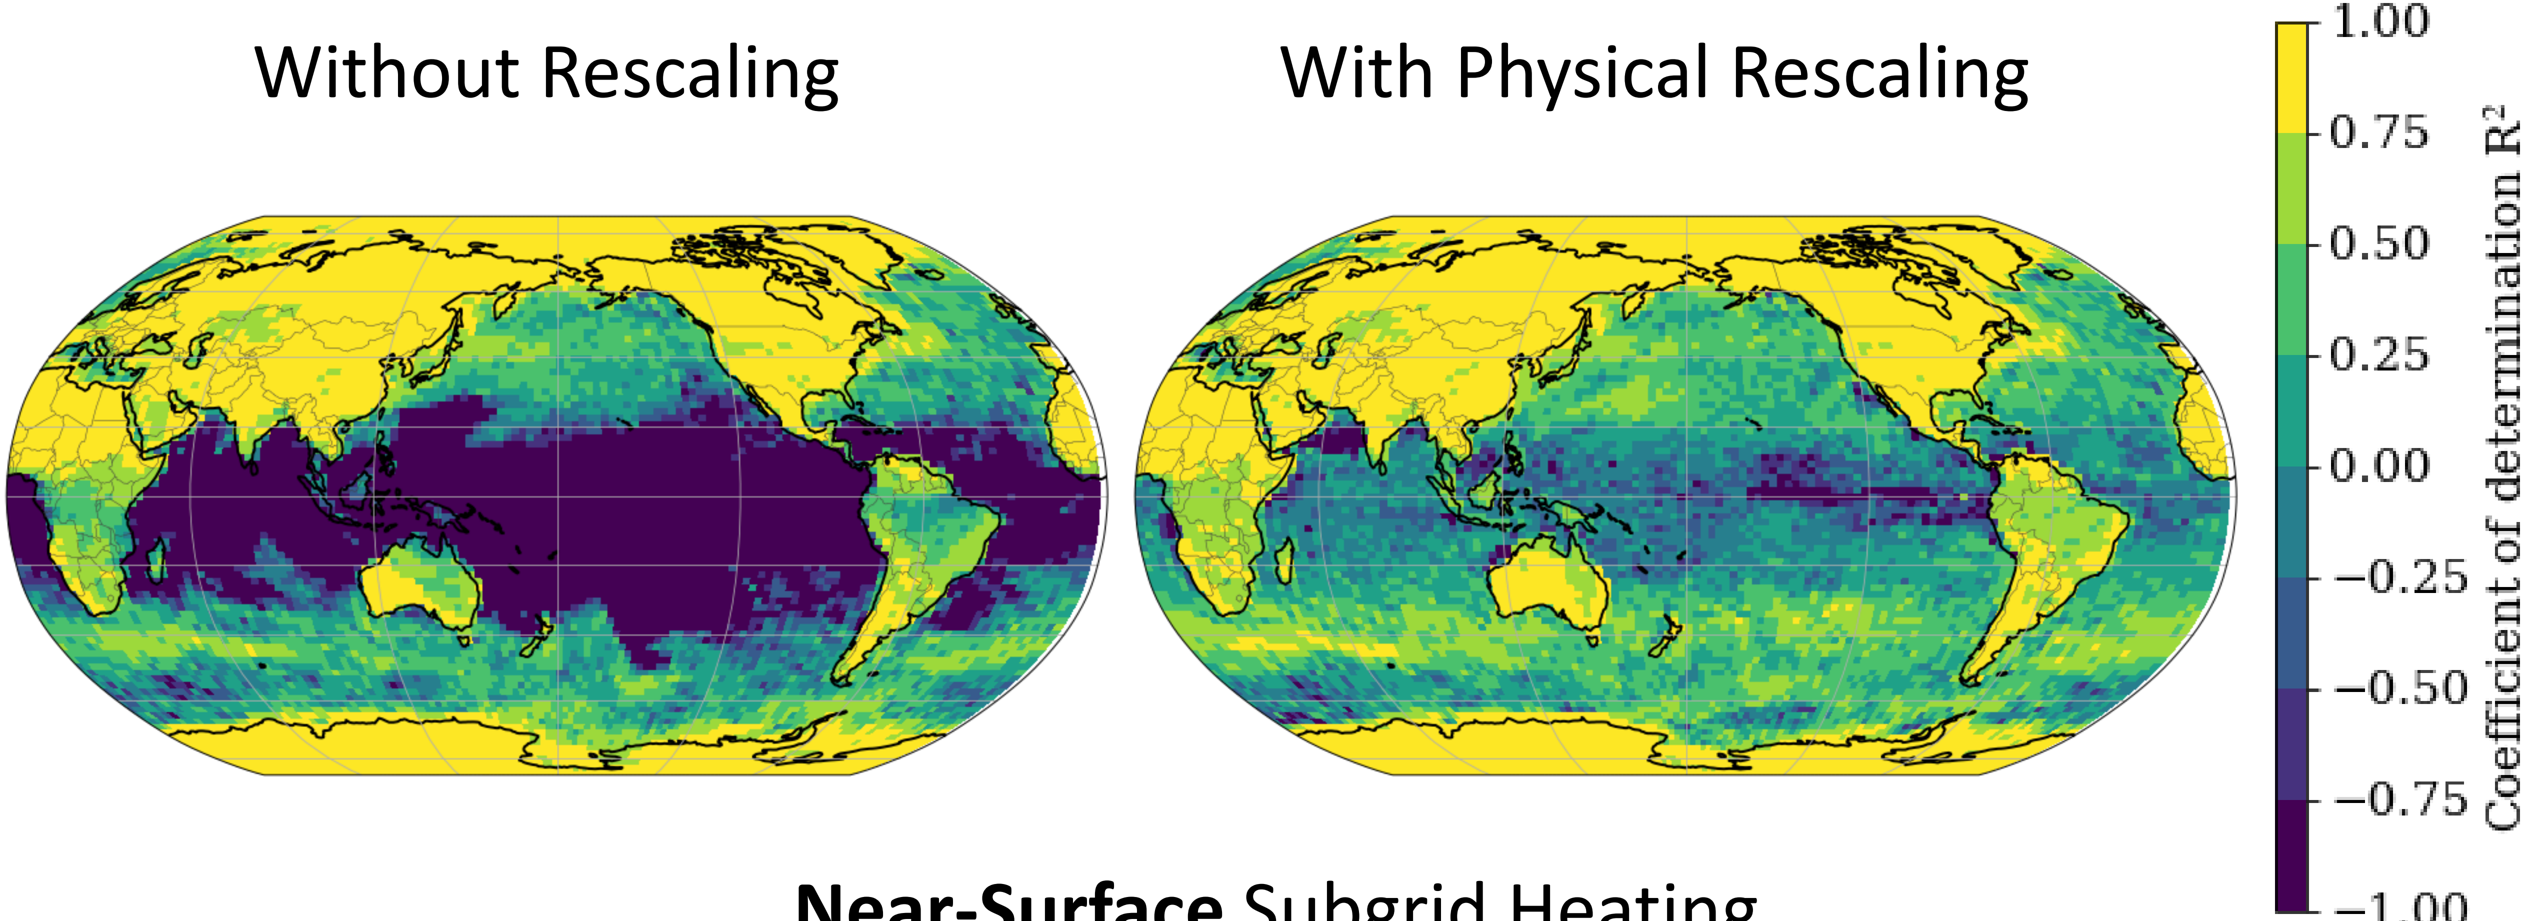

Near-Surface Subgrid Heating

<u>See:</u> Beucler et al. (2024)

Physically-Informed Neural Networks Generalize Better Across Climates **in Earth-like configurations**

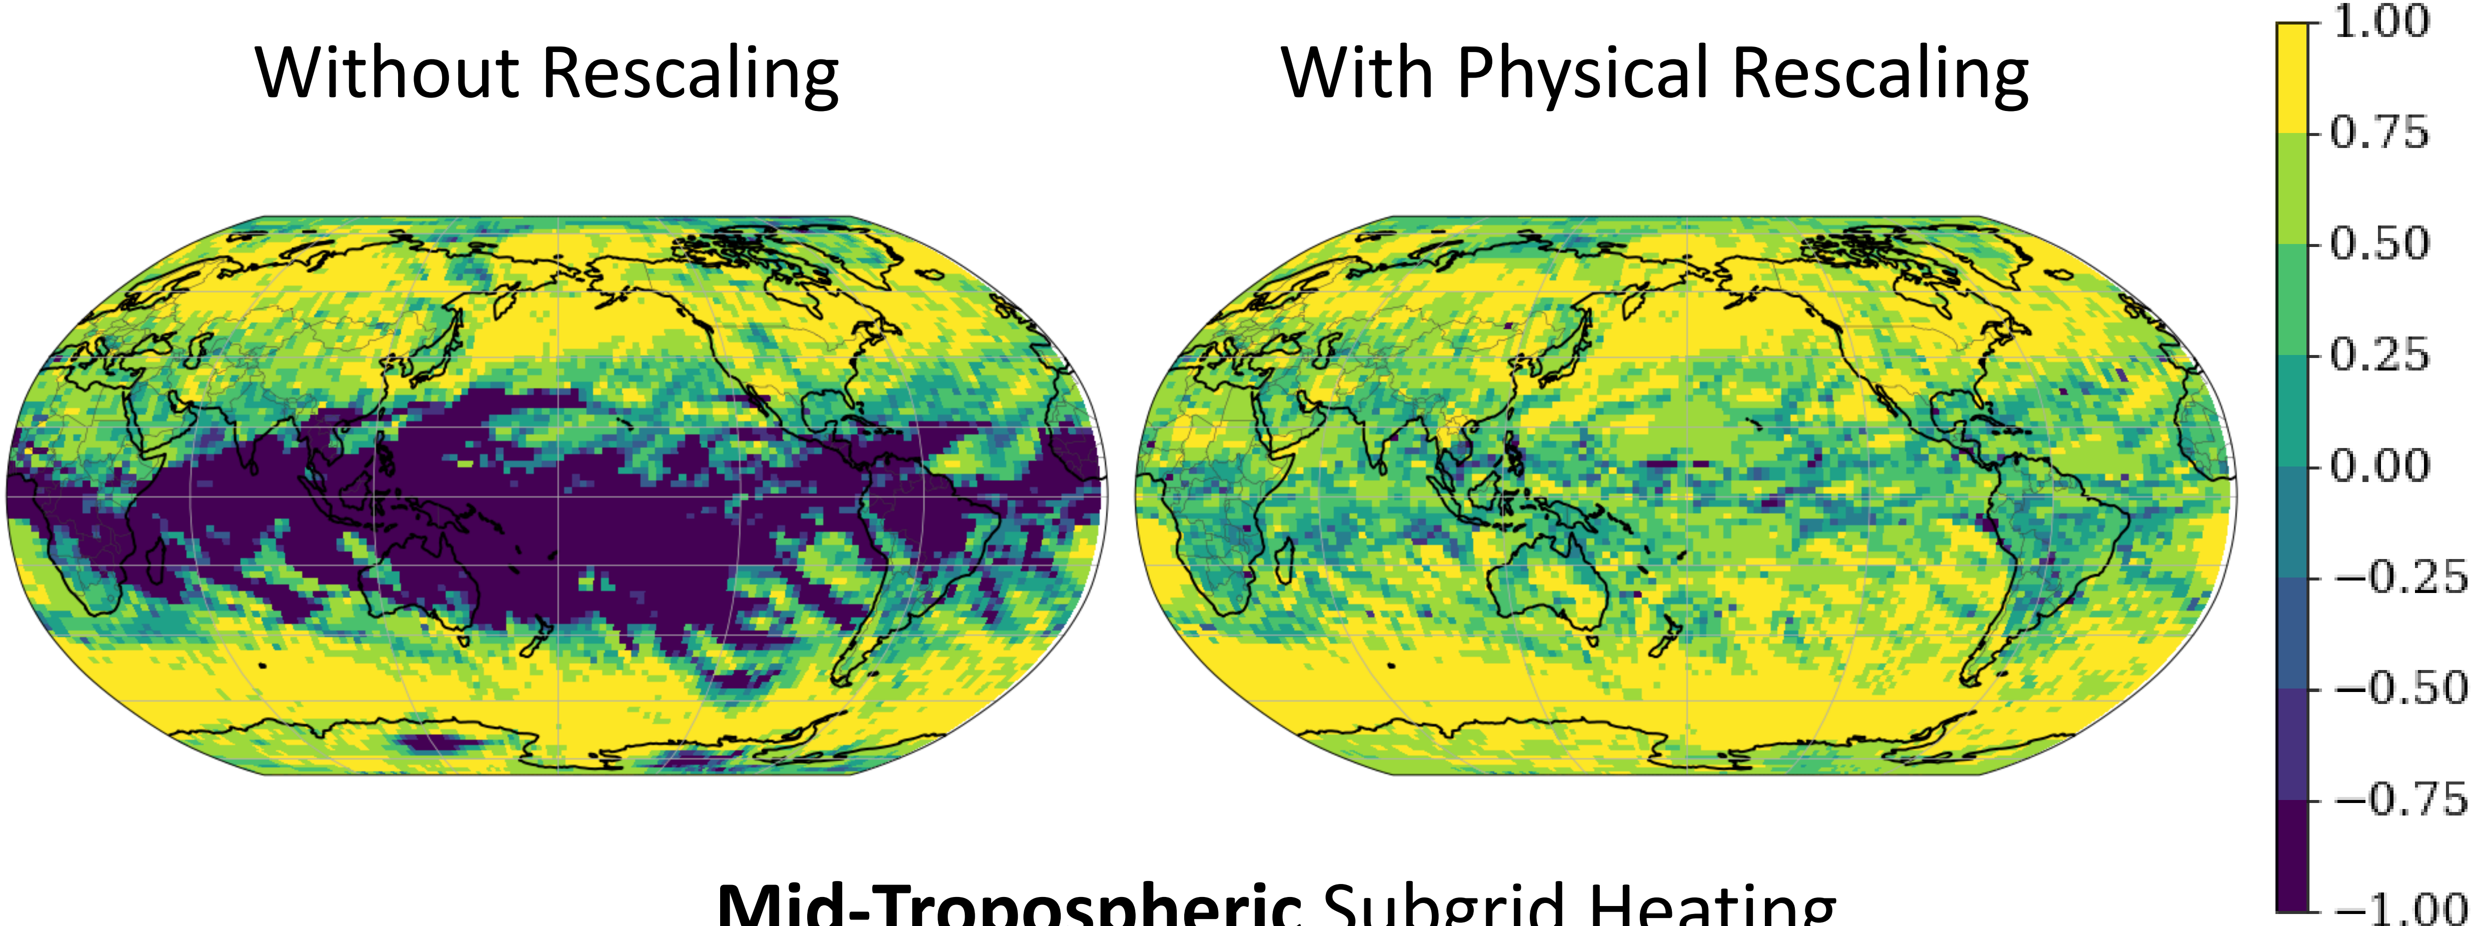

Mid-Tropospheric Subgrid Heating

<u>See:</u> Beucler et al. (2024)

Unexpected discovery: Climate-invariant NNs more local than Brute-Force NNs

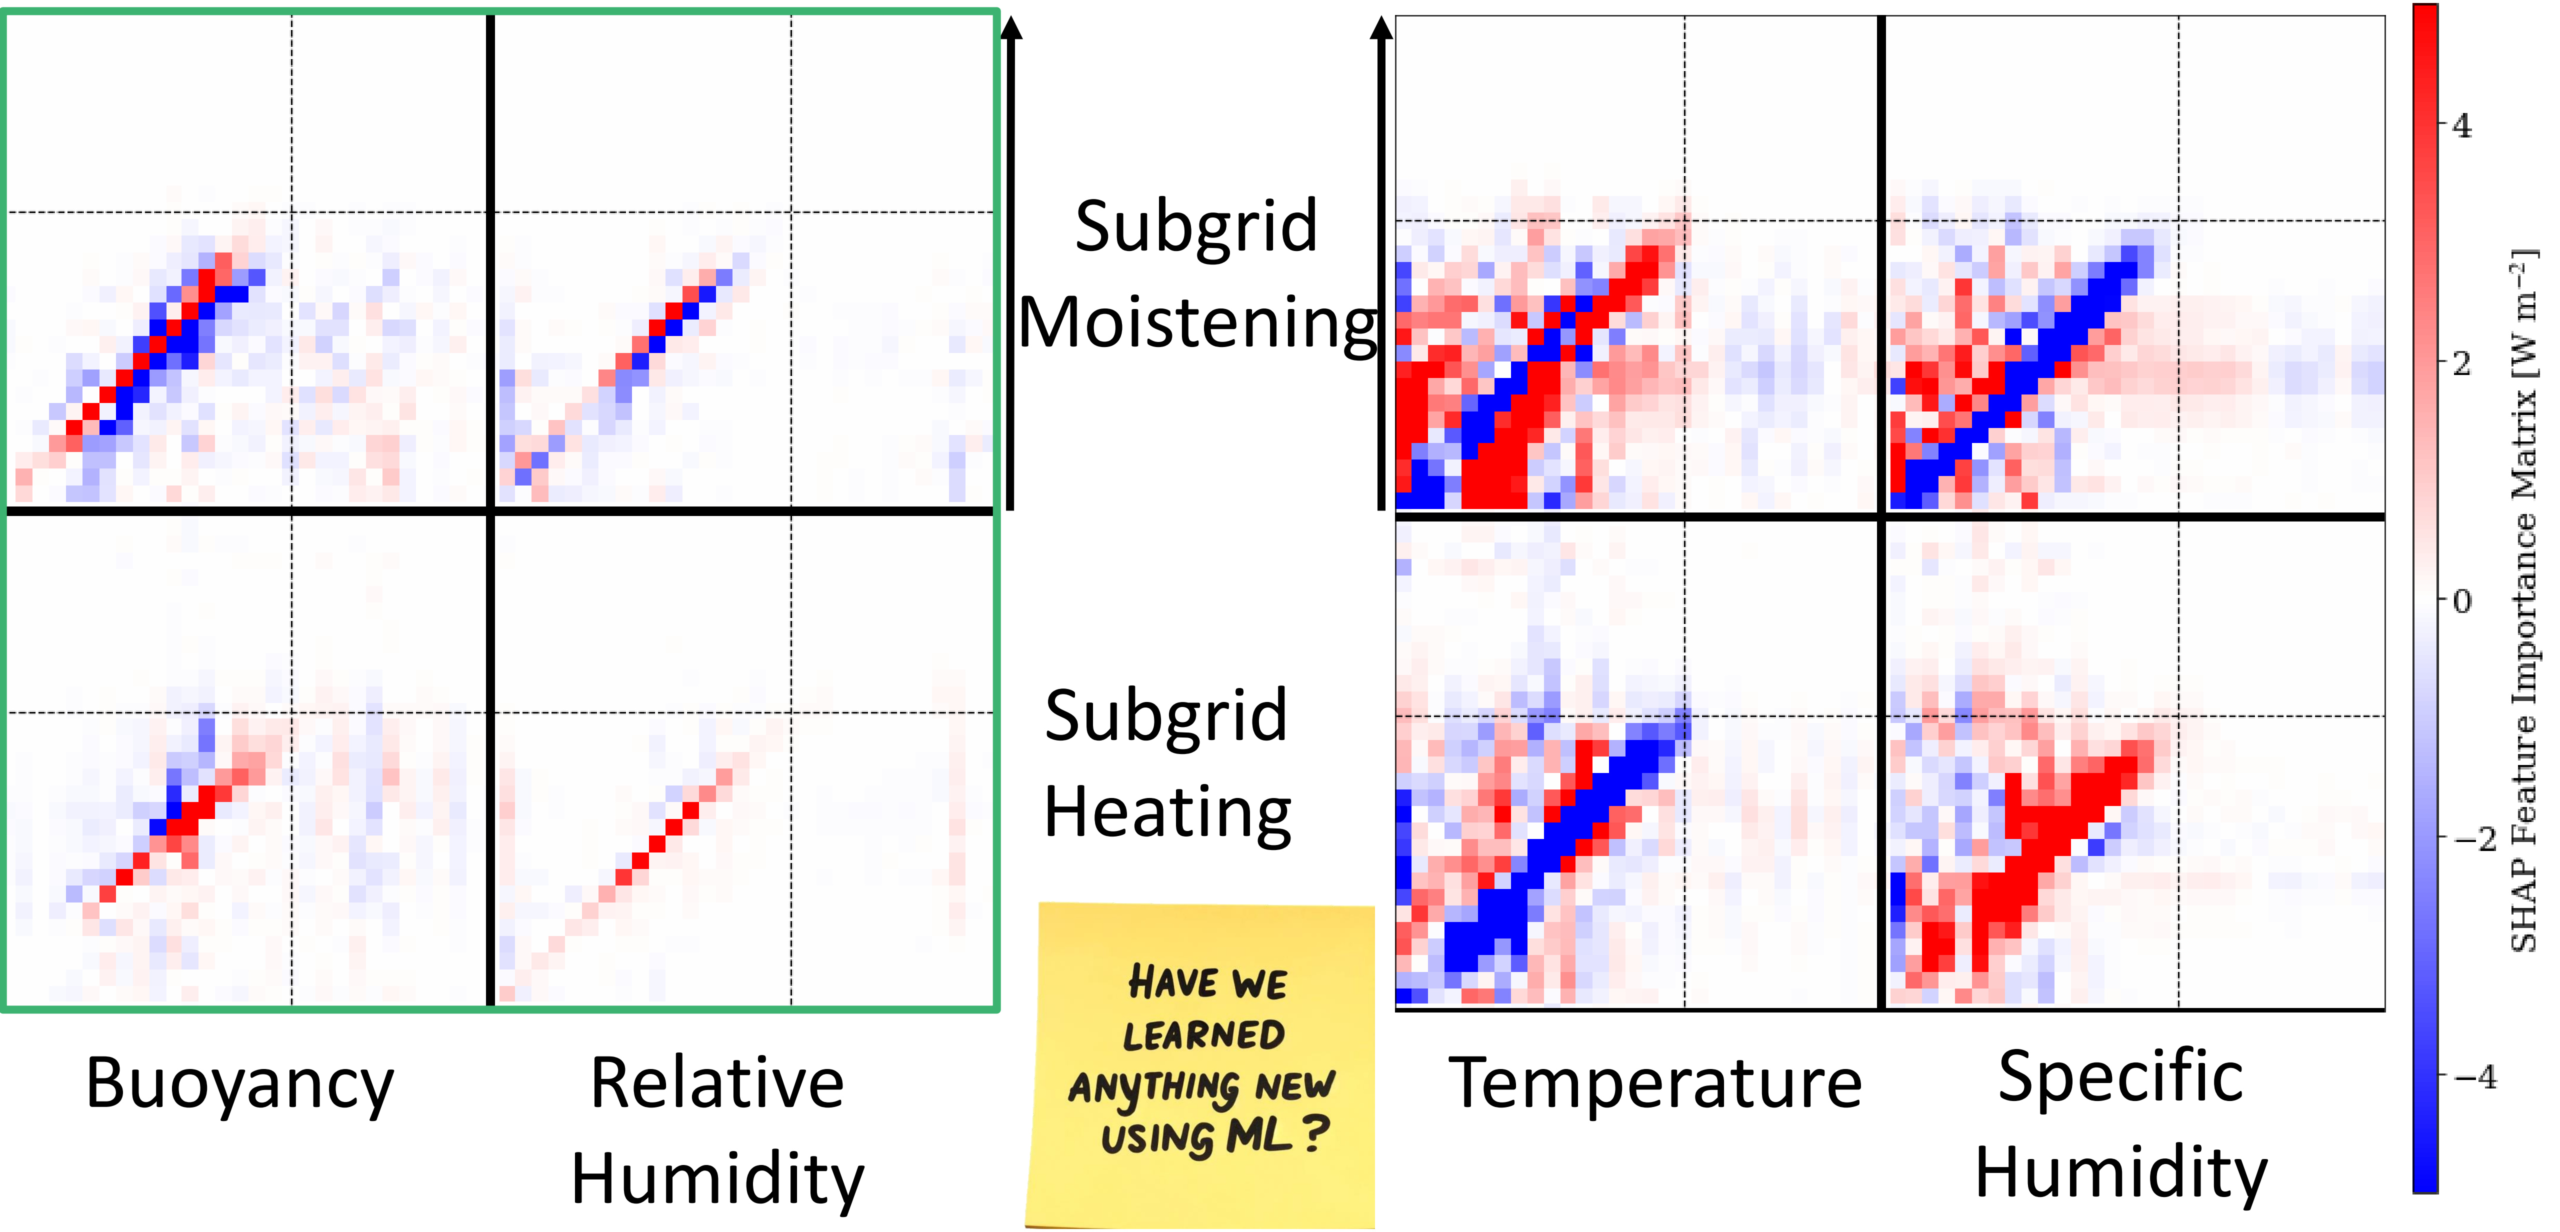

1) Data-driven parameterizations may not only accelerate, but also **improve** Earth System Models

2) They may lead to unexpected discoveries

3) They benefit from domain knowledge

4) Many challenges remain unsolved…

4) Many challenges remain unsolved:

1. Stability, extrapolation behavior, and recalibration of the host model

Journal of the Atmospheric Sciences

Volume 77: Issue 12

Sections

Editorial Type: Article

Article Type: Research Article

Interpreting and Stabilizing Machine-Learning Parametrizations of Convection

Noah D. Brenowitz, Tom Beucler, Michael Pritchard, and Christopher S. Bretherton

4) Many challenges remain unsolved:

1. **Stability**, extrapolation behavior, and recalibration of the host model

Why bother with XAI?

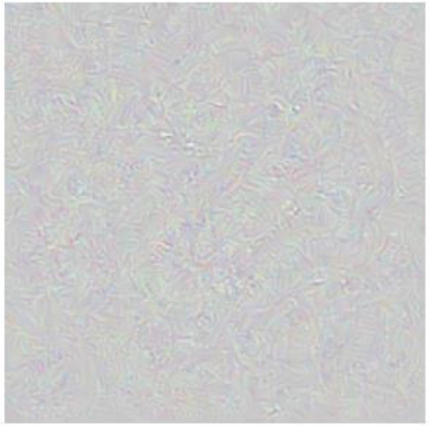

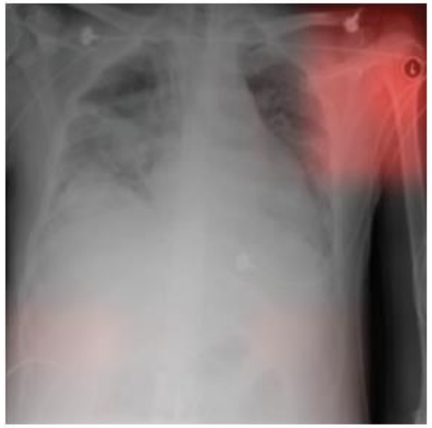

Article: Super Bowl 50

Paragraph: *"Peython Manning became the first quarterback ever to lead two different teams to multiple Super Bowls. He is also the oldest quarterback ever to play in a Super Bowl at age 39. The past record was held by John Elway, who led the Broncos to victory in Super Bowl XXXIII at age 38 and is currently Denver's Executive Vice President of Football Operations and General Manager. Quarterback Jeff Dean had a jersey number 37 in Champ Bowl XXXIV."*

Question: *"What is the name of the quarterback who was 38 in Super Bowl XXXIII?"*

Original Prediction: John Elway

Prediction under adversary: Jeff Dean

Task for DNN	Caption image	Recognise object	Recognise pneumonia	Answer question
Problem	Describes green hillside as grazing sheep	Hallucinates teapot if certain patterns are present	Fails on scans from new hospitals	Changes answer if irrelevant information is added

Source: Geirhos et al. (2020) arXiv:2004.07780.

Why bother with XAI? Because NNs take shortcuts!

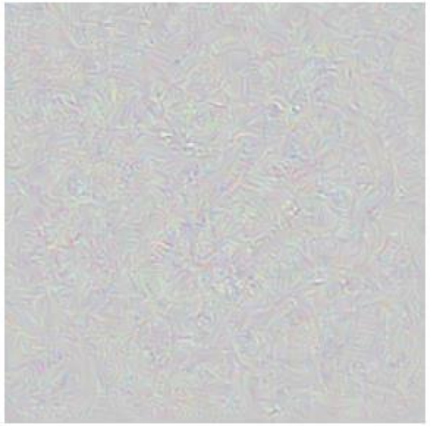

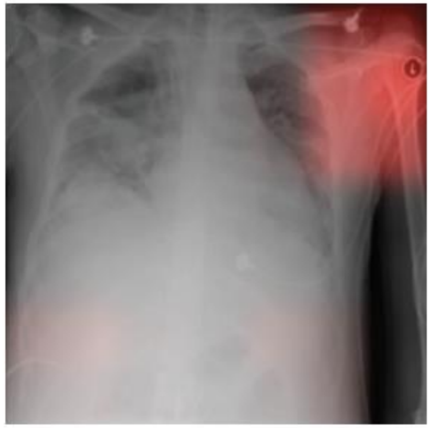

Article: Super Bowl 50

Paragraph: *"Peython Manning became the first quarterback ever to lead two different teams to multiple Super Bowls. He is also the oldest quarterback ever to play in a Super Bowl at age 39. The past record was held by John Elway, who led the Broncos to victory in Super Bowl XXXIII at age 38 and is currently Denver's Executive Vice President of Football Operations and General Manager. Quarterback Jeff Dean had a jersey number 37 in Champ Bowl XXXIV."*

Question: *"What is the name of the quarterback who was 38 in Super Bowl XXXIII?"*

Original Prediction: John Elway

Prediction under adversary: Jeff Dean

Task for DNN	Caption image	Recognise object	Recognise pneumonia	Answer question
Problem	Describes green hillside as grazing sheep	Hallucinates teapot if certain patterns are present	Fails on scans from new hospitals	Changes answer if irrelevant information is added
Shortcut	Uses background to recognise primary object	Uses features irrecognisable to humans	Looks at hospital token, not lung	Only looks at last sentence and ignores context

Source: *Geirhos et al. (2020) arXiv:2004.07780.*

Should you blindly trust hybrid AI-climate models?

