Mathematics and Artificial Intelligence for Weather and Climate

Week Two

Abstracts

10–13 November 2026

November 10, 2026

Machine Learning–driven Advances in Geophysical Data Assimilation

Marc Bocquet

CEREA, ENPC, EDF R&D, Institut Polytechnique de Paris, Île-de-France, France

Machine learning (ML), particularly deep learning (DL), is playing a growing role in geophysical data assimilation (DA), a cornerstone of numerical weather prediction that combines observations with model predictions to estimate the state of a system. ML can enhance existing methods, complement them, or potentially replace parts of the DA cycle.

This talk reviews recent developments and promising directions for integrating ML into DA to improve forecasting in the geosciences. ML can provide efficient approximations of complex physical processes, correct model errors, and help combine information from models and observations. In ensemble DA, which uses multiple forecasts to represent uncertainty, ML can improve how these forecasts are generated and how DA methods are tuned. It also opens the possibility of learning how to update forecasts with observations, or even learning the entire DA and forecast cycle.

I will illustrate these opportunities with two examples: using DL to discover new and efficient ways to combine forecasts with observations, and embedding generative AI within classical DA methods.

Merging Data Assimilation and Machine Learning at Various Degrees: Examples from Arctic Sea Ice and Radiative Transfer

Alberto Carrassi

University of Bologna, Italy

Data assimilation (DA) and machine learning (ML) share a common mathematical ground: both seek to infer the state of a system, or the rules that govern it, from incomplete and noisy data. Their convergence has become a defining shift across Earth-system science. In this talk I will discuss how DA and ML can be merged at increasing degrees of integration, with examples from Arctic sea ice, coupled atmosphere–ocean dynamics and atmospheric radiative transfer.

At a first level, DA trains ML. Analysis increments carry information about the systematic error of the forecast model, and a neural network trained on them can learn to correct it. I will show how this hybrid approach improves forecast skill and attractor reconstruction in a coupled atmosphere–ocean model, substantially reduces errors in the Icepack sea-ice column model, and, embedded in a 3D-Var global ocean–sea-ice

reanalysis, clearly improves sea-ice thickness in the central Arctic.

At a second level, ML replaces components of the DA system. I will present AI-neXtSIM, a fully data-driven sea-ice emulator, and discuss two known limitations of deterministic emulators: the damping of small, fast scales in the power spectrum and an under-dispersion that makes ensembles collapse. I will show how physics-like hyperparameters, playing the role of the air–ice drag coefficient, can be identified in the network and used to build ensembles. On the observation side, I will introduce an ML emulator of a radiative transfer model for ESA's forthcoming far-infrared FORUM mission, whose errors lie largely within the instrument noise at a fraction of the computational cost.

At a third level, ML reshapes the DA algorithm itself. An ensemble transform Kalman filter operating in the latent spaces of two variational autoencoders can respect the non-Gaussian, manifold-constrained statistics of sea-ice models, while transfer learning lets it adapt to non-autonomous conditions.

I will close with a perspective rooted in dynamical systems theory. To be trusted under a changing climate, AI emulators must be faithful not only to the mean state but to the chaotic properties of the system: Lyapunov spectra, power spectra, intermittency and bifurcations. These are precisely the properties that shape weather extremes, making them central to any use of AI for extreme-event prediction. Using the Lorenz-96 and Kuramoto–Sivashinsky models, I will argue that structural instability and rare regime shifts challenge extrapolation and transfer learning, and that dynamical precursors such as covariant Lyapunov vectors should inform how emulators are trained.

Downscaling Regional Forecasts to Operational Decisions: Weather Impact Analysis for Infrastructure Resilience

Ronald Chan

*A*STAR, Singapore*

Weather is a critical driver of risk and operational performance across infrastructure sectors, including aviation route planning, maritime port operations, the safe deployment of emerging alternative fuels, and urban flood resilience. In this talk, I will share how we combine numerical weather prediction, machine learning, and physics-informed models to transform weather data into actionable decision-support tools for both strategic planning and real-time operations. Through a series of case studies, I will highlight the opportunities and challenges of developing weather-informed decision-support systems that operate under uncertainty.

