Participant Welcome Packet

Unifying Innovations in Forecasting Capabilities Workshop

A Unified Forecast System (UFS) Collaboration Powered by the Earth Prediction Innovation Center (EPIC)

Dates:
Monday, July 20, 2026 - Friday, July 24, 2026

Time:
Monday: 9:00 AM – 4:00 PM EDT
Tuesday 9:00 AM – 7:00 PM EDT
Wednesday: 9:00 AM – 5:15 PM EDT
Thursday 9:00 AM - 5:30 PM EDT
Friday: 9:00 AM – 12:30 PM EDT

Location:
College Park Marriott Hotel & Conference Center
3501 University Boulevard E, Hyattsville, MD, 20783, and Online

Table of Contents

Goals

1. Share the current status of and future plans for the community-based Unified Forecast System.

2. Share successes and challenges related to contributing to the Unified Forecast System.

3. Identify ways academia, industry, and operations can work together to enhance the Unified Forecast System, accelerate their own contributions, and measure their success.

Objective

To enhance the collaboration between the public, private, and academic sectors within the Weather Enterprise. By highlighting the importance of community building and the shared goal of advancing the UFS, the workshop serves as a key event for fostering enhanced communication and collaboration. UIFCW26 is about engaging and uniting our efforts to advance forecasting capabilities for a more informed future. The theme for this year's workshop is *New Frontier: Optimizing New Technology in a Changing Weather Enterprise.*

In-Person, Virtual, and Hybrid Logistics

In-person

Address: 3501 University Boulevard E, Hyattsville, MD, 20783-7998

Parking: Parking is free and available in the hotel parking lot.

Lunch: Lunch will be provided to ALL in-person attendees by the hotel catering during the scheduled 1-hour lunch break Tuesday through Friday of the workshop. A lunch will be offered on Monday but only for those that signed up to participate in one of the training events offered that day. Additionally, federal employees are **REQUIRED** to pay for lunch if they intend to participate in the lunches. Lunch will be served through the honor system.