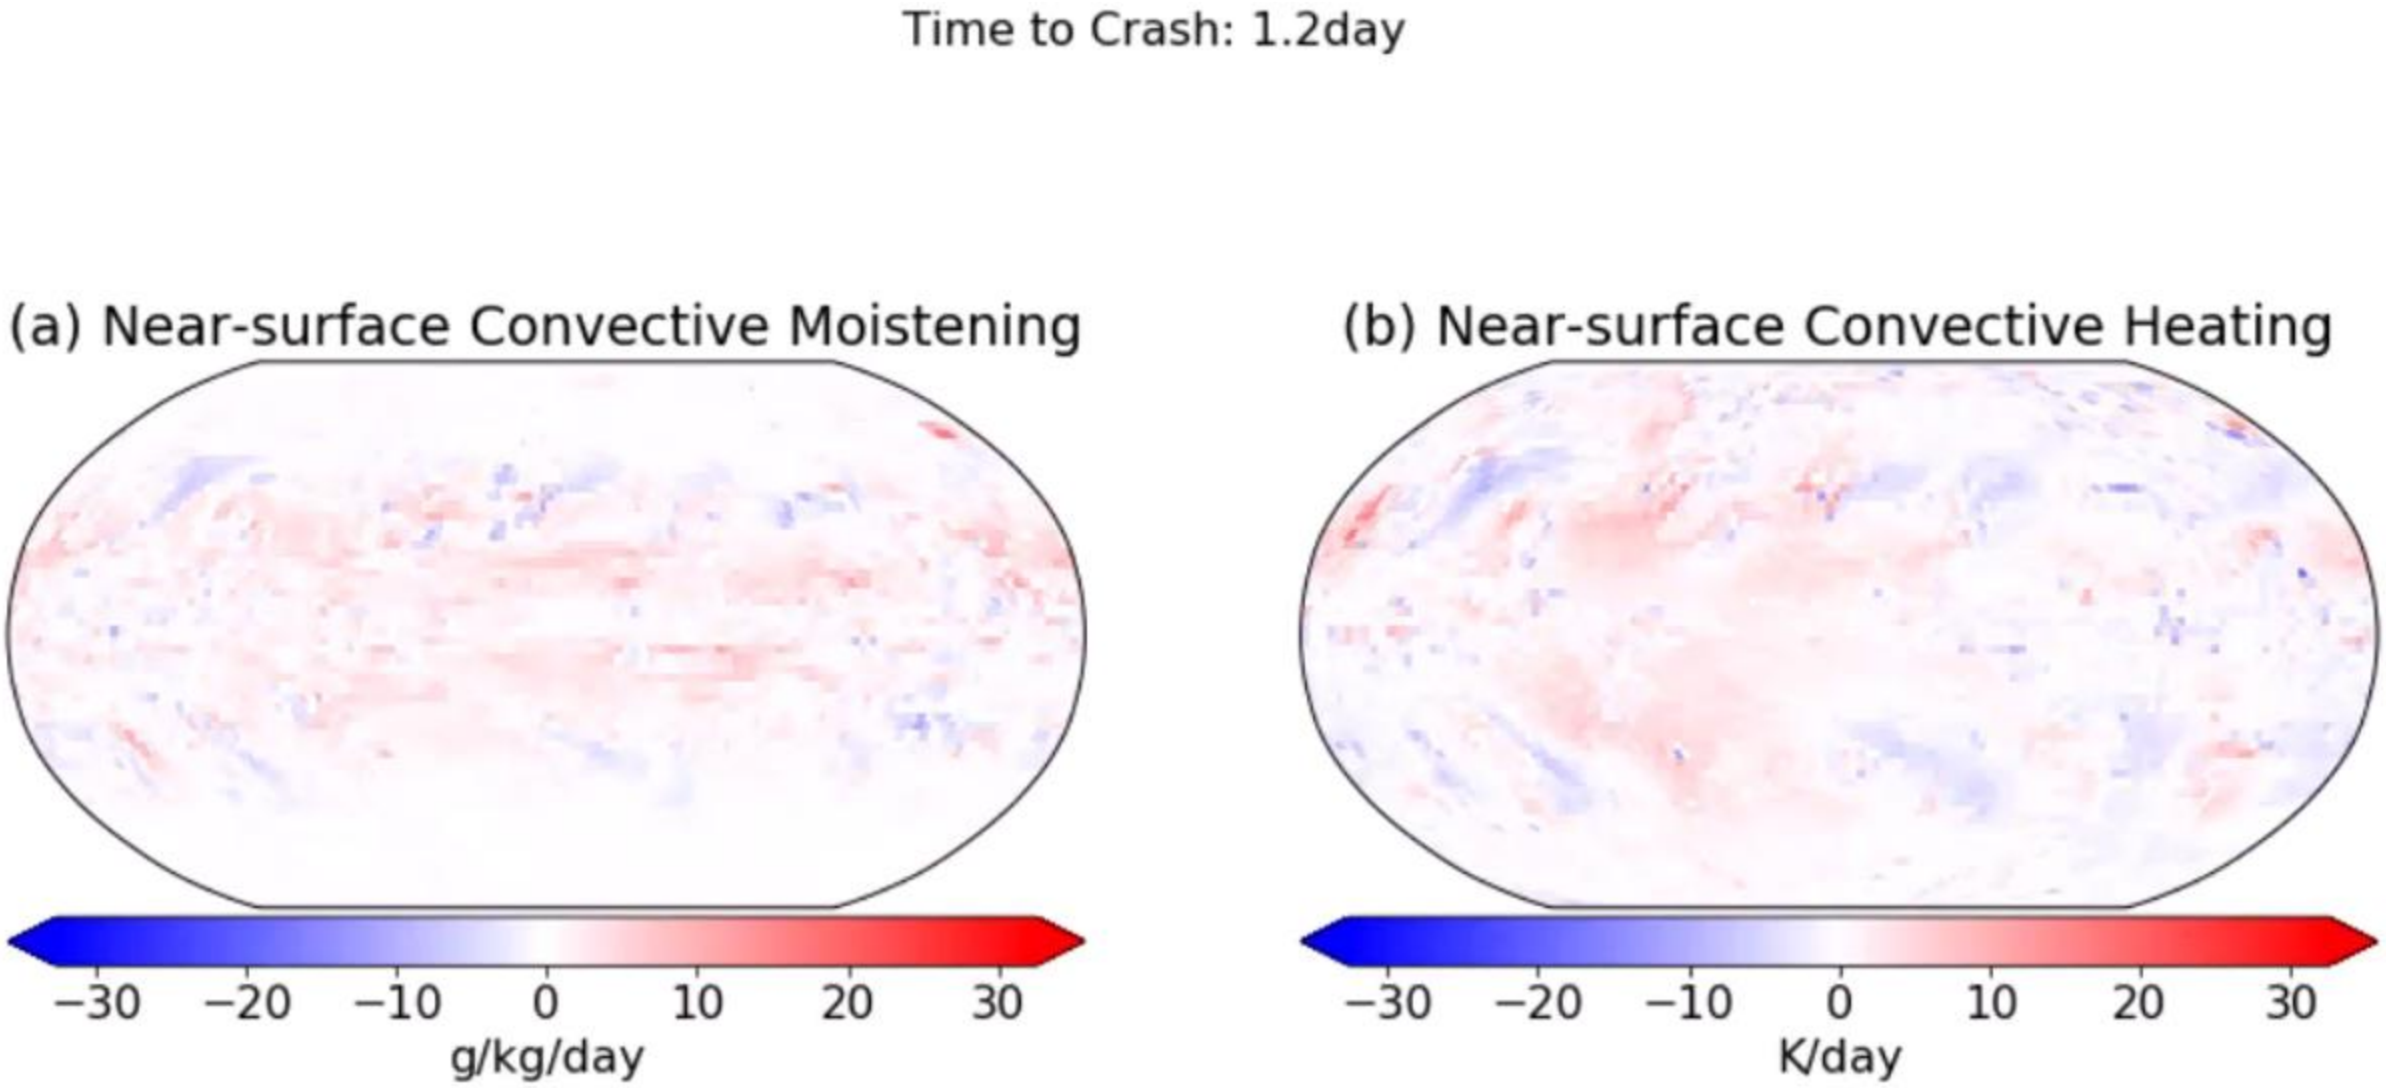

<u>See:</u> Brenowitz, Beucler et al. (2020)

Tailor NN interpretability techniques to parametrization task

Saliency Map

Squared Jacobian of the emulated mapping

Deep learning libraries efficiently calculate Jacobian via automatic differentiation

$$\boldsymbol{J} \stackrel{\mathrm{def}}{=} \left(\frac{\partial \mathbf{Output}}{\partial \mathbf{Input}}\right)_{\mathbf{Input_0}}$$

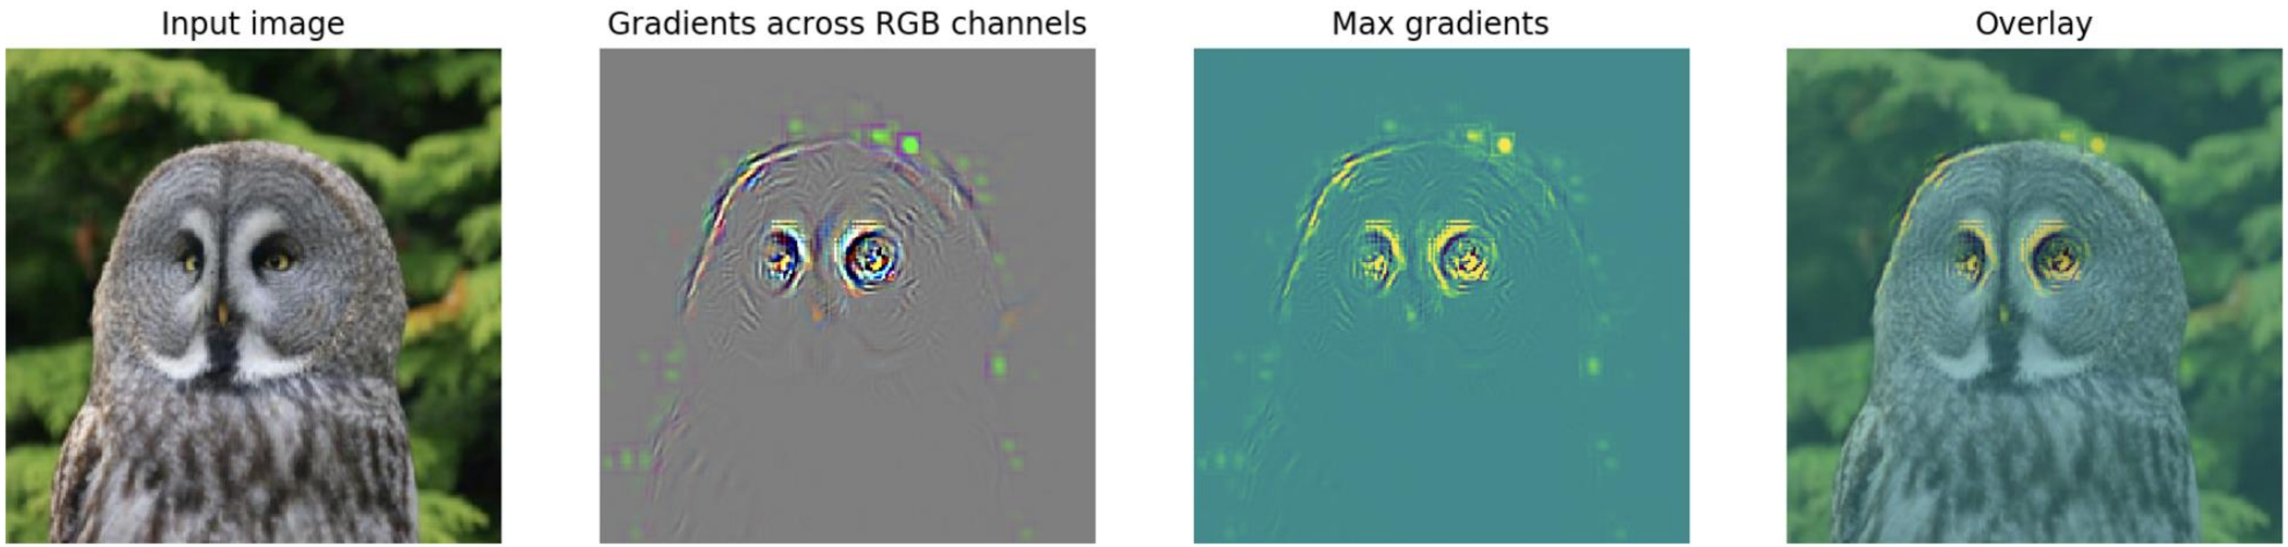

<u>Image source:</u> flashtorch (Github) <u>See</u>: Paszke et al. (2017), Springenbert et al. (2015)

Jacobian reveals linear response of convection

$$\left(\frac{\partial \mathbf{Output}}{\partial \mathbf{Input}}\right) = \frac{\partial\,(\text{Convective Moistening})}{\partial\,(\text{Moisture})} \quad [1/\text{day}]$$

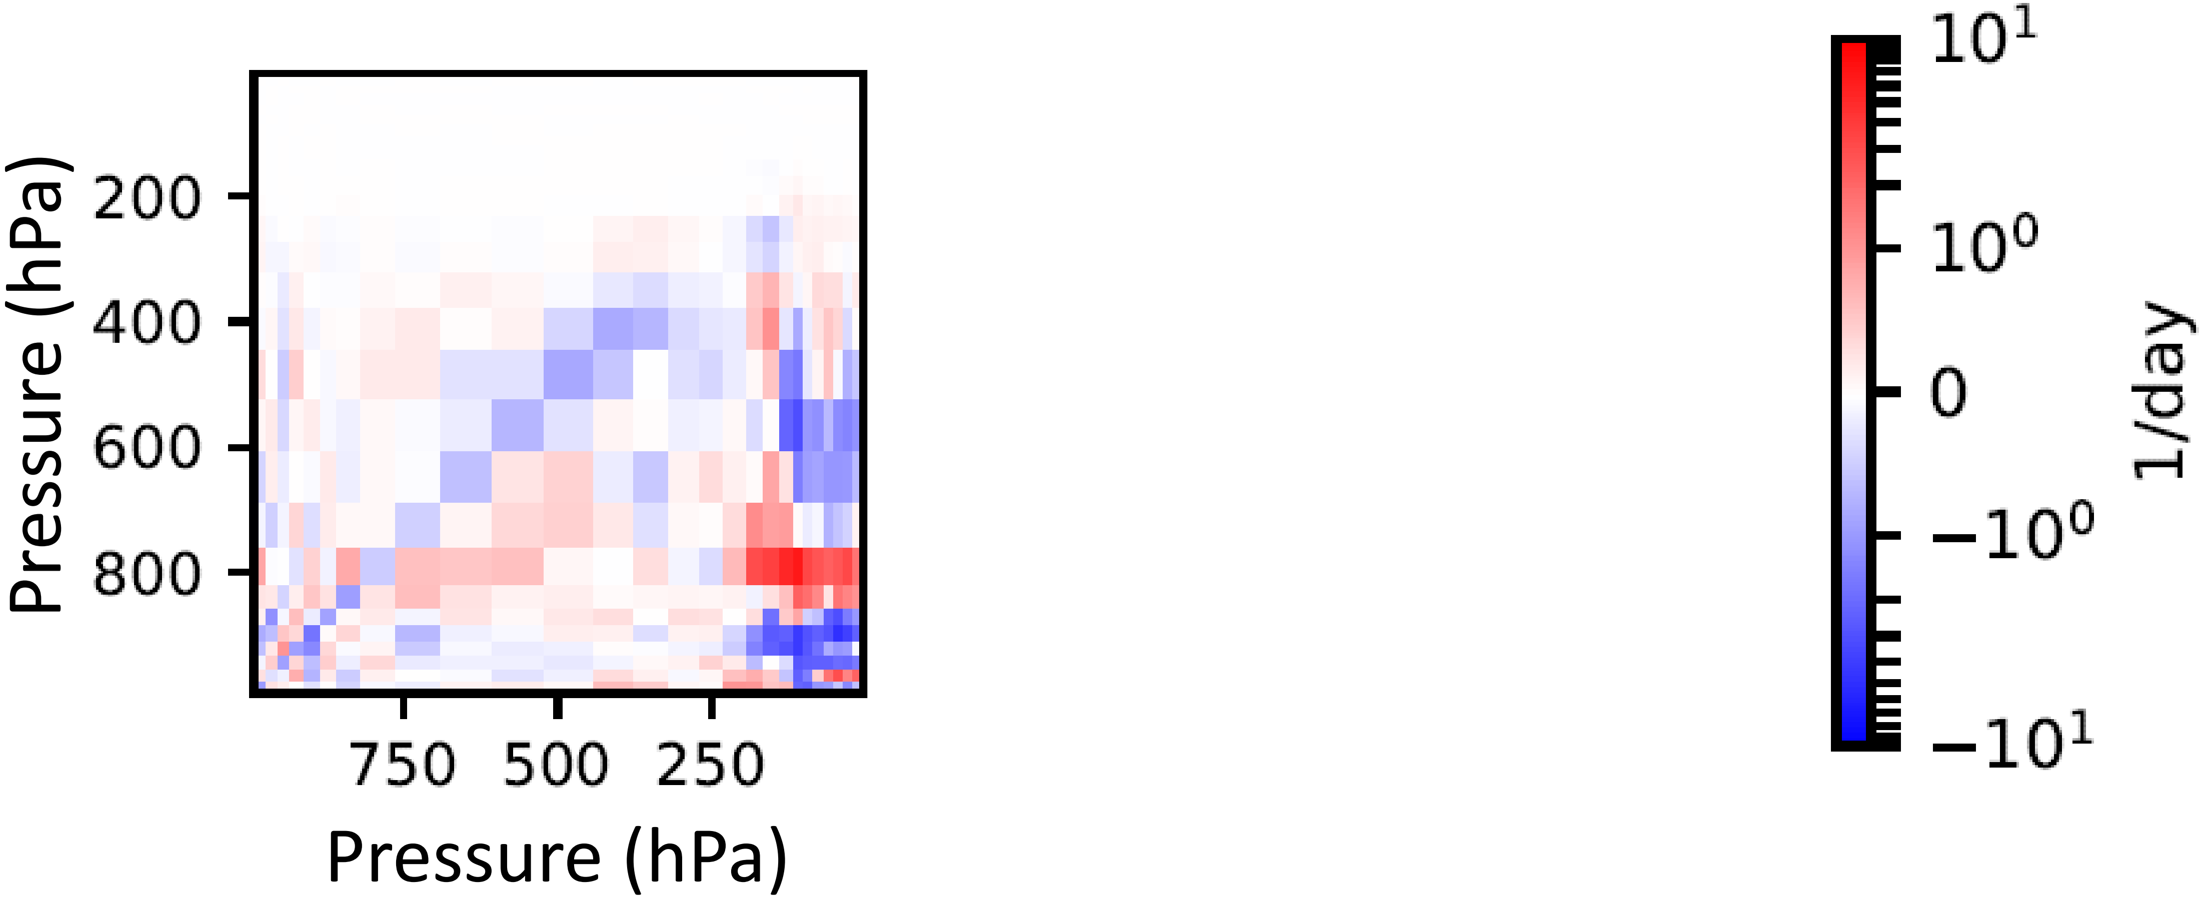

<u>See</u>: Kuang (2018, 2007), Herman and Kuang (2013), Beucler et al. (2018)

Jacobian reveals linear response of convection

$$\left(\frac{\partial \mathbf{Output}}{\partial \mathbf{Input}}\right) = \frac{\partial\,(\text{Convective Moistening})}{\partial\,(\text{Moisture})}\quad [1/\text{day}]$$

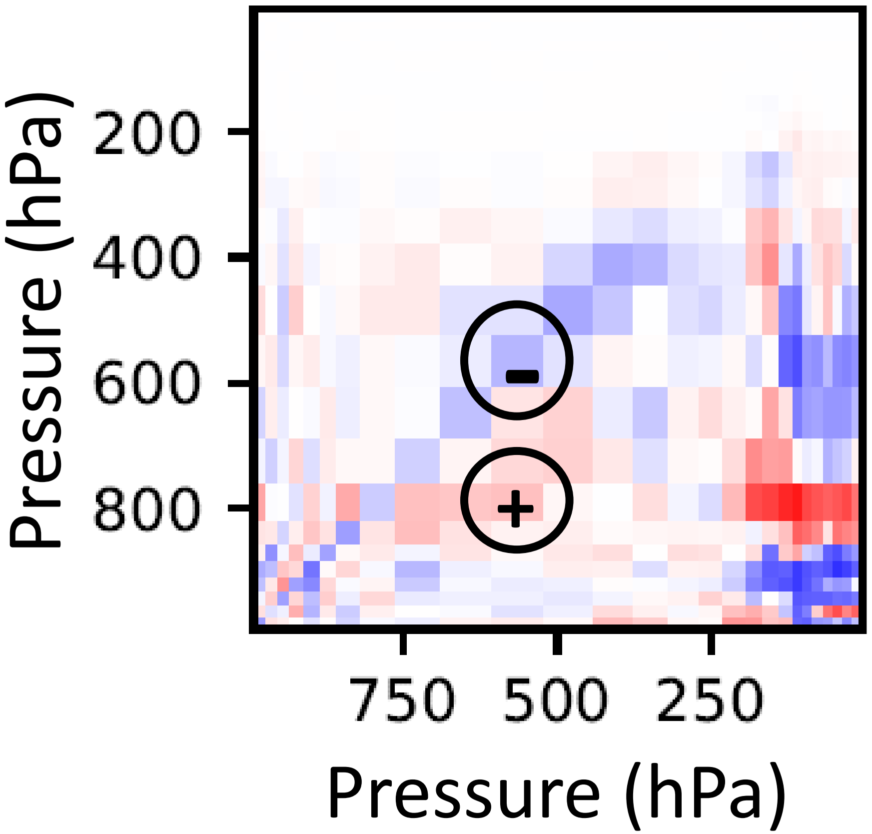

Local anomalies are removed

And redistributed in lower atmosphere

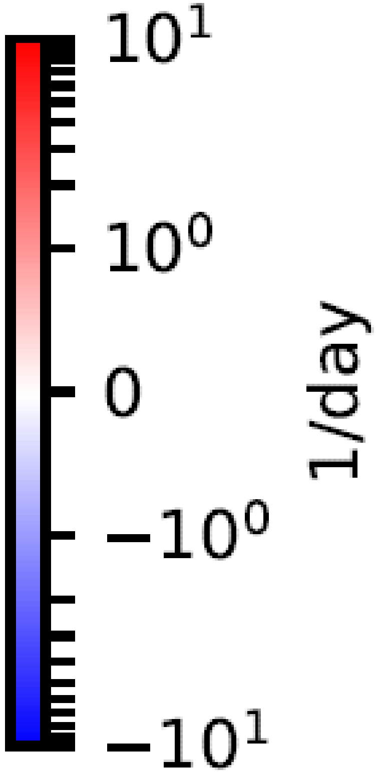

<u>See</u>: Kuang (2018, 2007), Herman and Kuang (2013), Beucler et al. (2018)

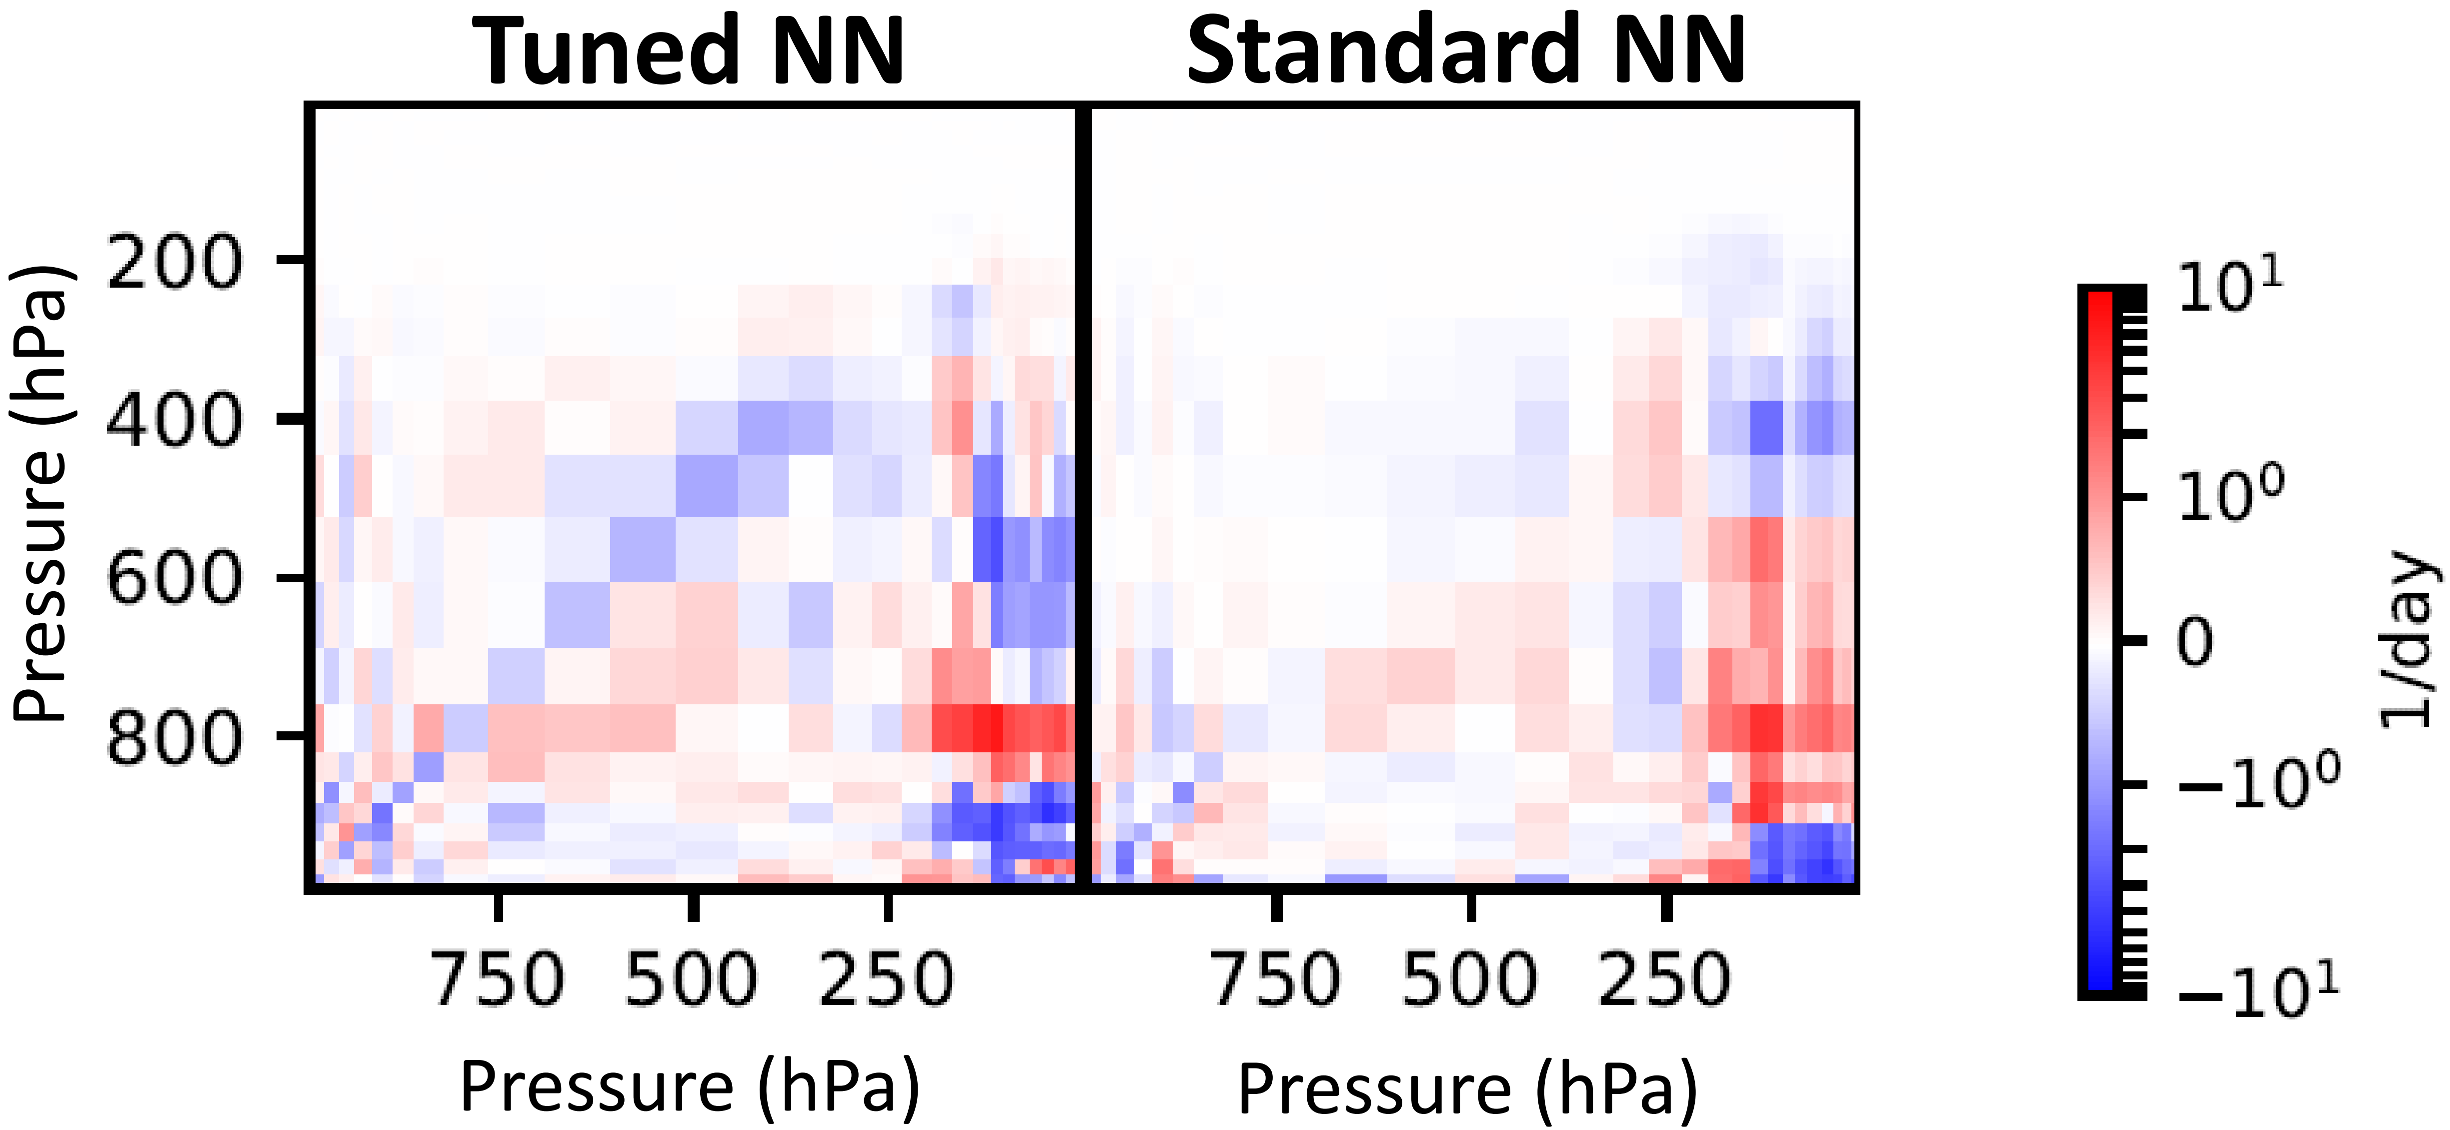

<u>See</u>: Kuang (2018, 2007), Herman and Kuang (2013), Beucler et al. (2018)

Challenge: Offline & online objectives often misaligned in hybrid ESMs

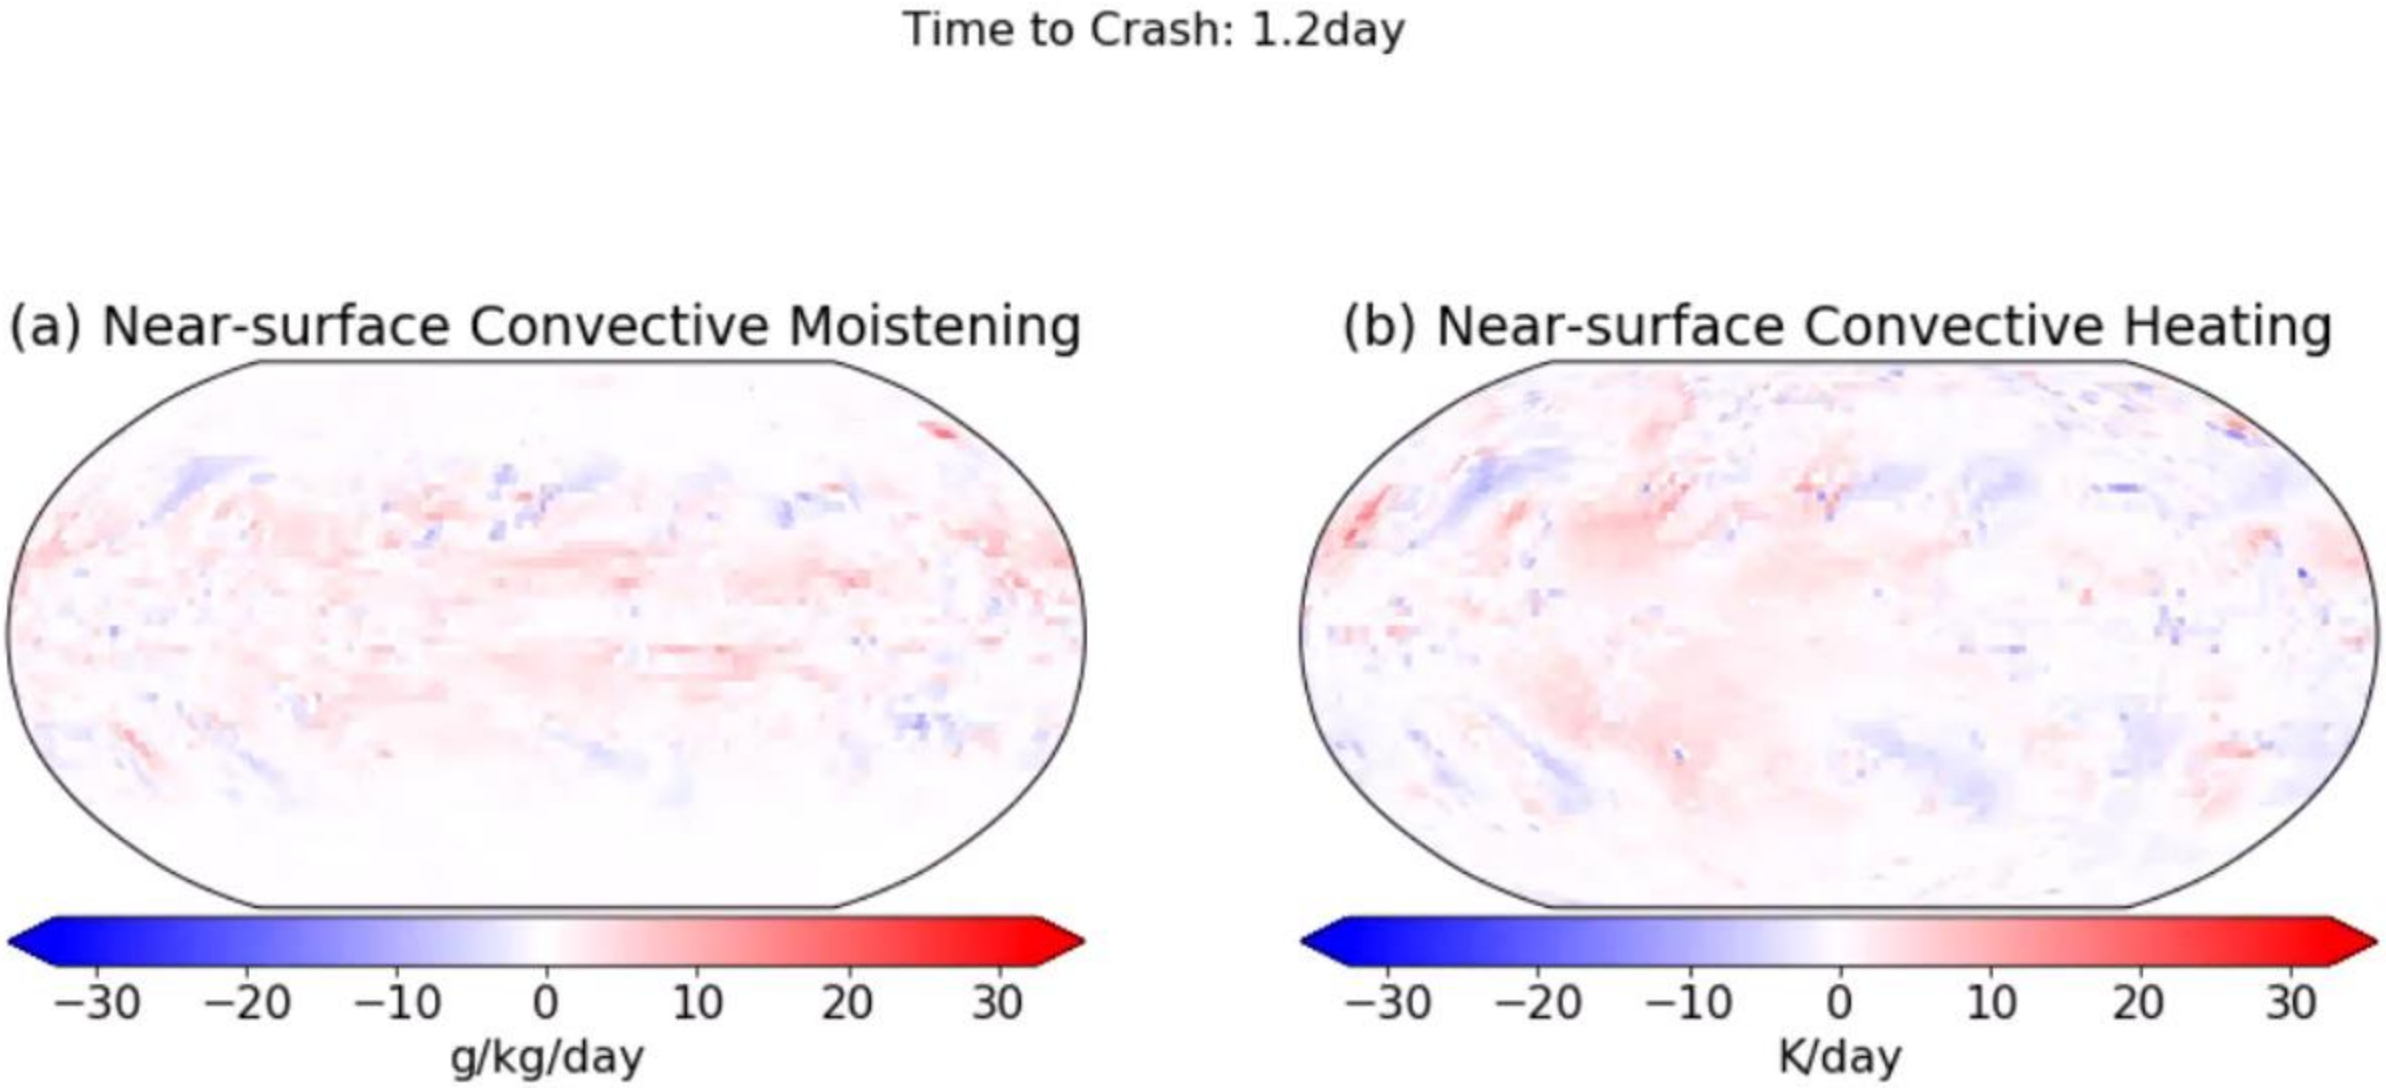

See: Brenowitz, Beucler et al. (2020)

Derivations in Brenowitz et al. (2020, JAS)

the authors upon request.

