Mathematics and Artificial Intelligence for Weather and Climate

Week One

Abstracts

02-06 November 2026

November 6, 2026

ATEMPO (Amortized Transfer Entropy Model Per-lag Over-time)

Emma Andrews

National University of Singapore, Singapore

Transfer entropy (TE) measures how much the past of one variable improves predictions of another. It is a natural tool for asking how parts of the climate system are coupled. In practice, TE is estimated one lag and one time window at a time, with one model fitted per configuration. That makes it expensive and noisy to ask the questions that matter in climate: which regions are coupled, at what delay, and whether the coupling depends on the state of the system.

I'll present ATEMPO (Amortized Transfer Entropy Model Per-lag Over-time), a diffusion-based estimator that learns a single model across source lag, conditioning depth and position in time. From one fit, ATEMPO gives the whole TE surface. Differencing that surface across lags gives a delay profile: how much information each lag adds that deeper lags don't already provide.

I then apply ATEMPO to boxes spanning the Pacific Walker circulation, identifying which regions are coupled and at what delay, and comparing the recovered delays with established methods such as PCMCI. I also show how the coupling changes with the state of the system, and how the approach extends to spatiotemporal fields.

Preparations for Next-Generational AI- and Physical Model-Based Regional Weather Prediction and Climate Projections at CCRS

Dale Barker

Centre for Climate Research Singapore, Singapore

Singapore is a small-island state situated in the deep tropics, approximately one degree north of the equator and 104 degrees longitude. Singapore's weather is characterized by generally warm, humid conditions with extreme weather hazards on hourly to daily timescales dominated by heavy rainfall events from localized convection, high heat stress, and transboundary haze.

The Centre for Climate Research Singapore (CCRS) is the research division of Singapore's National Meteorological Service - the Meteorological Service Singapore (MSS). Research at Singapore spans a wide range of timescales, from the past 50 years (regional reanalysis), through the next few hours (nowcasting), next 1-2 days (short-range local weather prediction), subseasonal to seasonal prediction, out to

regional climate change projections out to to 2100 and beyond.

CCRS has traditional employed the 'Unified Model' (UM)-based 'SINGV' system, initially for local km-scale NWP, but more recently for dynamically downscaled 'V3' regional climate change projections and digital twin scenarios of Singapore's future urban landscape. Since 2024, CCRS has been actively engaged with national and international partners to replace the physical UM-based SINGV system with next-generation 'SINGV_NG' modelling capabilities combining both physical and AI-based modelling systems. CCRS is investigating the use of local fine-scale data (e.g. observations, SINGV NWP, V3 projection data) to fine-tune global AI models for both AI-based local NWP and climate downscaling. In addition, a range of off-the-shelf AI-based global models (AIFS, Graphcast, ForecastNet, Aurora, etc.) are being evaluated within CCRS's regional 'WeatherBench'.

This talk with provide an update on the latest CCRS's efforts to develop and apply next generation physical and AI-based approaches suitable for local weather and climate science applications in the deep-tropics.

Understanding Predictability, Physical and Dynamical Processes to inform Data-driven Weather and Climate Prediction Models R&D

Gilbert Brunet

Environment and Climate Change Canada, Canada

Extreme weather events are increasingly producing unprecedented floods, droughts, wildfires, and ecosystem disruptions across spatial and temporal scales. Rapid urbanization and the evolving global economy have amplified the need for accurate climate-change projections and improved numerical weather prediction, particularly for disruptive—and occasionally beneficial—weather and climate phenomena. Meeting this demand requires continued progress in understanding atmospheric predictability, physical processes, and dynamical systems, especially as artificial intelligence (AI) becomes more deeply integrated into prediction science. A central challenge is to identify the physical and dynamical processes that govern the predictability of diverse weather and climate events so that AI-based prediction systems can be trained in a scientifically grounded and optimal manner. Our discussion will highlight several key examples, ranging from predictors of tropical-cyclone genesis to diagnostics of planetary-scale climate extremes.

