Generative AI Models
ECE 598 LV – Lecture 21

Lav R. Varshney

7 April 2022

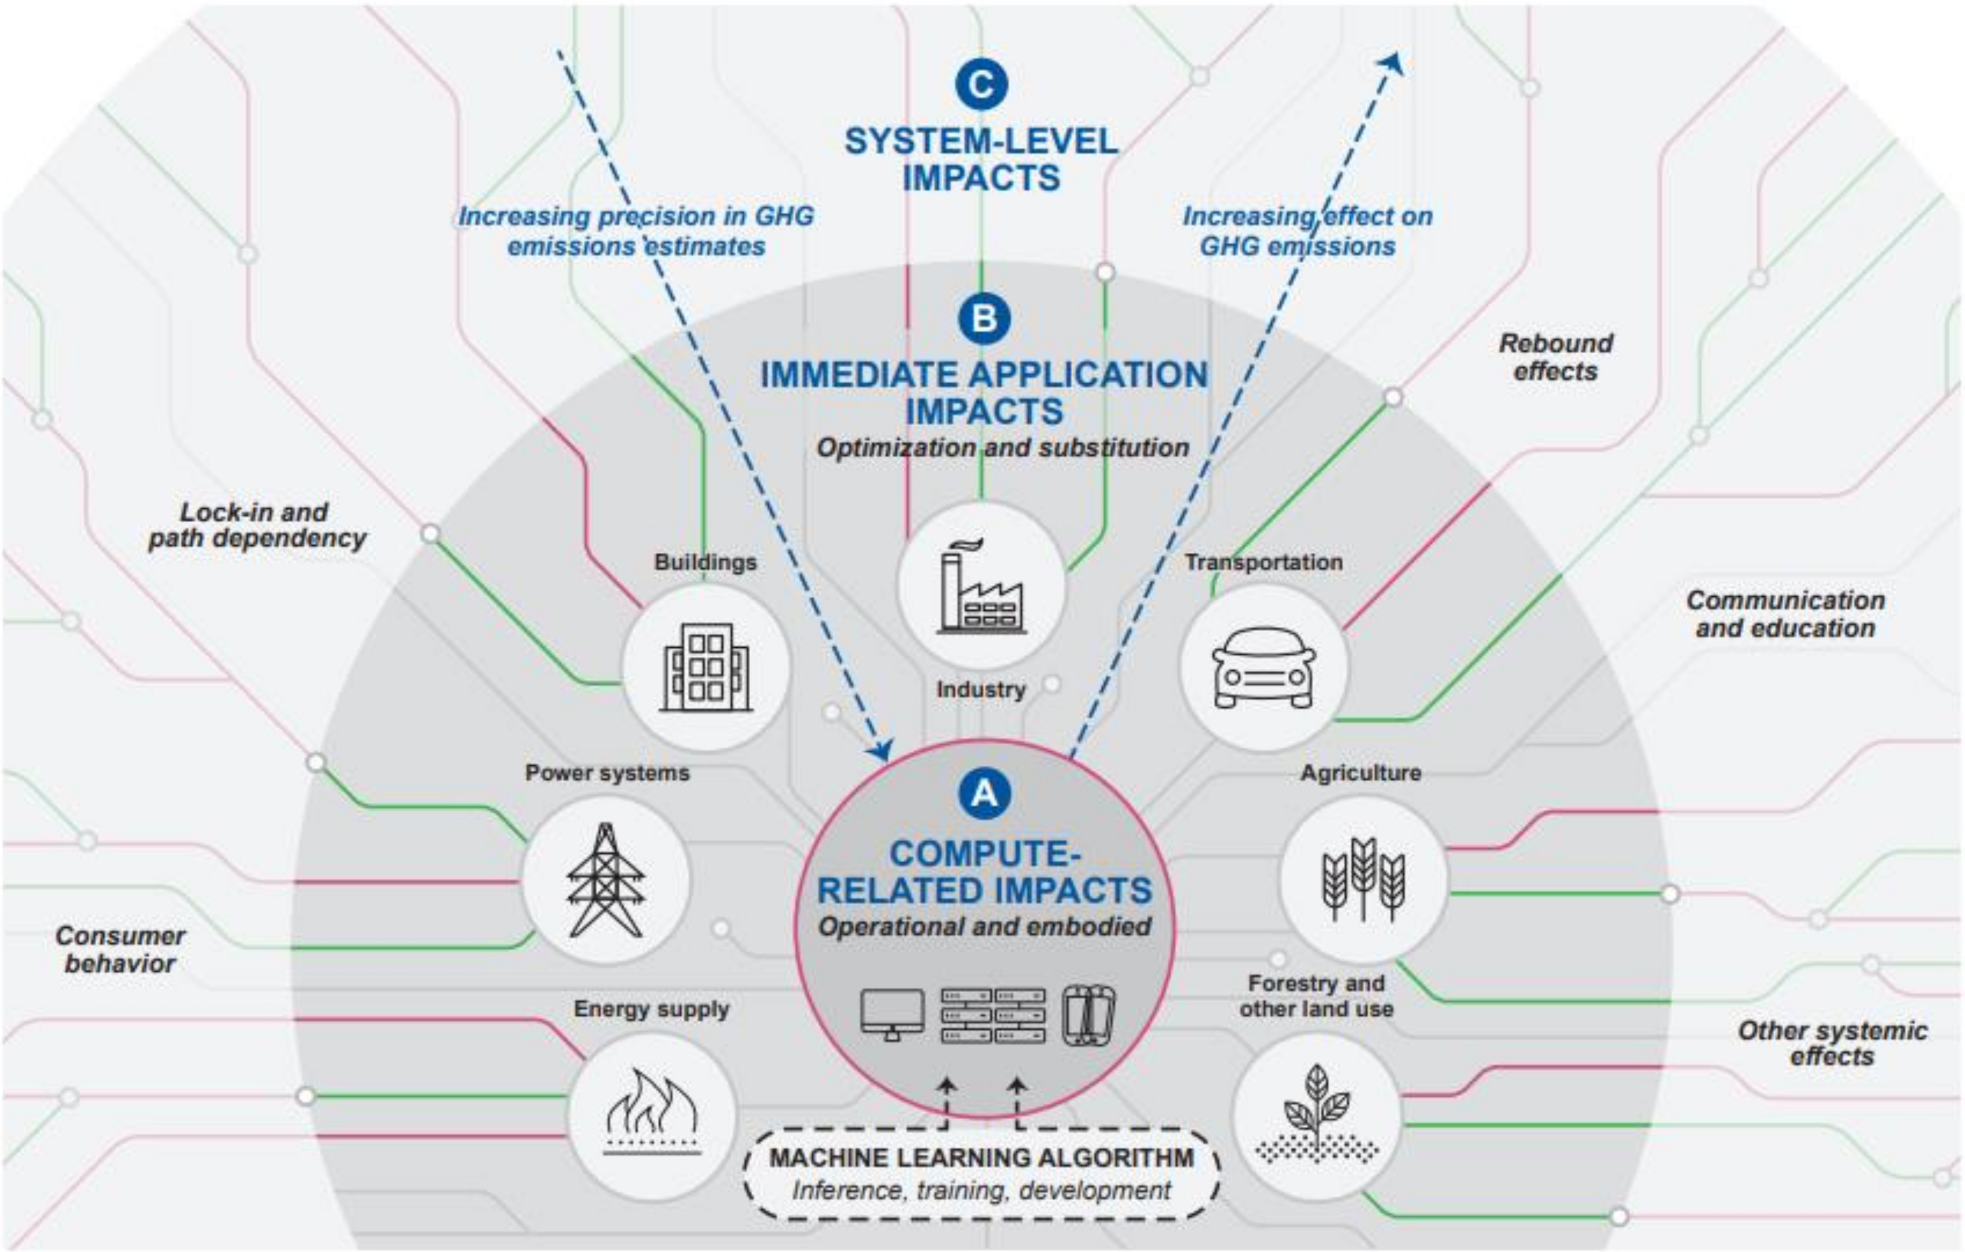

[Kaack et al, 2021]