In aviation, we enhance guidance on convective weather hazards by integrating numerical weather forecasts with radar and satellite observations, improving hazard prediction and supporting operational decision-making. In maritime port operations, we have developed AI-enhanced systems to assess the risk of extreme winds in Singapore's port areas, translating weather forecasts into container safety metrics and operational risk assessments. We are also developing rapid meteorology-enabled capabilities for real-time assessment of toxic plume dispersion risks associated with emerging alternative fuels. These capabilities leverage observations from automatic weather stations to downscale regional weather model outputs, enabling better representation of urban-scale weather variability and localised impacts. In addition, we are exploring the use of CCTV imagery to detect rain streaks and infer local rainfall conditions.

Together, these case studies demonstrate how advances in mathematics, meteorology, and AI can enhance the value of weather forecasts by translating meteorological information into operationally relevant insights, helping to build more resilient and weather-informed infrastructure systems across a range of critical sectors.

Topological Autoencoders for Nonlinear Dynamics: A Templex-Based Approach

Gisela D. Charó

Center for Marine and Atmospheric Research, Universidad de Buenos Aires, Argentina

Characterizing strange attractors from data requires coordinate-free topological invariants, as standard metric descriptors—such as Lyapunov exponents or fractal dimensions—cannot fully capture the system dynamics and may assign similar values to topologically inequivalent attractors. While knot-and-link theory provides a topological fingerprint for three-dimensional flows, it fails in higher dimensions. Recently, Fainstein et al. (2025) demonstrated that autoencoders (AEs) trained with a flow-direction regularization loss—which forces the latent vector field to preserve orbit orientation—can reconstruct topologically equivalent attractors. However, validating these reconstructions via linking numbers remains strictly limited to 3D systems. To bridge this gap, we introduce a validation framework based on templex theory (Charó et al., 2022) to assess the preservation of attractor topology in AE reconstructions, including higher-dimensional spaces. Templex theory generalizes the knot-holder formalism to arbitrary dimensions by combining a cell complex of the underlying branched manifold with a flow-induced digraph, characterizing attractors through directed cycles organized around joining loci (JL) where distinct branches merge. Generatexes encode the equivalence classes of these directed cycles, defining directed topological invariants applicable in any dimension. Our method bypasses the explicit construction of a cell complex by using flow-aware clustering that accounts for both

point positions and local flow directions. By varying the threshold on inter-cluster transition fractions, we build a filtration of directed graphs to identify persistent generatexes across a range of flow thresholds directly from data. We apply this framework to evaluate the AE reconstructions proposed by Fainstein et al. (2025) for the 3D Lorenz and Rössler systems, as well as a 4D seasonally forced Lorenz model (Vallis, 1986, 1988) as a toy model for El Niño–Southern Oscillation (4D-ENSO). By comparing persistent generatexes extracted from original phase-space trajectories and their latent representations, we demonstrate that the AE successfully preserves the directed topological organization across all cases. These findings establish persistent generatexes as a robust, dimension-agnostic tool for validating AE-based attractor reconstructions beyond the constraints of classical 3D knot-and-link invariants.

REFERENCES

1. Fainstein, F., Mindlin, G. B., & Groisman, P. (2025). Reconstructing attractors with autoencoders. Chaos: An Interdisciplinary Journal of Nonlinear Science, 35(1).

2. Charó, G. D., Letellier, C., & Sciamarella, D. (2022). Templex: A bridge between homologies and templates for chaotic attractors. Chaos: An Interdisciplinary Journal of Nonlinear Science, 32(8), https://doi.org/10.1063/5.0092933.

Challenges and Strategies for Data-driven Modeling of Multiscale Turbulent Systems: Focus on Climate Modeling, Forced Responses and Causality

Fabrizio Falasca

New York University, USA

Current advances in data-driven weather forecasting are motivating the development of autoregressive neural emulators for the climate system. Climate emulation, however, poses distinct challenges: useful models should be numerically stable, capture stationary statistics, and predict responses to external perturbations. Accurately capturing forced responses is essential for advancing scientific understanding as it allows us to probe causal mechanisms via "what if?" experiments.