At intermediate time scales, notably sub-seasonal to seasonal (S2S) prediction, we show how data-driven weather and climate models can be significantly optimized through targeted training strategies. The S2S variability of the global atmosphere has

been extensively characterized over recent decades. A prevailing paradigm decomposes this variability into two normal-mode components—quasi- or discrete modes and continuous spectra—interpreted as perturbations to a time-mean zonal flow (Brunet 1994; Zadra et al. 2002a,b). Discrete modes represent transient phenomena such as storms and teleconnections, while the continuous spectrum acts as an atmospheric "skeleton" that modulates large-scale circulation. Under this framework, the discrete spectrum encompasses familiar S2S variability patterns such as the Arctic Oscillation and the Madden–Julian Oscillation.

We demonstrate that this atmospheric decomposition, known as empirical normal modes (ENMs), forms a complete and unique Hamiltonian canonical basis. This canonical property enables substantial dimensionality reduction of the atmospheric evolution operator, offering a powerful pre-conditioning strategy for data-driven S2S prediction. We outline how ENMs can be computed using spatial covariance methods and discuss their theoretical equivalence to POP/LIM approaches based on time-lag covariances. Their predictive potential is illustrated through statistical regression and AI-based techniques following Brunet and Vautard (1996).

Finally, we will address atmospheric waviness, whose dynamical interpretation has remained ambiguous. Using potential vorticity (PV) as a materially conserved tracer, we formulate waviness in terms of the Rossby restoring-force mechanism and derive a PV-based measure of meridional displacement. Applied to 40 winters of ERA-5 data, this framework reveals robust trends in the restoring-force parameter γ_0 and links its variability to atmosphere angular momentum (AAM), ENSO, the Arctic Oscillation, and the Pacific–North American pattern. A weakened background PV gradient at high latitudes explains increased trend in meridional air-mass excursions and recent temperature extremes (Brunet, Lin and Davies, 2026).

S2S-JEPA: Predicting the Predictable at Subseasonal-to-Seasonal Timescales

Chenyu Dong

National University of Singapore, Singapore

Subseasonal-to-seasonal (S2S) prediction, from two weeks to two months ahead, remains the 'predictability desert' of weather and climate forecasting. Recent AI weather models excel up to two weeks ahead but deteriorate beyond, largely because they are trained to predict fine-scale details that are neither predictable nor essential at S2S timescales. We argue that a more physically grounded objective is to forecast only the slowly varying components that remain predictable. Computer vision reached the same conclusion with the Joint-Embedding Predictive Architecture (JEPA), which predicts in latent space, discarding unpredictable details. In this talk, I will present

S2S-JEPA, which brings the JEPA paradigm to S2S forecasting, tailored to the task through design elements from state-of-the-art AI weather models. Evaluated against the ECMWF extended-range ensemble, S2S-JEPA achieves comparable skill at weeks 5 to 6 and surpasses it on multiple metrics, while remaining behind at shorter leads. I will close with what the ablations tell us about where the skill comes from, and what this implies for AI-based S2S prediction.

Complexity science and AI Methods for Earth System Tipping Points

Jingfang Fan

Beijing Normal University, China

Critical phenomena are common in nature and society, from phase transitions and extreme weather to sea-ice loss, ecosystem degradation, and financial fluctuations. When complex systems approach a critical threshold, they often show early-warning signals such as critical slowing down, stronger fluctuations, expanding correlations, and reduced resilience. The Earth system is a coupled complex system involving the atmosphere, ocean, land, biosphere, and cryosphere. Local disturbances may be amplified and trigger cascading changes across regions and components. This talk will introduce key ideas from statistical physics, nonlinear dynamics, complex systems science, and artificial intelligence, including phase transitions, percolation, self-organized criticality, critical slowing down, complex networks, and data-driven prediction. Typical examples such as extreme weather and climate tipping points will be used to show how these methods help reveal mechanisms, improve early warning, and support climate-risk management.