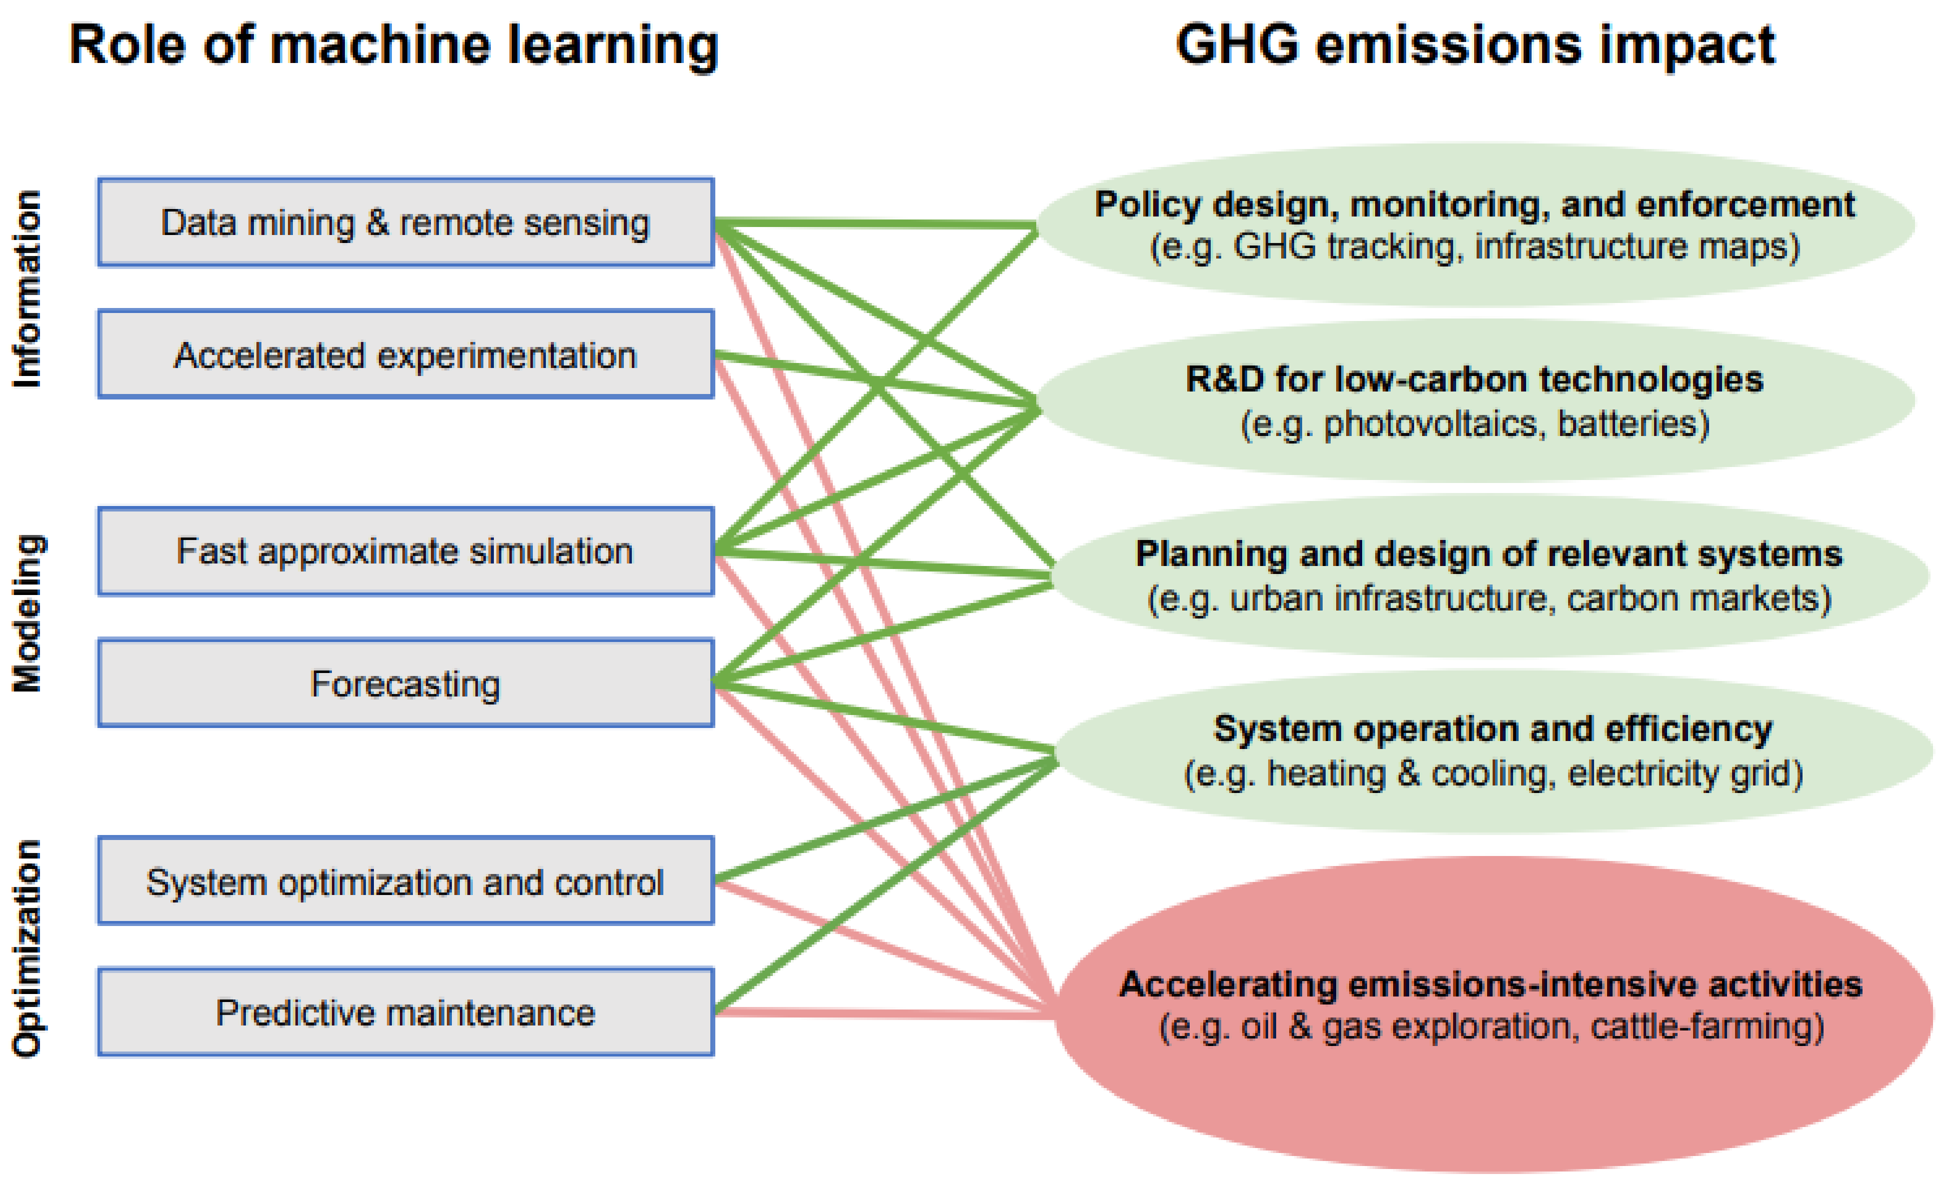

[Kaack et al, 2021]

	Causal inference	Computer vision	Interpretable models	NLP	RL & Control	Time-series analysis	Transfer learning	Uncertainty quantification	Unsupervised learning
1 Electricity systems									
Enabling low-carbon electricity		•	•		•	•		•	•
Reducing current-system impacts		•				•		•	•
Ensuring global impact		•					•		•
2 Transportation									
Reducing transport activity		•				•		•	•
Improving vehicle efficiency		•			•				
Alternative fuels & electrification					•				•
Modal shift	•	•				•		•	
3 Buildings and cities									
Optimizing buildings	•				•	•	•		
Urban planning		•				•	•		•
The future of cities				•			•	•	•
4 Industry									
Optimizing supply chains		•			•	•			
Improving materials									•
Production & energy		•	•		•				
5 Farms & forests									
Remote sensing of emissions		•							
Precision agriculture		•			•	•			
Monitoring peatlands		•							
Managing forests		•			•	•			
6 Carbon dioxide removal									
Direct air capture									•
Sequestering CO_2		•						•	•
7 Climate prediction									
Uniting data, ML & climate science		•	•			•		•	
Forecasting extreme events		•	•			•		•	

[Rolnick et al, 2019]

	Causal inference	Computer vision	Interpretable models	NLP	RL & Control	Time-series analysis	Transfer learning	Uncertainty quantification	Unsupervised learning
8 Societal impacts									
Ecology		•					•		
Infrastructure					•	•		•	
Social systems		•				•			•
Crisis		•		•					
9 Solar geoengineering									
Understanding & improving aerosols						•		•	
Engineering a planetary control system					•			•	
Modeling impacts						•		•	
10 Individual action									
Understanding personal footprint	•			•	•	•			
Facilitating behavior change				•					•
11 Collective decisions									
Modeling social interactions			•		•				
Informing policy	•	•		•				•	•
Designing markets					•	•			•
12 Education				•	•				
13 Finance				•		•		•	

[Rolnick et al, 2019]

Forming New Combinations: Concrete that has half as much embodied carbon and is much stronger

- 5% of worldwide CO_2 emissions caused by cement production
- Reduce environmental impacts of construction materials while complying with product specifications

- UCI ML repository concrete strength dataset + environmental impact evaluated using the Cement Sustainability Initiative's Environmental Product Declaration tool:
  - 1030 instances
  - 8 input variables (composition)
  - 1 (compressive strength)
  - 12 (environmental impact) output variables

- Train a conditional generative neural network model to be able to create novel formulations of concrete

Conditional Variational Autoencoder (CVAE)

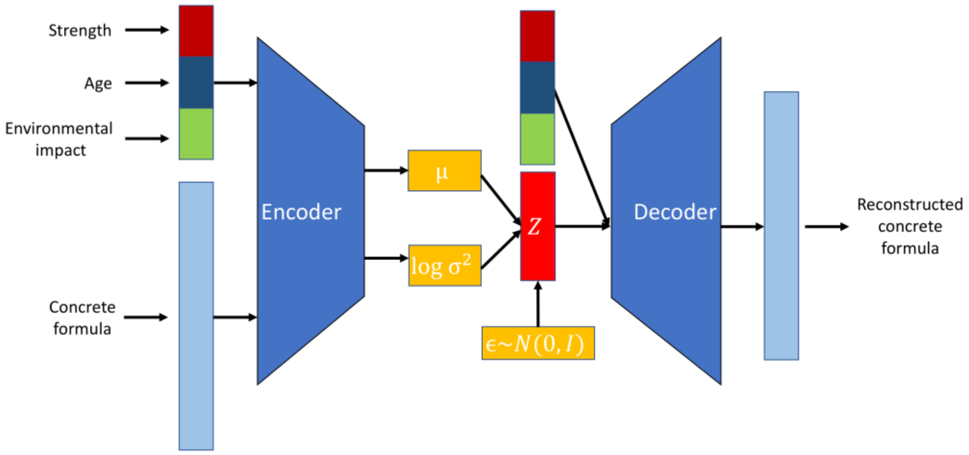

Strength	$[0,1]$
Age	$\{0,1\}^6$
Environmental Impact	$[0,1]^{12}$
Concrete formula	$[0,1]^7$

Forming New Combinations: Concrete that has half as much embodied carbon and is much stronger