In the first part of this talk, I will discuss conceptual limitations that arise when learning models of multiscale dynamical systems such as the Earth's climate from observational data. In particular, I will clarify how the ability of a climate emulator to reproduce perturbed statistics depends critically on (a) the choice of an appropriate coarse-grained representation and (b) stochastic parameterizations of unresolved processes. These considerations will motivate reduced-order stochastic models tailored to specific scientific questions, processes, and scales as a complementary approach to

general-purpose climate emulators.

I will then present novel strategies for constructing and validating such low-order stochastic models. First, I will introduce a discrete-time formulation of turbulent dynamical systems with exact energy-conserving nonlinearities that can be parameterized by neural networks, yielding stable physics-constrained emulators. Second, I will leverage the fluctuation-dissipation theorem (FDT), together with advances in generative modeling, to estimate responses to perturbations from unperturbed data alone, providing a new benchmark for model validation. Third, I will show that the FDT also provides causal constraints that can be optionally incorporated into neural networks loss functions. After testing the framework on idealized models of geophysical turbulence, I will present a real-world application to tropical climate dynamics using reanalysis data. The resulting response-validated stochastic model is used to investigate the long-term causal drivers of El Niño–Southern Oscillation variability by probing its response to perturbations across tropical ocean basins and variables. Taken together, these results provide a modular framework for stable reduced-order models capable of probing causal mechanisms in realistic, partially observed turbulent dynamical systems.

Attribution Science in the AI Era

Davide Faranda
CNRS, France

Every major storm, heatwave or flood now raises the same question: did climate change cause it, and by how much? Extreme-event attribution answers this by comparing the observed event in today's climate with how the same weather situation would have looked in a world without human-driven warming. This talk reviews how the field has developed and how artificial intelligence is reshaping it. We first introduce ClimaMeter, an observation-based rapid attribution framework developed at IPSL-CNRS. It builds factual and counterfactual worlds directly from ERA5 reanalysis by comparing circulation analogues from past and present periods, without a climate model in the loop. ClimaMeter produces reports within 12 to 48 hours of an event and has analysed more than 100 extremes worldwide. It also assesses the role of natural variability, changes in seasonality, the rarity of the event and population exposure. For the June 2026 western European heatwave, the analysis shows that human-driven climate change made similar conditions up to 2.5°C warmer, with about 327 million people exposed. We then move from hazards to impacts, using the August 2023 heatwave in France as a case study. Climate change raised electricity demand and peak load while reducing wind and nuclear output, and solar gains were too small to compensate. The result was a compound stress on the power grid that the hazard alone would not reveal. Finally, we discuss where machine learning enters attribution:

emulation, downscaling, generative counterfactuals and the diagnosis of hazards that physical models do not resolve. As an example, we show that a U-Net trained on ERA5 predictors and lightning observations outperforms simpler statistical models. We conclude that AI can bring speed, local relevance and new hazards into attribution, but only interpretable and physically consistent models should be part of an attribution chain that must remain transparent and trustworthy.

Multivariate Extremes in Dynamical Systems

Ana Cristina Moreira Freitas
University of Porto, Portugal

We present a theory for multivariate extreme value analysis of dynamical systems. Namely, we provide conditions adapted to the dynamical setting that enable the study of dependence between extreme values of the components of d-dimensional observables evaluated along the orbits of the systems. We study this cross-sectional dependence, which results from the combination of spatial and temporal dependence structures. We give several illustrative applications, where concrete systems and sources of dependence are introduced and analysed.

Acknowledgements: The author was partially supported by CMUP, member of LASI, which is financed by national funds through FCT – Fundação para a Ciência e a Tecnologia, I.P., under the projects with reference UID/00144/2025 and associated DOI given by https://doi.org/10.54499/UID/00144/2025.