APPENDIX

Derivation of 2D Anelastic Wave Dynamics

a. Continuous equations

The linearized hydrostatic anelastic equations in the horizontal direction x and height z are given by

$$q_t + \overline{q}_z w = Q_2',$$

$$s_t + \overline{s}_z w = Q_1', \quad \text{and}$$

$$u_t + \phi_x = -du.$$

The prognostic variables are humidity q, dry static energy $s = T + (g/c_p)z$, horizontal velocity u, and vertical velocity w. These are assumed to be perturbations from a large-scale state denoted by an overbar. The anelastic geopotential term is given by $\phi = p'/\rho_0$, where $\rho_0(z)$ is a reference density profile specified for the full nonlinear model.

These prognostic equations are completed by assuming hydrostatic balance and mass conservation. Hydrostatic balance is given by

$$(Aw)_k = a_k w_{k-1} + b_k w_k + c_k w_{k+1}, \quad \text{where}$$

$$a_k = \frac{\rho_{k-1}}{(z_k - z_{k-1})(z_{k+1/2} - z_{k-1/2})\rho_{k-1/2}},$$

$$b_k = -\frac{\rho_k}{(z_{k+1/2} - z_{k-1/2})} \times \left[\frac{1}{(z_{k+1} - z_k)\rho_{k+1/2}} + \frac{1}{(z_k - z_{k-1})\rho_{k-1/2}}\right], \quad \text{and}$$

$$c_k = \frac{\rho_{k+1}}{(z_{k+1} - z_k)(z_{k+1/2} - z_{k-1/2})\rho_{k+1/2}}.$$

The index k ranges from 1 to N, the number of vertical grid cells, and z is the height.

The rigid-lid boundary conditions are satisfied by: $w_0 = -w_1$ and $w_{n+1} = -w_n$. It is not simply w_0 because the vertical velocity should be located at the cell center. These boundary conditions can be implemented by modifying the matrix representation of A to satisfy

$$(Aw)_1 = -a_1 w_1 + b_1 w_1 + c_1 w_2,$$
$$(Aw)_n = a_n w_{n-1} + b_n w_n - c_n w_n$$

at the lower and upper boundaries.

Coupling Linear Response Function to Gravity Waves reveals the unstable NN **offline**

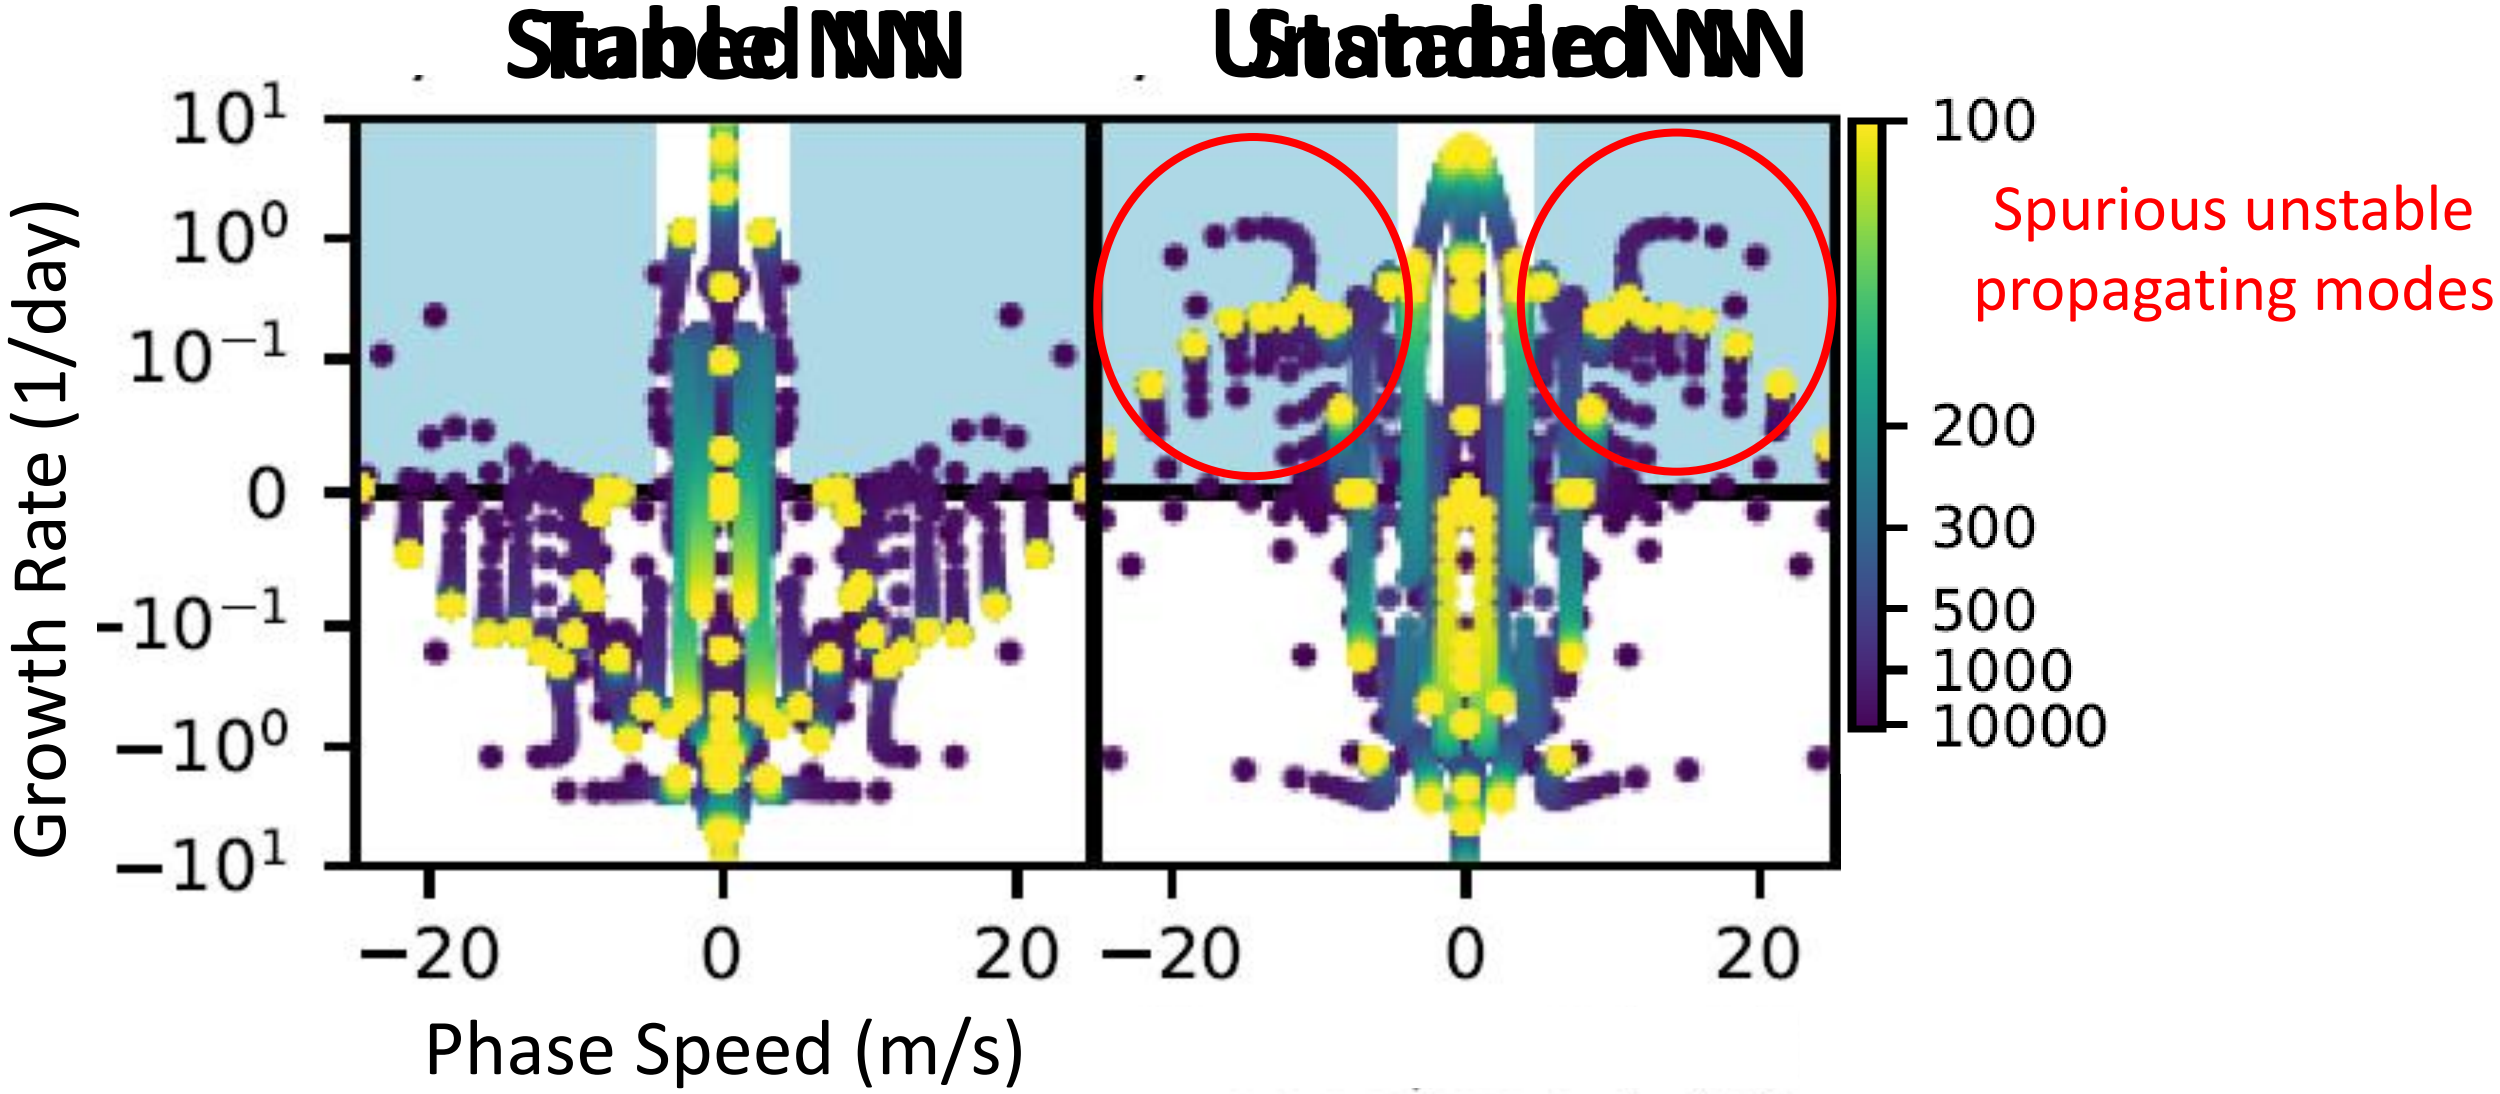

<u>See</u>: *Kuang (2018), Brenowitz, Beucler et al. (2020)*

Stability diagram helped stabilize NNs **offline**

Super Parametrized

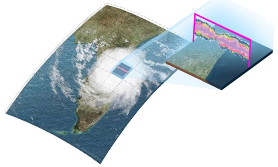

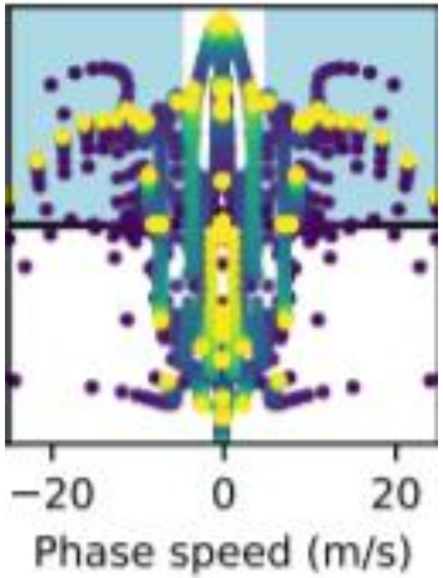

“Regularize” Inputs
by adding
Gaussian noise

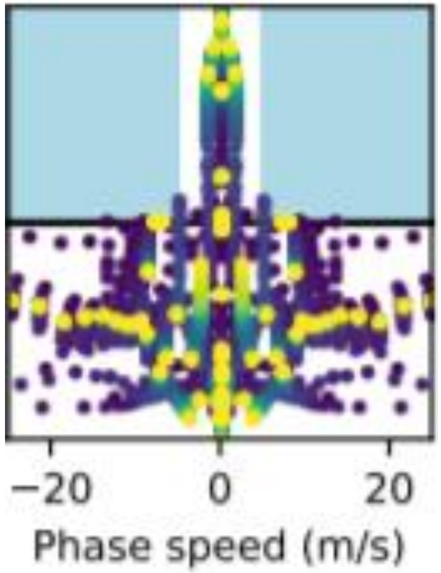

Global Cloud-Resolving

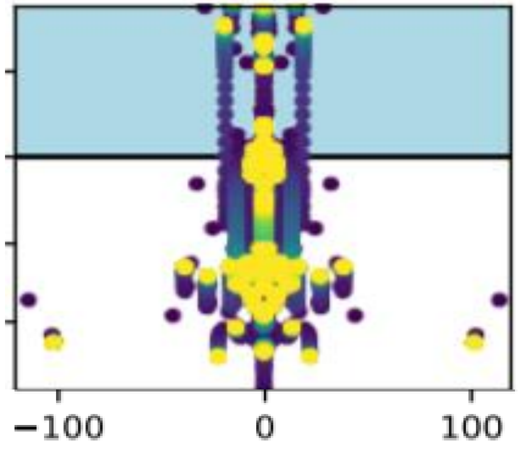

Remove
upper-atmospheric
Inputs

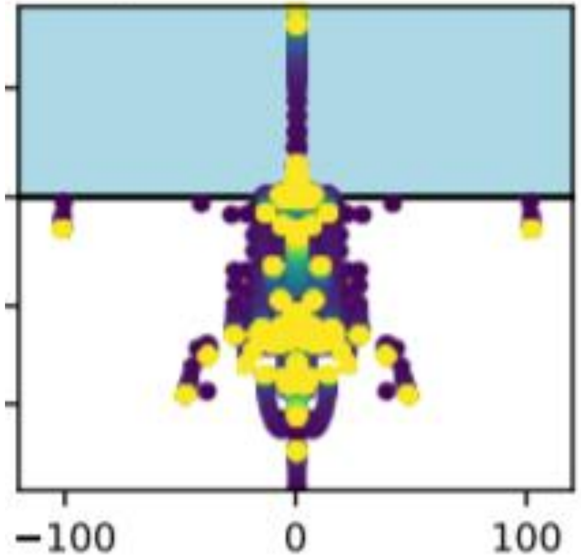

Both stabilized NN ran without crashing for **1month+** when coupled to climate models

Super Parametrized

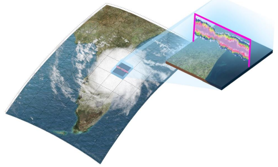

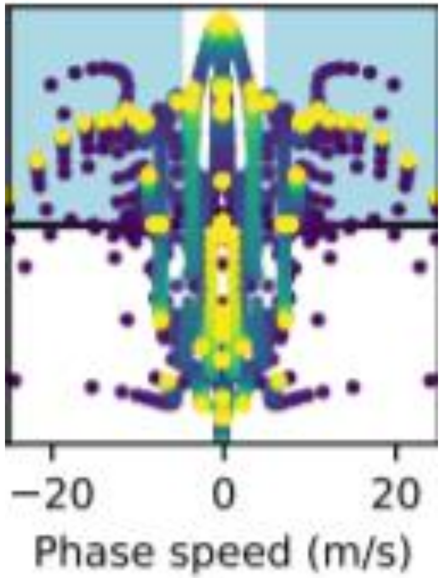

“Regularize” Inputs
by adding
Gaussian noise

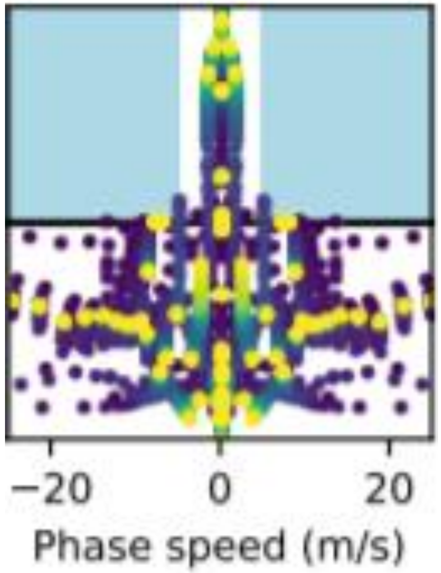

Global Cloud-Resolving

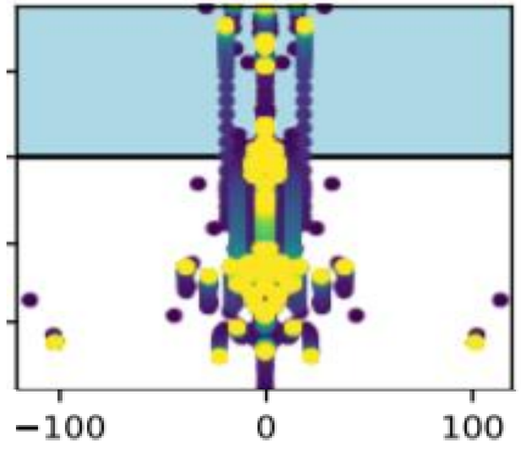

Remove
upper-atmospheric
Inputs

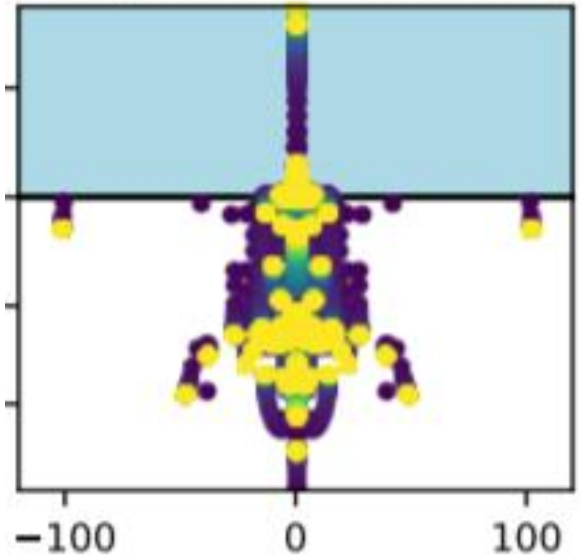

Journal of the Royal Statistical Society Series B: Statistical Methodology, 2024, **00**, 1–25
https://doi.org/10.1093/jrsssb/qkae108

ROYAL STATISTICAL SOCIETY
DATA | EVIDENCE | DECISIONS
Journal of the Statistics Society
Series B
Statistical Methodology
B

Original Article

Engression: extrapolation through the lens of distributional regression

Xinwei Shen and Nicolai Meinshausen

Seminar für Statistik, Department of Mathematics, ETH Zürich, Zürich, Switzerland
Address for correspondence: Xinwei Shen, Seminar für Statistik, Department of Mathematics, ETH Zürich, Rämistrasse 101, 8092 Zürich, Switzerland. Email: xinwei.shen@stat.math.ethz.ch

Progression: an extrapolation principle for regression

Gloria Buriticá
Université Paris-Saclay, AgroParisTech, INRAE, UMR MIA Paris-Saclay, 22 place de l'Agrono 91123 Palaiseau, France
Sebastian Engelke
Research Institute for Statistics and Information Science, University of Geneva, Boulevard du d'Arve 40, 1205 Geneva, Switzerland.

Journal of Advances in Modeling Earth Systems / Volume 14, Issue 9 / e2022MS003219

Research Article Open Access

Correcting a 200 km Resolution Climate Model in Multiple Climates by Mac Learning From 25 km Resolution Simulations

Spencer K. Clark, Noah D. Brenowitz, Brian Henn, Anna Kwa, Jeremy McGibbon, W. Andre Perkin Oliver Watt-Meyer, Christopher S. Bretherton, Lucas M. Harris

First published: 02 September 2022
https://doi.org/10.1029/2022MS003219

4) Many challenges remain unsolved:

1. Stability, **extrapolation** behavior, and recalibration of the host model

4) Many challenges remain unsolved:

1. Stability, extrapolation behavior, and **recalibration of the host model**

Offline Learning = Best fit given the collected data
Online Learning = Integrating dynamical model during training

$$
\begin{tikzcd}
\mathbf{y}(t) \arrow[r, dash] \arrow[d, dash] & f(\mathbf{y}(t)) \arrow[r] & \mathbf{y}(t+\Delta t) & \cdots & \mathbf{y}(t+N\Delta t) \arrow[r, dash] & f(\mathbf{y}(t+N\Delta t)) \arrow[r] & \mathbf{y}(t+(N+1)\Delta t) \\
\mathcal{T}(\mathbf{y}(t)) \arrow[d] & g(\bar{\mathbf{y}}(t)) \arrow[d] & \color{blue}{\mathcal{L}_{\text{post}}} \arrow[u, dashed, blue] \arrow[d, dashed, blue] & \color{blue}{\cdots} & \color{blue}{\mathcal{L}_{\text{post}}} \arrow[u, dashed, blue] \arrow[d, dashed, blue] & g(\bar{\mathbf{y}}(t+N\Delta t)) \arrow[d] & \color{blue}{\mathcal{L}_{\text{post}}} \arrow[u, dashed, blue] \arrow[d, dashed, blue] \\
\bar{\mathbf{y}}(t) \arrow[ur, bend left] \arrow[dr, bend right] & + \arrow[r] & \bar{\mathbf{y}}(t+\Delta t) & \cdots & \bar{\mathbf{y}}(t+N\Delta t) \arrow[ur, bend left] \arrow[dr, bend right] & + \arrow[r] & \bar{\mathbf{y}}(t+(N+1)\Delta t) \\
 & \mathcal{M}(\bar{\mathbf{y}}(t)|\theta) \arrow[u] & & & & \mathcal{M}(\bar{\mathbf{y}}(t+N\Delta t)|\theta) \arrow[u] &
\end{tikzcd}
$$

Figure source: *Frézat et al. (2022)*

Challenge: Offline & online objectives often misaligned in hybrid ESMs

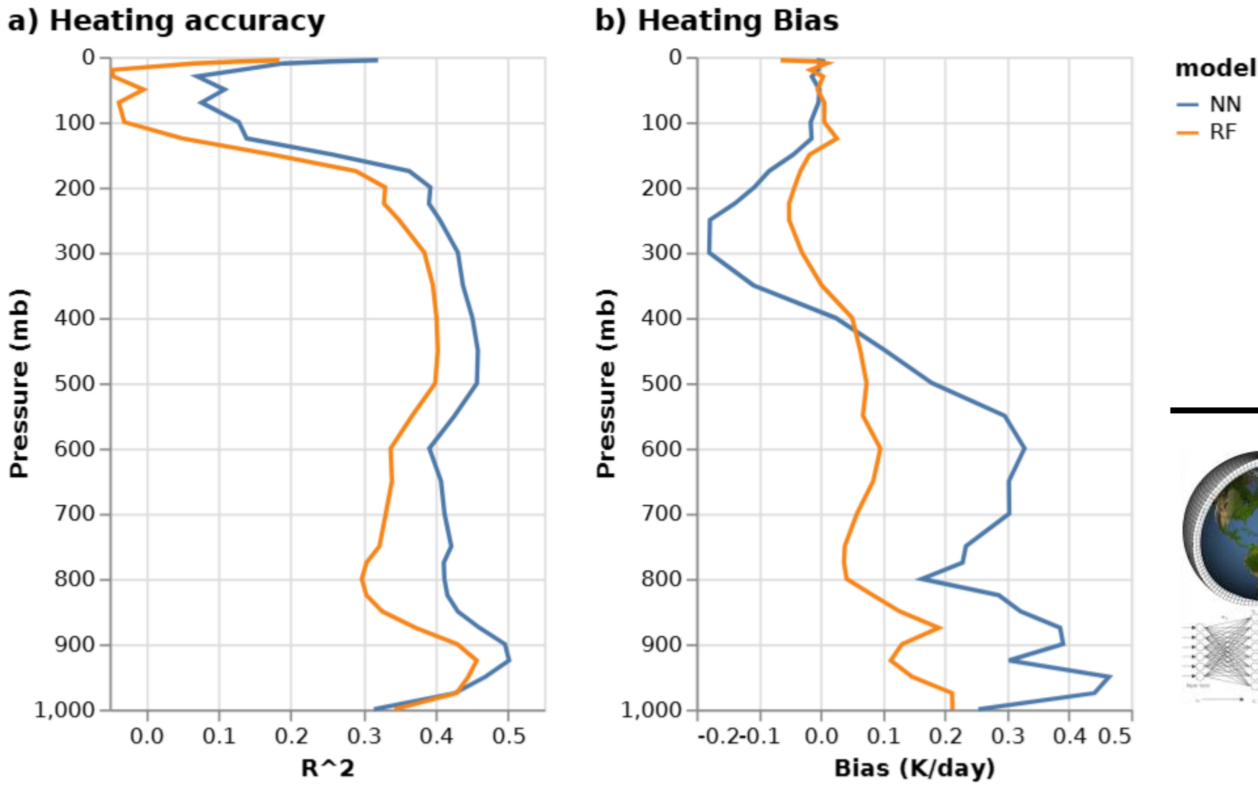

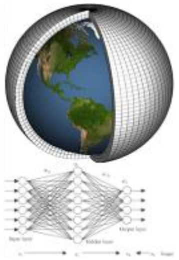

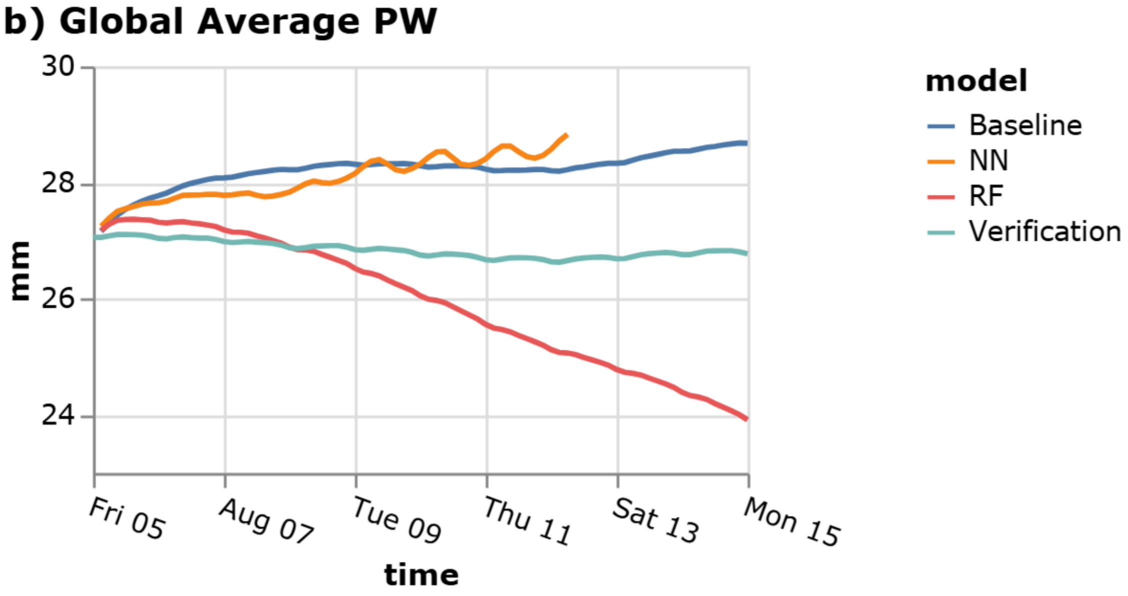

Source: Brenowitz et al. (2020)

Challenge: Offline & online objectives often misaligned in hybrid ESMs

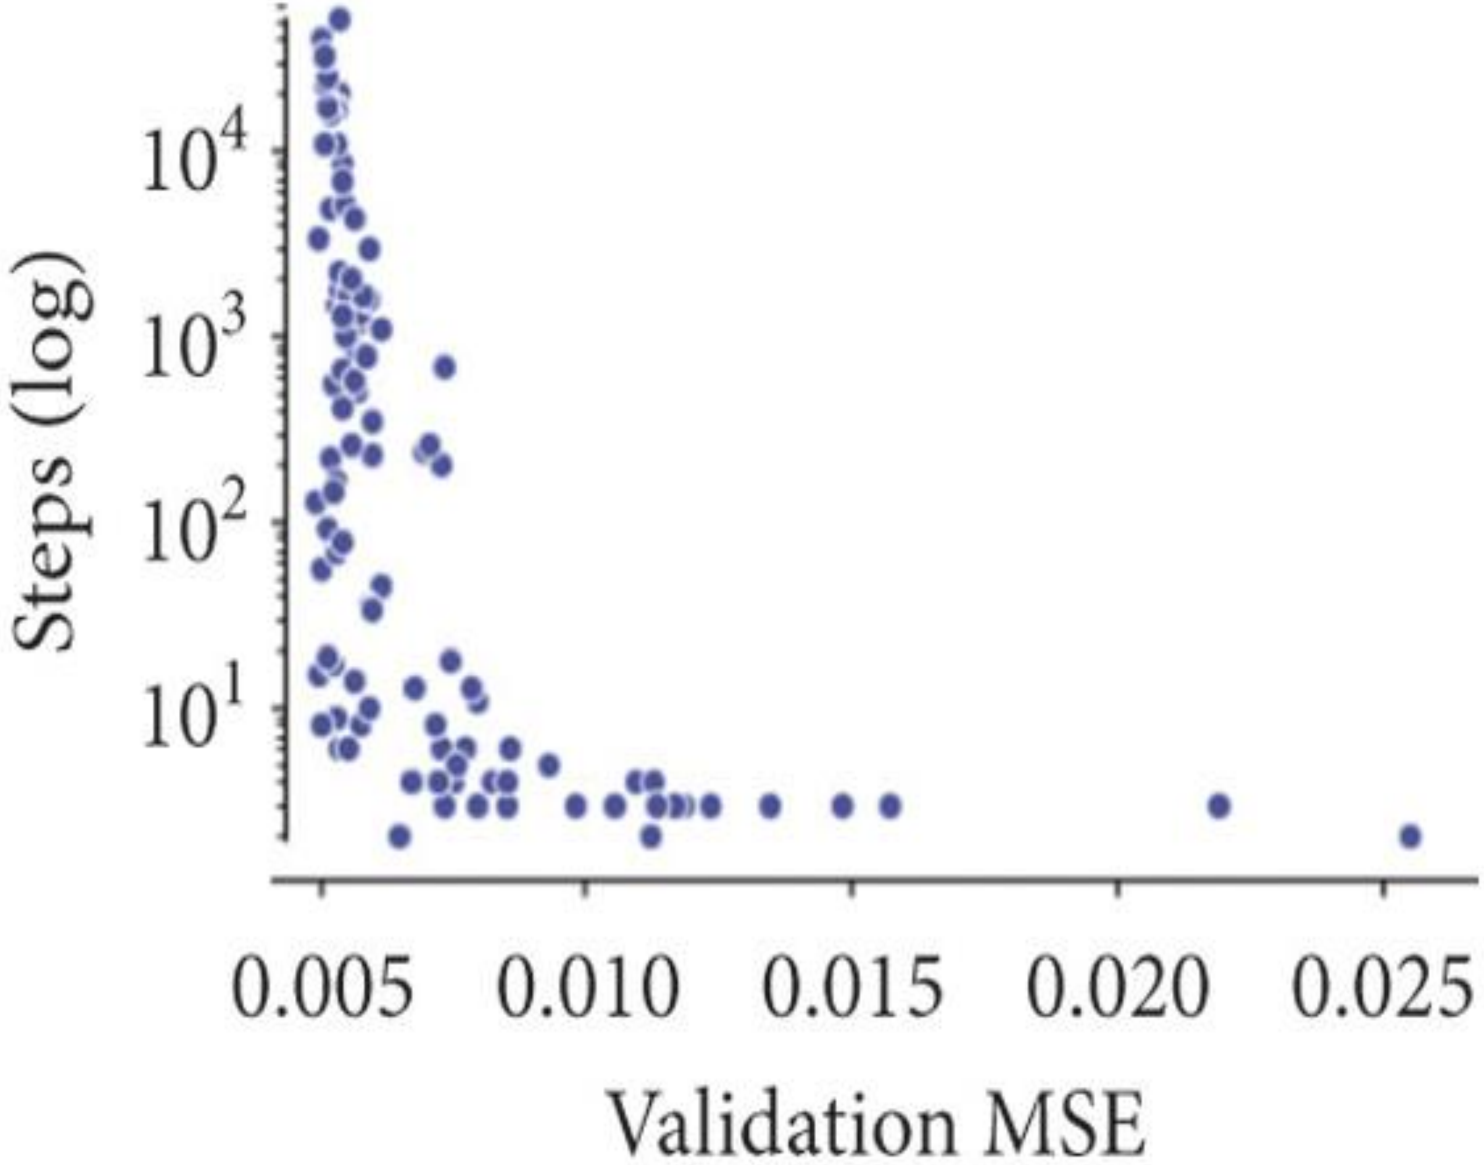

Sources: Ott et al. (2020)