Critical Transitions in Complex Dynamical Systems: Theory and Applications to Ecosystems and Climate

Ulrike Feudel

University of Oldenburg, Germany

Many systems in nature are characterized by the coexistence of different stable states for a given set of environmental parameters and external forcing. Examples for such behavior can be found in different fields of science ranging from mechanical or chemical systems to ecosystem and climate dynamics. Perturbations, applied to those natural systems can lead to a critical transition from one stable state to another. Such critical transitions—also called tipping phenomena—can happen in various ways: (1) due to bifurcations, i.e. changes in the dynamics when external forcing or parameters

are varied extremely slow, (2) due to fluctuations which are always inevitable in natural systems, (3) due to shocks or extreme events, and (4) due to rate-induced transitions, i.e. when external forcing changes too fast compared to the ability of the forced system to follow the changes.

We discuss these critical transitions and their characteristics and illustrate them with examples from mechanical and natural systems. Moreover, we relate them to the concept of resilience developed in ecology. Special emphasis is given to non-autonomous systems, in which we highlight the interplay between different time scales, like the dissipative time scale and the time scale of the variation of parameters or forcing. Moreover, we discuss the role of unstable states, that are not directly observable in nature, but nevertheless act as the organizing centers of the dynamics.

Thermodynamically Consistent Modeling of Geophysical Fluids: A Variational Perspective

François Gay-Balmaz

Nanyang Technological University, Singapore

We present a Lagrangian variational framework for the modeling of single- and multicomponent geophysical fluids, encompassing both reversible (entropy-conserving) and irreversible (entropy-generating) processes. The framework naturally incorporates moisture and extends to semi-compressible fluid models relevant to atmospheric and oceanic dynamics. By construction, the resulting equations satisfy the first and second laws of thermodynamics, providing a systematic route to thermodynamically consistent models. We also discuss recent developments in using this variational structure to construct thermodynamically consistent numerical discretizations for geophysical fluid systems.

Global warming & global weirding: A nonautonomous dynamical systems perspective

Michael Ghil

University of California, Los Angeles, USA

The phenomenon of global warming is by now scientifically well established. The important aspects of "how, how much and when" are being explored quantitatively, as well as qualitatively, and extra-scientific objections persist. This aspect of the climate sciences refers, however, mostly if not exclusively, to mean global near-surface atmospheric temperatures and to those other climatic and environmental phenomena

that can be linked in a more-or-less straightforward manner to it. A large number of phenomena associated with regional extrema of temperature and precipitation, such as strong and persistent floods and droughts or cold and warm extrema, are still quite puzzling and often exploited in extra-scientific objections to global warming.

The theory of nonautonomous dynamical systems (NDSs) has by now been given considerable play in studying the effects of anthropogenic forcing on global warming. More recently, NDS theory has been shown to be a powerful methodology in going beyond the effect of time-dependent forcing on means and on to full distributions of equivariant measures. This review talk will cover some of the results in the above-mentioned directions for deterministic as well as random forcing. A possible role for AI in elucidating these phenomena will be mentioned as well.

The theory of nonautonomous dynamical systems (NDSs) has by now been given considerable play in studying the effects of anthropogenic forcing on global warming. More recently, NDS theory has been shown to be a powerful methodology in going beyond the effect of time-dependent forcing on means and on to full distributions of equivariant measures. This review talk will cover some of the results in the above-mentioned directions for deterministic as well as random forcing. A possible role for AI in elucidating these phenomena will be mentioned as well.

References

1. Ashwin, P., U. Feudel, M. Ghil, 2026: Introduction to the Double Focus Issue: Nonautonomous dynamical systems in the sciences, *Chaos: An Interdisciplinary Journal of Nonlinear Science*, **36**, chosen by the Editors as Featured Article, 080401, https://doi.org/10.1063/5.0333561.

2. Chekroun, M. D., E. Simonnet, and M. Ghil, 2011: Stochastic climate dynamics: Random attractors and time-dependent invariant measures, *Physica D*, **240**(21), 1685–1700, https://doi.org/10.1016/j.physd.2011.06.005.

3. Ghil, M., and V. Lucarini, 2020: The physics of climate variability and climate change, *Rev. Mod. Phys.*, **92**(3), 035002, https://doi.org/10.1103/RevModPhys.92.035002.

4. Ghil, M., A. Haraux, J. Roux, and N. Bodnariuk, 2026: *Mathematics for the Environment*, Springer Science & Business Media, 900 pp., in final stages of preparation, to appear in 2 vols., *Mathematical Tools & Applications to Climate and Ecology*.