Clustering Profiles of Extremes in Dynamical Models

Jorge Freitas
University of Porto, Portugal

Dynamical systems are very useful for modelling meteorological phenomena, where heavy-tailed extremes may exhibit different clustering profiles. We consider specially designed decorated point processes, capable of describing and recording these clustering profiles. These tools are then used to gain further statistical insight, including the study of enriched functional limit theorems and large deviations for sums of heavy-tailed observables.

Detecting Transient Metastability and Optimising System Perturbations from Observations

Gary Froyland

UNSW Sydney, Australia

In the first part of this talk I describe how to identify the emergence (and disappearance) of transient almost-invariant sets using a time-expanded Fokker-Planck generator. The leading eigenvectors of this inflated generator capture the full spacetime evolution of the almost-invariant sets. I will illustrate this with the tracking of the onset and decay of atmospheric blocks, and the continual emergence and disappearance of 3D turbulent superstructures in Rayleigh-Benard convection. In the second part of the talk I describe how to apply the operator-based methodology of optimal linear response to systems that are only partially observable. Given a user-specified objective that the user desires to maximally increase or decrease, one may explicitly determine the optimal perturbation to maximally move this objective in the desired direction. Crucially, these optimal perturbations may be obtained in the original observation space or in any other desired observed coordinates. An initial illustration is given for the Lorenz system (with full observation), and then we apply our techniques to optimally perturb the decorrelation time and cycle period of the El-Nino Southern Oscillation.

How Surprising are Rare Climate Extremes?

Melinda Galfi

Vrije Universiteit Amsterdam, Netherlands

From a human perspective, extreme weather and climate events are always surprising. However, the statistical and physical characteristics of different types of extremes can vary substantially, as reflected in the three distinct tail behaviours of limiting extreme-value distributions. In the climate system, extremes on large spatial scales frequently exhibit strong similarities. For example, intense heatwaves in the mid-latitudes tend to share common spatio-temporal evolution and underlying circulation patterns. We refer to this property as the "typicality of extremes". Notably, this resemblance tends to increase with event intensity, a feature that can be interpreted within the framework of large deviation theory. At smaller spatial scales, by contrast, extremes display greater heterogeneity and complexity, often associated with fat-tailed distributions. In this presentation, I first discuss these characteristics in a steady-state climate. I then examine how global warming may alter the behaviour and typicality of extremes and discuss the implications of these properties for the training of AI models.

Multimodal Spatiotemporal Atmospheric Data Assimilation with Latent Video Flow-matching

Romit Maulik

Purdue University, USA

Data assimilation (DA) uses Bayesian inference to update the state of a numerical forecast model with observed data. In this study, we propose a fundamentally different, unified approach to atmospheric data assimilation. We use latent video flow-matching to sample temporally consistent trajectories from a prior trained using ERA5 reanalysis (69 variables over an 8-day window). We also use posterior sampling to assimilate real observation sources, such as those from the NOAA Integrated Global Radiosonde Archive and the Integrated Surface Database. Because the prior generates a continuous trajectory, it naturally propagates information between observed and unobserved frames. Therefore, we can perform various DA tasks, such as filtering and smoothing, simply by changing the observed frames. Moreover, we generate full-state ensemble forecasts directly from sparse observations, achieving performance competitive with state-of-the-art observation-to-forecast models.

Data-driven Meteorological Forecasts of Extremes in a Changing Climate

Gabriele Messori

Uppsala University, Sweden

Weather forecasts are often evaluated on average meteorological skill. However, what is most valuable to society are skillful forecasts of extreme events. Such events are by definition rare, and our rapidly changing climate also makes unprecedented events more likely than they would otherwise be. Both these aspects make the data-driven forecasting of extreme events challenging. Moreover meteorological skill, even when restricted to extreme events, is not necessarily a good indicator of the value of a forecast for any given user. In this presentation, I discuss the evaluation of data-driven weather forecasts for extremes, the challenge posed by unprecedented events and ways to go beyond meteorological skill when evaluating whether data-driven forecasts bring added user value beyond numerical weather forecasts.