Challenge: Offline & online objectives often misaligned in hybrid ESMs

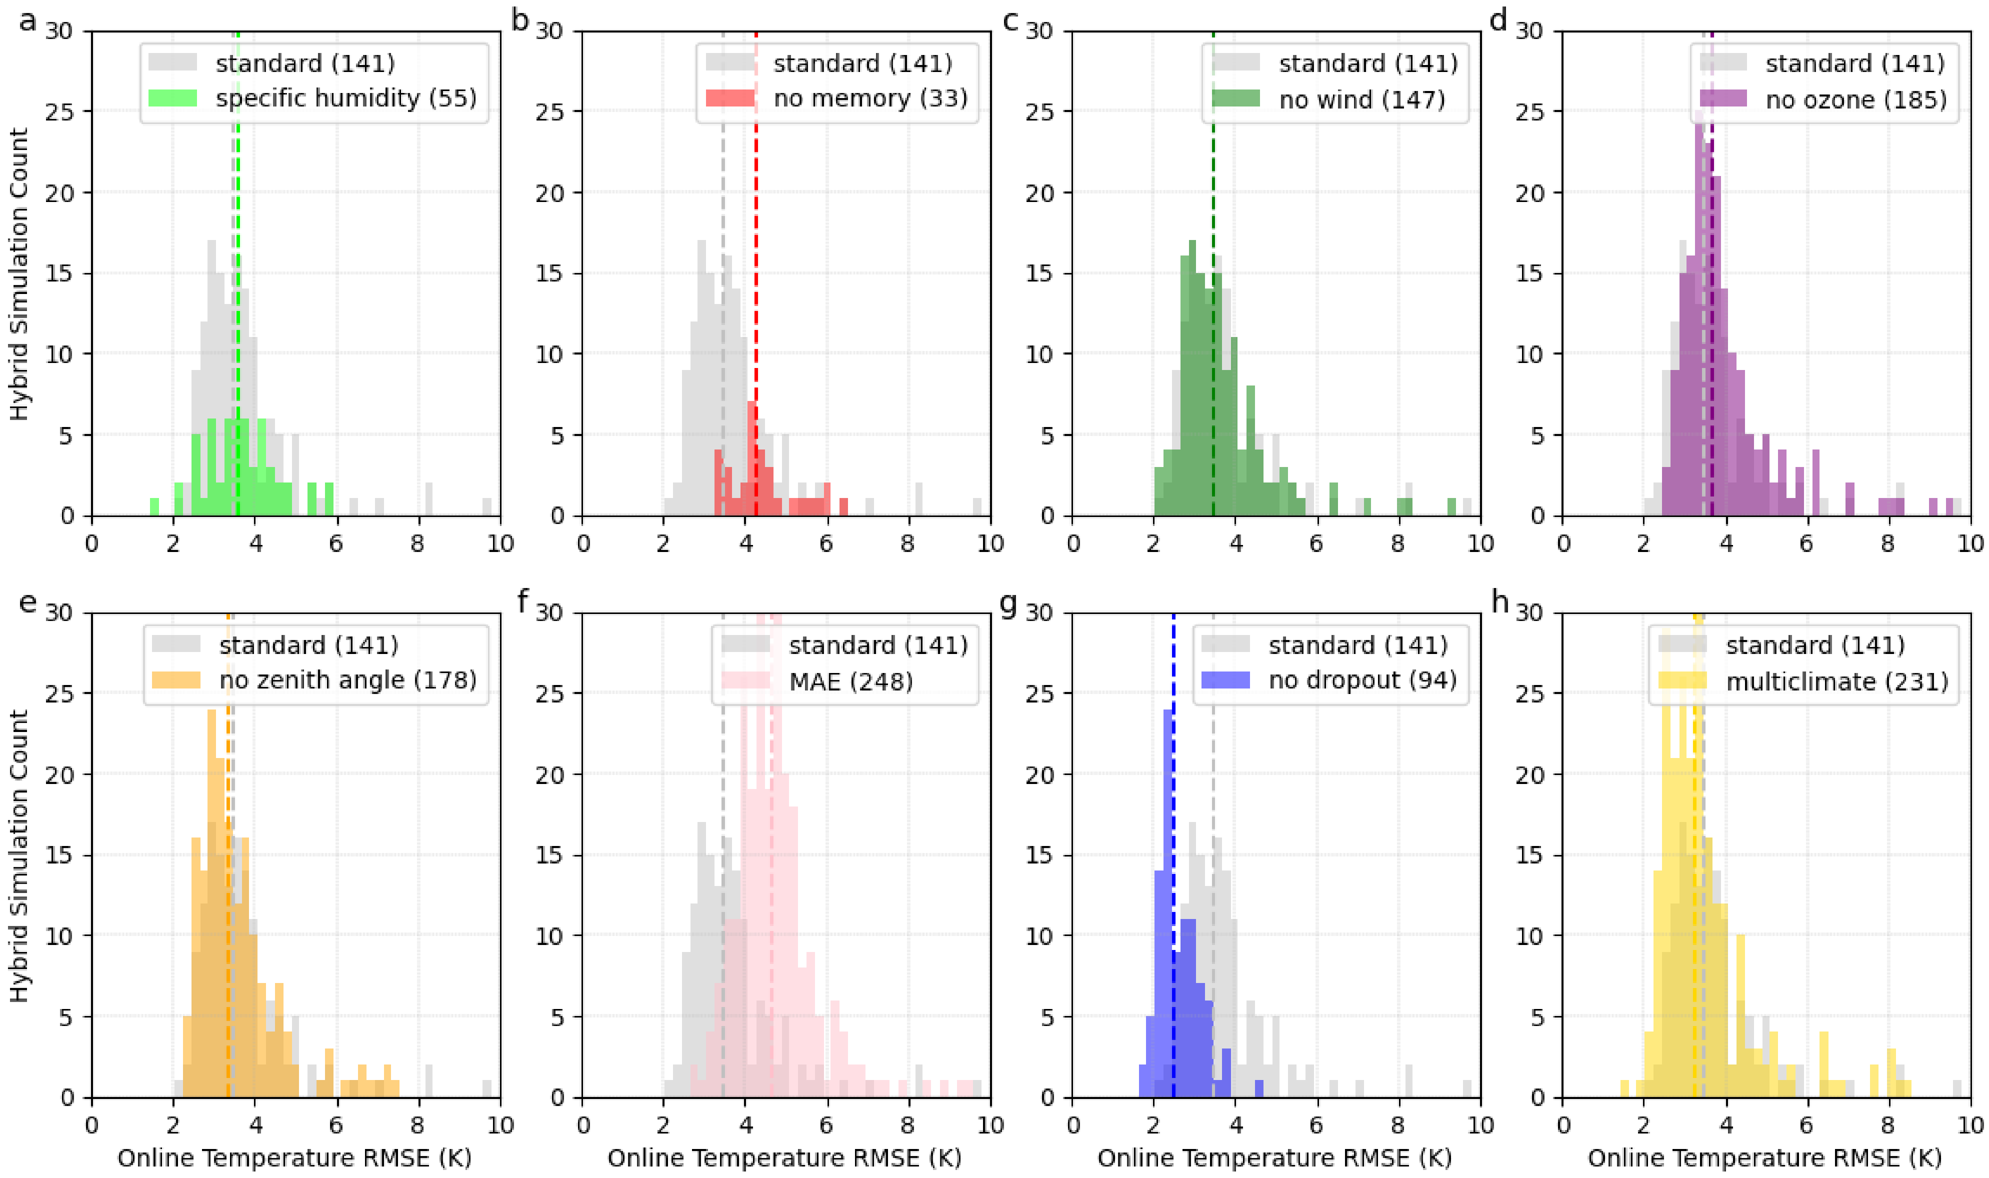

Source: Lin et al. (2024)

Journal of Advances in Modeling Earth Systems / Volume 16, Issue 11 / e2024MS004485

Research Article | Open Access

Online Learning of Entrainment Closures in a Hybrid Machine Learning Parameterization

Costa Christopoulos, Ignacio Lopez-Gomez, Tom Beucler, Yair Cohen, Charles Kawczynski, Oliver R. A. Dunbar, Tapio Schneider

First published: 14 November 2024
https://doi.org/10.1029/2024MS004485

4) Many challenges remain unsolved:

1. Stability, extrapolation behavior, and **recalibration of the host model**

Strategy: Machine Learn Parameters of Physical Models

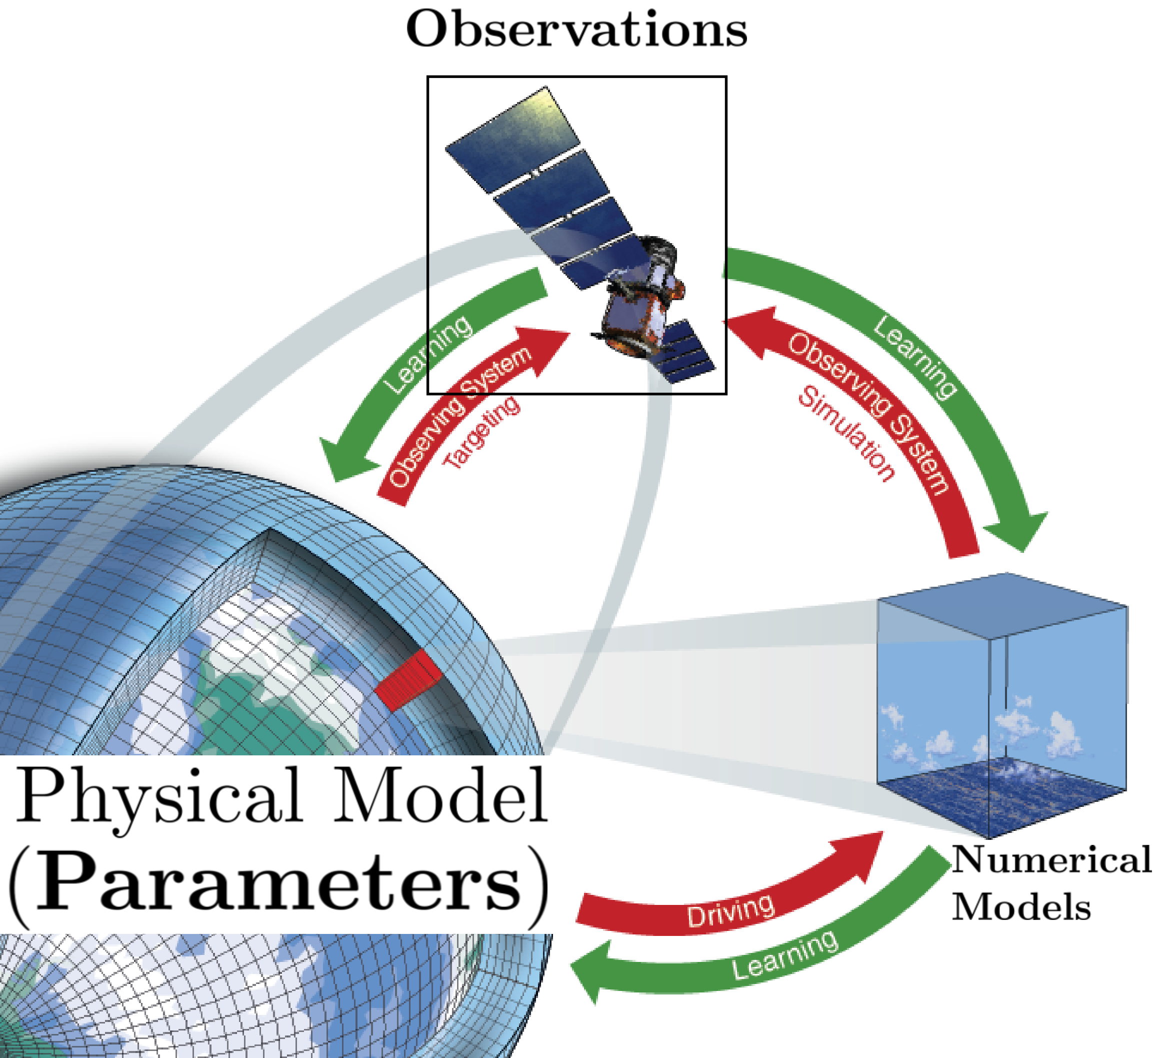

Parameter estimation/calibration problem

Subfield of data assimilation = Optimal state estimation given real-time data
See ECMWF resources on 4D-Var

$$p\left(\text{Parameters}|\text{Obs}\right)$$

ML helps maximize the likelihood

See: Brajard et al. (2021), ECMWF Fact sheets; Image Source: CliMA

Learn Parameters of Physical Models
ML-based frameworks: CES & ESEm

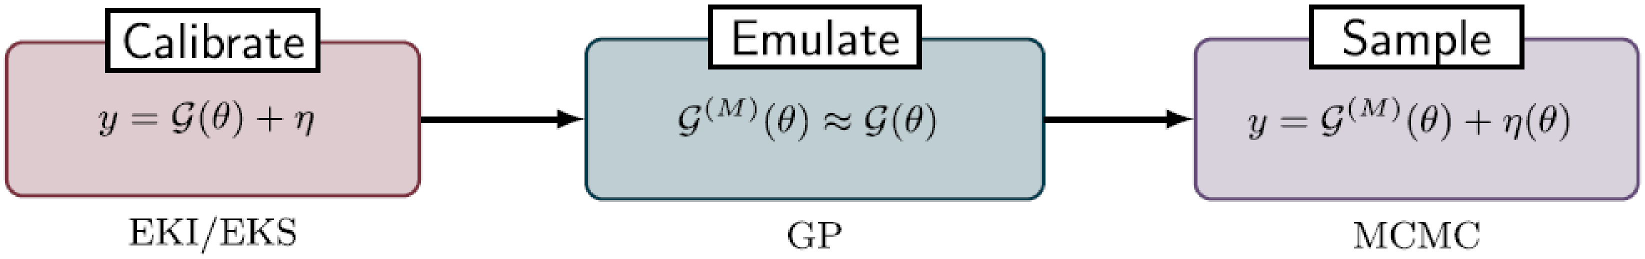

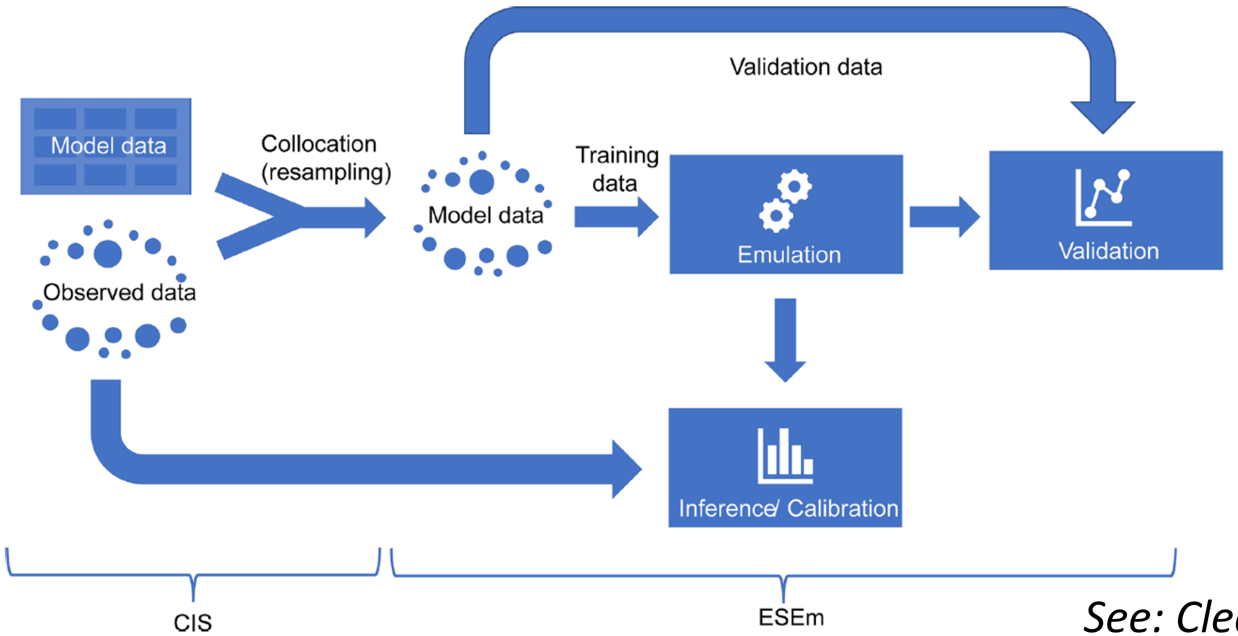

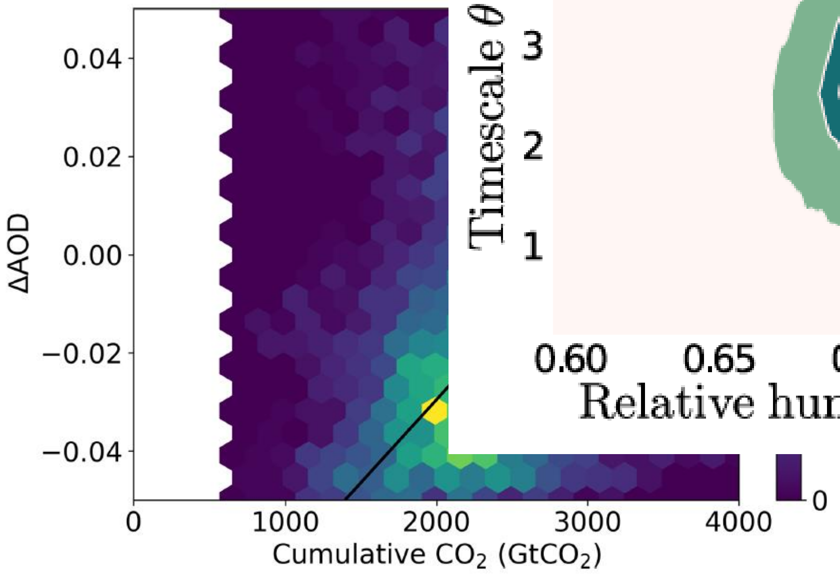

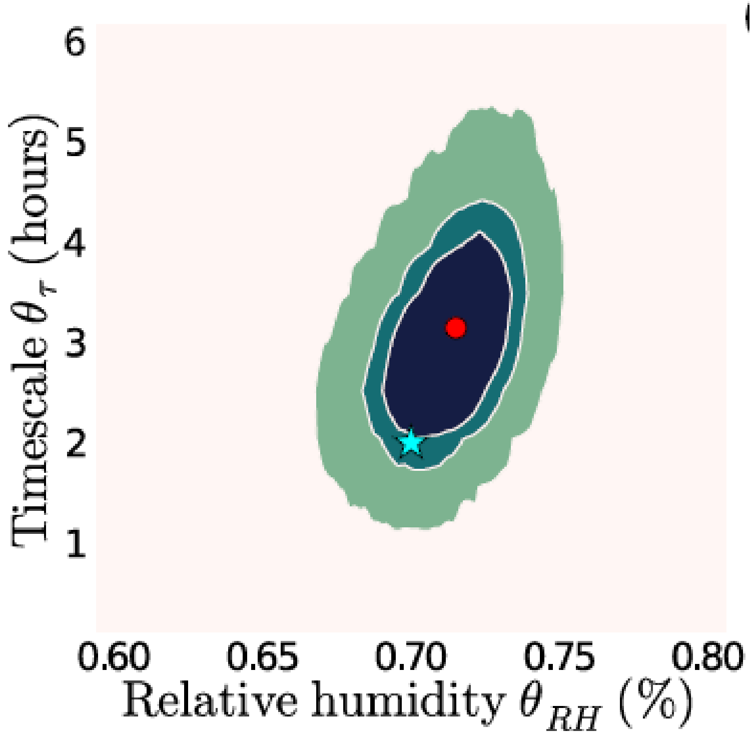

See: Cleary et al. (2021), Howland et al. (2022), Watson-Parris et al. (2021)

Ensemble Kalman Inversion = Gradient-free, data-assimilation technique
Inverse problem framing: Direct learning from climate statistics

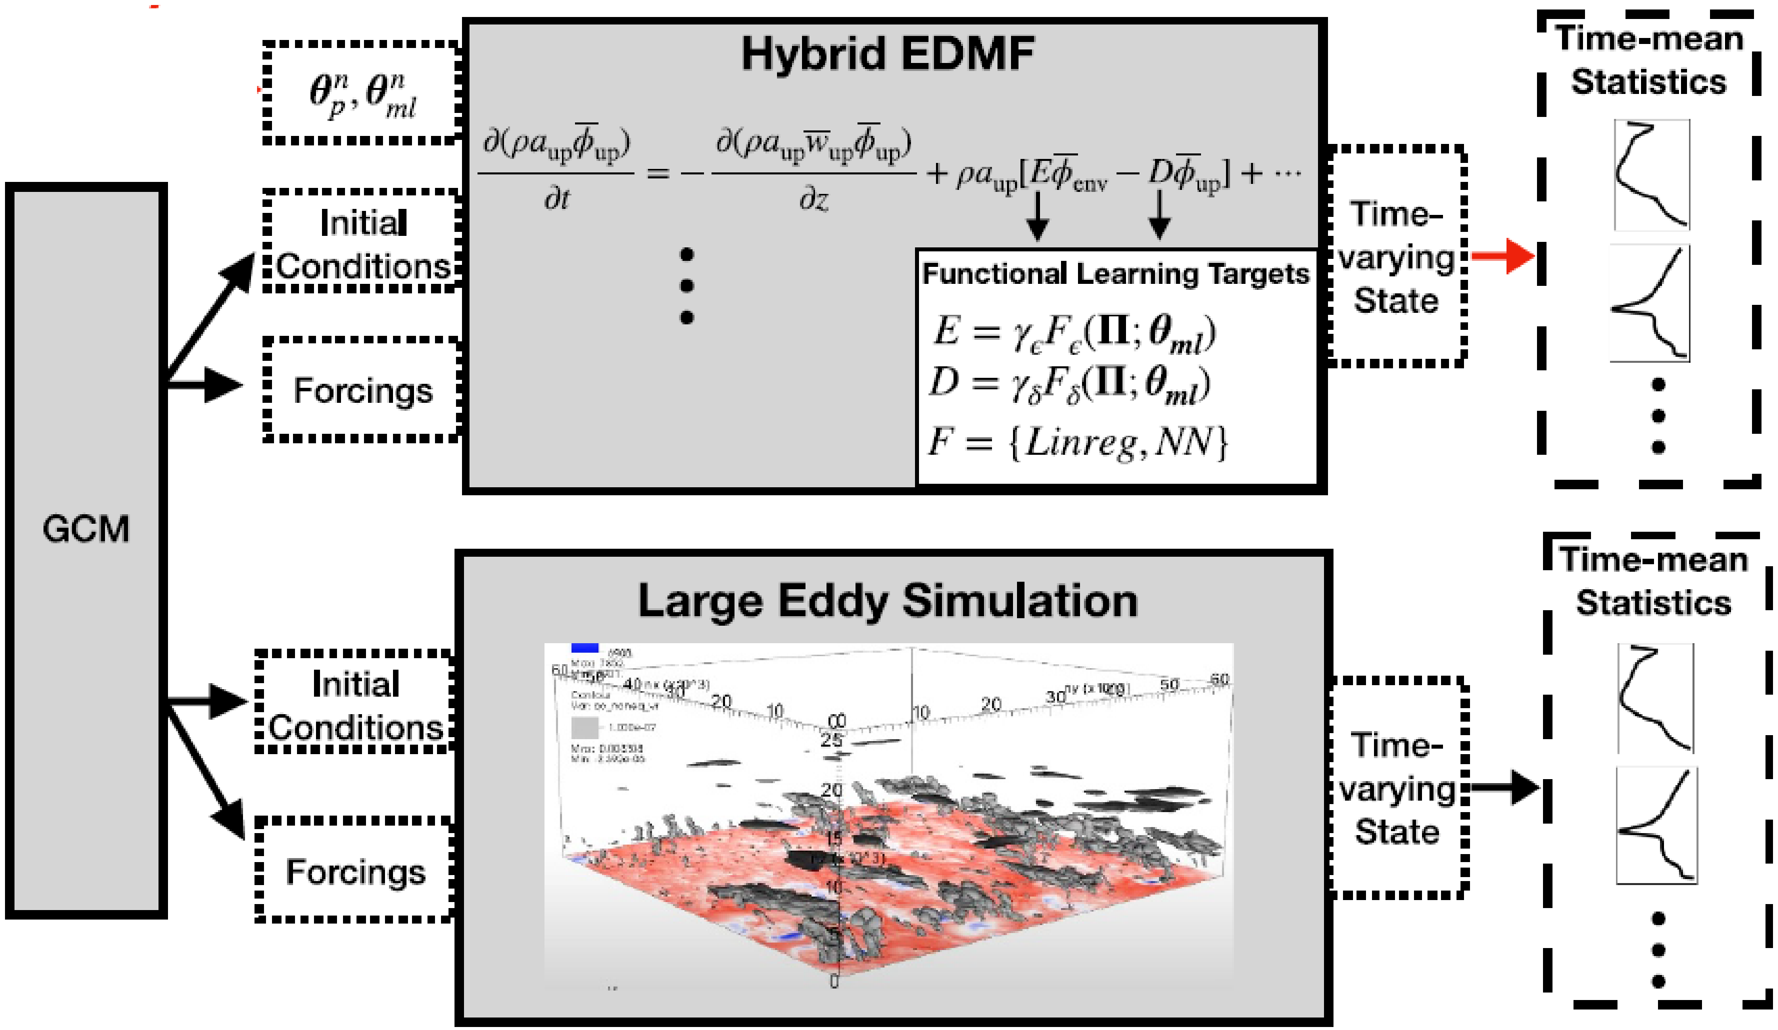

<u>See</u>: Christopoulos et al. (2024); Lopez-Gomez et al. (2022)

Advantages: Guaranteed stability once trained, +4K generalization

Challenges: Cost, instability and host model issues during training

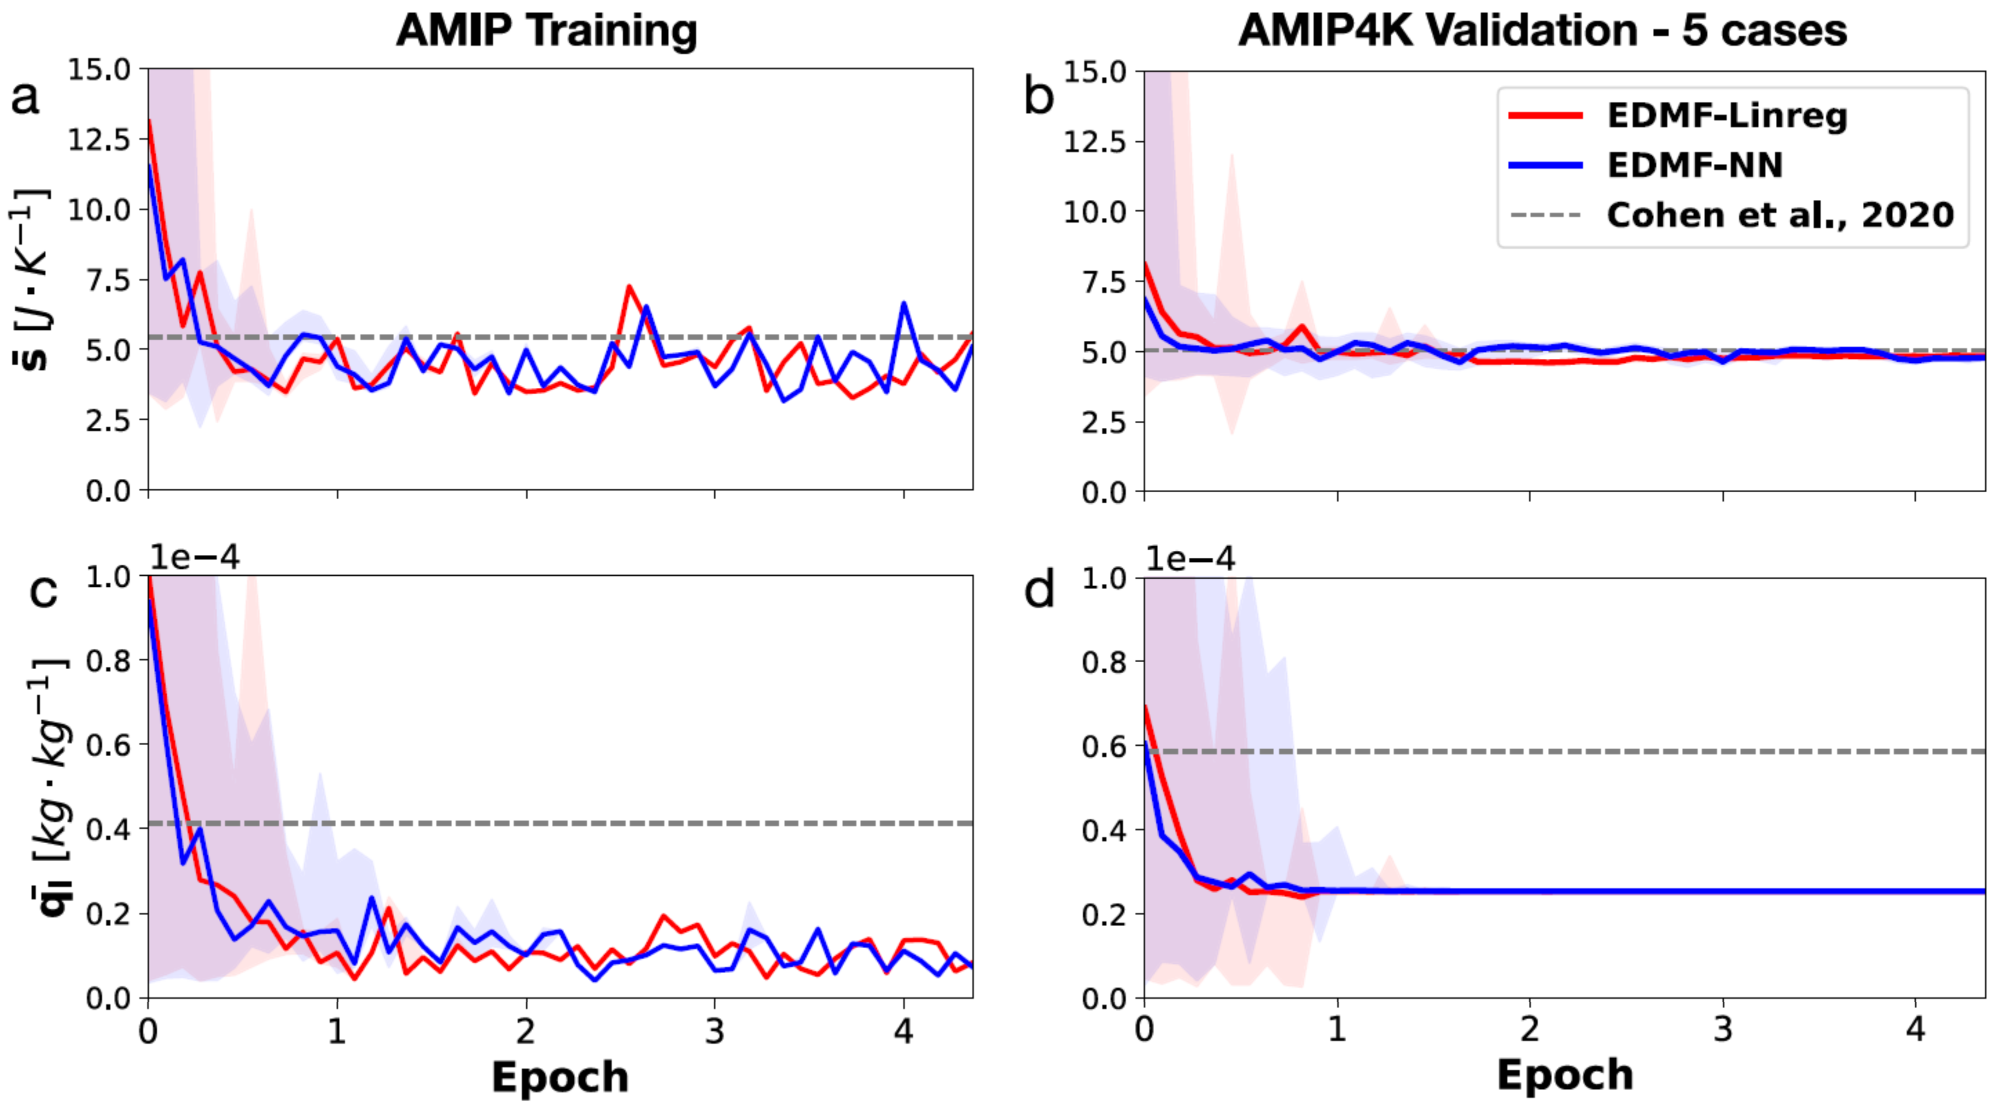

See: Christopoulos et al. (2024)

JGR Atmospheres

Research Article | Open Access

Causally-Informed Deep Learning to Improve Climate Models and Projections

Fernando Iglesias-Suarez, Pierre Gentine, Breixo Solino-Fernandez, Tom Beucler, Michael Pritchard, Jakob Runge, Veronika Eyring

First published: 19 February 2024 | https://doi.org/10.1029/2023JD039202 | Citations: 10

4) Many challenges remain unsolved:

1. Stability, extrapolation behavior, and recalibration of the host model
2. Best way to incorporate causality?

Can Causal Discovery Improve Parameterizations?