Stronger and more than 50% reduction in carbon emissions

DeKalb data center currently under construction

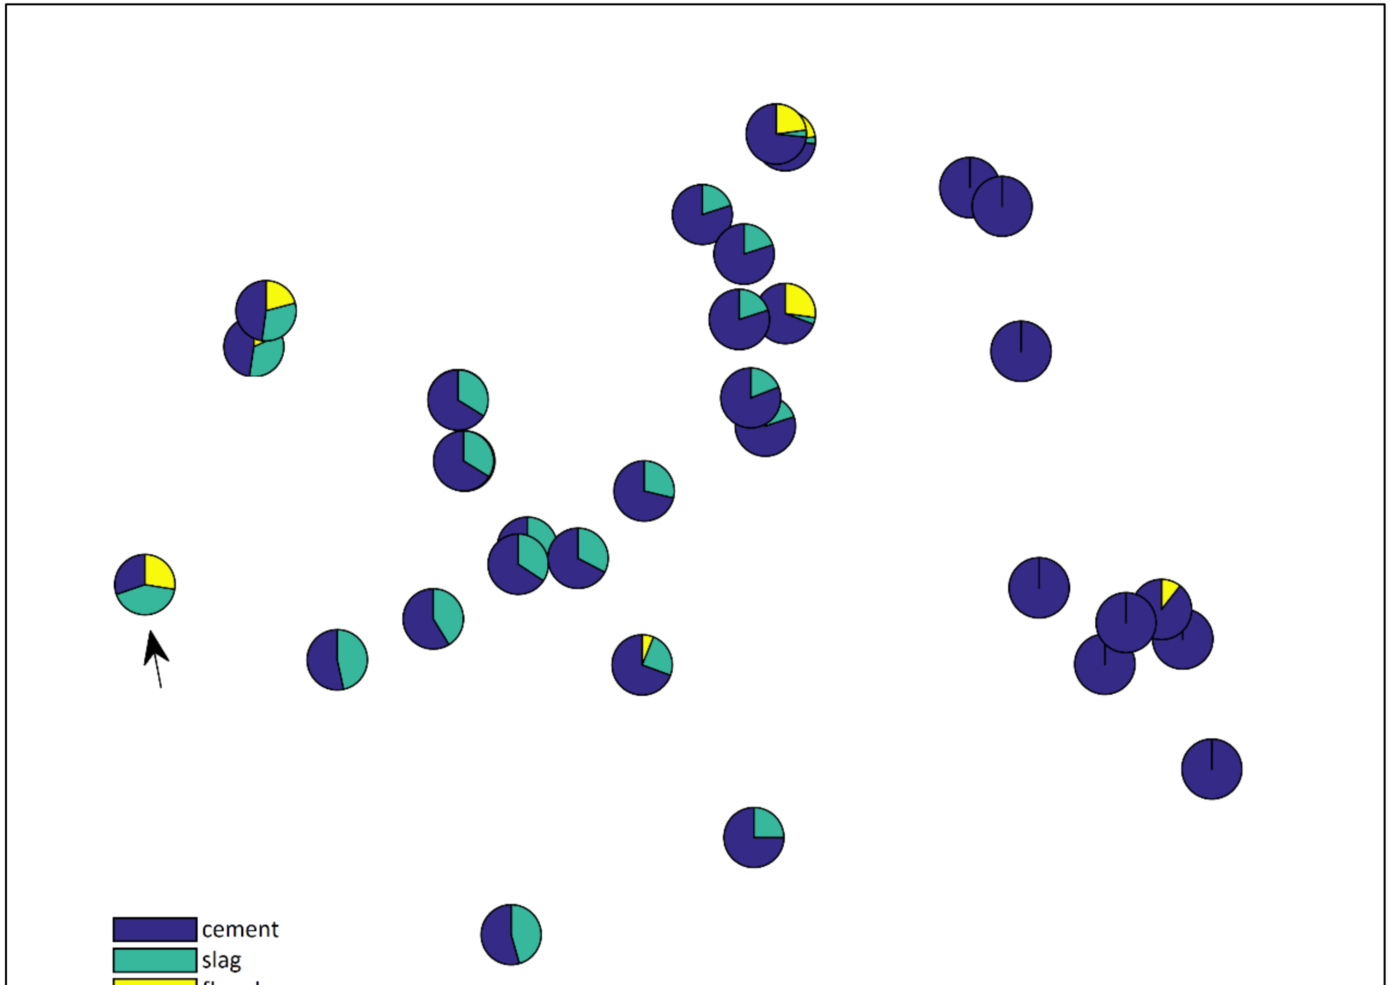

Isomap embedding of concrete formulations

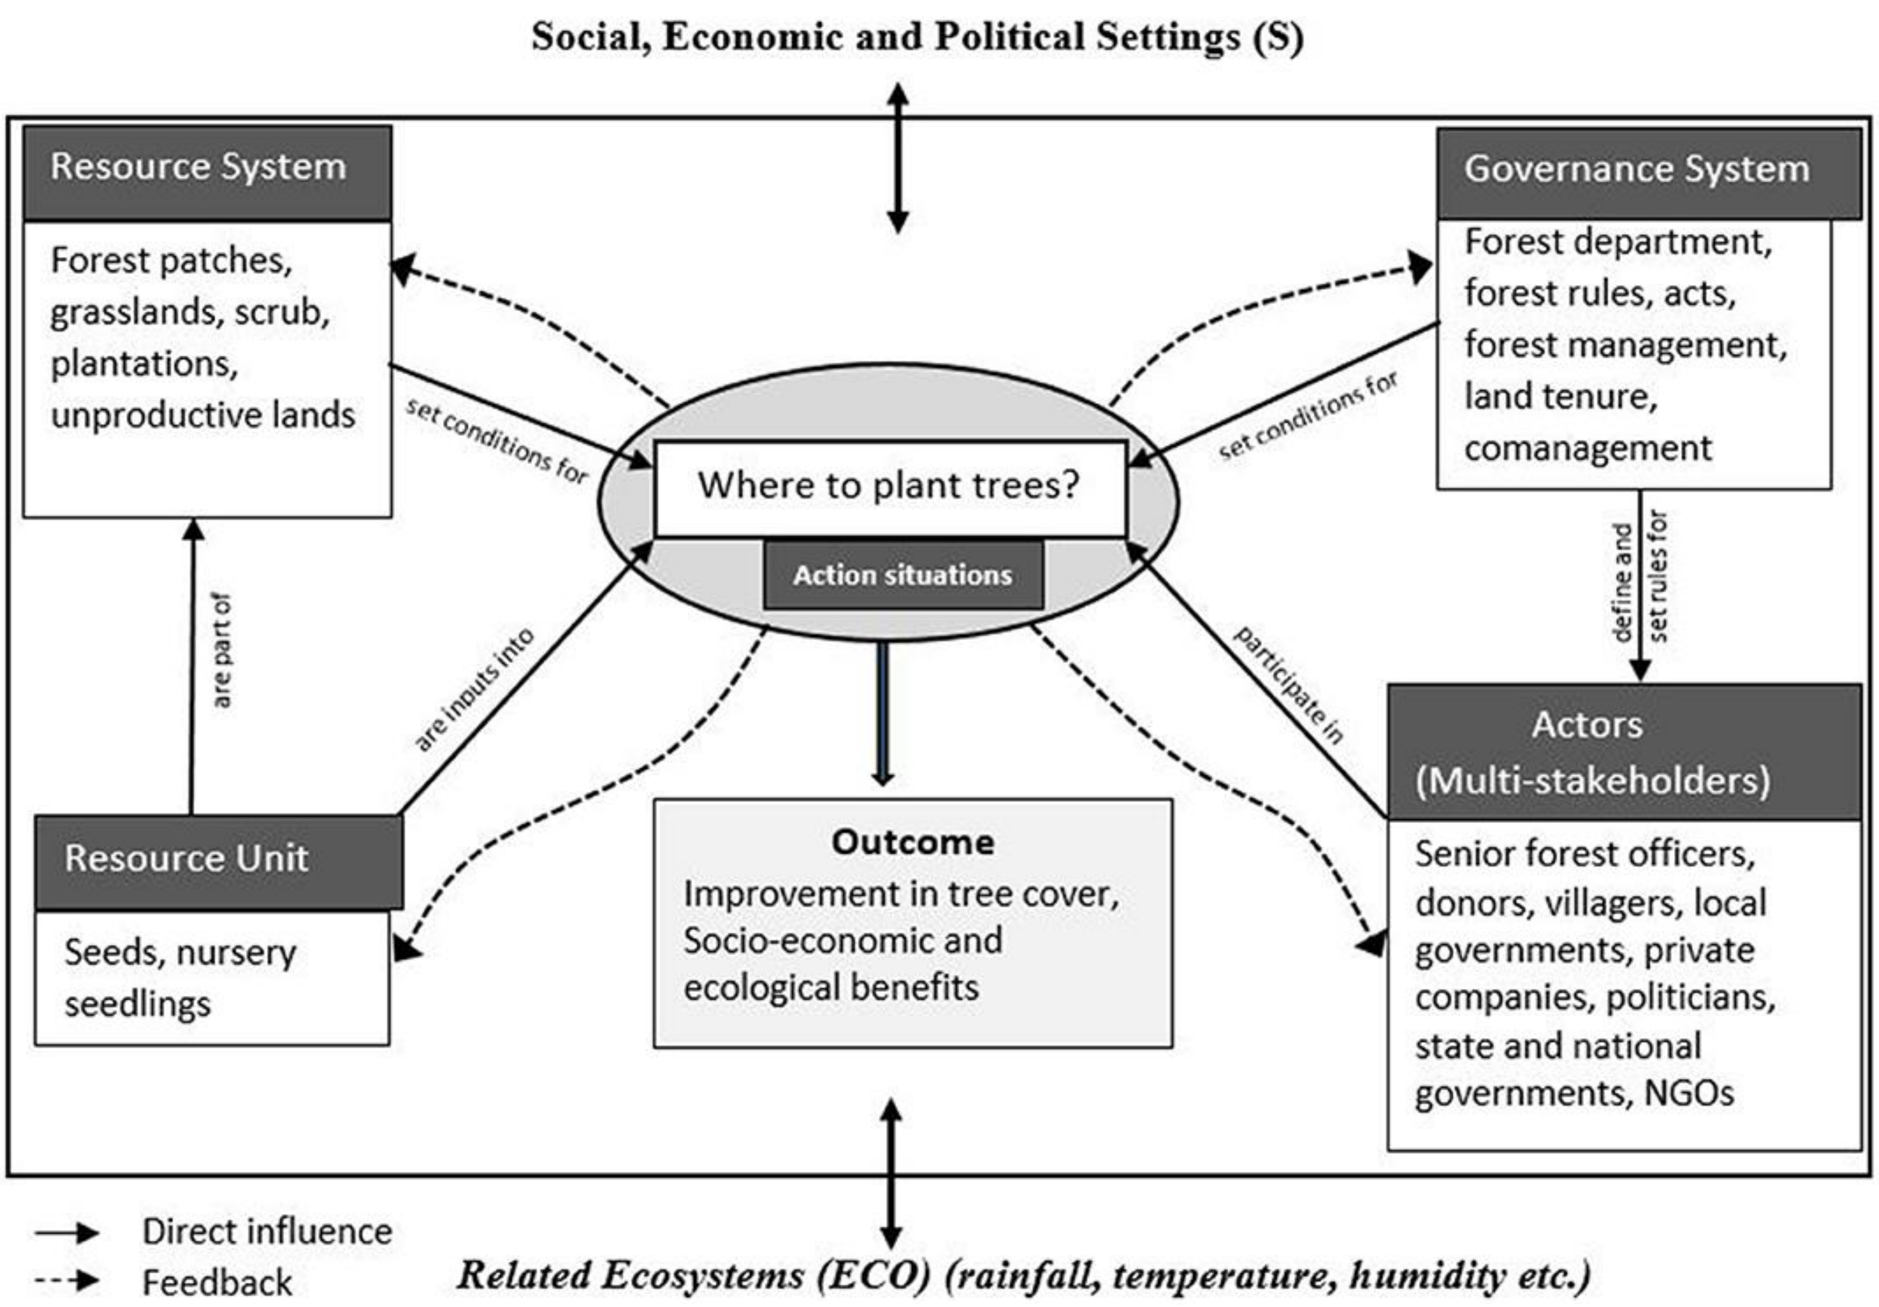

[Rana and Varshney, 2021]