5. Maraldi, B., H. A. Dijkstra, and M. Ghil, 2025: Intraseasonal atmospheric variability under climate trends, *Chaos*, **35**(5), 053136, https://doi.org/10.1063/5.0253103.

6. Walsh, B., 2014: Forget global warming and get ready for global weirding. *TIME Magazine*.

Learning Macroscopic and Mesoscopic Dynamics

Qianxiao Li

National University of Singapore, Singapore

We discuss some recent work on constructing stable and interpretable macroscopic and mesoscopic dynamics from trajectory data using deep learning. We adopt a hypothesis-driven approach: contrary to using generic neural networks as functional approximators for the vector fields driving a dynamical system or operator based learning approaches, we start with a structured hypothesis class in which well-posedness and numerical stability are a priori ensured. Learning on trajectory data then identifies a specific member of this family suitable for capturing the observed dynamical process, from which scientific predictions and analysis can be made. We demonstrate this approach on a variety of scientific problems, including the dynamics of polymers, magnetic systems and diffusion/mixing problems.

In-Context Operator Networks (ICON): A Research Program for Numerical Intelligence

Yang Liu

National University of Singapore, Singapore

Intelligence is commonly understood as the ability to acquire knowledge, adapt to unfamiliar situations, and solve new problems. Large language models exhibit this capacity by inferring task-relevant knowledge from textual context and applying it to new tasks. Yet intelligence need not be confined to language. Scientific and social systems often reveal themselves numerically before we can fully describe them in words. We call the ability to acquire and apply knowledge from numerical context "numerical intelligence" and view it as a foundational pillar of artificial intelligence alongside linguistic intelligence. We proposed "In-Context Operator Learning" paradigm and the corresponding model "In-Context Operator Networks" (ICON) for numerical intelligence. We will show how a single ICON model (without fine-tuning) manages multiple families of PDE problems. In our latest work, we trained Unified ICON (UNICON) on real-world datasets across different disciplines, including weather, hydrology, power, and traffic. And it generalizes to completely novel disciplines, e.g. web activity. Apart from the model itself, we can also design "harness" for ICON due to the flexibility of the context. We will introduce building blocks, including "chain of

operators" and "contextual ensemble learning". With such techniques, we outperformed state-of-the-art (SOTA) expert models in disciplines unseen in training. See more details on the website: https://scaling-group.github.io/

From Physics to AI and back

Gianmarco Mengaldo

National University of Singapore, Singapore

In this talk, we explore how we can use pre-existing knowledge (e.g., physics) to improve AI systems, and how we can possibly extract some knowledge from AI systems. On the first topic, we present a novel physics-enhanced deep learning hybrid method, namely CondensNet, for resolving cloud physics in climate models. On the second topic, we present some results on the use of explainable AI for climate applications. We conclude with an overall perspective bridging the two topics.

From Trajectories to Climate Learning and Forecasting with Reservoir Computing

Juan-Pablo Ortega

Nanyang Technological University, Singapore

Reservoir computing provides a state-space approach to learning dynamical processes from observations. This talk reviews the echo state property, fading memory, and universality, and explains how generalized synchronization and embedding results connect reservoir representations with the underlying dynamics.

We then consider autonomous forecasting, distinguishing individual trajectory prediction from the evolution of distributions and long-term statistics. Recent results relate finite-horizon trajectory errors to approximation accuracy and the instability of the learned dynamics, and identify conditions for accurate forecasts of a system's statistical climate. Numerical examples illustrate how trajectories can separate while their distributions remain close. The talk emphasizes the main mathematical ideas and their relevance to weather and climate modeling.

This talk is based on joint work with Lukas Gonon, Lyudmila Grigoryeva, Allen Hart, James Louw, and Florian Rossmannek.