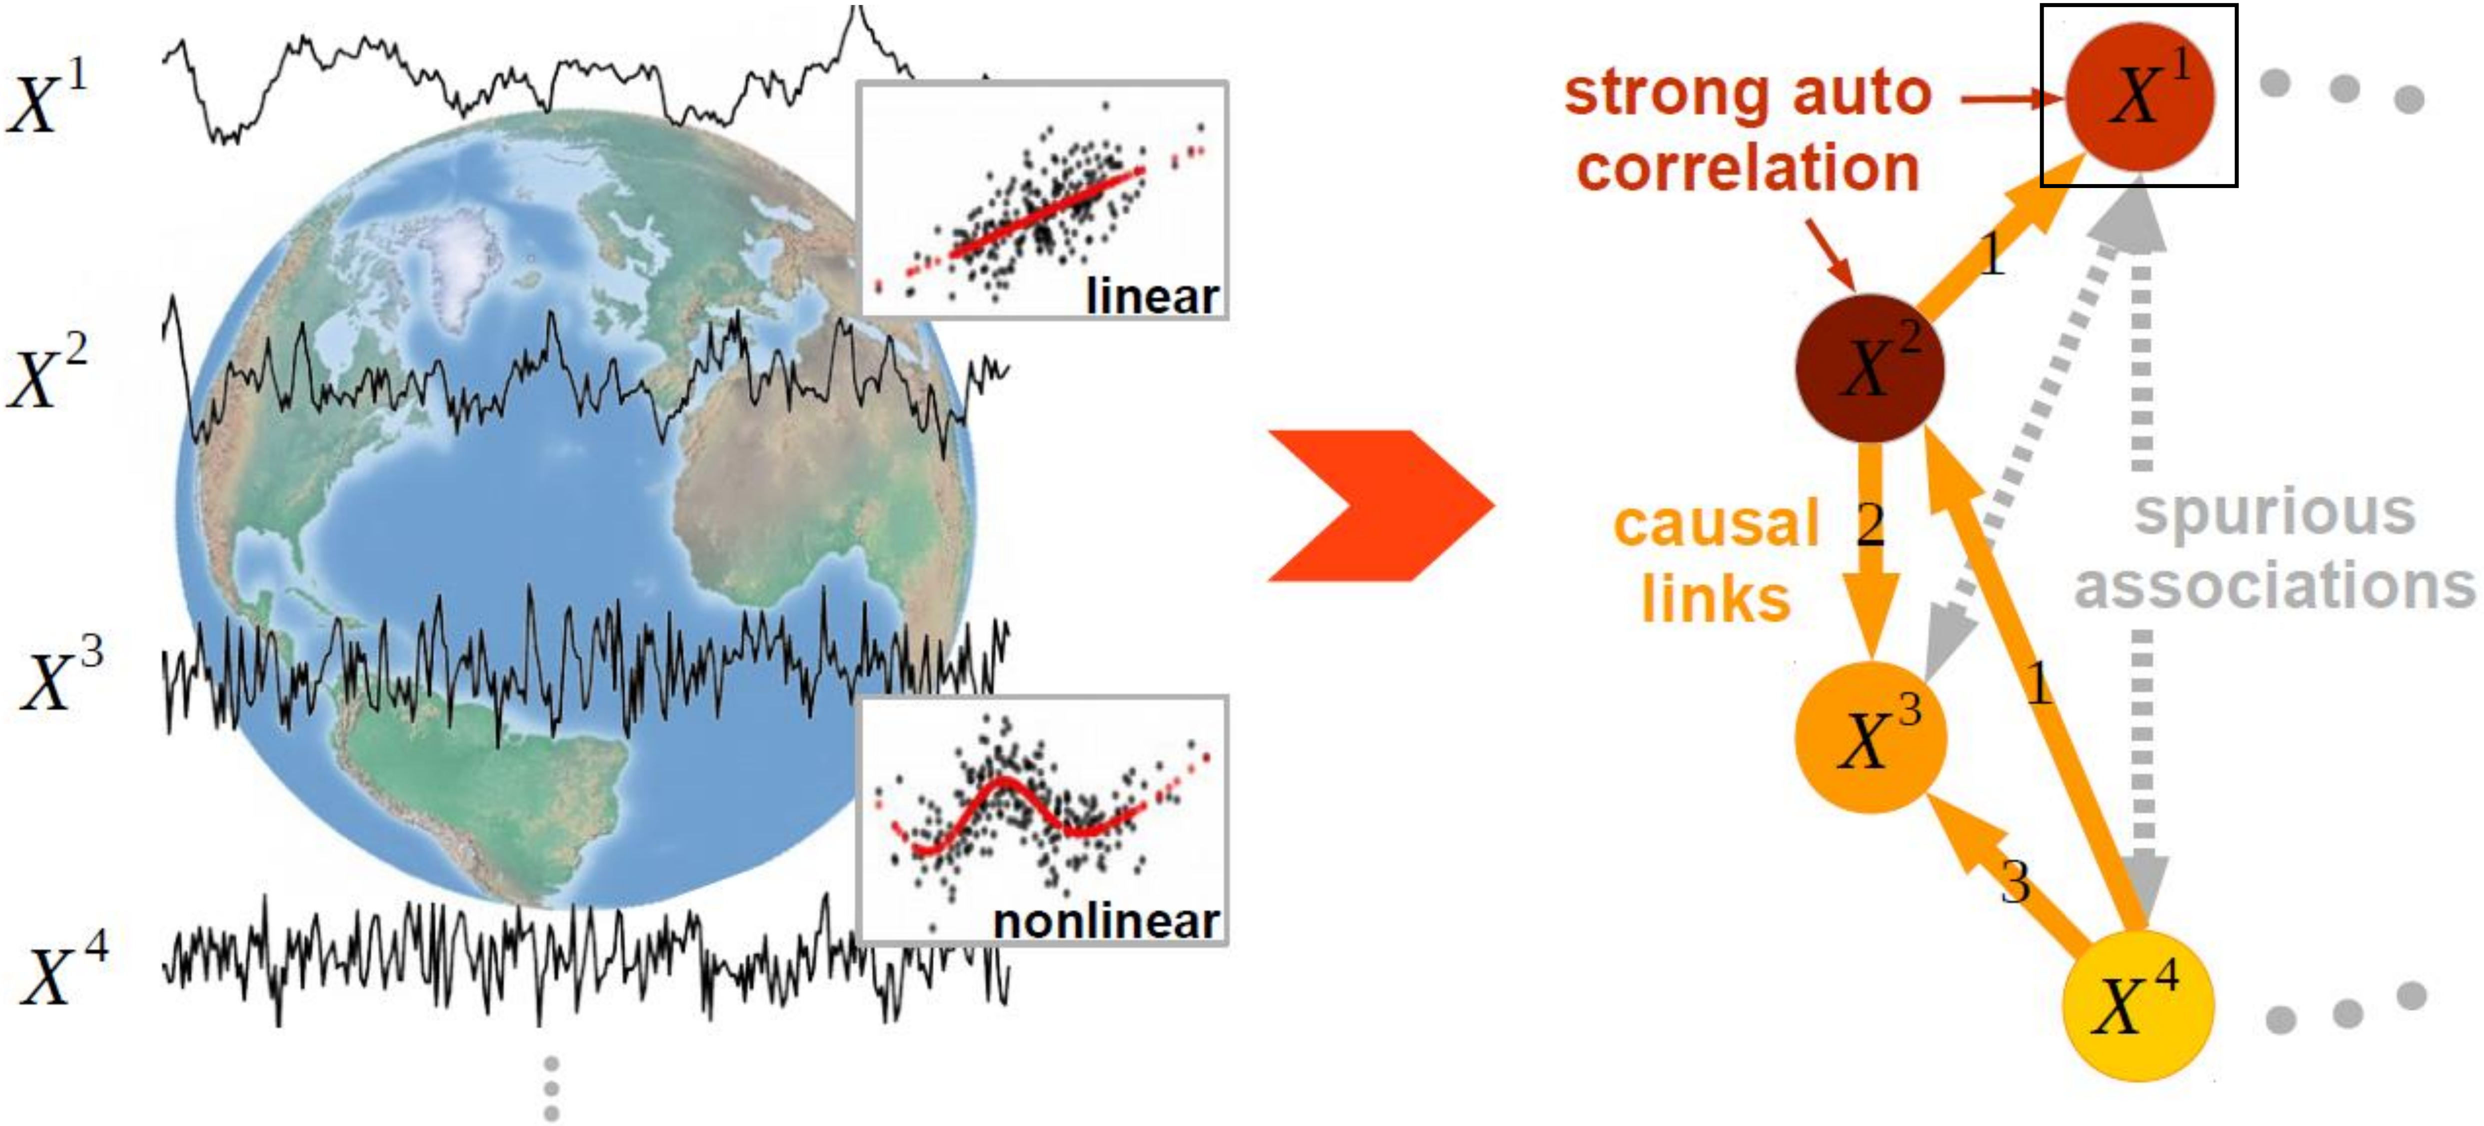

<u>Source</u>: Runge et al. (2019), <u>See</u>: Kretschmer et al. (2016), Runge et al. (2019), Spirtes & Glymour (1991)

Causal feature selection = Eliminating non-causal predictors

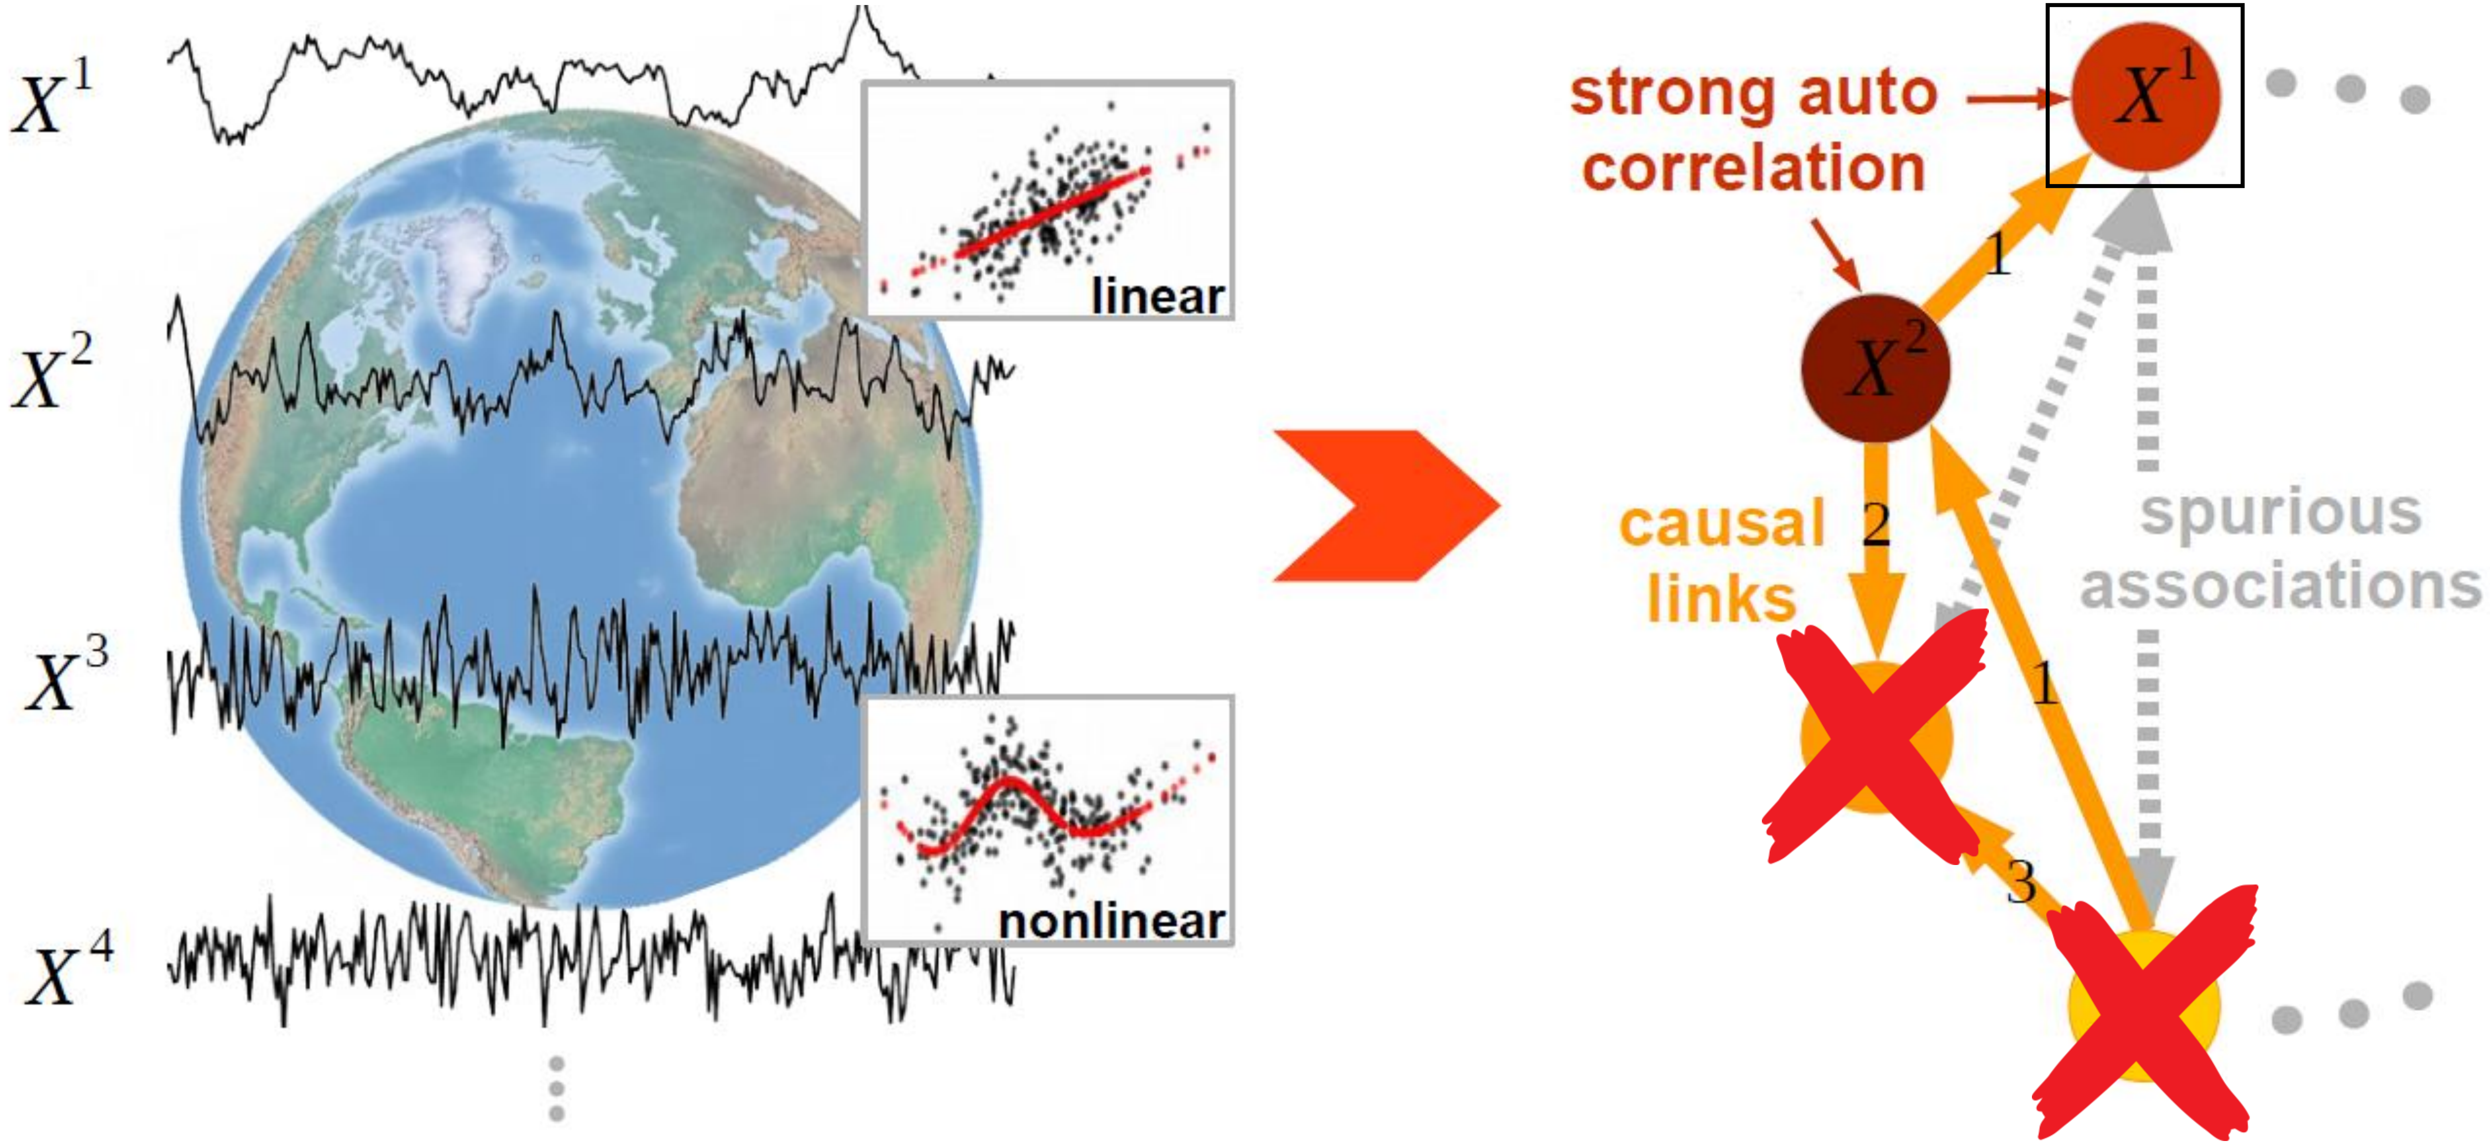

See: Geiger et al. (1990), Pena et al. (2007), Gao and Ji (2017); Image source: Res

Causal feature selection improves the robustness & stability of hybrid climate-AI simulations

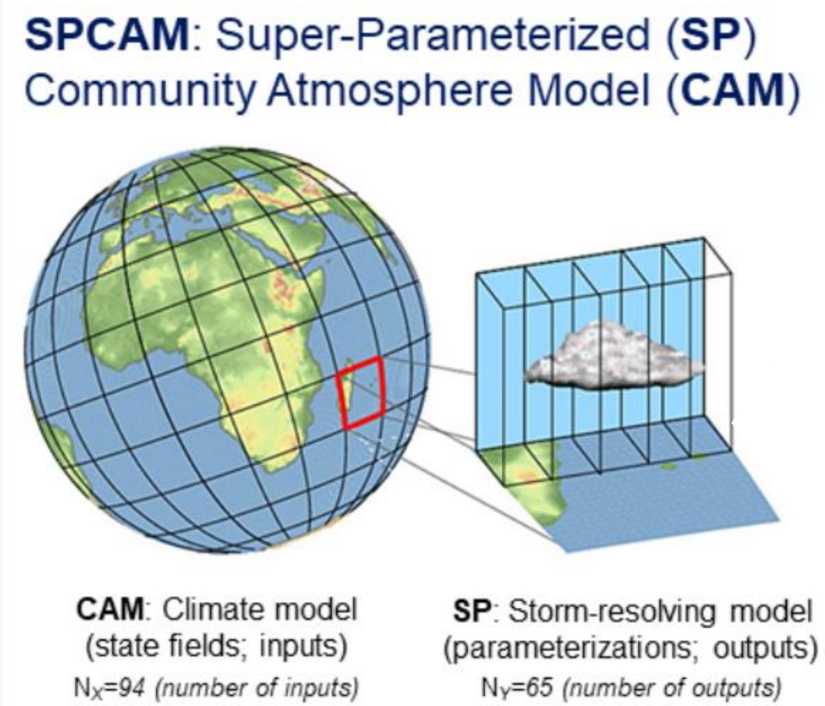

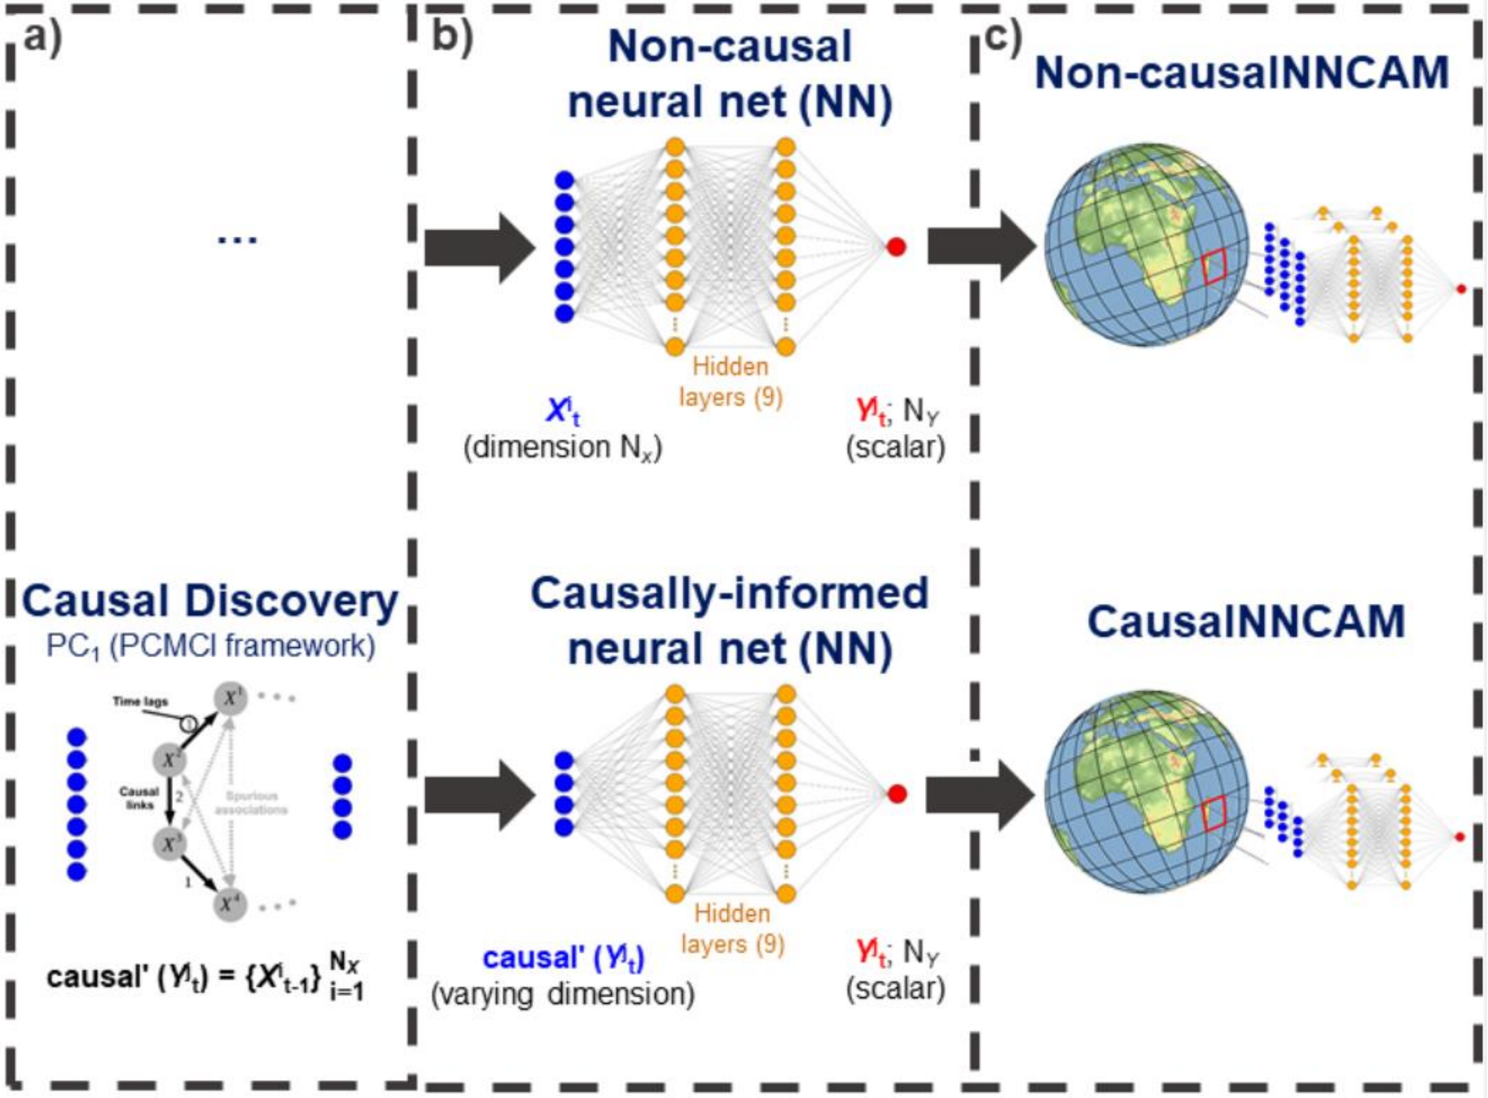

See: Iglesias-Suarez et al. (2024)

4) Many challenges remain unsolved:

1. Stability, extrapolation behavior, and recalibration of the host model
2. Best way to incorporate causality?
3. Best way to incorporate and evaluate stochasticity?

When given an input, **generative** models predict **distributions** of outputs rather than a single output vector

A **generative model** is a statistical model of the joint probability distribution $P(X, Y)$ on given observable variable X and target variable Y;[1]

Uncertainty Quantification

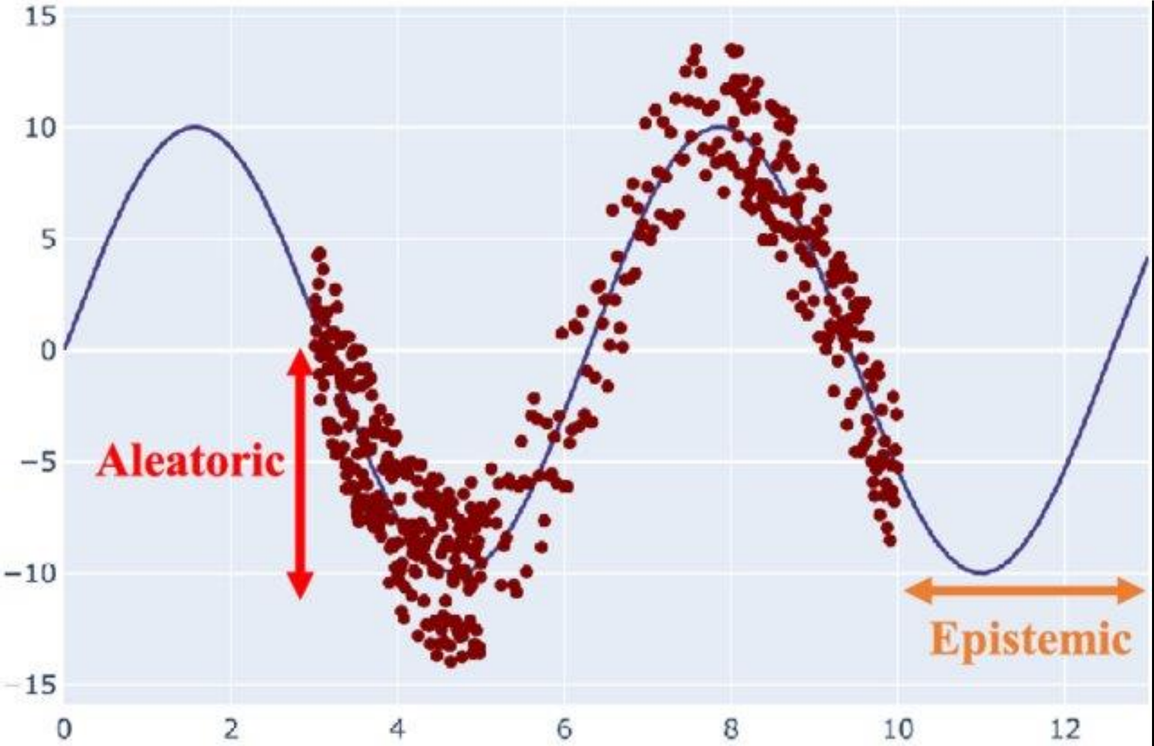

Abdar et al. (2021)

Stochastic Modeling

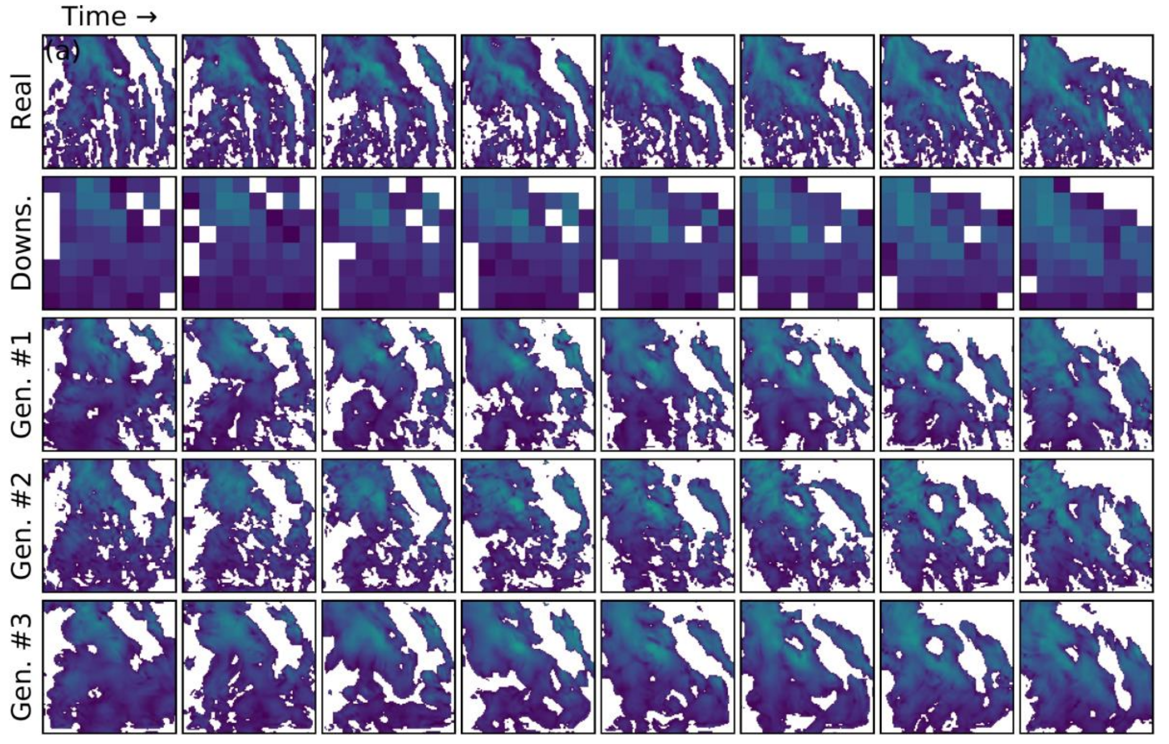

In practice, there are many ways of adding UQ!

Ensemble Prediction (EP)

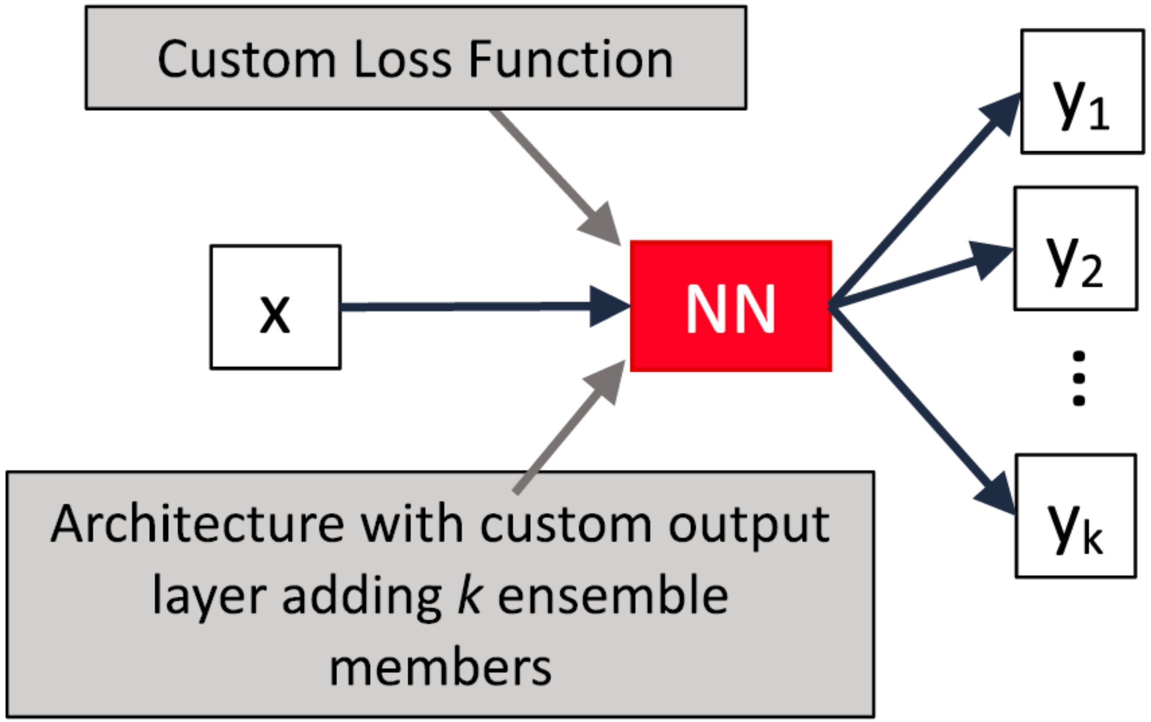

Multi-Model (MM)

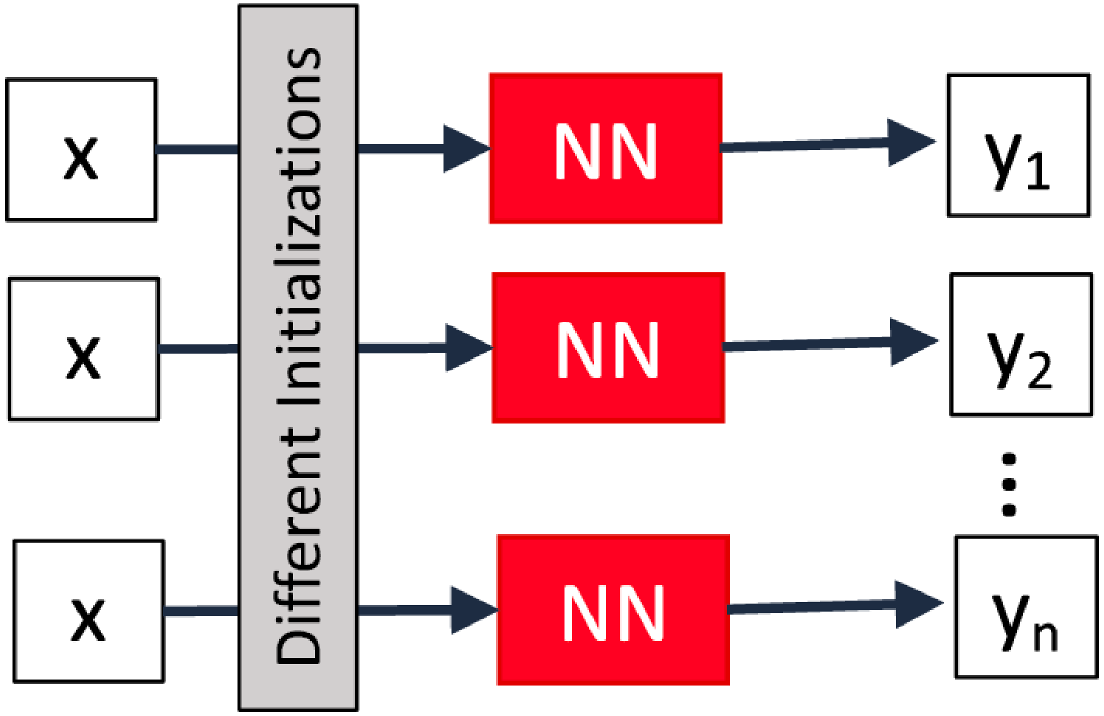

Haynes et al. (2023)

In practice, there are many ways of adding UQ!