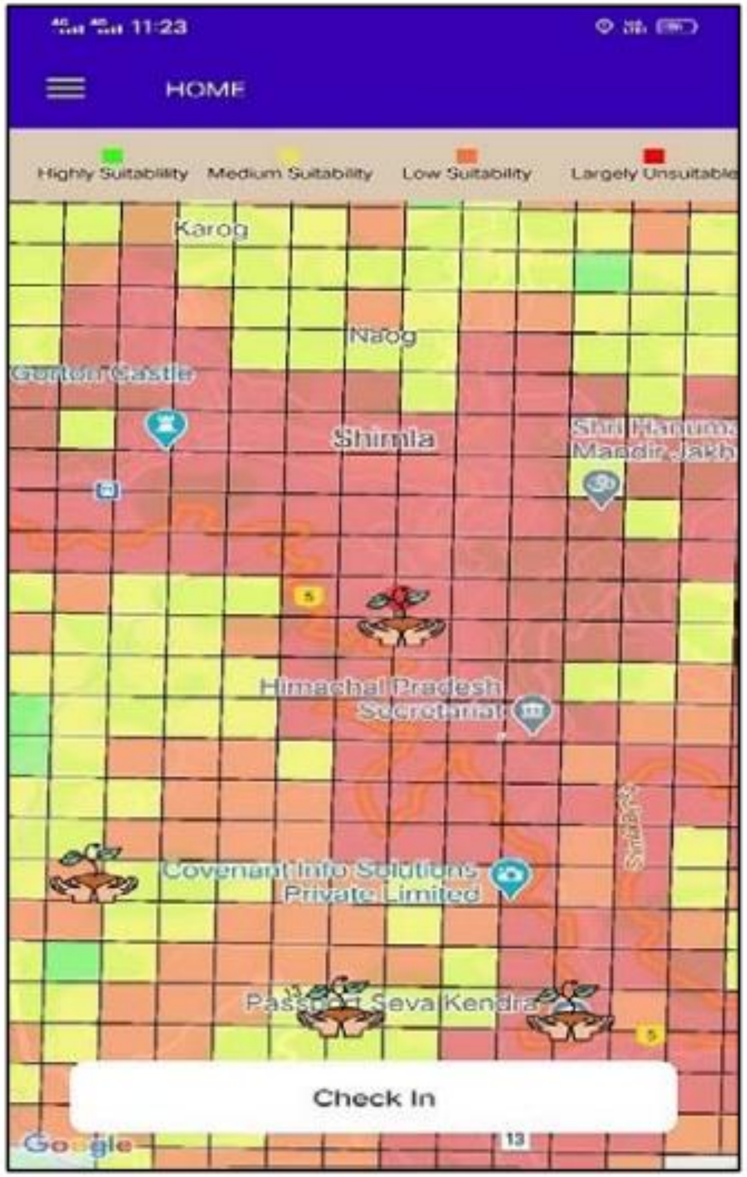

Fig. 2. ePSA (mobile app) showing site suitability classes

Planting trees at the right places: Recommending suitable sites for growing trees using algorithm fusion

Pushpendra Rana[1] and Lav R. Varshney[2]

[1]Indian Forest Service, Shimla, India
Affiliated Researcher, University of
Illinois, Urbana Champaign
pranaifs27@gmail.com

[Rana and Varsney, 2021]

The Day After Tomorrow

https://www.youtube.com/watch?v=Ku_IseK3xTc

Atlas of a Changing World

https://www.youtube.com/watch?v=YUIXp7uoZVc

A 20-Year Community Roadmap for Artificial Intelligence Research in the US

CCC
Computing Community Consortium
Catalyst

AAAI
Association for the Advancement of Artificial Intelligence

Vignette 16

Aishwarya is a climate scientist trying to make predictions of future climate at the local and regional scale. It is essential that such predictions correctly quantify uncertainty. She chooses a climate model that is based on mathematical models of atmospheric physics, solar radiation, and land surface-atmosphere interactions. Unfortunately, running the model at the required fine level of detail is not possible due to the computational cost and the lack of sufficient observation data. Fortunately, recent advances in ML research have produced new physics-aware ML approaches that learn from data while incorporating knowledge about the underlying physics. These approaches run efficiently and produce models at a much finer level of detail than the original climate models. This makes it easy to run multiple models efficiently, which in turn allows Aishwarya to provide clear uncertainty bounds for the resulting predictions.

The results of these models are then used by Jia, who works at FEMA. Using machine learning methods, she combines climate predictions under different policy scenarios (no change versus reduced carbon emissions, etc.) to identify regions that are most vulnerable to extreme weather events such as hurricanes, floods, droughts, heat waves, and forest fires. With the aid of these causal models, she can plan appropriate responses. For example, her physics-aware ML model produces inundation maps in response to extreme meteorological events (hurricane, heavy rain) to identify areas of flooding, which is fed into an AI system for smart cities to perform evacuation planning, emergency relief operations, and planning for long-term interventions (e.g., building a sea wall to ward off storm surge).

Thanks to the 20 years of research investment in AI, physics-aware machine learning techniques are available that can process multimodal, multi-scale data and also handle heterogeneity in space and time, as well as quantify uncertainty in the results. The combination of physics-based climate and hydrological models with machine learned components allows Jia to produce more accurate predictions than would be possible with pure physics-based or pure machine learned models alone. This hybrid approach also generalizes better to novel scenarios, identifying new threats that could result in injury or death. In 2035, these models are applied to revise flood maps, saving many lives in the floods caused by hurricane Thor in North Carolina.

An AI-based Framework for Accelerated Discovery of Climate Impacts

Q: What is the impact of climate change on <insert sector>?

1: **Understand** the driving processes
2: **Identify** the impacts of a changing climate
3: **Design** effective solutions to mitigate impacts
4: **Enable** key stakeholders

An AI-based Framework for Accelerated Discovery of Climate Impacts

A state-of-the-art process-based land and food production model

Nature Food 2021:
"CO2, CH4 and N2O emissions from the production and consumption of all plant- and animal-based foods"

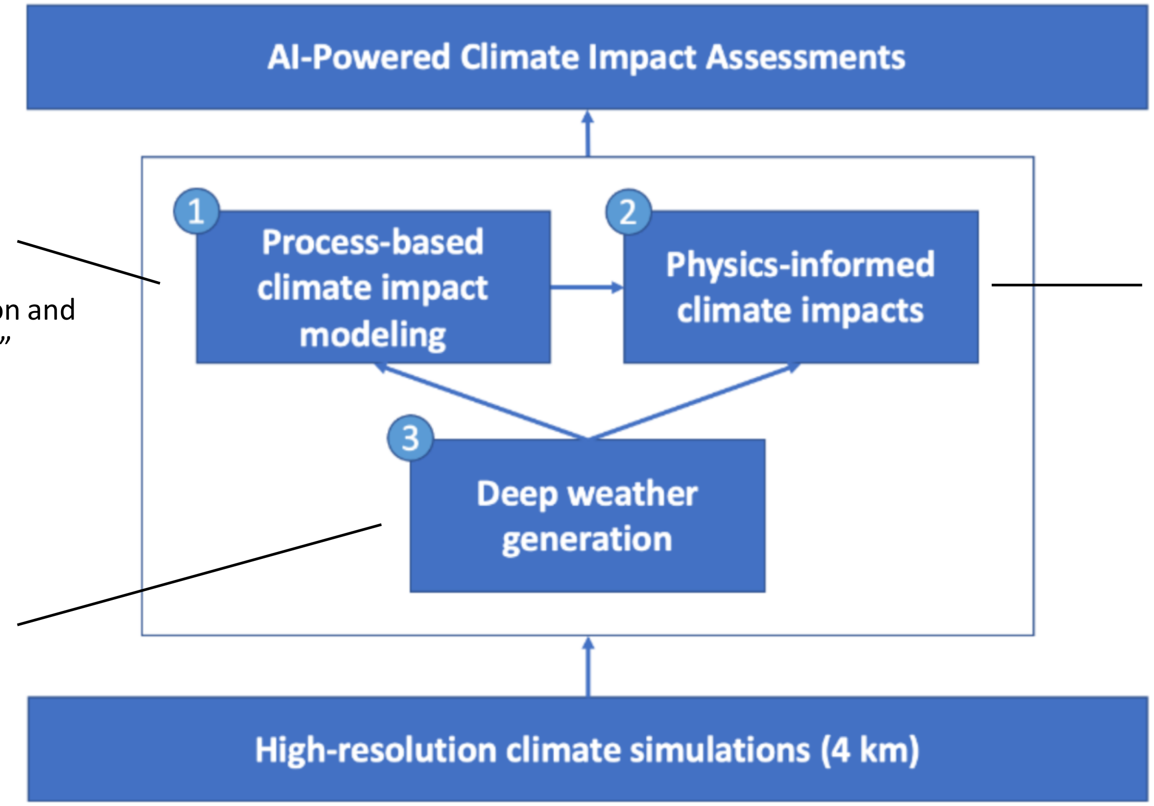

Novel physics-informed AI methods to accelerate discovery of climate impacts on food productivity and GHG emissions

Combine objectives of data fidelity and satisfaction of physical constraints, with symmetries enforced via group-equivariance and parameter sharing

Generative AI techniques for rapid climatic scenarios creation

ICML 2021:
"Mapping complex weather data to a known distribution implements an efficient control for weather field synthesis towards more (or less) extreme scenarios"

State-of-the-art dynamically downscaled high-resolution climate simulations for North America

Unprecedented climate projections developed by a consortium of researchers at UIUC and Argonne https://s3.us-east-1.amazonaws.com/climate-change-ai/papers/icml2021/39/slides.pdf