Persistent anticyclones, AI and quantum computing: Why We Need a 'CERN for Climate Change'

Talk is delivered virtually

Tim Palmer

University of Oxford, UK

At the thermodynamic level, climate change is simple 19th Century physics. Predictions of warmer wetter winters and hotter drier summers draw on such thermodynamic ideas. However, these do not explain what happened this summer, where drought, heatwaves and fire spread across large parts of Europe. The proximate cause of such meteorological conditions was a persistent anticyclone positioned over Western Europe. Were such anticylonic conditions made more likely by climate change? To answer this question requires us to understand not just the thermodynamic aspects of climate change, but the dynamic feedbacks too. This changes climate change prediction from a simple thermodynamic problem to a highly complex, nonlinear 21st Century problem. Indeed I argue that it is so complex that current generation climate models are simply not good enough to provide reliable answers. I conclude we need a "CERN for Climate Change" (where nations pool human and computational resources) to understand whether the summer of 2026 is going to be the new year-on-year normal for Europe, or whether it is a summer we are unlikely to experience again for many years. Along the way, I will try to discuss the potential role of AI and quantum computing in helping answer these questions.

Tracing Uncertainties of Extreme El Niño Prediction to Mesoscale Ocean Currents

Bo Qin

Tongji University, China

The predictability of the El Niño-Southern Oscillation (ENSO) has been extensively investigated by tracing monthly averaged or large-scale initial errors, such as heat content and coupled atmosphere-ocean perturbations. Yet these slowly evolving initial errors may originate from smaller-scale ocean dynamics, which exist in current numerical model and ocean observing system. The relations and energy transition processes between small scale processes and basin wide event remain poorly investigated. In this investigation, mesoscale ocean-current errors are revealed as a hidden source of ENSO forecast uncertainty. A differentiable ocean AI model is constructed, trained on high-resolution GLORYS reanalysis product and coupled with conditional nonlinear optimal perturbation (CNOP) method to identify the most influential current errors for extreme El Niño forecasts. The resulting CNOP-type initial errors are strongly event dependent. For the 1997/98 El Niño, the current error

located in the eastern Pacific results in the largest forecast error through locally suppressing original subsurface warming. For the 2015/16 event, the dominate error originates in the western Pacific, reorganizes upper ocean temperature structure and affects the eastern Pacific through the equatorial waveguide. Further analysis reveals that these current errors reshape mesoscale eddy count and eddy kinetic energy, then triggers large-scale temperature errors that weaken El Niño development. The monthly growth rate of the obtained CNOP initial error is aligned with the spring predictability barrier, and their structures are reflected in simulated analysis errors caused by unresolved mesoscale currents. These findings extend ENSO predictability beyond large-scale thermal signals, reveal a mesoscale route of ENSO forecast failure and provide a physical basis for targeted ocean observations and data assimilation.

Analysis of a Two-Layer Diffusive Energy Balance Model

Cristina Urbani

Universitas Mercatorum, Italy

Energy Balance Models (EBMs) provide a simple yet effective framework for describing the Earth's climate. Introduced by Budyko and Sellers in the late 1960s, classical EBMs describe the evolution of the zonally averaged surface temperature through a one-dimensional diffusive equation.

A natural extension of these models is to include vertical resolution by accounting for the energy exchange between the Earth's surface and the atmosphere. In this talk, I will present the mathematical analysis of a two-layer diffusive EBM. I will first briefly review the main results obtained for the spatially homogeneous model, including well-posedness, positivity, and long-time behaviour [1].

I will then focus on the spatially dependent system [2]. The main ingredient is the establishment of new maximum and comparison principles for coupled degenerate cooperative parabolic systems. These results allow us to prove global well-posedness, preservation of positivity, and the existence of distinguished equilibria. Finally, I will discuss the asymptotic dynamics of the model by proving the existence of a global attractor.

I will conclude with a brief glimpse into the stochastic version of the model, highlighting some of the mathematical results obtained so far in [3].

References

1. P. Cannarsa, V. Lucarini, P. Martinez, C. Urbani, and J. Vancostenoble, *Analysis of a Two-Layer Energy Balance Model: Long-Time Behaviour and Greenhouse Effect*, *Chaos*, **33** (2023), 113111.

2. P. Cannarsa, V. Lucarini, P. Martinez, C. Urbani, and J. Vancostenoble, *Comparison Principles and Long-Time Behaviour for a Diffusive Energy Balance Model with Vertical Resolution*, (2026), submitted.