Parametric Distributional Prediction (PDP)

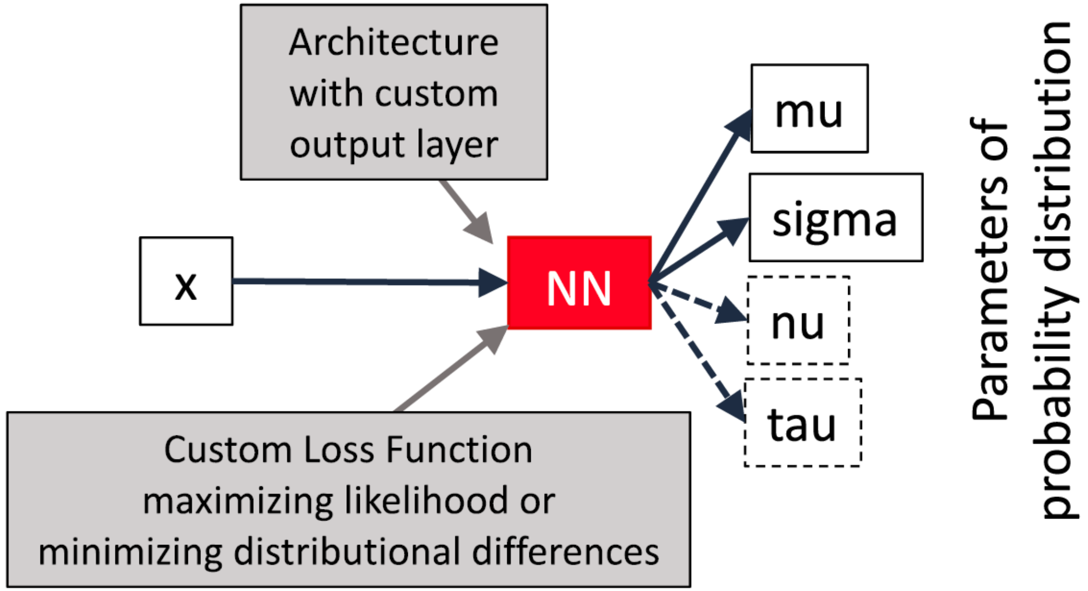

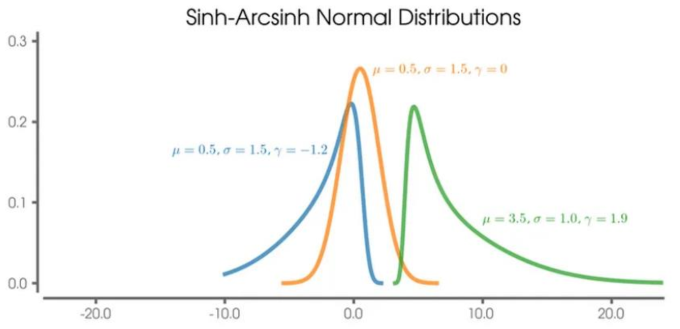

Non-Parametric Distributional Prediction (NPDP)

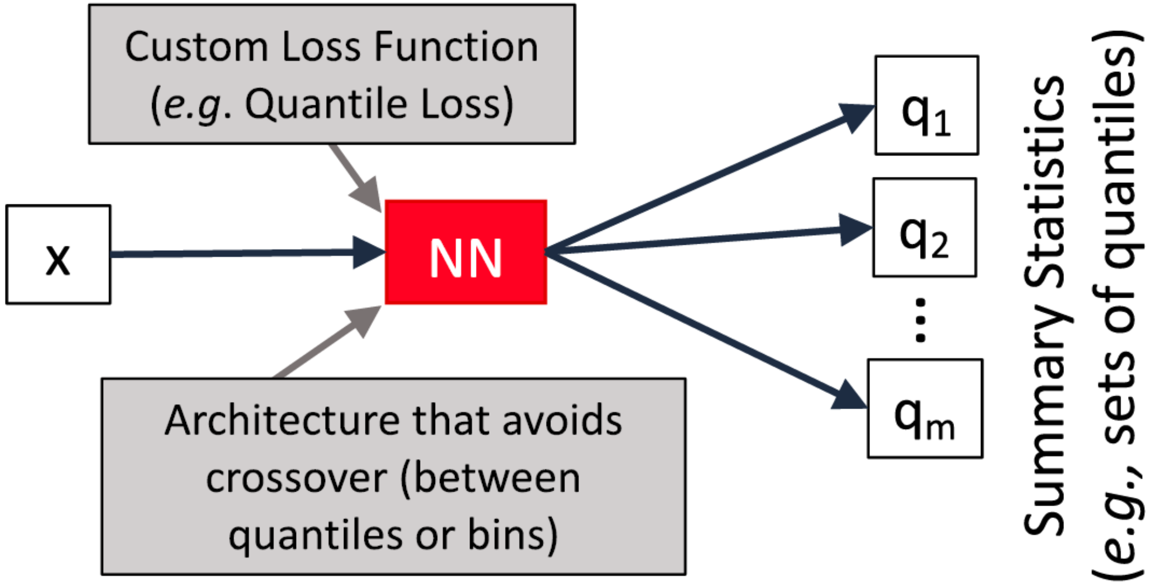

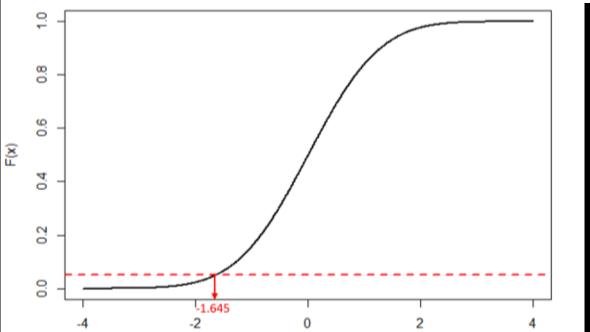

Haynes et al. (2023). Kujawska et al. (2021)

MC-Dropout randomly drops neural connections (Training) Regularization & (Test time) UQ

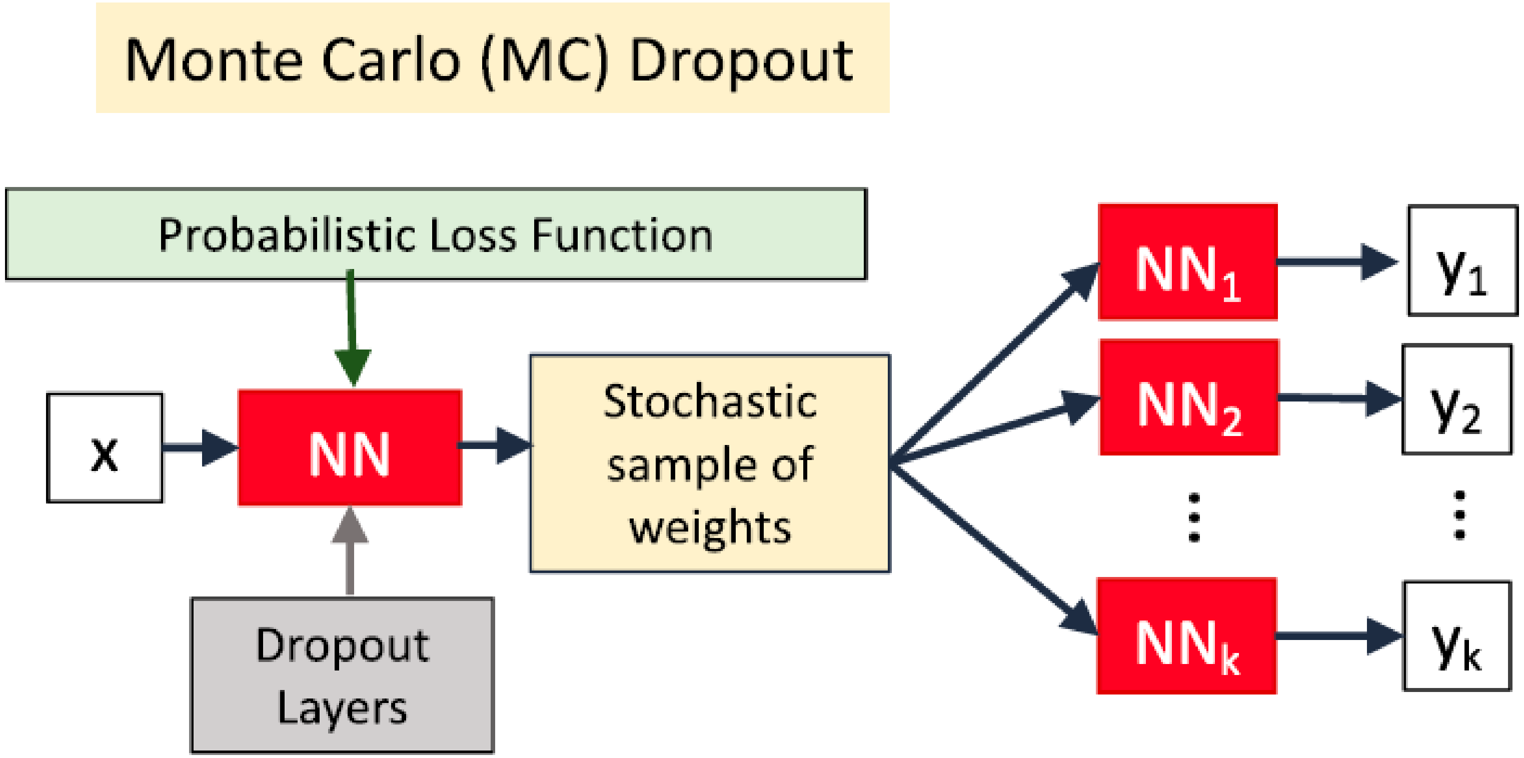

Ensemble of deterministic predictions [probabilistic predictions]

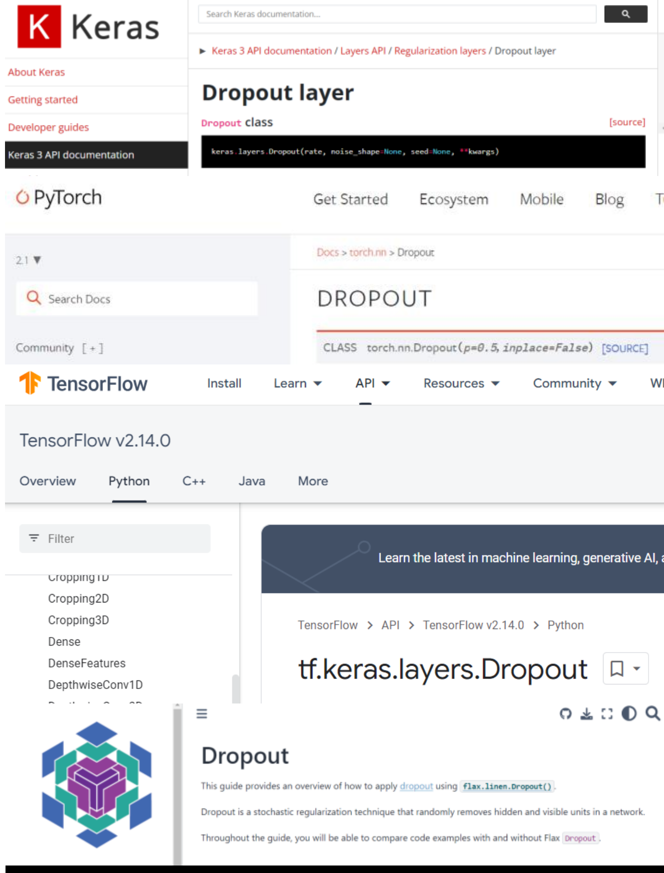

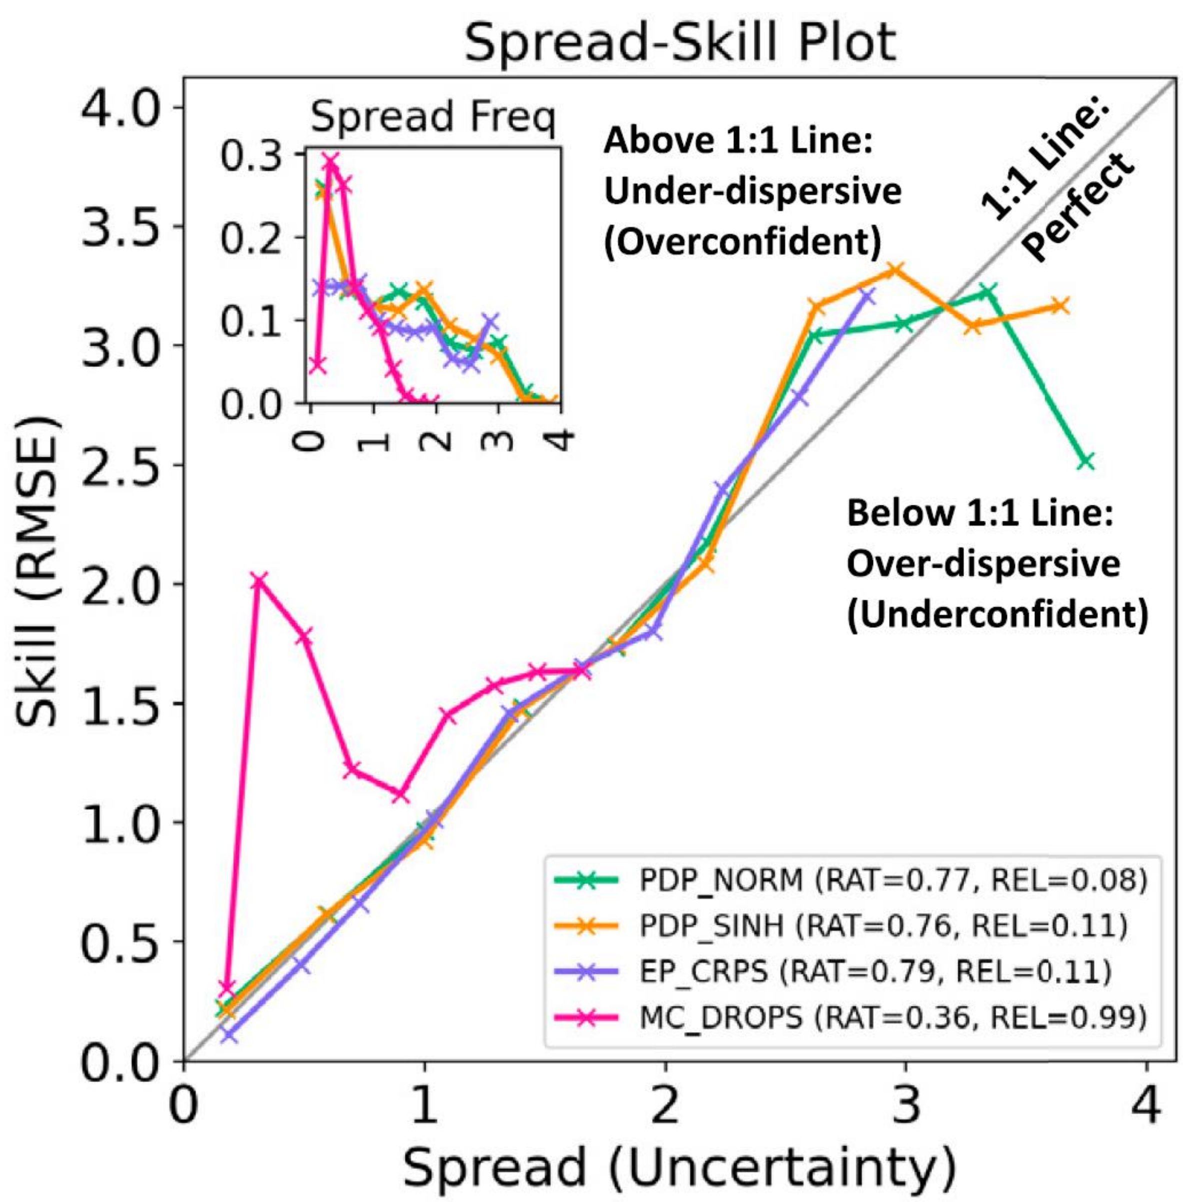

Spread-Skill Plot and CRPS* evaluate the generated distribution
*(Continuous Ranked Prob. Score)

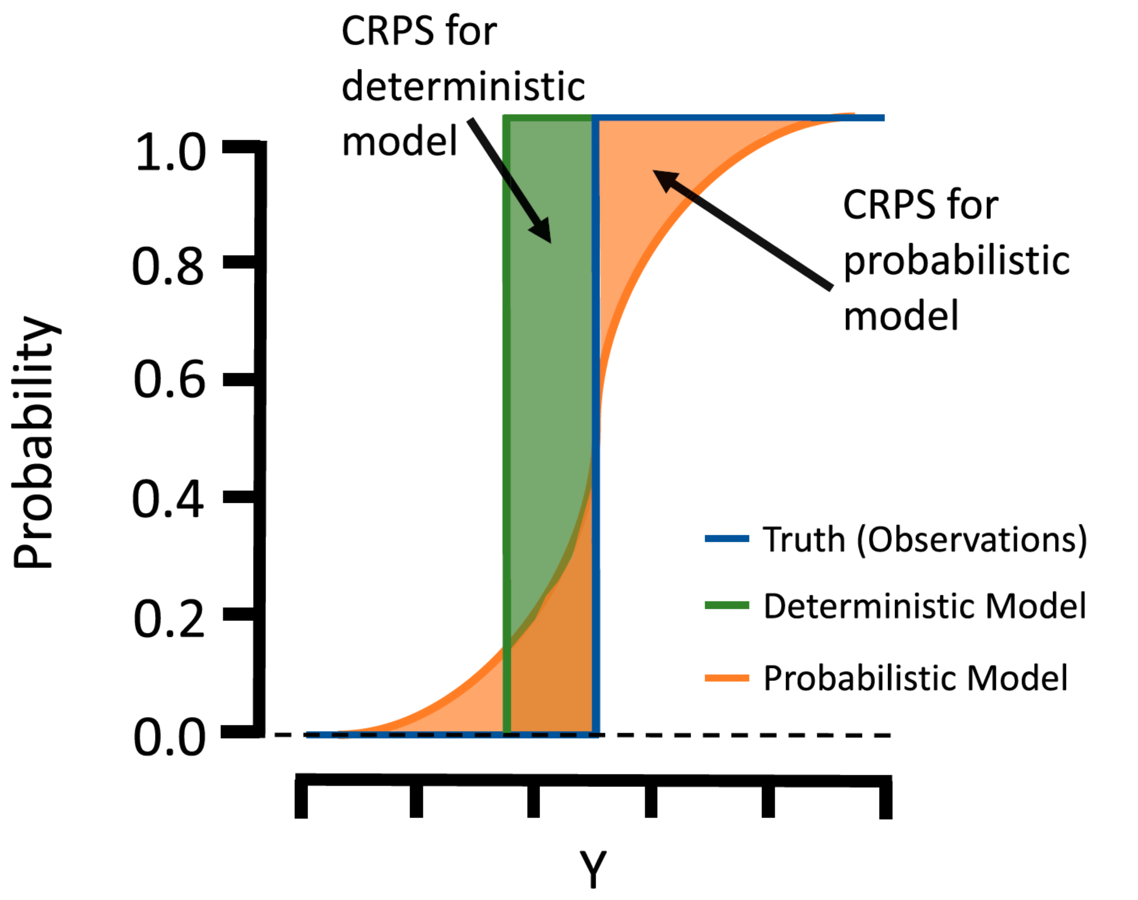

1) Generative Adversarial Nets pit a generator (fed noise) against a discriminator (fed the fake or real images)

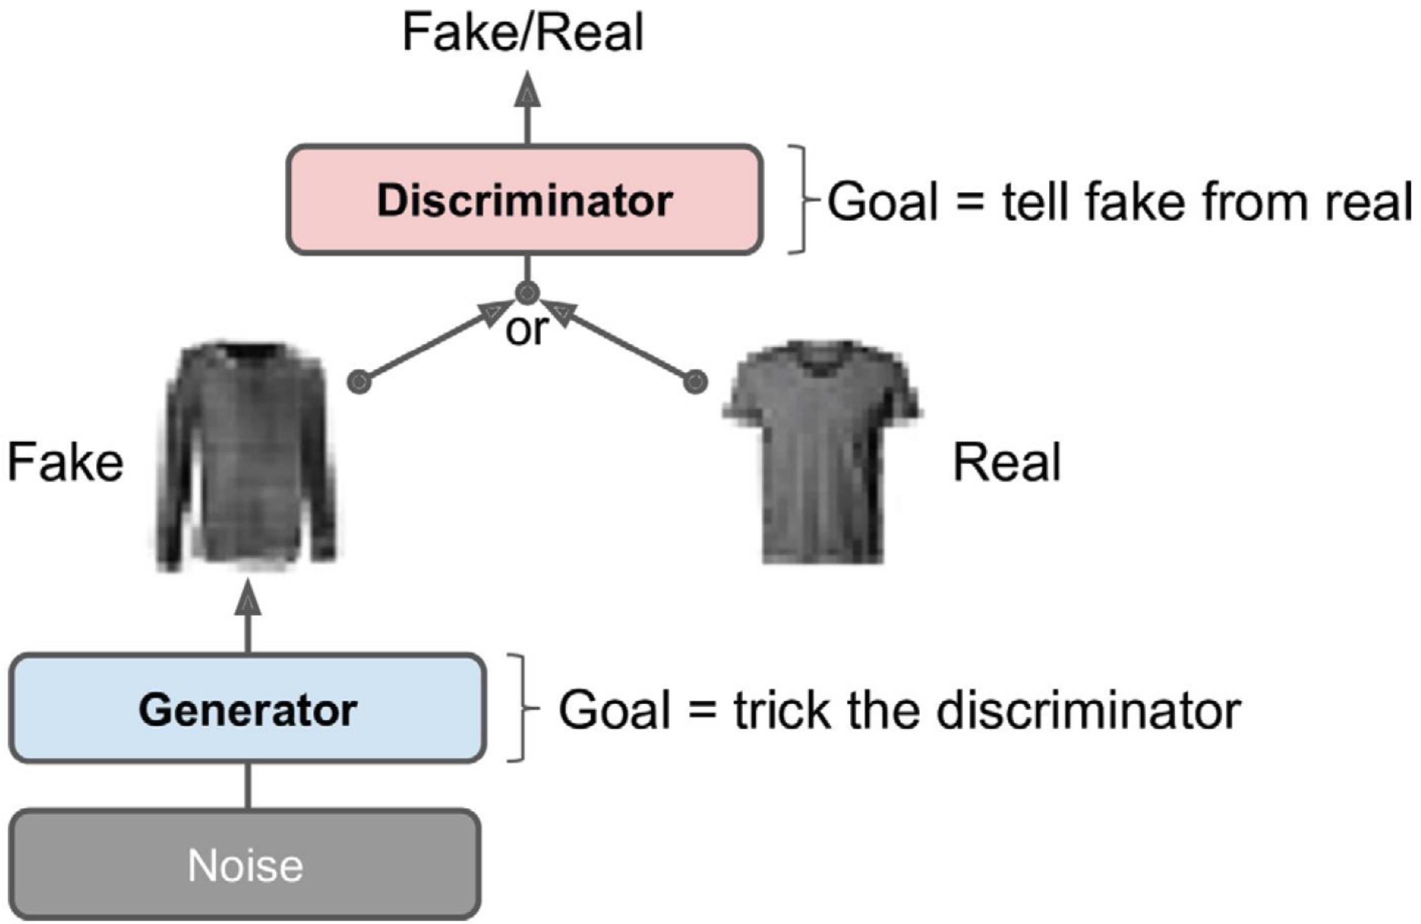

Géron textbook, Wikipedia (GAN), Bodla et al. (2018)

2) Diffusion Probabilistic Models smoothly perturb data by adding noise, then reverse this process to generate new data from noise.

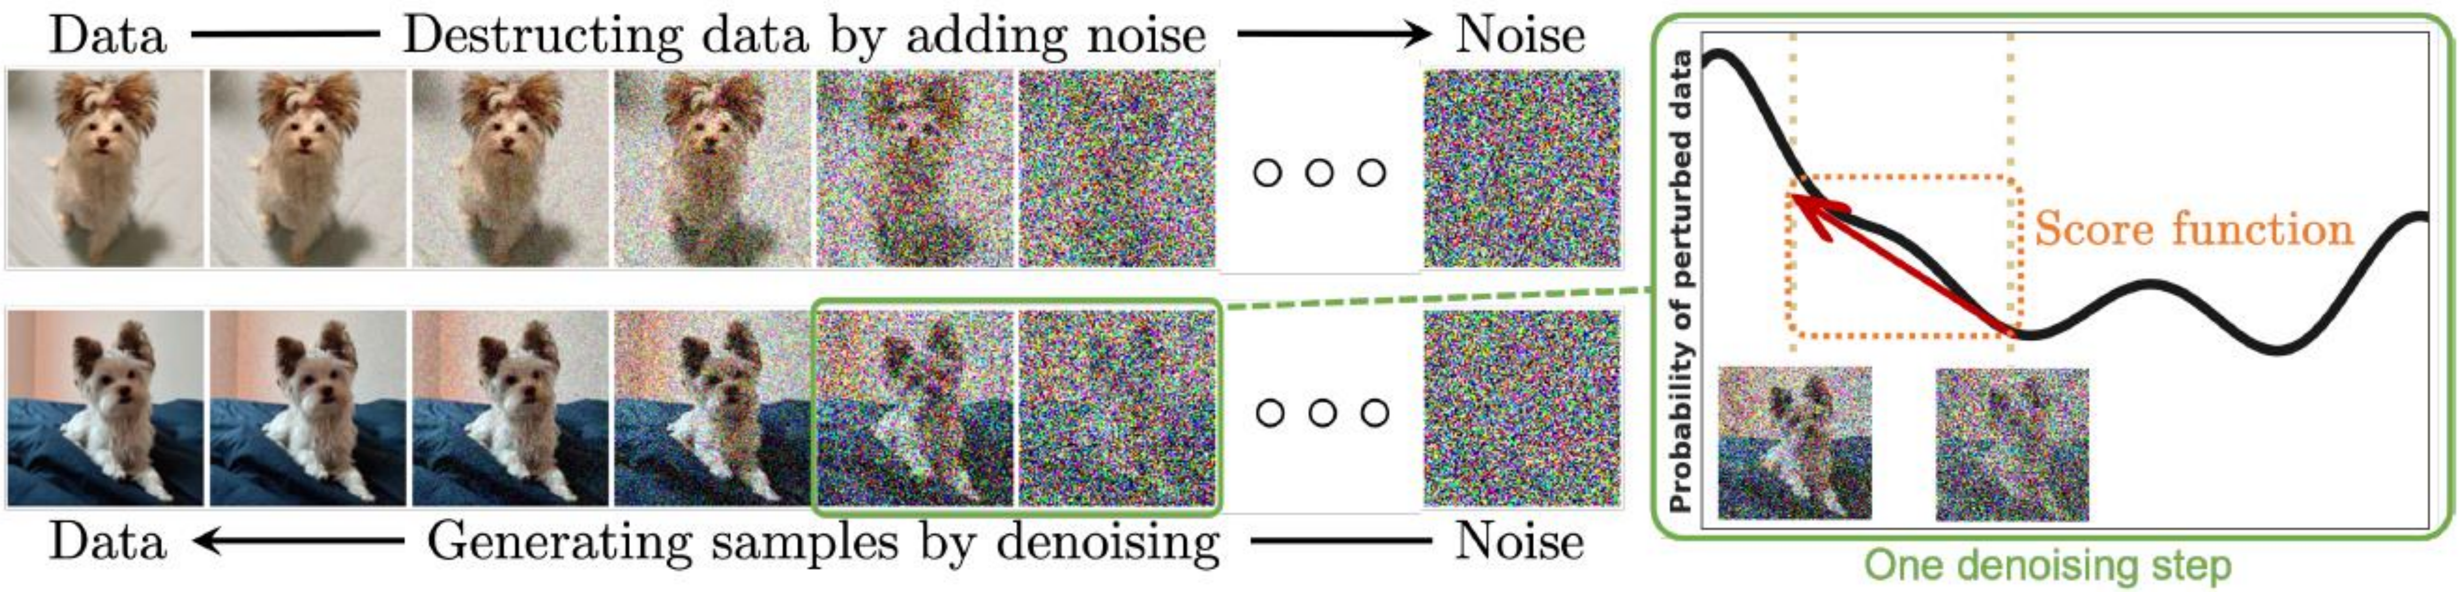

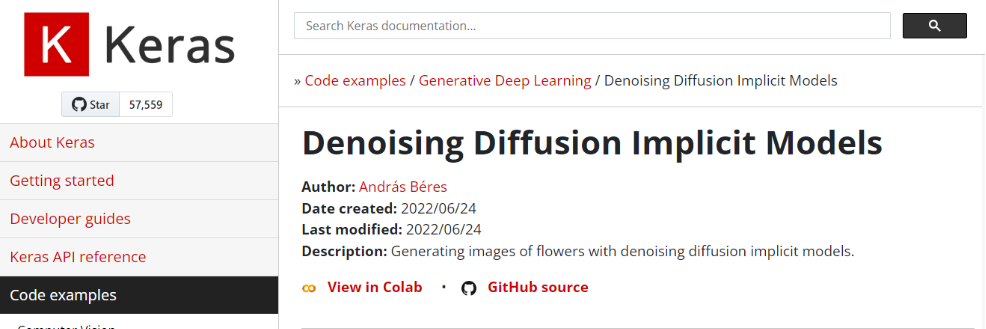

Diffusers

license Apache-2.0 release v0.14.0 Contributor Covenant 2.0

Yang et al. (2022)

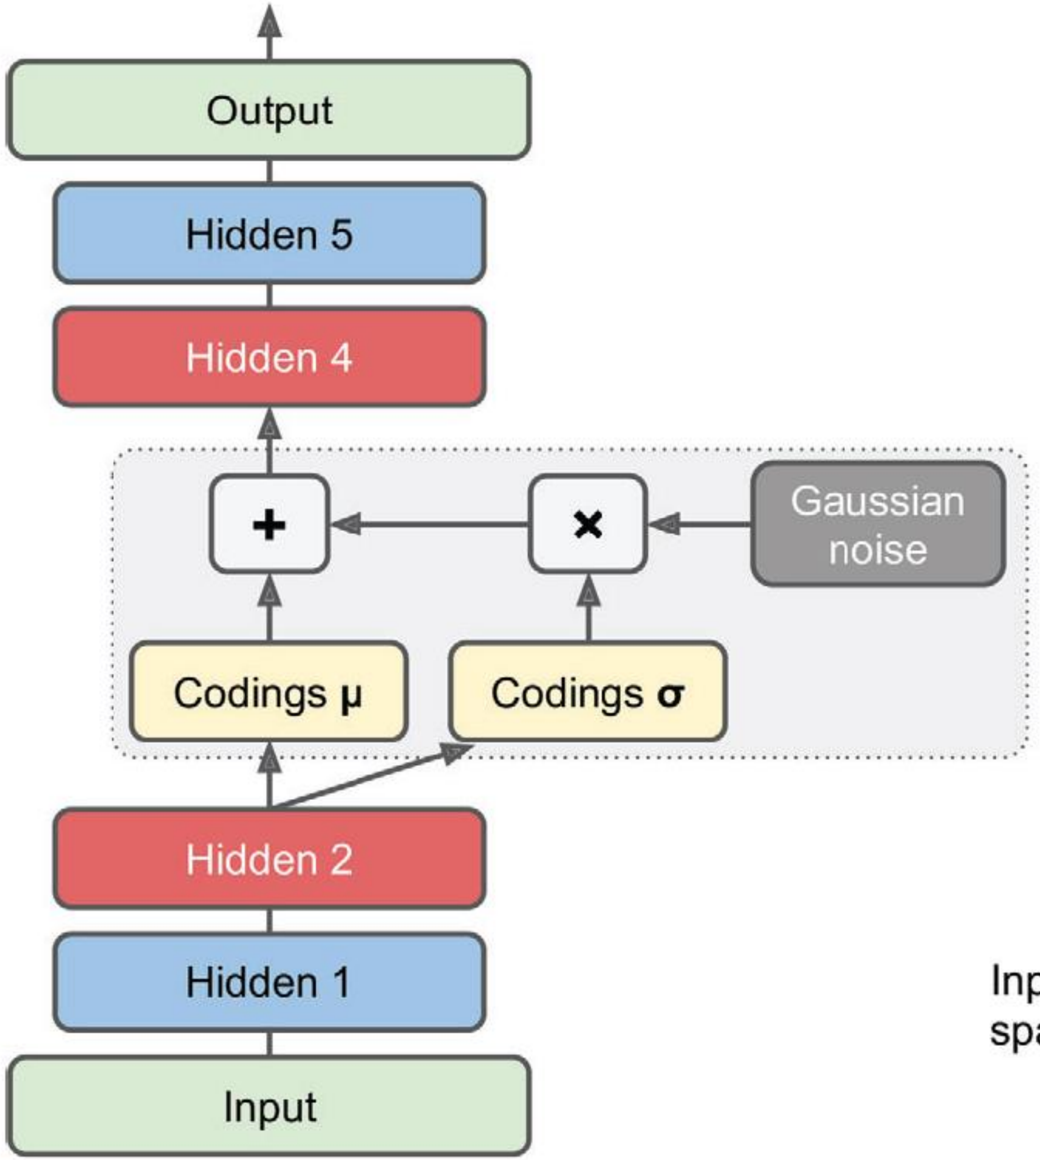

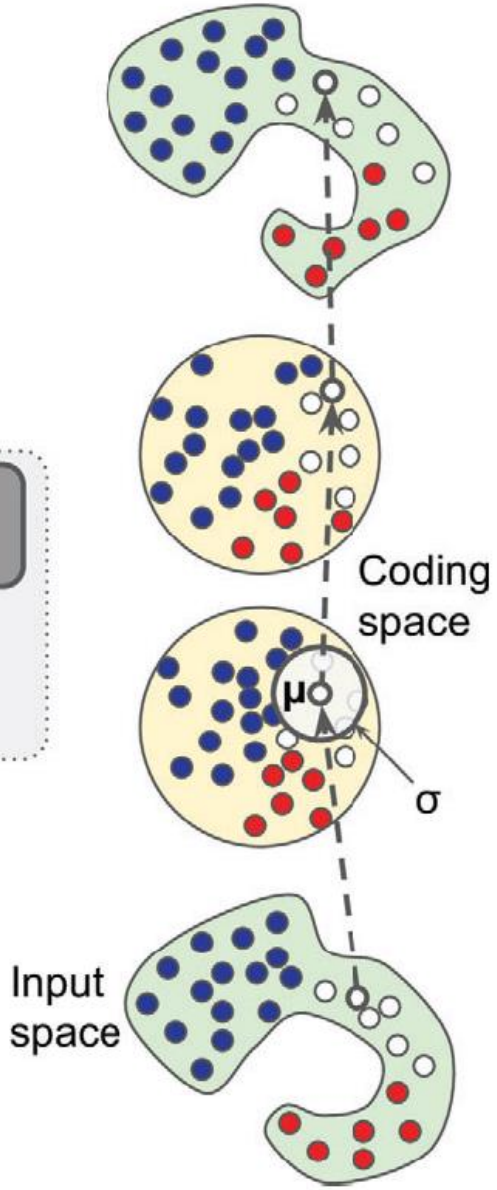

3) Variational Autoencoders probabilistically encode/decode data from *latent representations*