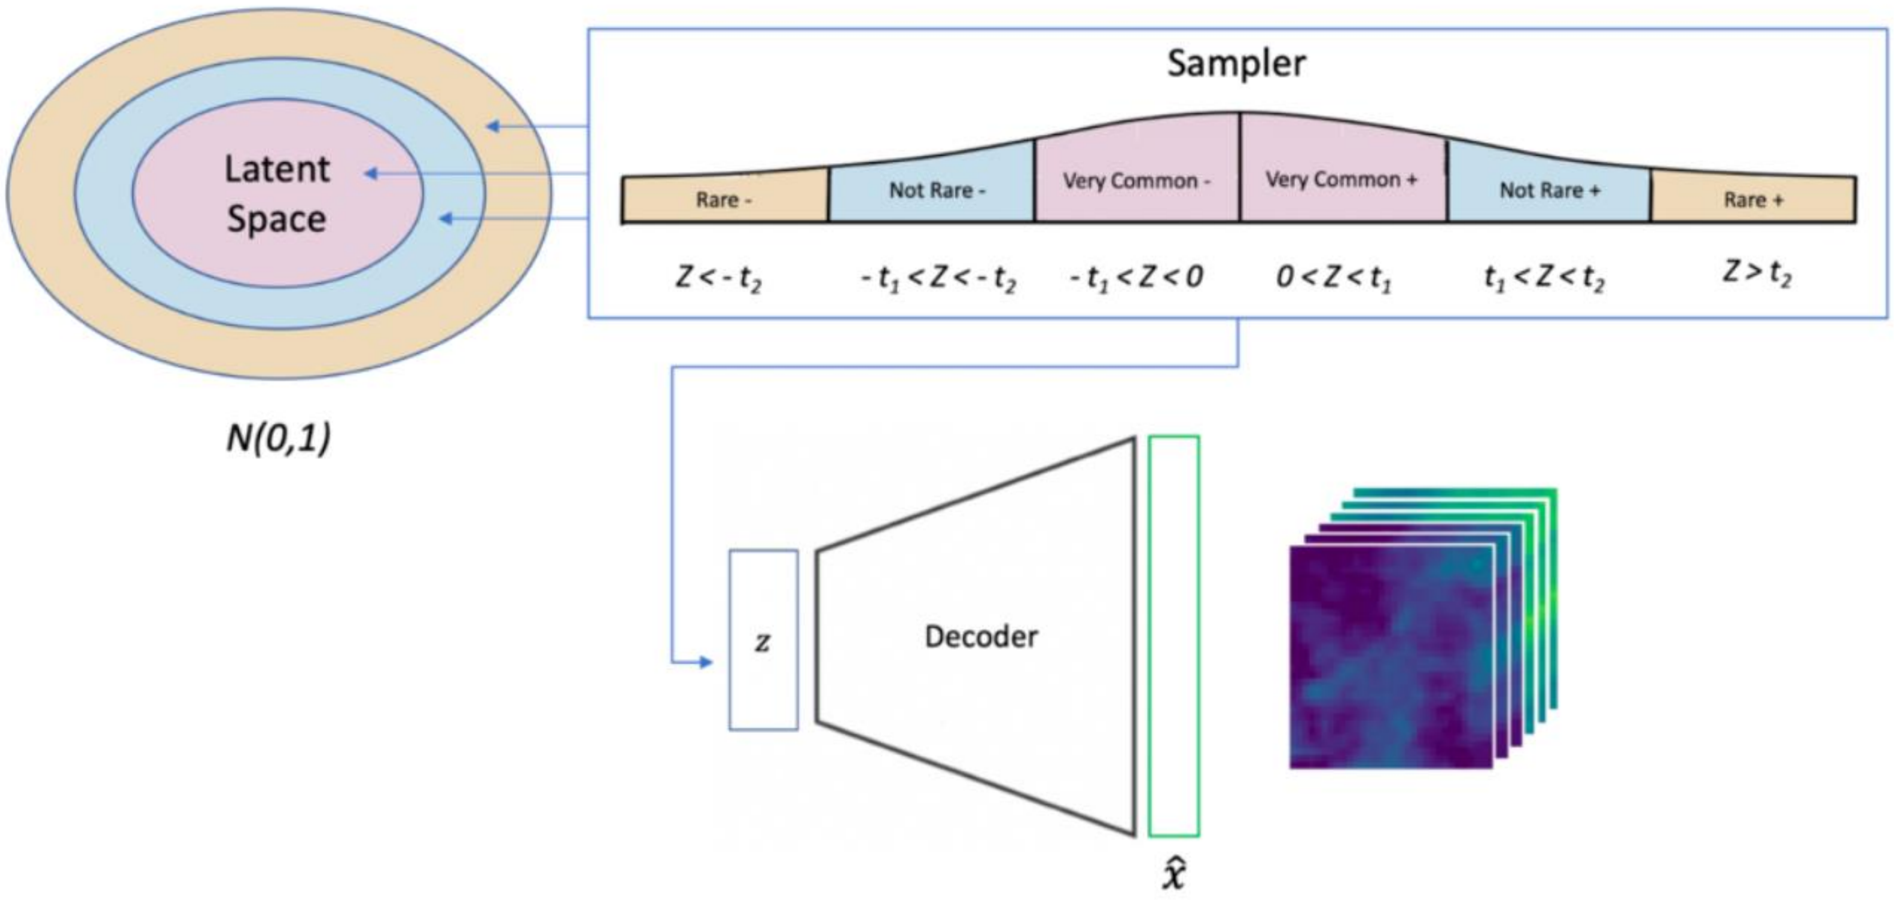

PHILOSOPHICAL TRANSACTIONS A

royalsocietypublishing.org/journal/rsta

Review

Check for updates

Cite this article: Kashinath K *et al.* 2021 Physics-informed machine learning: case studies for weather and climate modelling. *Phil. Trans. R. Soc. A* **379**: 20200093. https://doi.org/10.1098/rsta.2020.0093

Accepted: 24 November 2020

Physics-informed machine learning: case studies for weather and climate modelling

K. Kashinath[1], M. Mustafa[1], A. Albert[1,2], J-L. Wu[1,3], C. Jiang[1,4], S. Esmaeilzadeh[5], K. Azizzadenesheli[6], R. Wang[1,7], A. Chattopadhyay[1,8], A. Singh[1,2], A. Manepalli[1,2], D. Chirila[9], R. Yu[7], R. Walters[10], B. White[2], H. Xiao[11], H. A. Tchelepi[5], P. Marcus[4], A. Anandkumar[3,12], P. Hassanzadeh[8] and Prabhat[1]

2. Physics-informed machine learning: objectives, approaches, applications

(a) Objectives of physics-informed machine learning

By incorporating physical principles, governing laws and domain knowledge into ML models, the rapidly growing field of PIML seeks to:

- Build physically consistent and scientifically sound predictive models.
- Increase data efficiency, i.e. train models with fewer data points.
- Accelerate the training process, i.e. help models converge faster to optimal solutions.
- Improve the generalizability of models to make reliable predictions for unseen scenarios, including applicability to non-stationary systems, e.g. a changing climate.
- Enhance transparency and interpretability to make models more trustworthy.

Section and case study reference	PIML application/ modelling task from §2c	Physical processes	Datasets	PIML approaches from §2b	ML model	PIML objectives achieved from §2a
§3ai: Wu *et al.* [18]	emulation	Rayleigh-Bénard convection	direct numerical simulation (DNS)	custom loss, stochasticity, multi-scale, spectral	GAN	physically consistent, accelerated training
§3aii: Manepalli *et al.* [102]	emulation	mountain snowpack melting (hydro-climate)	hydro- meteorological observational product	custom loss, stochasticity, UQ	conditional GAN	physically consistent
§3aiii: Daw *et al.* [22]	emulation	lake temperature dynamics	observational lake characteristics data	custom architecture, UQ	LSTM	physically consistent, data efficient, interpretable
§3aiv: Beucler *et al.* [14]	emulation	atmospheric convection and clouds	super-parametrized community atmosphere model	custom loss, custom architecture	NN	physically consistent, generalizable
§3bi: Singh *et al.* [103]	downscaling/ super-resolution	atmospheric winds	weather research and forecasting model	custom loss, stochasticity, spectral	GAN	physically consistent, generalizable
§3bii: Vandal *et al.* [89]	downscaling/ super-resolution	precipitation	reanalysis product (PRISM)	custom architecture, stochasticity, UQ	Bayesian NN and CNN	physically consistent, data efficient
§3biii: Jiang *et al.* [104]	downscaling/ super-resolution	Rayleigh-Bénard convection	DNS	custom loss, custom architecture, multi-scale, spatio-temporal coherence	encoder-decoder NN	physically consistent, generalizable, scalable
§3ci: Wang *et al.* [105]	forecasting	Rayleigh-Bénard convection	DNS	custom loss, custom architecture, multi-scale, physics-based structure	encoder-decoder NN	physically consistent, generalizable, interpretable
§3cii: Wang *et al.* [30]	forecasting	Rayleigh-Bénard convection, Ocean currents	DNS, ocean reanalysis product (ORAS5)	custom architecture, equivariant, spatio-temporal coherence	residual network	physically consistent, data efficient, generalizable
§3ciii: Chattopadhyay *et al.* [31]	forecasting	geophysical turbulence	DNS	custom loss, custom architecture, equivariant, spatio-temporal coherence	spatial transformer network	physically consistent, data efficient, stable