Climate Dynamics: from AI to Topology and Back

Denise Sciamarella

Centre National de la Recherche Scientifique, France

AI models now reproduce many statistical properties of weather and climate. Whether they correctly capture the mechanisms that generate those statistics remains an open question. Two datasets can agree in their means, variances and spectra while organizing their trajectories in phase space in different ways. This talk shows that directed algebraic topology, through the templex formalism, offers a language to pose this question, to intercompare models and data, and to assess whether a model can be understood in terms of the main nonlinear mechanisms underlying a given dynamics.

The talk examines from this perspective the AI Model Intercomparison Project (AIMIP), which brings together AI weather and climate models under a common protocol: in its first phase, AI models trained on reanalysis simulate the atmosphere over 1979–2024 under observed sea surface temperatures and sea ice. We focus on processes related to tropical phenomena that classical models have historically failed to predict or simulate accurately.

But the story does not end here: there is also the way back. Machine learning can help construct the topological objects that yield directed invariants from time-varying experimental, numerical or observational data. The templex and its stochastic counterpart, the random templex, provide a non-metric notion of chaos in finite time. These objects allow decomposing a dynamics into fundamental units held together by topological bonds, as well as defining topological modes of variability as metric manifestations of those units. Where topological tools struggle, AI can help, in a way reminiscent of Hasselmann's view of fast weather acting on slow climate. The goal of this two-way trip is clear: building a solid bridge between machine learning and human understanding of time and structure.

REFERENCES

1. Brunton, S. L., Noack, B. R., & Koumoutsakos, P. (2020). Machine learning for fluid mechanics. Annual review of fluid mechanics, 52(1), 477-508, https://doi.org/10.1146/annurev-fluid-010719-060214.

2. Charó, G. D., Letellier, C., & Sciamarella, D. (2022). Templex: A bridge between homologies and templates for chaotic attractors. Chaos: An Interdisciplinary Journal of Nonlinear Science, 32(8), https://doi.org/10.1063/5.0092933.

3. Ghil, M., & Sciamarella, D. (2023). Dynamical systems, algebraic topology and the climate sciences. Nonlinear Processes in Geophysics, 30(4), 399-434, doi: https://doi.org/10.5194/npg-30-399-2023.

4. Charó, G. D., Ghil, M., & Sciamarella, D. (2023). Random templex encodes topological tipping points in noise- driven chaotic dynamics. Chaos: An Interdisciplinary Journal of Nonlinear Science, 33(10), https://doi.org/10.1063/5.0140660.

5. Sciamarella, D., & Charó, G. D. (2024). New elements for a theory of chaos topology. In Topological Methods for Delay and Ordinary Differential Equations: With Applications to Continuum Mechanics (pp. 191-211). Cham: Springer International Publishing, doi: https://doi.org/10.1007/9.

6. Henn, B., Bretherton, C. S., Koldunov, N., Lessig, C., Molina, M. J., Arcomano, T., ... & Lopez-Gomez, I. (2026). AIMIP Phase 1: systematic evaluations of AI weather and climate models. arXiv preprint, doi: https://doi.org/10.48550/arXiv.2605.06944.

7. Sciamarella, D. (2026). Functorial invariants for chaos topology from data. Physical Review E, 114(2), 024205. https://doi.org/10.1103/zx76-5115.

Turning Weather Uncertainty into Better Airport and Airspace Decisions

Paul Tan

International Centre For Aviation Innovation, Singapore

Asia-Pacific aviation is entering a period where traffic growth and tightly coupled operations will make weather uncertainty costly. Passenger demand has moved above 2019 levels and is projected to reach 4.1 billion journeys by 2044. Yet airport capacity is not determined by the runway alone. It depends on the most constrained link across airspace, approach, runway, taxiway, apron, gates and turnaround. A convective or lightning event can therefore create holding, diversions, stopped ramp work, gate conflicts and a recovery tail that lasts far longer than the weather itself.