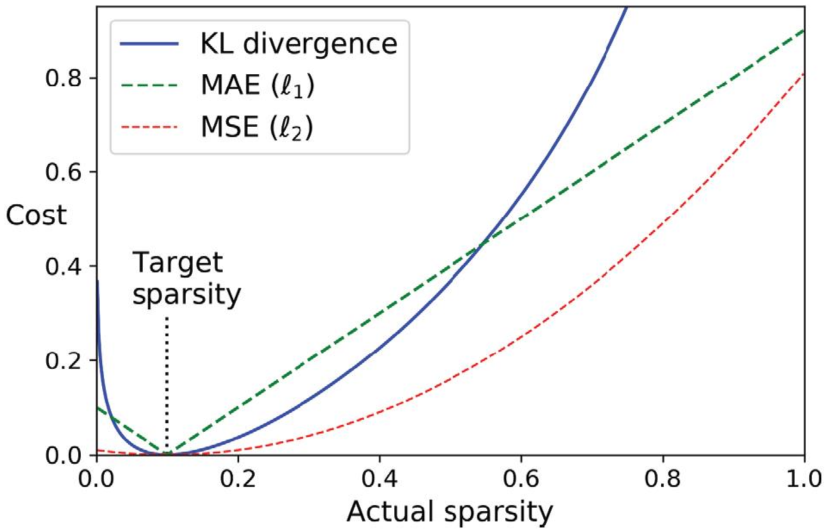

Latent manifold or "Latent space"

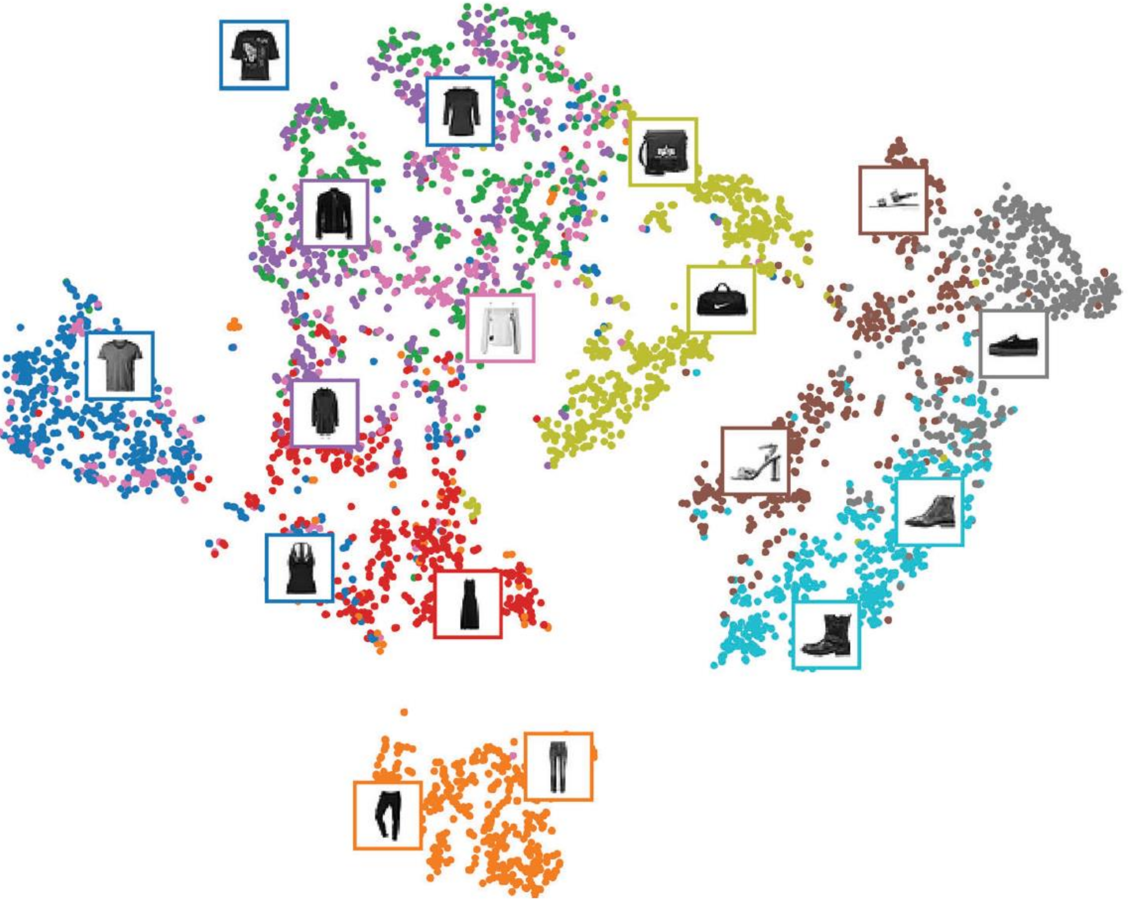

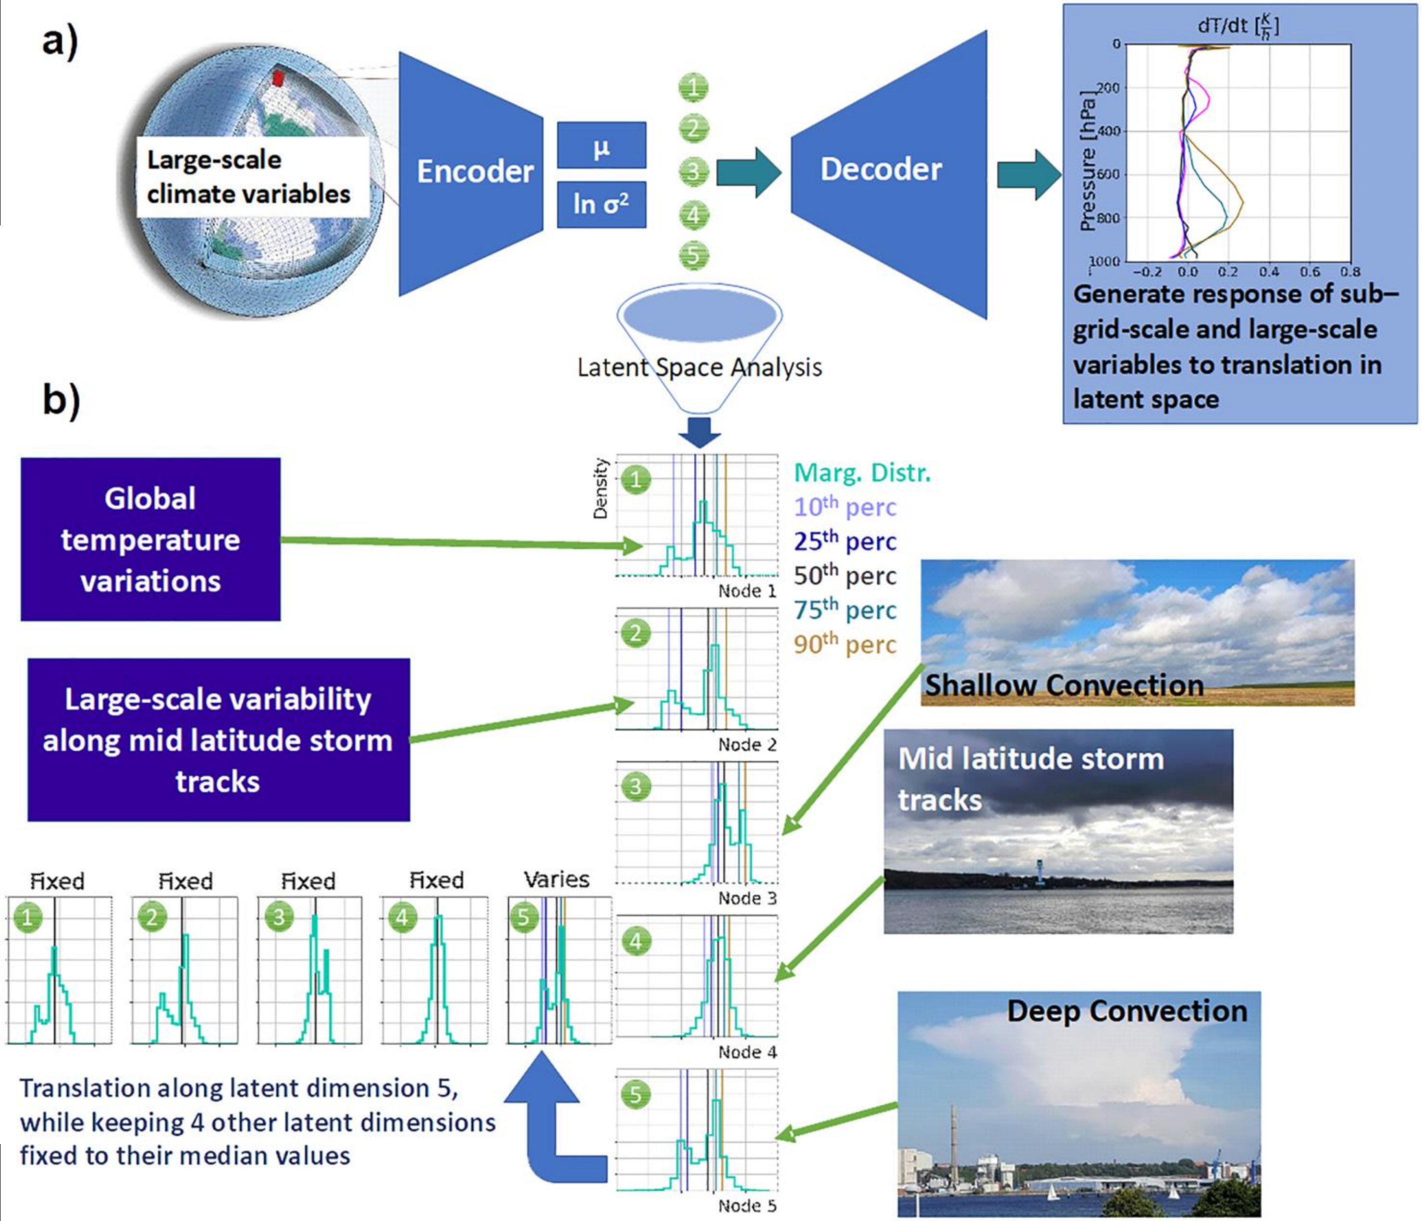

Géron textbook, Behrens et al. (2022)

Latent manifolds can be used for e.g., scientific analysis and semi-supervised learning

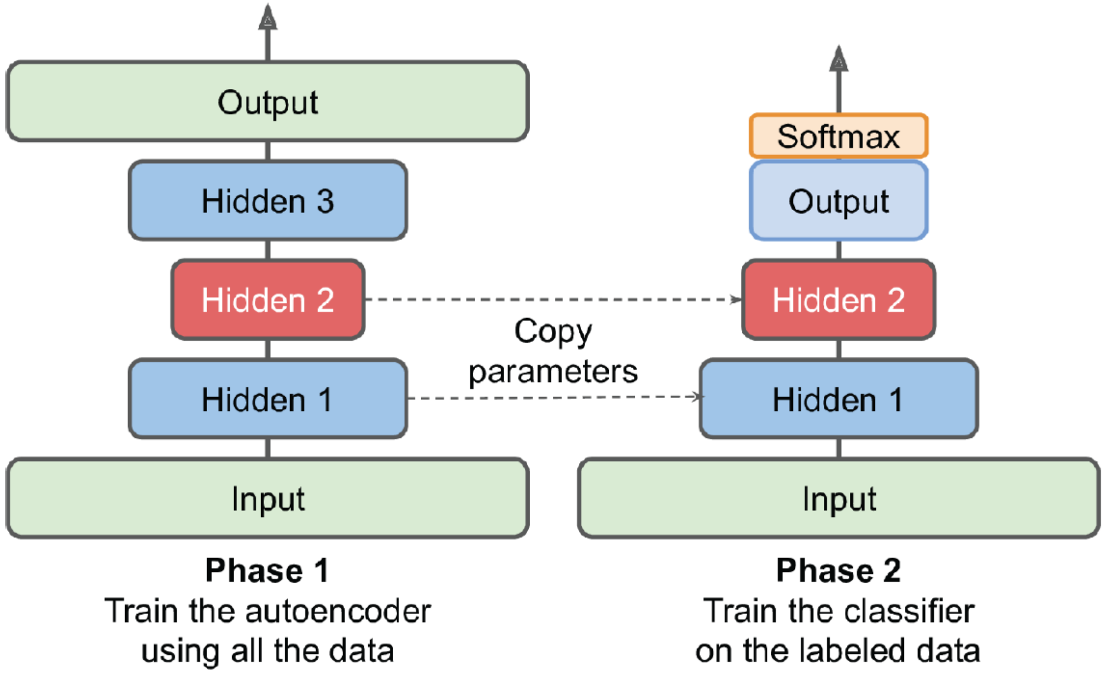

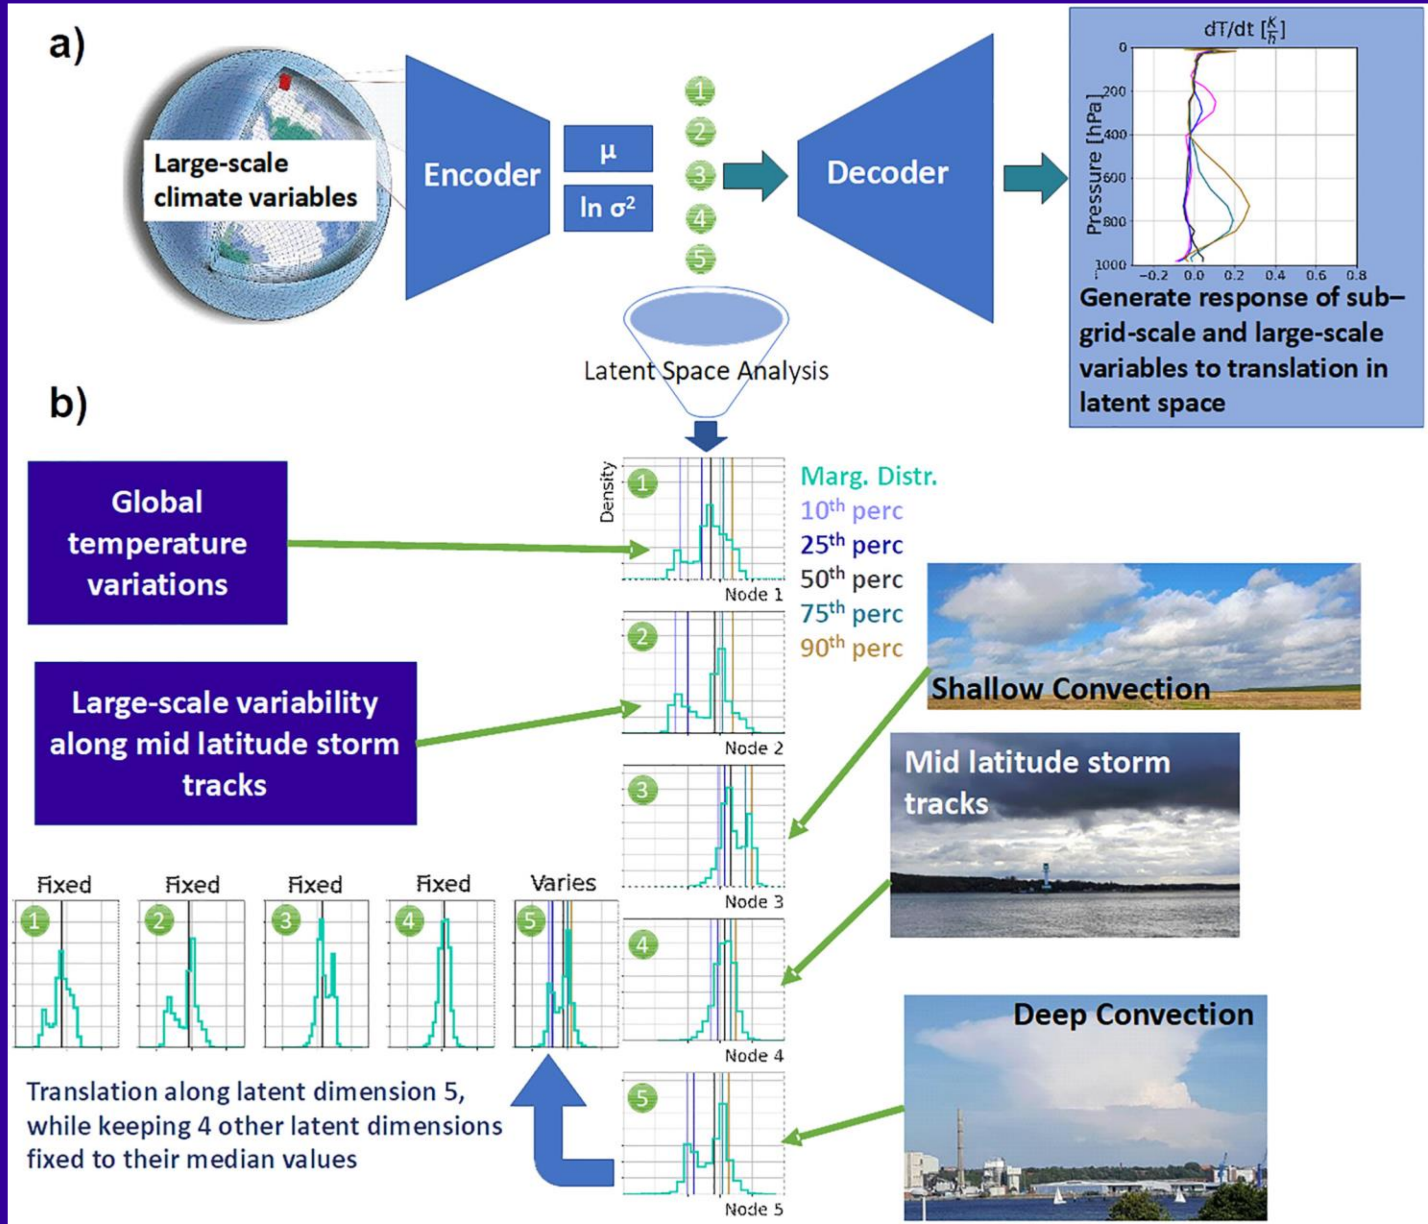

Journal of Advances in Modeling Earth Systems / Volume 17, Issue 4 / e2024MS004272

Research Article | Open Access

Simulating Atmospheric Processes in Earth System Models and Quantifying Uncertainties With Deep Learning Multi-Member and Stochastic Parameterizations

Gunnar Behrens, Tom Beucler, Fernando Iglesias-Suarez, Sungduk Yu, Pierre Gentine, Michael Pritchard, Mierk Schwabe, Veronika Eyring

First published: 13 April 2025
https://doi.org/10.1029/2024MS004272

4) Many challenges remain unsolved:

1. Stability, extrapolation behavior, and recalibration of the host model
2. Best way to incorporate causality?
3. Best way to incorporate and evaluate stochasticity?

ICTP25: Hybrid AI-Climate Modeling

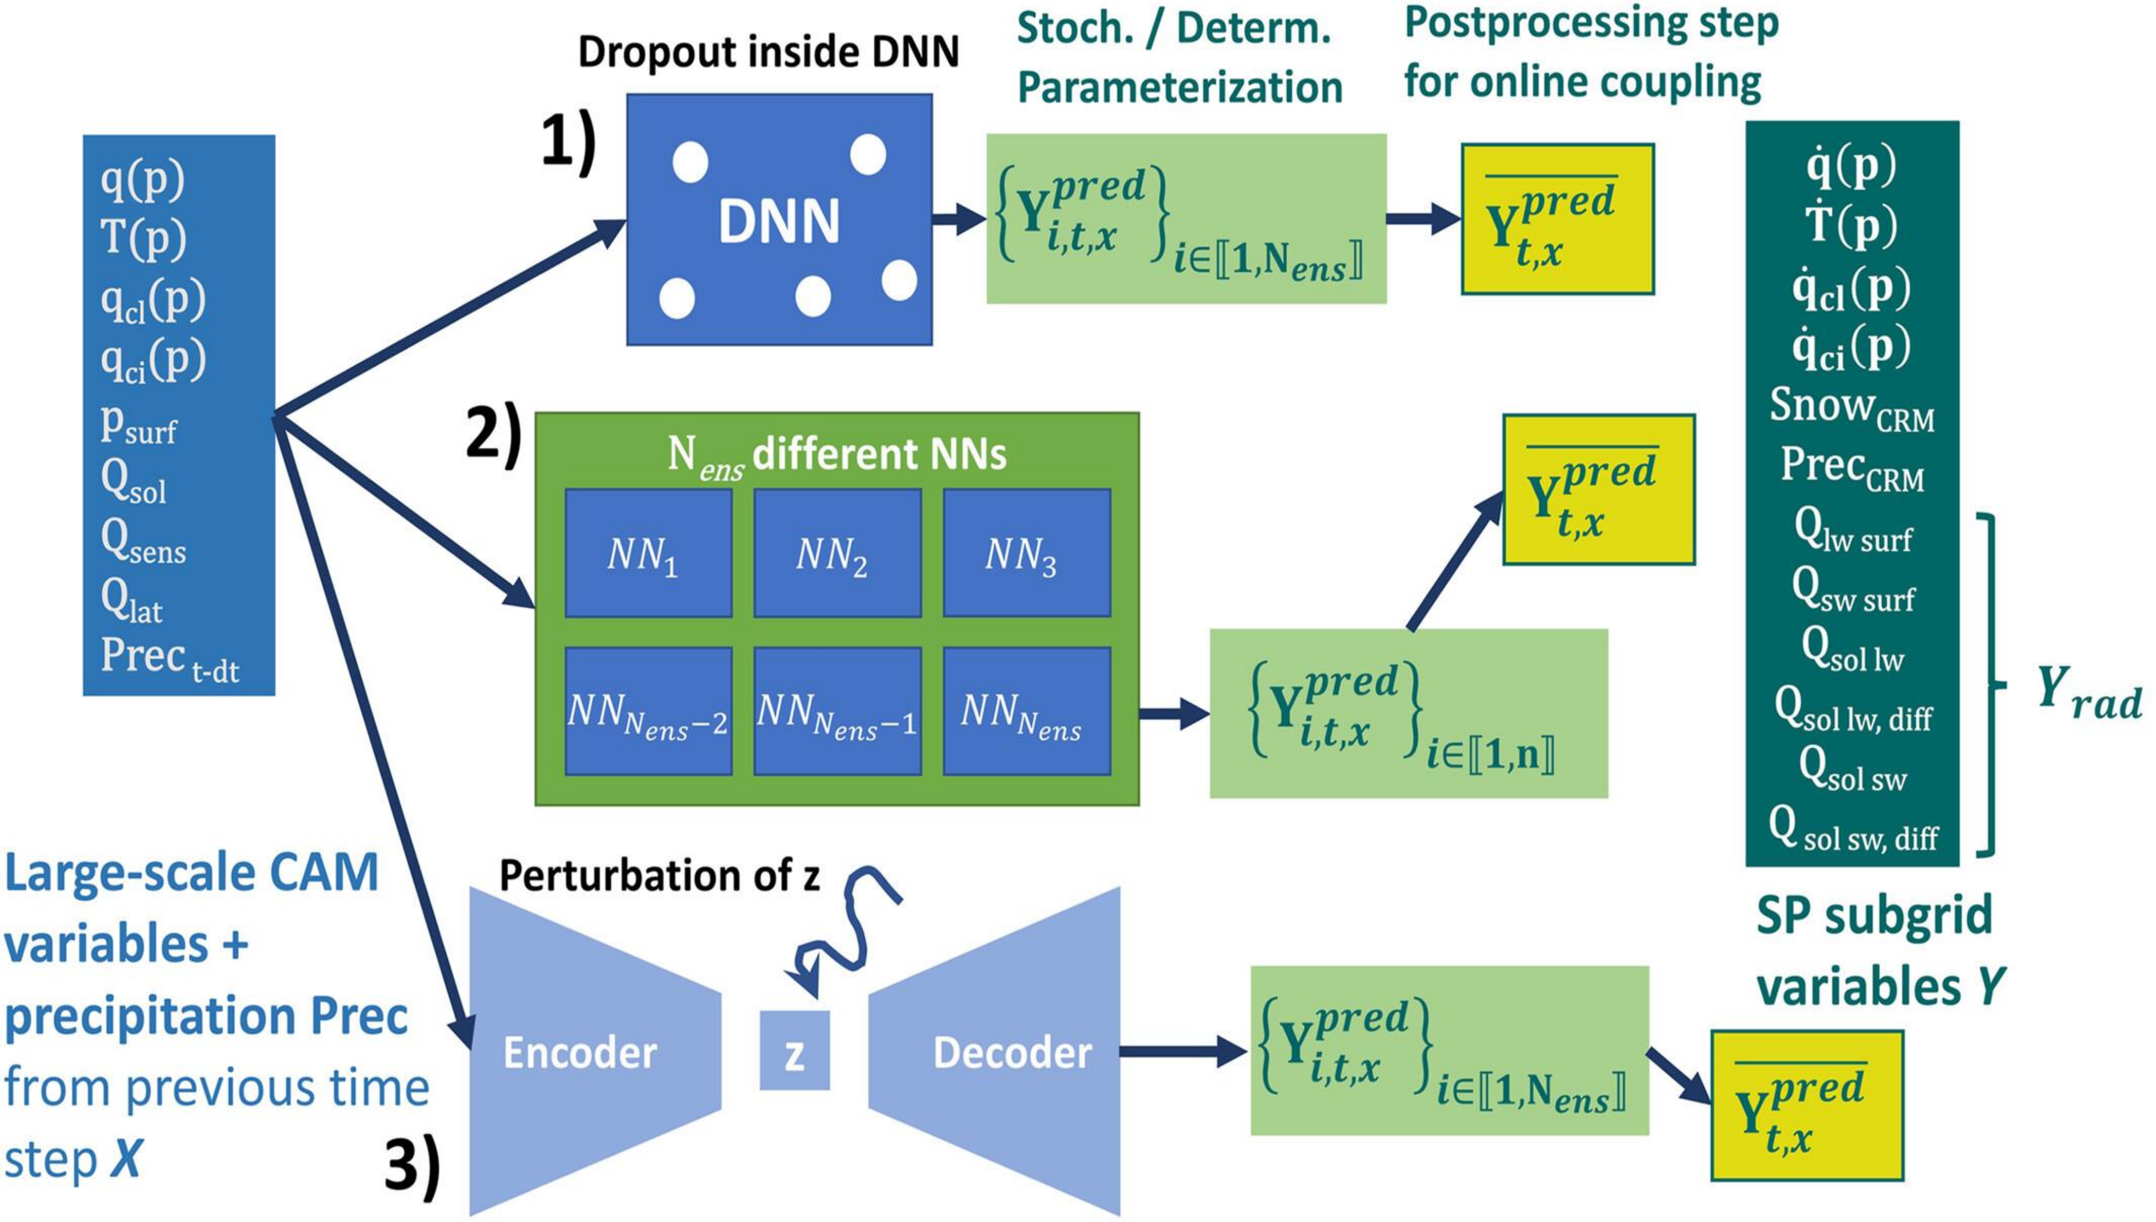

Figure: Behrens et al. (2025)

ANN ensembles and latent space perturbations lead to well-calibrated uncertainty **offline**

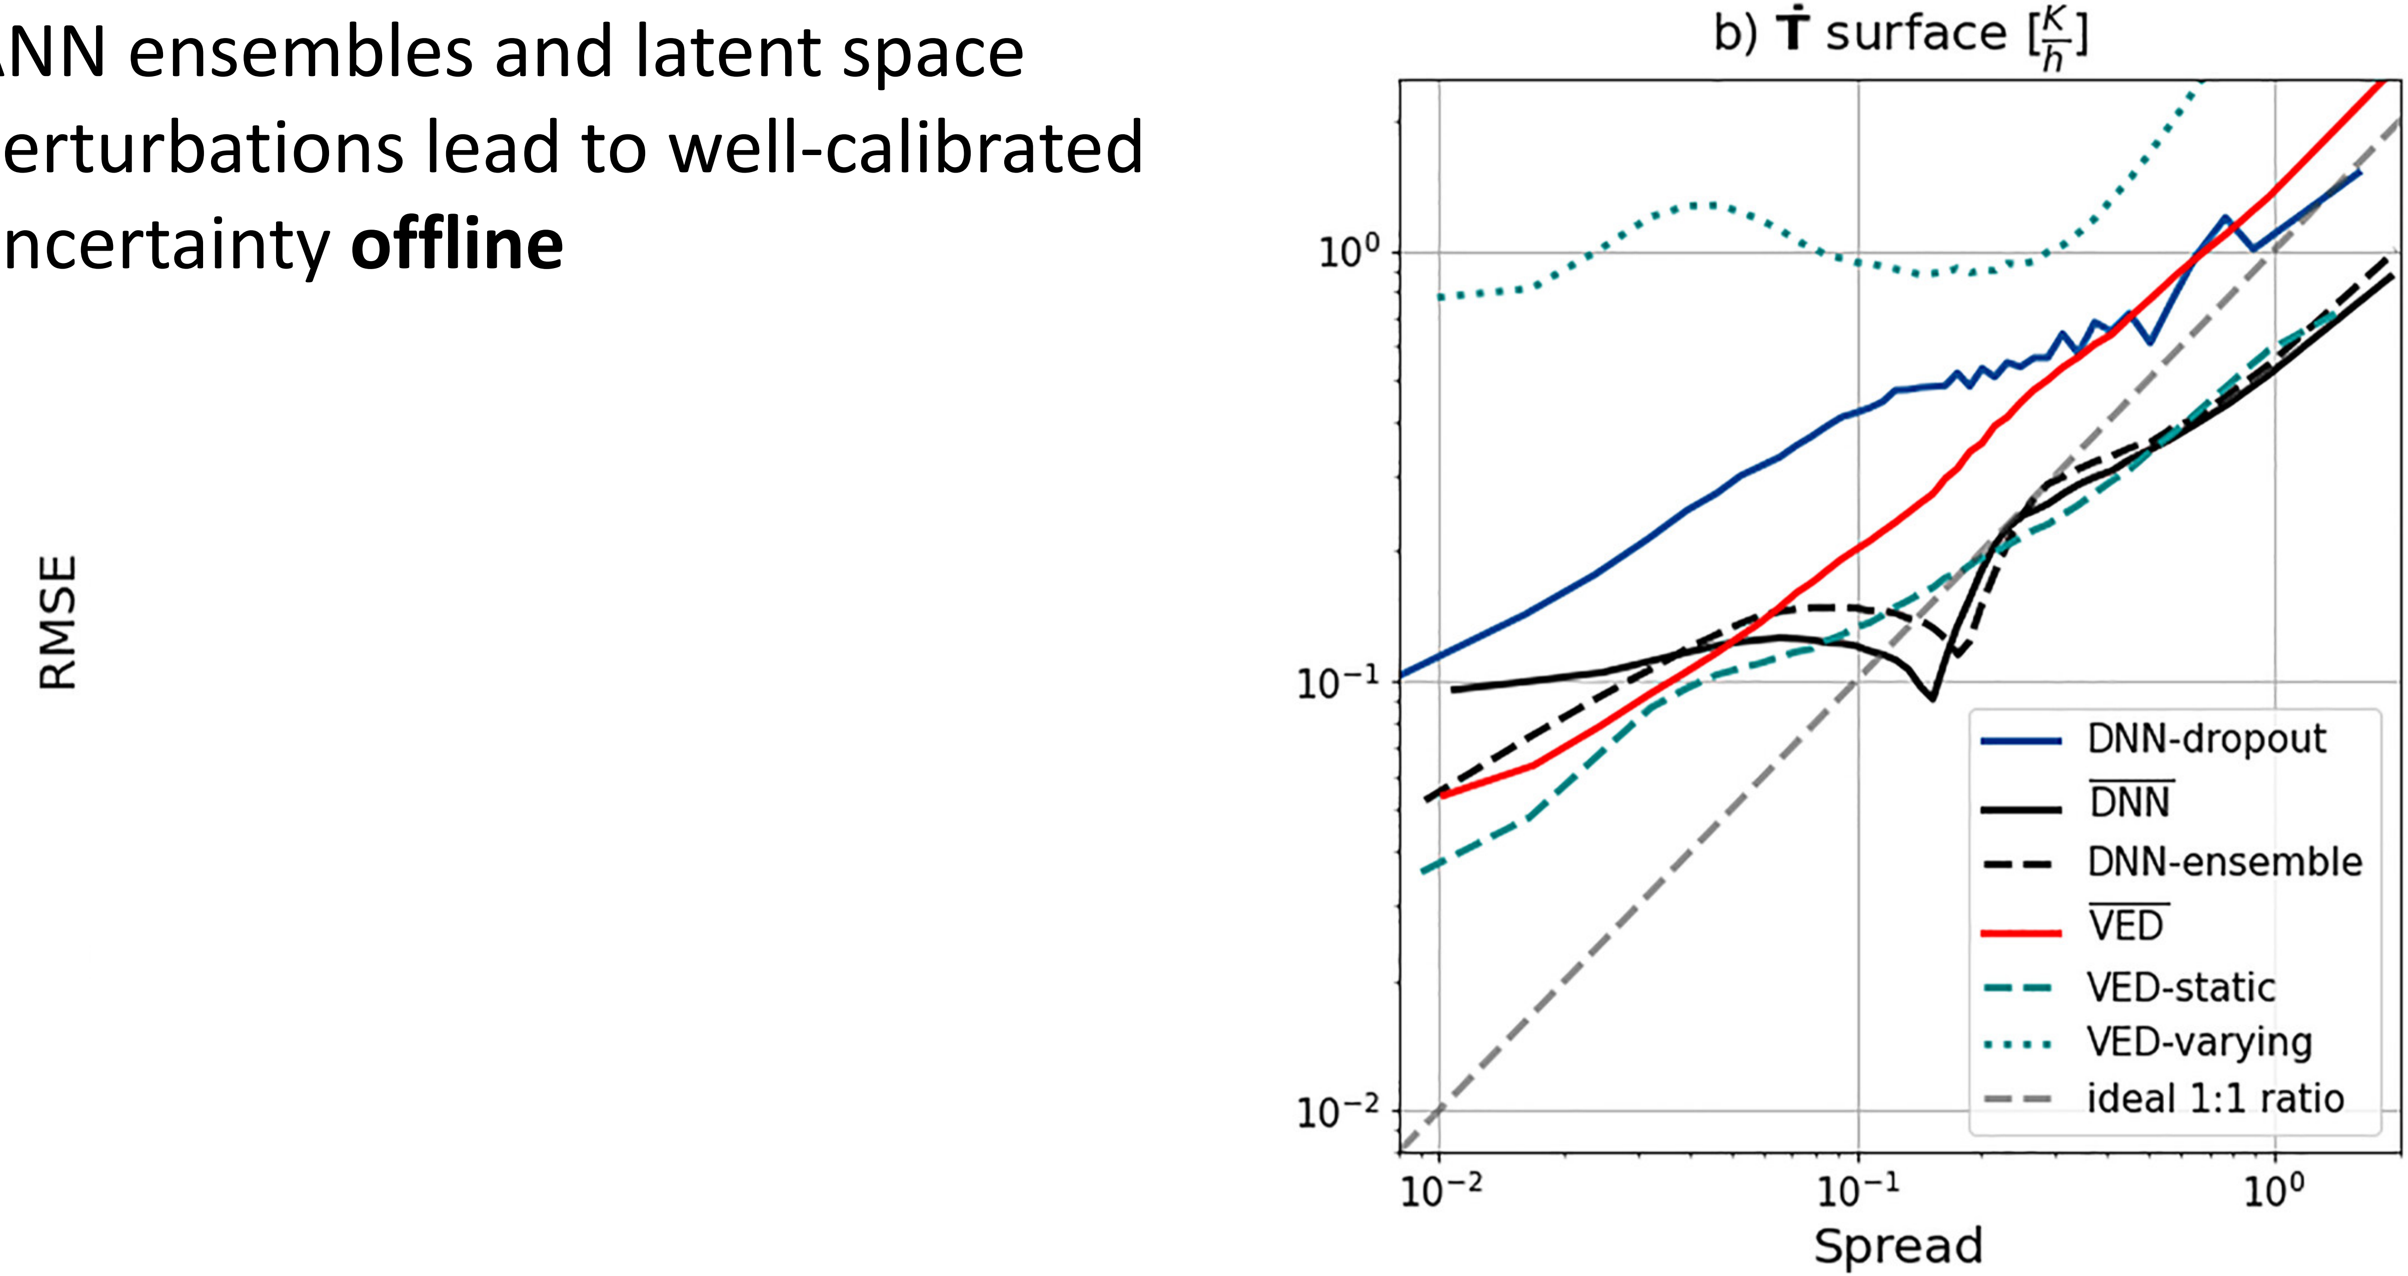

Fig. source: Behrens et al. (2025), See also: Guillaumin & Zanna (2021), Mansfield & Sheshadri (2024), Schneider et al. (2021)