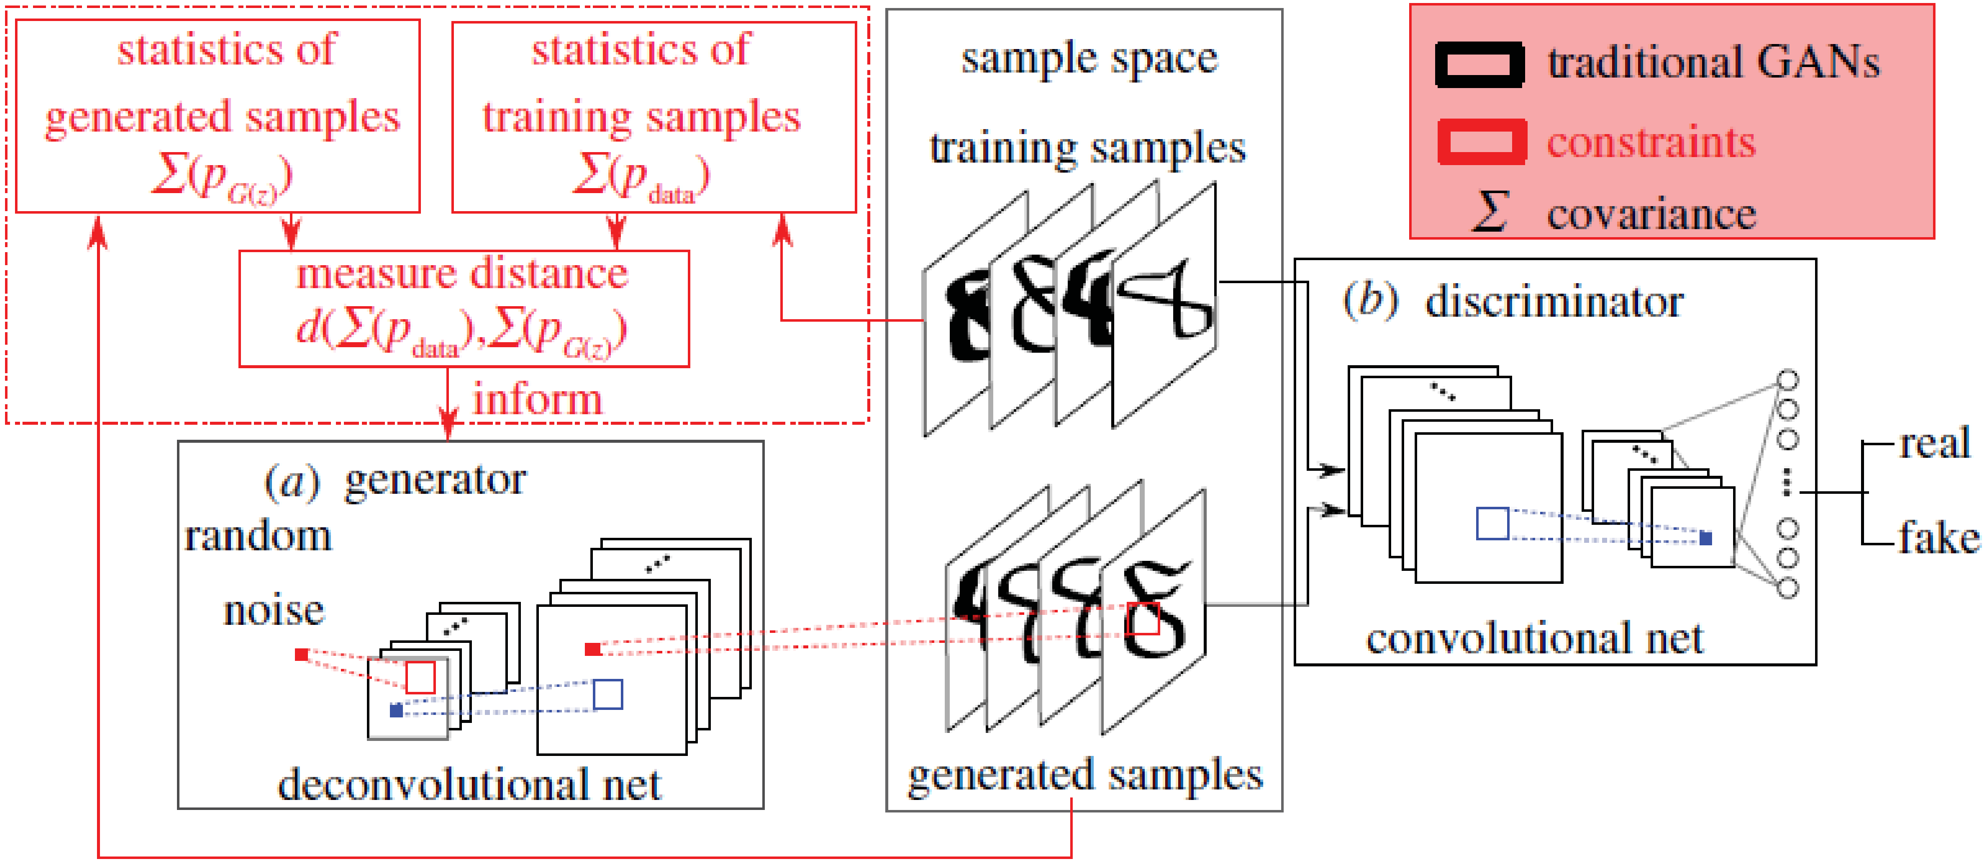
statistics of generated samples
Σ(p_G(z))
statistics of training samples
Σ(p_data)
measure distance
d(Σ(p_data),Σ(p_G(z)))
inform
sample space
training samples
generated samples
traditional GANs
constraints
Σ covariance
(a) generator
random
noise
deconvolutional net
(b) discriminator
convolutional net
real
fake

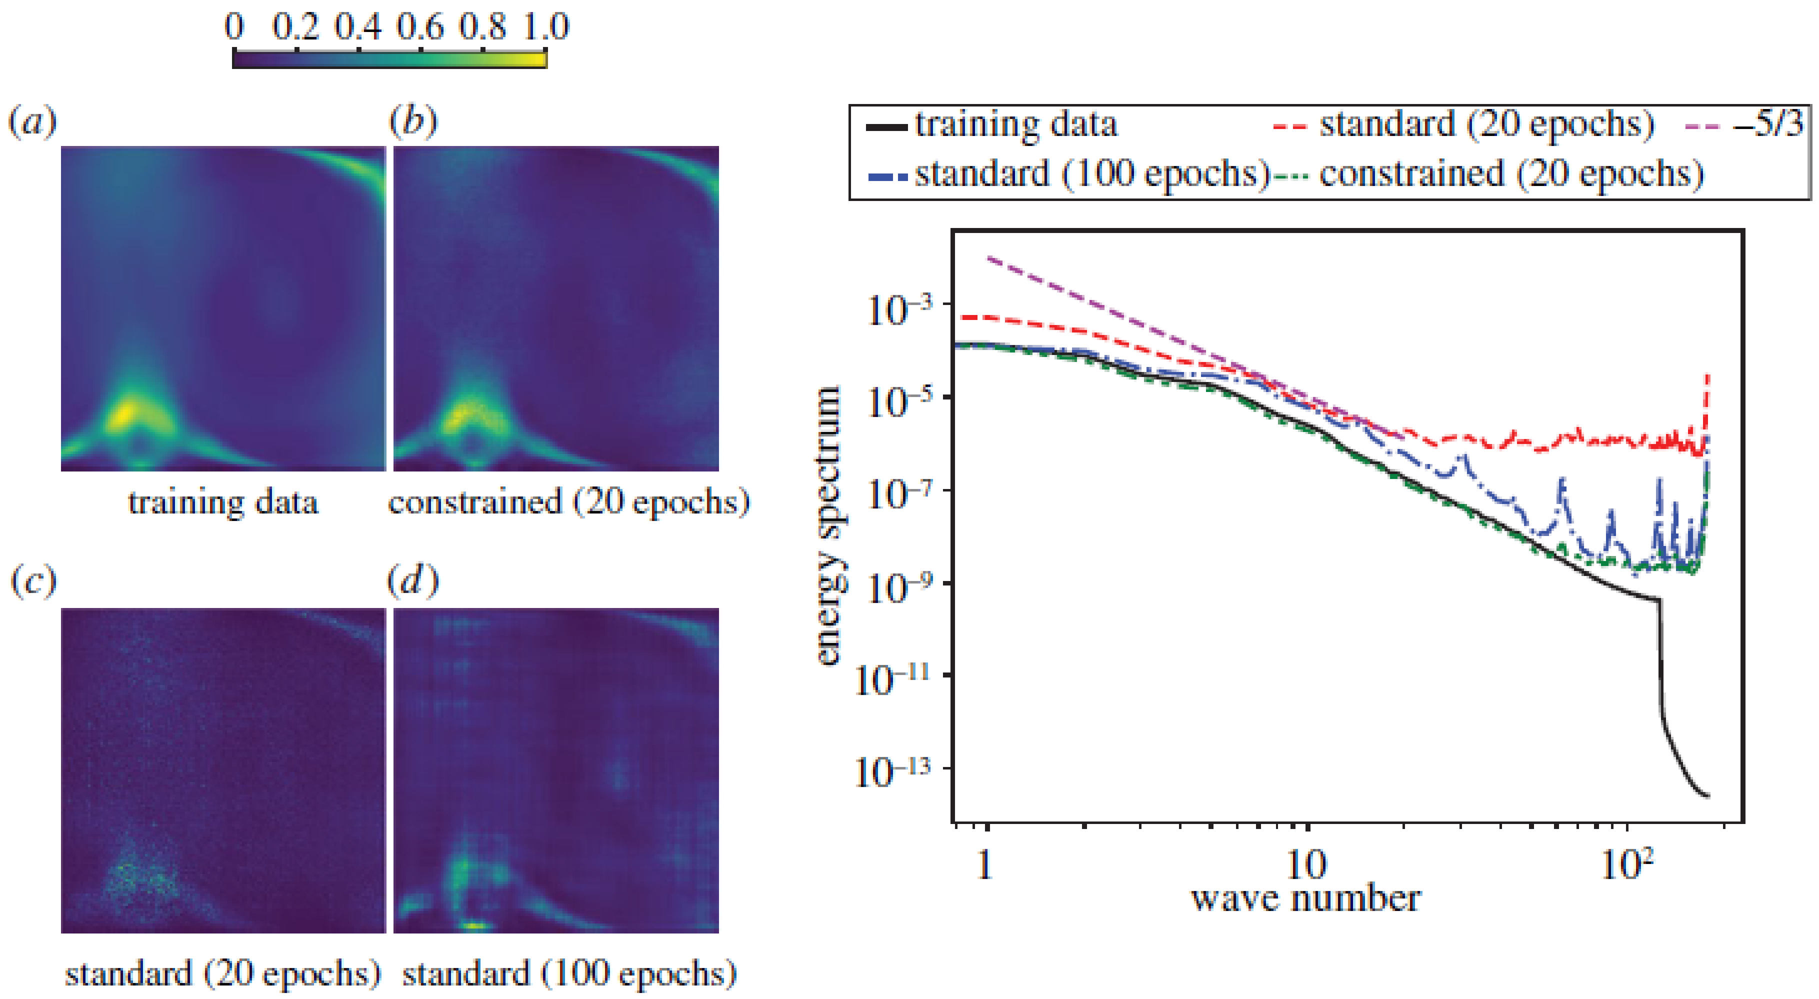

Figure 2. A comparison between the training data (truth), a standard GAN trained up to 20 epochs and 100 epochs, and the constrained GAN trained up to 20 epochs. Left: (*a*–*d*) time-averaged turbulent kinetic energy fields over a square spatial domain of size 256 × 256. Right: turbulent kinetic energy spectra. The −5/3 line is predicted by theory. The constrained GAN captures the spectrum at all except the highest wavenumbers, i.e. the finest scales of the flow. Figure reproduced from [18]. (Online version in colour.)