Singapore illustrates both the urgency and the opportunity. Lightning is a routine operational risk, occurring on nearly half the days of the year, while perimeter-based alerts and fixed cessation timers can impose broad restrictions without representing how exposure changes across space and time. Existing initiatives already move beyond general weather products toward operational impact, but a gap remains: forecasts do not translate uncertainty into usable infrastructure, sustainable movement rate, total cost and a defensible action.

The proposed shift is from forecasting the weather to forecasting the airport. This requires uncertainty to be carried through five linked questions: what the weather will do, what infrastructure remains usable, how much traffic can be handled, what the

consequences will cost and which action is most robust. The challenge is compounded by sparse observations over surrounding seas, where storms often form, and by the limits of satellite imagery in resolving the lower atmosphere and internal storm state. Better decisions will depend on integrating existing observations, targeted new sensing, operational data and models that preserve multiple plausible outcomes rather than collapsing them into one confident answer.

Success must be measured in operational terms: useful warning time, calibrated uncertainty, fewer wasted stop-minutes, narrower restrictions, reduced diversions and faster recovery. Delivering this capability will require a tested weather-to-capacity model, event-level decision records, safe replay and shadow-testing environments, and one shared picture across airports, airlines, handlers and air traffic control. The invitation to researchers and industry is practical: see disruption earlier, quantify its consequences better and test responses before the system loses flexibility, while preserving human authority and safety.

AI for Actionable Convective Weather Nowcasting

Shili Xiang
*A*STAR, Singapore*

Forecasting convective weather can provide valuable advance notice for flood preparedness, transport operations, infrastructure management and outdoor activities. AI is advancing short-term weather prediction, but operational value depends on more than forecast accuracy: predictions should capture hazardous developments, convey uncertainty and provide useful information despite gaps in observation coverage. This talk brings together four recent contributions from our group towards actionable convective weather nowcasting. First, SimCast uses short-to-long-term knowledge distillation to improve precipitation forecasts, particularly for heavy rainfall, without additional inference overhead. Second, deep evidential learning estimates forecast uncertainty in a single forward pass, while revealing limitations in representing rare, intense events. Third, GeoRad-3D separates storm motion and deformation from uncertain changes in radar echo intensity, enabling more skilful forecasts of evolving three-dimensional storm structures. Finally, SRDiff translates satellite observations into radar-like sequences that can feed existing nowcasting models, offering a pathway to extending forecasting capabilities beyond radar coverage. Drawing on benchmark evaluations and tropical weather examples, I will discuss how these advances could support more informed operational decisions and examine the remaining challenges in translating predictive skill into useful decision support solutions.

REFERENCES

1. Yin, Y., Chen, S., Li, Y., Wang, L., Jin, R., Cui, W., & Xiang, S. (2025). SimCast: Enhancing precipitation nowcasting with short-to-long term knowledge distillation. In 2025 IEEE International Conference on Multimedia and Expo (ICME) (pp. 1–6). IEEE. https://doi.org/10.1109/ICME59968.2025.11209905.

2. Wang, L., Xu, C., Yin, Y., Zhou, C. X., Heng, P., & Xiang, S. (2025). Uncertainty analysis in deep evidential learning for radar-based weather nowcasting. In 2025 IEEE Radar Conference (RadarConf25) (pp. 1–6). IEEE. https://doi.org/10.1109/RadarConf259087.2025.11205055.

3. Li, Y., Zhou, Z., Chen, S., Zhang, J., Liang, J., Chui, K. W., Chen, C., Liu, F., Yin, Y., & Xiang, S. (2026). GeoRad-3D: Factorized geometry transport and residual radiometry for 3D radar nowcasting. In the Fortieth Annual Conference on Neural Information Processing Systems (NeurIPS 2026; accepted for publication). https://neurips.cc/virtual/2026/poster/152938.