3. J. Broeker, P. Cannarsa, G. Carigi, T. Kuna, and C. Urbani, *A Stochastic Two-Layer Energy Balance Model: Well-Posedness and Exponential Ergodicity*, (2026), preprint.

Statistics of Open Systems and Metastability

Sandro Vaienti

CNRS, France

We review a few recent results on open dynamical systems, eventually randomly perturbed. Two main points will be addressed. Limit theorems characterizing convergence to equilibrium on the surviving set and metastability whenever "openess" allows jumps of the trajectories among different regions in phase space.

Predictability Hierarchies of Extreme Events Revealed by AI Model

Yuxuan Yang

National University of Singapore, Singapore

Predicting extreme events in chaotic dynamical systems remains a fundamental challenge due to their sensitivity to initial conditions and the high computational cost of ensemble forecasting. Conventional approaches often characterize predictability through system-wide measures or individual case studies, leaving systematic differences among extreme events largely unexplored. In this talk, we present a data-driven framework for quantifying event-wise predictability horizons using AI-based ensemble forecasts. By evaluating probabilistic forecast skill against climatology, we determine the maximum lead time over which individual extreme events remain predictable. We first demonstrate this framework in two-dimensional turbulent Kolmogorov flow, where autoregressive diffusion models generate ensembles of possible future trajectories. Our results reveal a pronounced hierarchy of predictability horizons across extreme events. Further analysis identifies the importance of large-scale flow structures in sustaining predictability, while comparisons with numerical ensembles and different model architectures support the robustness of the observed hierarchy. We then extend the framework to European heatwaves using AI-based weather ensemble forecasting, revealing substantial variations in predictability both across individual events and

geographical regions. Comparisons with numerical weather forecasts further assess the consistency of these patterns across forecasting systems. Together, these findings demonstrate that extreme events possess distinct predictability limits and establish AI-based ensemble forecasting as a promising tool for systematically investigating predictability hierarchies in complex dynamical systems.

Learning a Recurrence-based Latent Space of North Atlantic-European Weather Regimes with Joint Embedding Predictive Architectures

Marco Zanchi

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

European winter weather is largely organised into a small number of recurrent and persistent large-scale flow configurations of the North Atlantic European circulation, known as weather regimes. These regimes influence cold spells, windstorms, and prolonged anomalies in energy demand and renewable energy production. They are widely used in sub-seasonal prediction, climate model evaluation, and the interpretation of extreme events. Weather regimes are typically identified by clustering in a truncated Empirical Orthogonal Function (EOF) space, whose dimensionality is fixed a priori according to a variance criterion. Yet variance is not the fundamental property that defines a weather regime: regimes are atmospheric states that recur frequently and in which the atmosphere tends to persist.

We propose to learn a representation of the circulation in which recurrence itself provides the training signal, using a Joint Embedding Predictive Architecture (JEPA). JEPAs learn without labels and without reconstructing the input: a convolutional encoder is trained to map different views of the same underlying state to nearby latent representations. The choice of views therefore determines what the representation retains (what the views have in common) and what it discards. In computer vision, these views are typically generated through hand-crafted augmentations, which have no obvious physical justification for geopotential fields. Here, we derive them directly from the dynamics. For each winter day in ERA5 daily 500-hPa geopotential height anomalies from 1979 to the present, we identify the 15 closest circulation analogues, excluding temporally adjacent days, and treat them as different views of the same dynamical state. The resulting latent space therefore retains the features shared by recurrences of a flow configuration, while VICReg regularisation prevents representation collapse.

We evaluate the recurrence-based latent space against EOF and beta-variational autoencoders representations of equal dimension on a held-out period. We first ask whether the classical regimes (NAO+, NAO-, Atlantic Ridge and Scandinavian

Blocking) emerge as distinct regions of the latent space without being imposed, using externally defined circulation indices as linear probes. We then test whether the latent geometry encodes the local dimension and persistence of the flow, and whether residence times and transition pathways are more coherent than in variance-based coordinates. The approach is designed to be transferred to climate model simulations, where it can address whether anthropogenic forcing alters this organisation and not only the mean state.