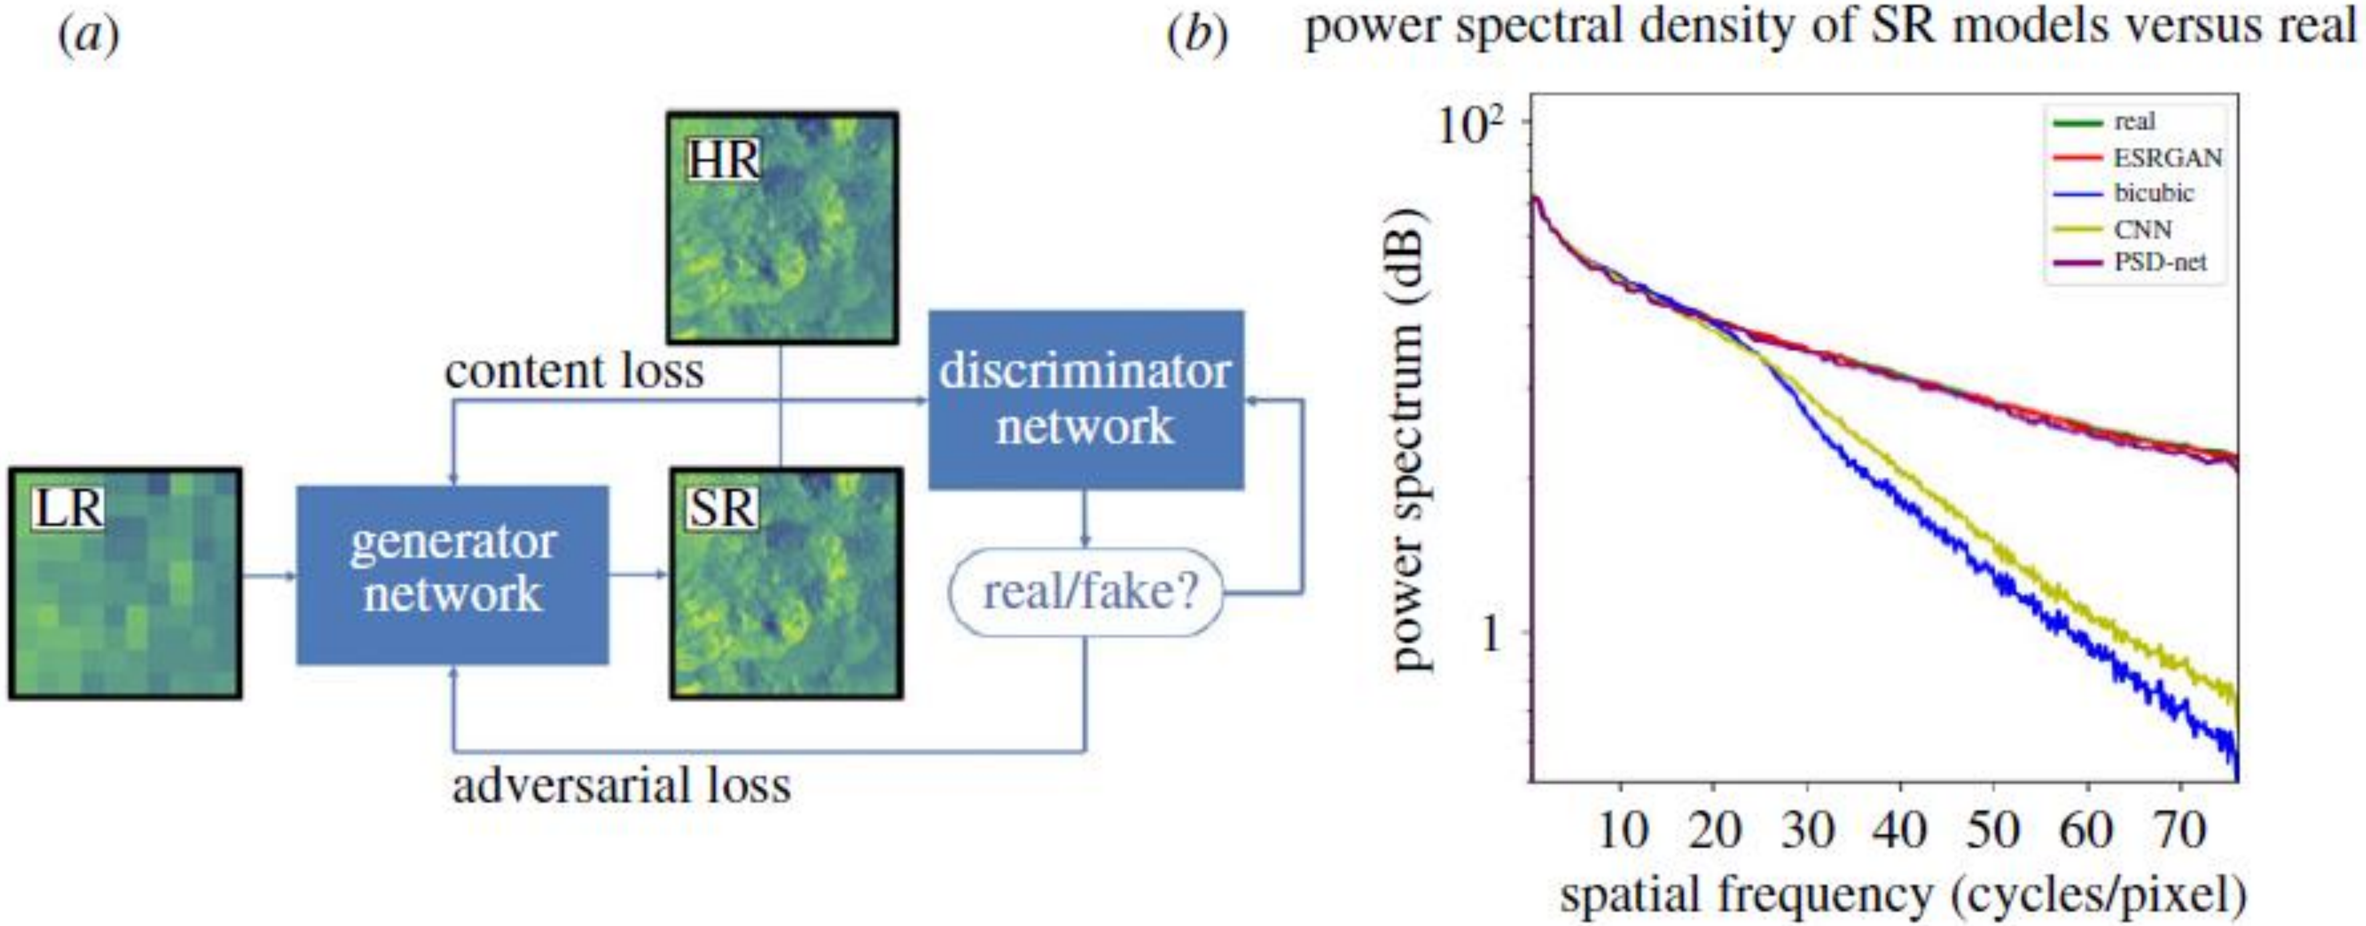

Figure 6. (*a*) Schematic of GAN for SR. Figure reproduced from [97]; (*b*) Power spectral density (PSD) plot of SR methods compared. Figure reproduced from [103]. (Online version in colour.)