4. Qin, Y., Cao, J., Wang, T., Yin, Y., Li, L., Xiang, S., Zhang, Y., & Zimmermann, R. (2026). SRDiff: A cross-modal diffusion model for satellite-to-radar translation in precipitation nowcasting. IEEE Transactions on Geoscience and Remote Sensing, 64, Article 4106416. https://doi.org/10.1109/TGRS.2026.3686188.

Beyond Extremes: a Few Thoughts on Black Swans and Aliens, where Statistics and Physics Meet

Pascal Yiou

Laboratoire des Sciences du Climat et de l'Environnement, France

Climate extremes are intensifying globally. Our ability to anticipate unprecedented events is limited by the computational burden of sampling distribution tails of variables and the limited number of observations. Solving this fundamental discrepancy has led to several research alleys. This presentation will explore some of them. The overarching goal is to understand the gap between known observed records, and potential (unobserved) most extremes. The notion of "black swans" has been popularized by Taleb (2010). Black swans can be defined by observations that escape our empirical understanding. I propose a pragmatic statistical way of defining black swans, and argue that such unexpected events can be fairly frequent. The simulation of climate black swans can be achieved through the use of rare event algorithms (Noyelle, 2024). Such approaches seem to allow extrapolating empirical scaling laws (Lancelin et al., 2026; Miloshevich et al., 2024). But nothing guarantees that the obtained extremes are the most extremes. If we assume that climate dynamical variables (i.e., the variables of the primitive equations (Peixoto and Oort, 1992)) evolve

around a compact (strange) attractor (Temam, 1988), then those variables are bounded. I will illustrate why this is true for temperature, and how to determine an upper bound from physical principles. I will discuss the implication of the existence of bounds and the fact that records are far from those bounds (for temperature). I will introduce "aliens", which are events that have never been reached, but are physically possible. This offers a new paradigm that complements black swans.

REFERENCES

1. Lancelin, A., Wikner, A., Dubus, L., Le Priol, C., Abbot, D. S., Bouchet, F., Hassanzadeh, P., and Weare, J.: AI-boosted rare event sampling to characterize extreme weather, *Physical Review Letters*, 137, 064201, 2026.

2. Miloshevich, G., Lucente, D., Yiou, P., and Bouchet, F.: Extreme heat wave sampling and prediction with analog Markov chain and comparisons with deep learning, *Environmental Data Science*, 3, e9, https://doi.org/10.1017/eds.2024.7, 2024.

3. Noyelle, R.: Statistical and dynamical aspects of extreme heatwaves in the mid-latitudes, PhD Thesis, Univ. Paris-Saclay, 289 pp., 2024.

4. Peixoto, J. P. and Oort, A. H.: Physics of climate, American Institute of Physics, New York, xxxix, 520 p. pp., 1992.

5. Temam, R.: Infinite-dimensional dynamical systems in mechanics and physics, Springer-Verlag, New York, xvi, 500 pp., 1988.

Climate Downscaling in Southeast Asia Using Generative AI Methods

Xiaoling Zhou

Centre for Climate Research Singapore, Singapore

Artificial Intelligence (AI) has emerged as a promising approach for enhancing regional climate projections through efficient, high-resolution climate downscaling. This presentation introduces our AI downscaling strategy for Southeast Asia (SEA), focusing on the development of generative AI models for regional climate applications. An overview of the datasets, model architecture, and evaluation framework will first be presented, followed by the implementation of AI downscaling models based on the NVIDIA CorrDiff framework. The trained models are developed to downscale coarse-resolution global climate models (GCMs) to finer spatial resolution, with a particular focus on daily precipitation and daily maximum temperature. Model performance is evaluated using a range of statistical and spatial verification metrics to assess the representation of daily total precipitation, spatial structures and climate

change signals. Initial results demonstrate that the AI models can effectively reconstruct fine-scale climate features while maintaining good agreement with the reference datasets. Finally, the presentation discusses the remaining challenges for AI-based climate downscaling, particularly the temporal consistency and transferability of AI models across different GCMs. Future work on improving model generalization and robustness to support regional climate projection applications in Southeast Asia will also be outlined.