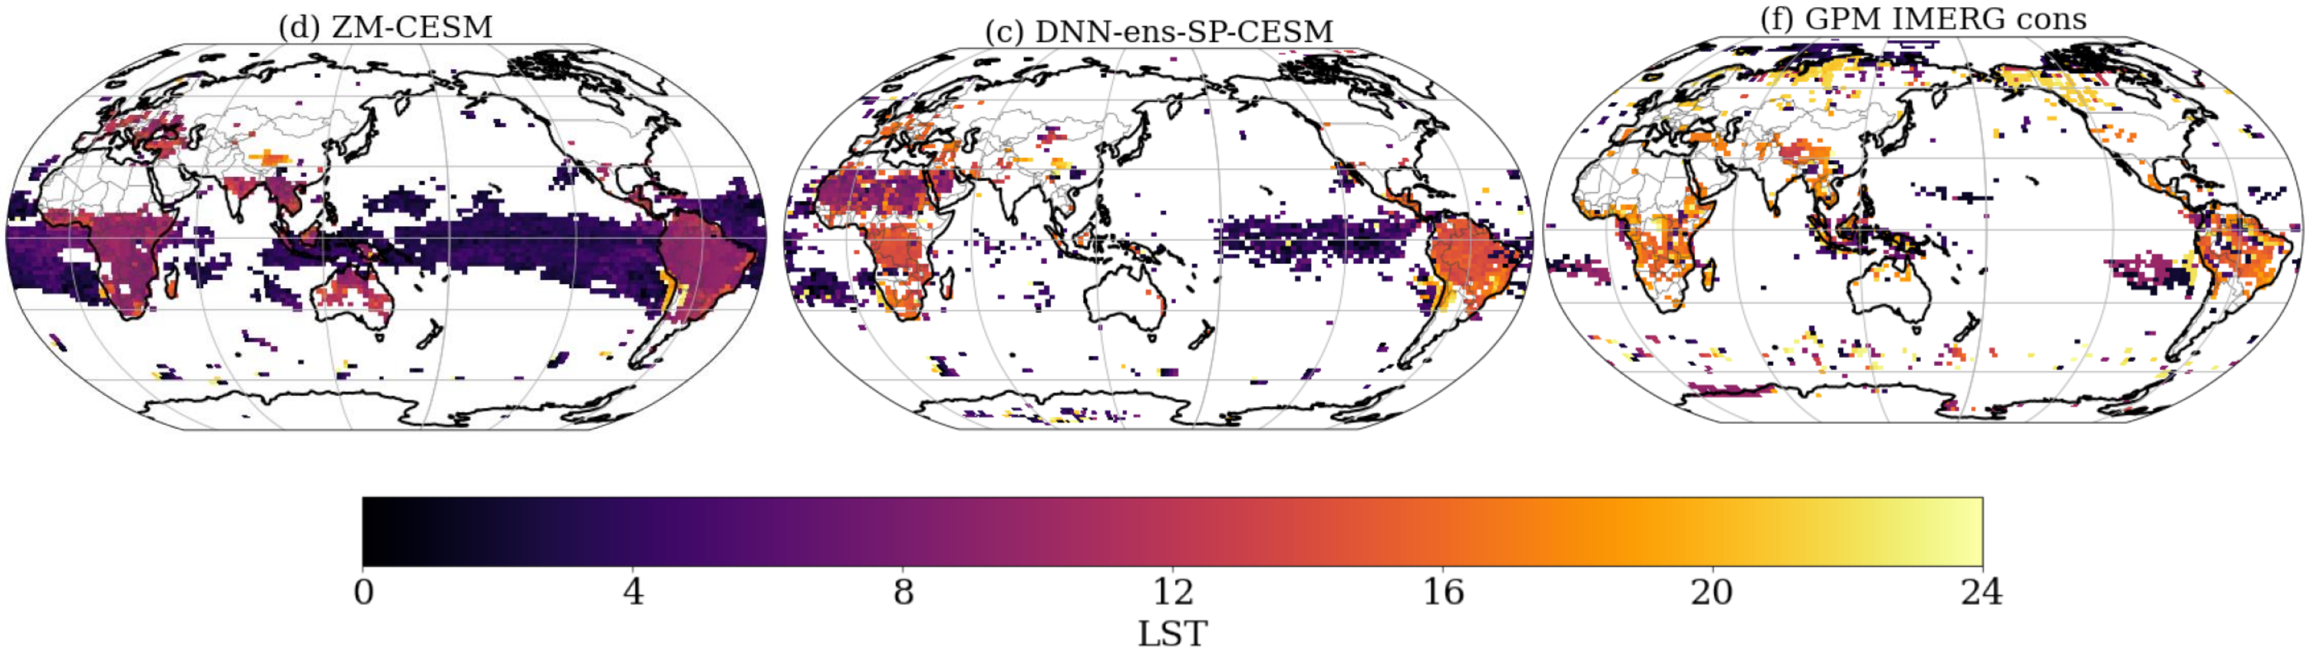

4) Many challenges remain unsolved:

1. Stability, extrapolation behavior, and recalibration of the host model
2. Best way to incorporate causality?
3. Best way to incorporate and evaluate stochasticity?
4. Grid-independence? Scale awareness? Transferability?

At hourly timescales, convection is non-local!
How can we keep our equation simple?

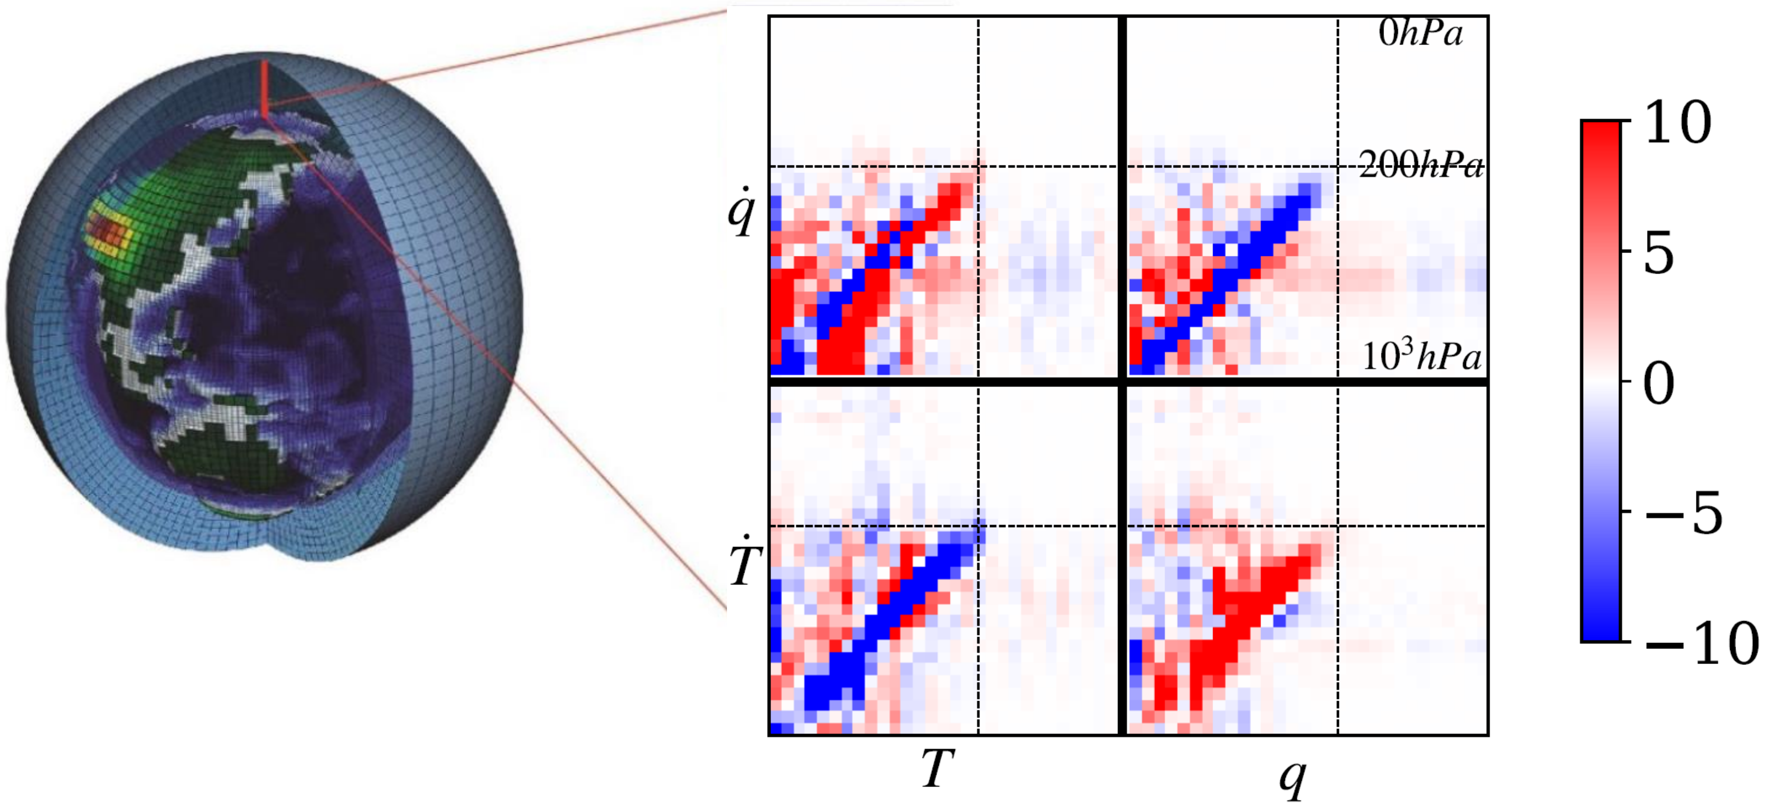

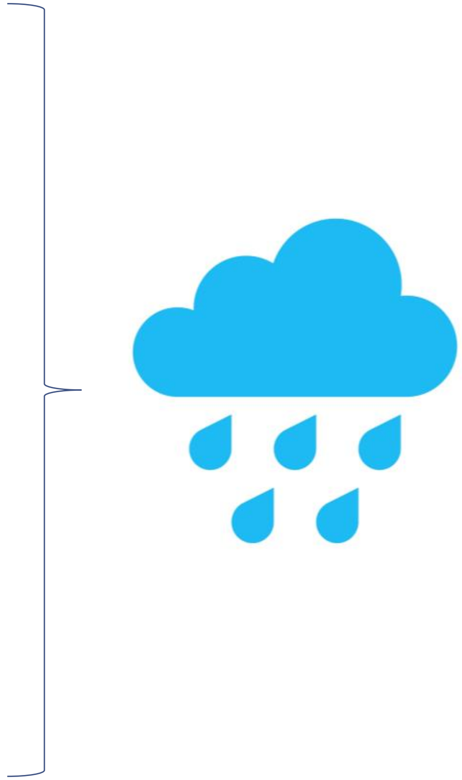

<u>Image source</u>: Gentine, Eyring & Beucler (2020); <u>Figure source</u>: Beucler et al. (2024)

<u>Idea</u>: 1. Learn a vertical integration kernel from data
2. Parameterize this integration kernel analytically

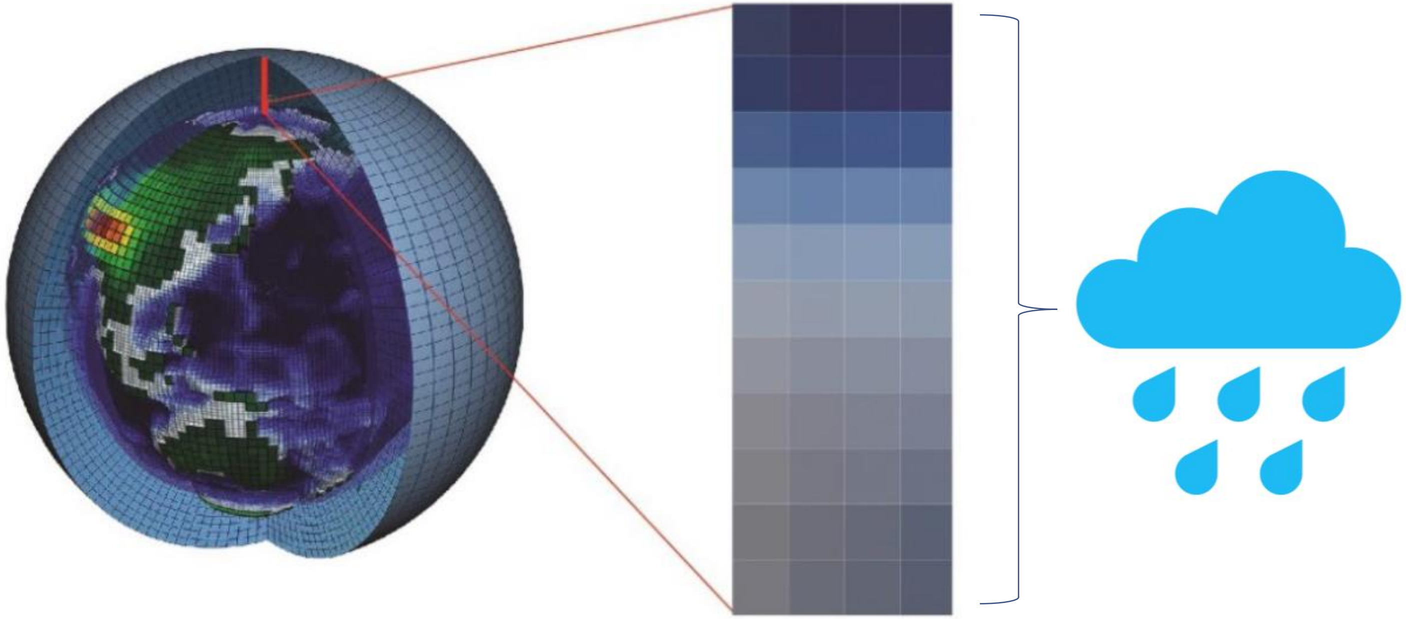

$$y = f\left[\boldsymbol{X}(\boldsymbol{p})\right] \approx f\left[\int_0^{p_s} \frac{dp}{g} \underbrace{k(p)}_{\text{Learned}} X(p)\right]$$

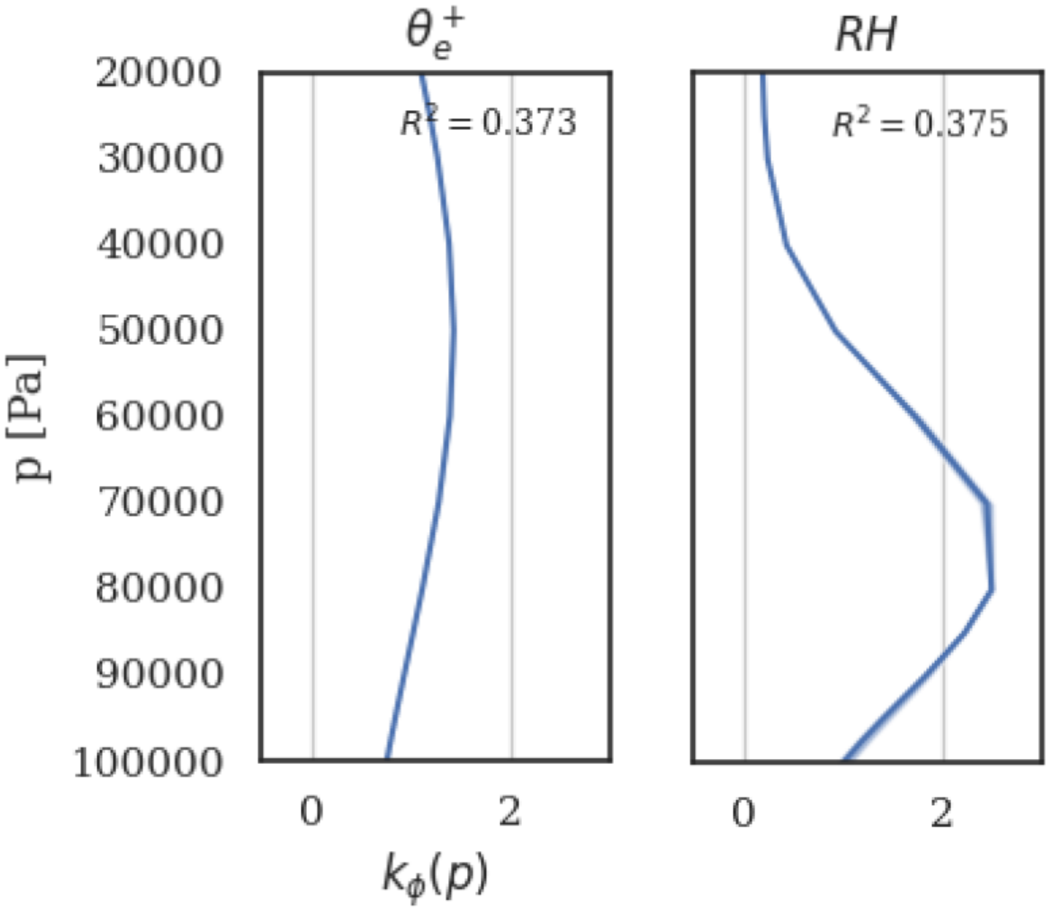

<u>See</u>: Beucler et al. (2024, AMS Tropical Meteorology)

Learning kernels shares analogy with neural operators

NeuralOperator Install User Guide API Examples Developer's Guide

Search the doc Go

Installing NeuralOperator

User Guide

API reference

Examples

NeuralOperator Developer's Guide

Limitation of Fixed Discretization

PDEs are, unfortunately, hard. Instead of learning the operator, people usually discretize the physical domain and cast the problem in finite-dimensional Euclidean space. Indeed, hundred years of effort has been made to develop numerical solvers such as the finite element method and finite difference method.

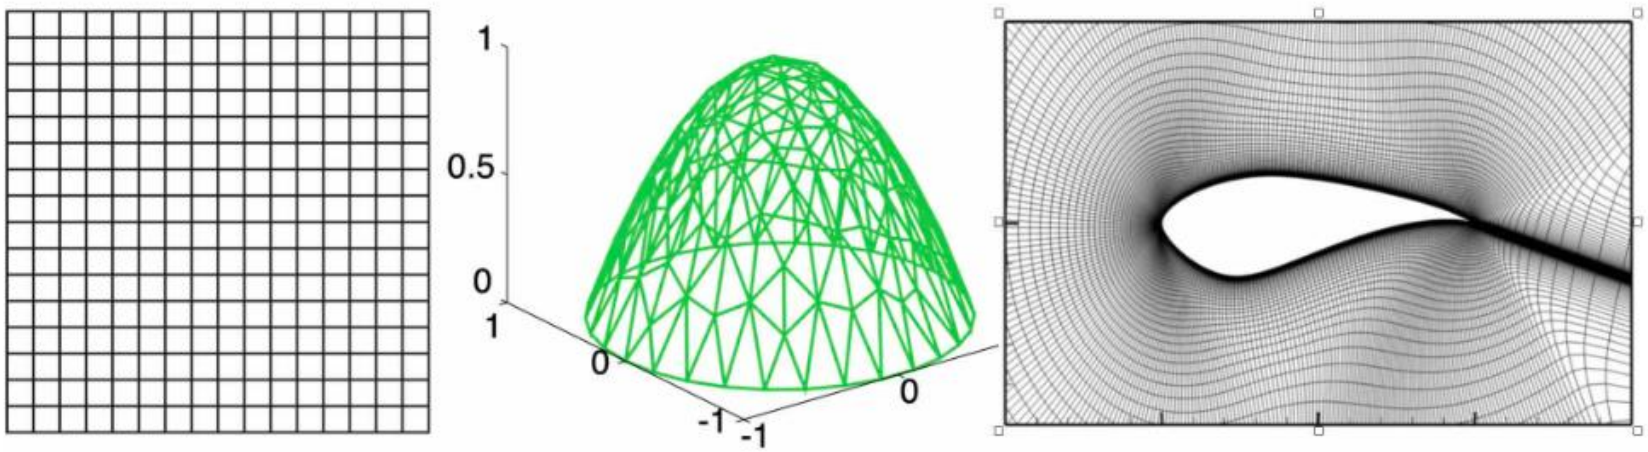

Source: https://neuraloperator.github.io/

ACE2 is a SOA climate model emulator based on Spherical Fourier Neural Operators

ACE2-SOM: Coupling an ML atmospheric emulator to a slab ocean and learning the sensitivity of climate to changed CO_2

Spencer K. Clark[1,2], Oliver Watt-Meyer[1], Anna Kwa[1], Jeremy McGibbon[1], Brian Henn[1], W. Andre Perkins[1], Elynn Wu[1], Lucas M. Harris[2], and Christopher S. Bretherton[1]

[1]Allen Institute for Artificial Intelligence, Seattle, WA, USA
[2]NOAA/Geophysical Fluid Dynamics Laboratory, Princeton, NJ, USA

Key Points:

- The Ai2 Climate Emulator coupled to a slab ocean accurately emulates temperature and precipitation CO_2 sensitivity in a physics-based model
- Inference in an out-of-sample scenario with gradually increasing CO_2 is also accurate except for regime shifts in its stratosphere
- Abrupt $4xCO_2$ inference reaches the correct equilibrium climate but the atmosphere warms too fast due to energy non-conservation

Group activity

4) Many challenges remain unsolved:

1. Stability, extrapolation behavior, and recalibration of the host model
2. Best way to incorporate causality?
3. Best way to incorporate and evaluate stochasticity?
4. Grid-independence? Scale awareness? Transferability?
5. **Questioning the entire climate model formulation** 🤯

README   MIT license

Towards Hybrid Earth System Modeling: A Living Review

This page reviews and organizes emerging hybrid Earth System Models (ESMs), which combine Machine Learning (ML) and physics-based components, alphabetically. Hybrid ESMs retain essential components for physical consistency (e.g., the dynamical core) while using ML to enhance parameterizations for small-scale processes (e.g., clouds). These models hold promise for improving long-term projections of Earth's physical climate and biogeochemical cycles.

If you notice any errors, omissions, or outdated information, please feel free to submit a pull request.

Author: Tom Beucler (UNIL); written in the context of AI4PEX.

Table of Contents

https://github.com/tbeucler/HybridESM

If you want to learn more: Lit. reviews are listed at https://github.com/tbeucler/ML_for_Environmental_Science

README   MIT license

Literature Reviews on Machine Learning for Environmental Science

- Stephan Rasp (Living Review): *State-of-the-art in AI-based weather forecasting.*
- Tom Beucler (Living Review): *State-of-the-art in hybrid Earth system modeling.*
- Eyring et al. (2024): *Pushing the frontiers in climate modeling and analysis with machine learning.*
- Beucler et al. (2021): *Machine Learning for Clouds and Climate.*
- Ullrich et al. (2025): *Recommendations for comprehensive and independent evaluation of machine learning-based Earth system models.*
- Reichstein et al. (2019): *Deep Learning and Process Understanding for Data-Driven Earth System Science.*
- Lai et al. (2025): *Machine Learning for Climate Physics and Simulations.*
- Zhu et al. (2023): *Machine Learning in Environmental Research: Common Pitfalls and Best Practices.*
- Reichstein et al. (2025): *Early warning of complex climate risk with integrated artificial intelligence.*
- Bergen et al. (2019): *Machine Learning for Data-Driven Discovery in Solid Earth Geoscience.*
- Rampal et al. (2024): *Enhancing Regional Climate Downscaling through Advances in Machine Learning.*
- Beucler et al. (2024): *Next-Generation Earth System Models: Towards Reliable Hybrid Models for Weather and Climate Applications.*
- Camps-Valls et al. (2025): *Artificial intelligence for modeling and understanding extreme weather and climate events.*
- Bracco et al. (2024): *Machine learning for the physics of climate.*
- Eyring et al. (2024): *AI-empowered next-generation multiscale climate modelling for mitigation and adaptation.*
- Rolnick et al. (2019): *Tackling Climate Change with Machine Learning.*
- Willard et al. (2020): *Integrating Physics-Based Modeling with Machine Learning: A Survey.*
- Sonnewald et al. (2021): *Bridging observations, theory and numerical simulation of the ocean using machine learning*

Hands-on exercises:

https://wp.unil.ch/dawn/teaching/

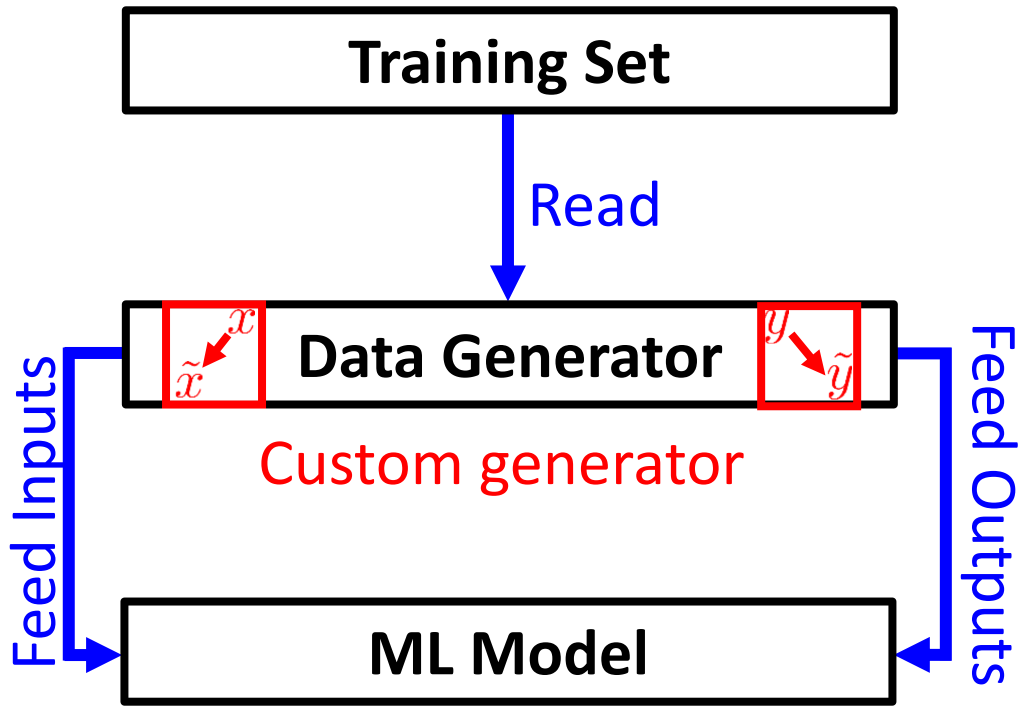

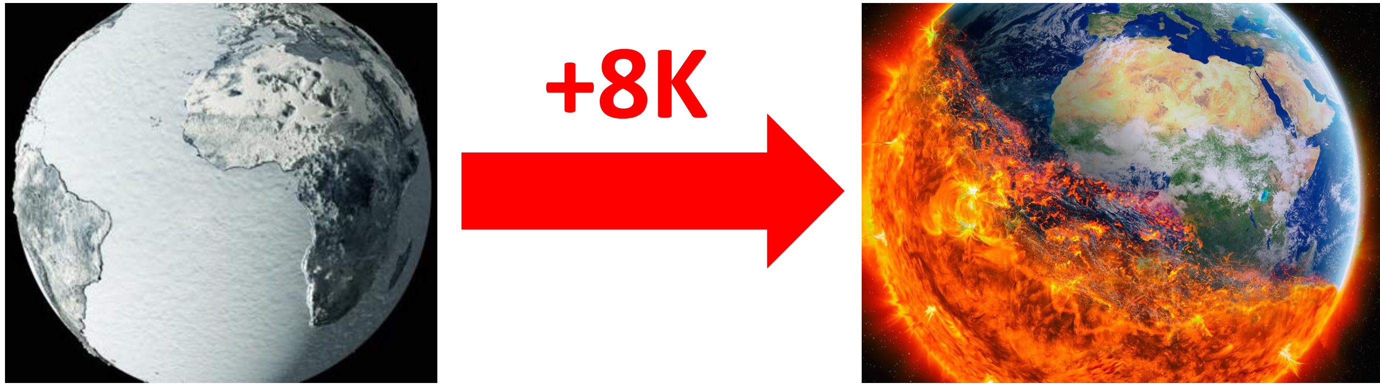

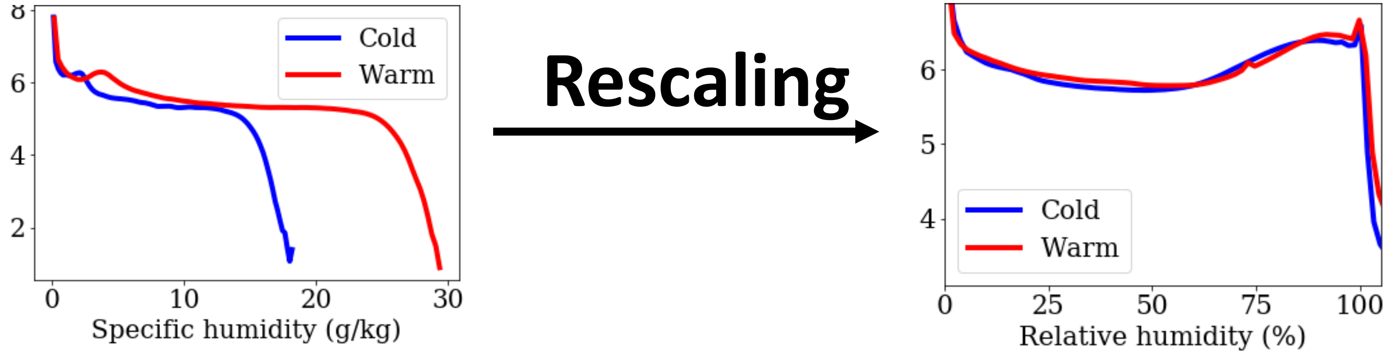

See: Beucler et al. (2021)