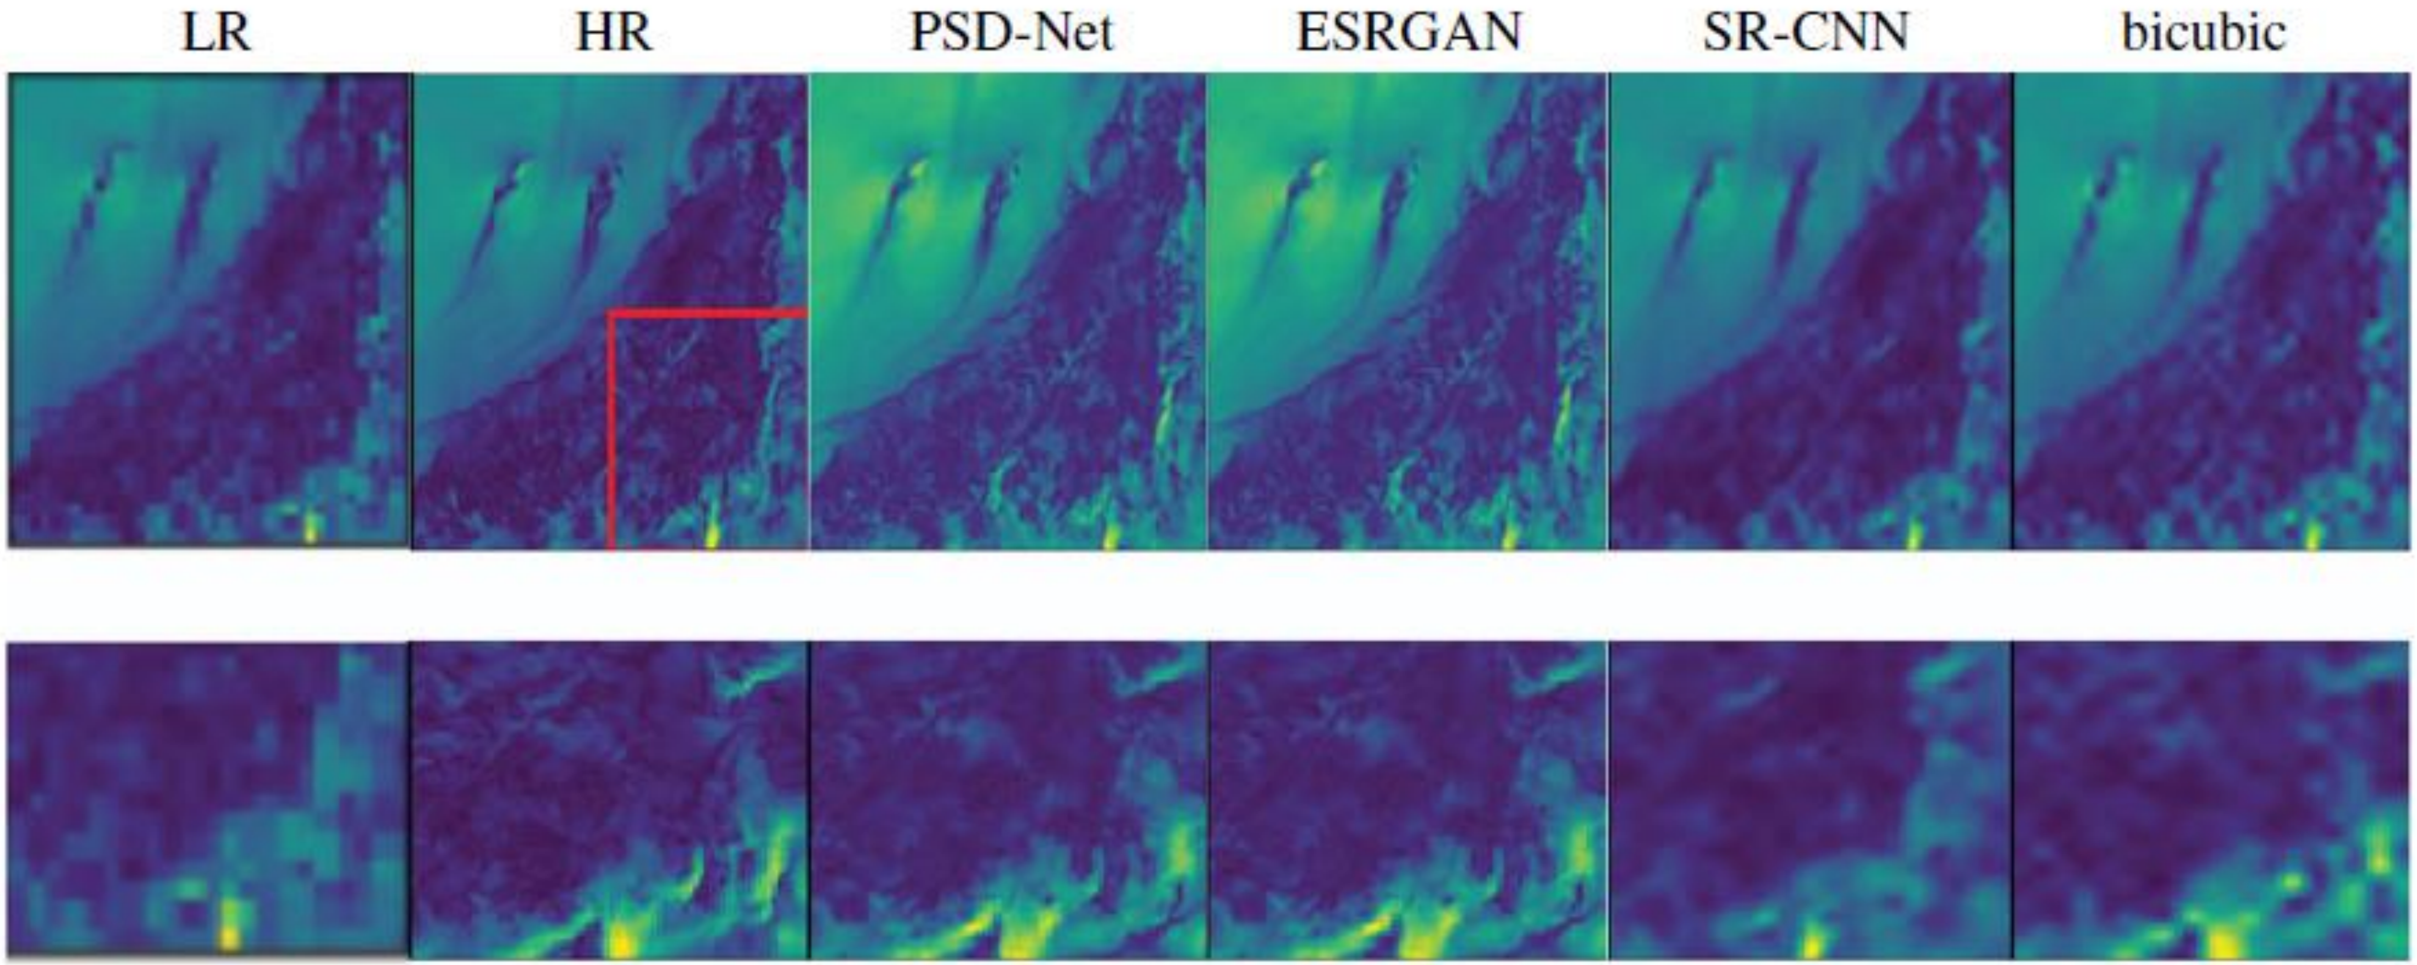

Figure 7. Comparison of low-resolution (LR), high-resolution ground truth (HR), and generated SR outputs from PSD-Net, ESRGAN, SRCNN and bicubic upsampling. The lower panel corresponds to the area of the red box in the upper panel. Figure reproduced from [103]. (Online version in colour.)

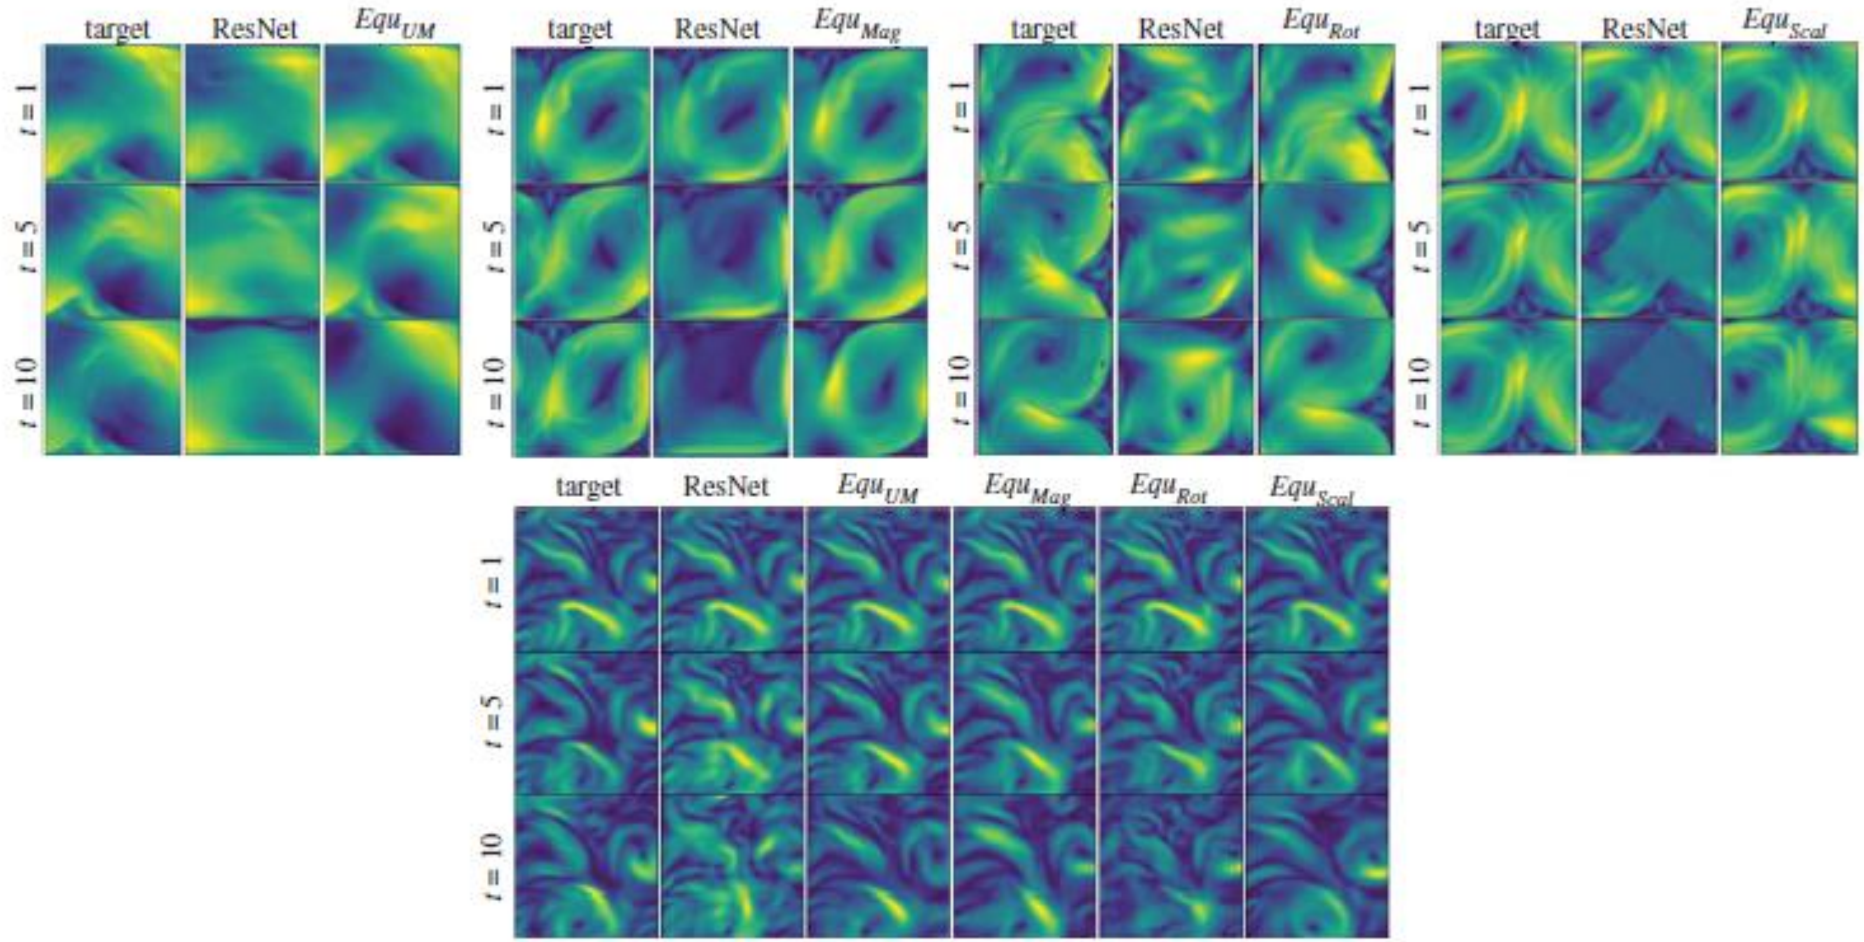

Figure 13. Comparison between performance of equivariant (Equ) and non-equivariant ResNet models for RBC velocity fields. From left to right are equivariant models under uniform motion, magnitude, rotation, and scale equivariance transformations. Tests are on future times, $t = 1, 5$ and 10. Bottom: Comparison between performance of equivariant (*Equ*) and non-equivariant ResNet models for ocean currents. *Equ* columns are equivariant models under uniform motion, magnitude, rotation, and scale equivariance transformations. No single equivariant model captures the target accurately; however, all equivariant models perform better than the non-equivariant baseline. Figure reproduced from [30]. (Online version in colour.)