Hotel Map

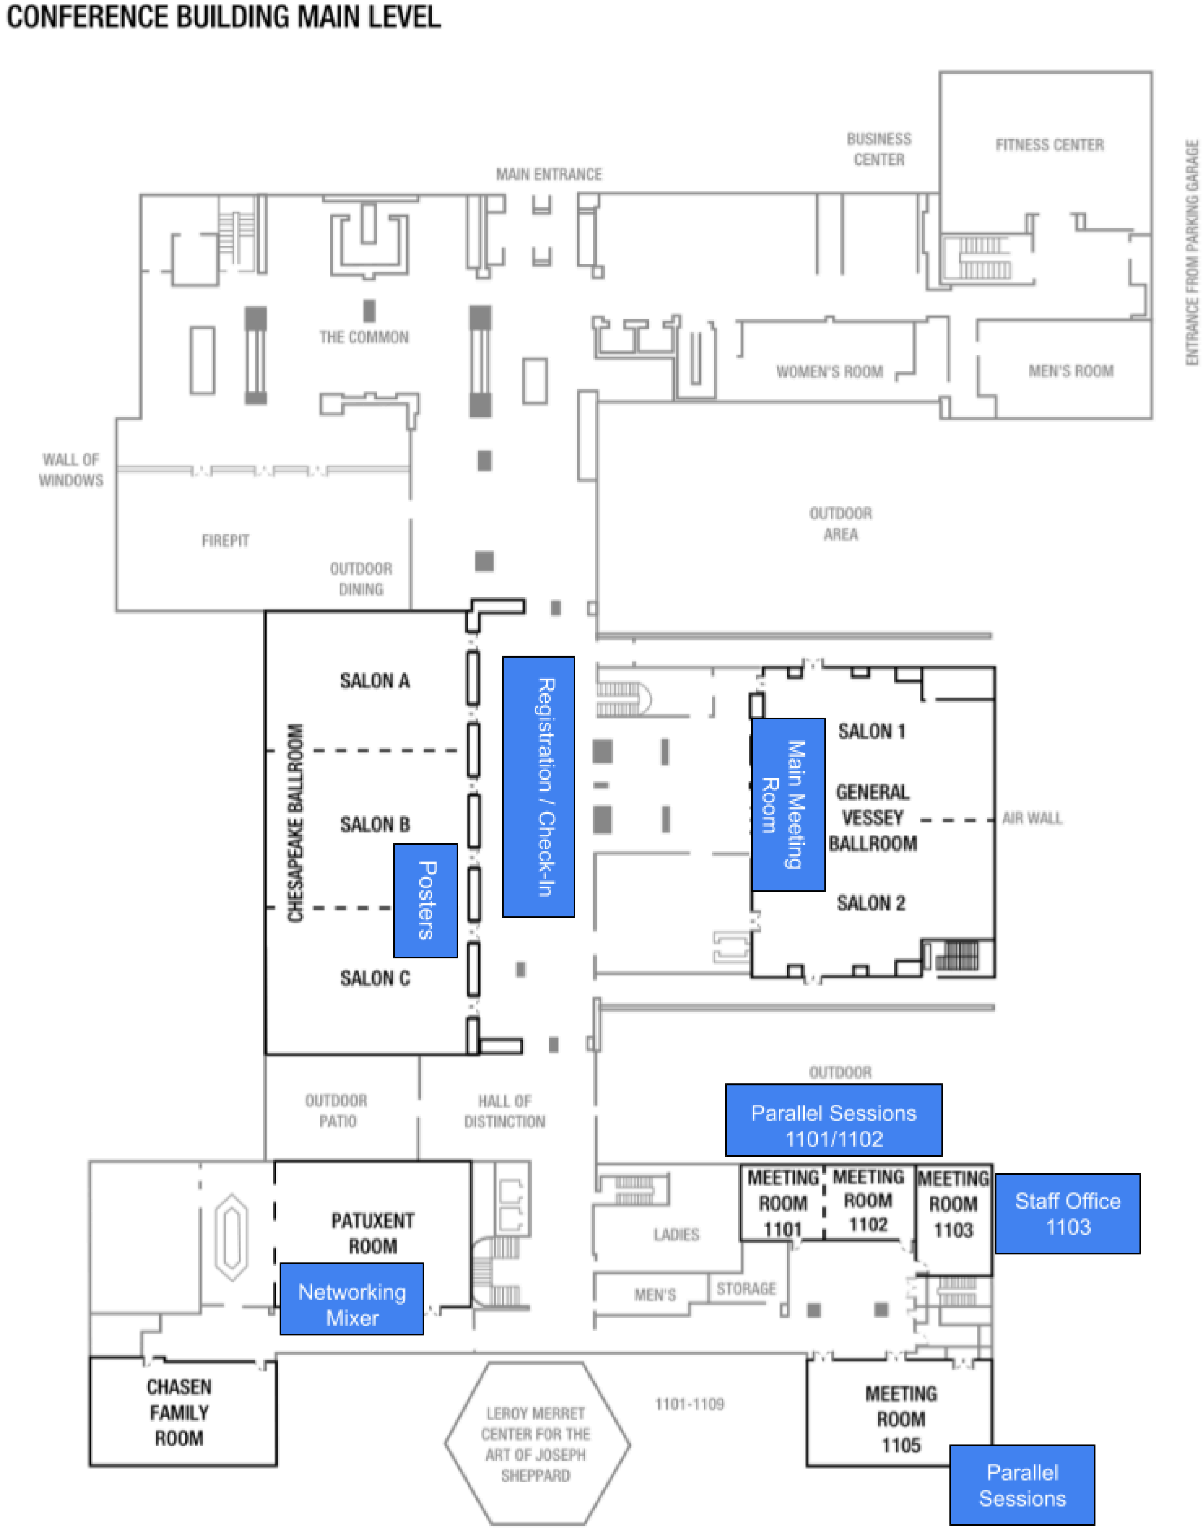
CONFERENCE BUILDING MAIN LEVEL
MAIN ENTRANCE
BUSINESS CENTER
FITNESS CENTER
ENTRANCE FROM PARKING GARAGE
THE COMMON
WOMEN'S ROOM
MEN'S ROOM
WALL OF WINDOWS
FIREPIT
OUTDOOR DINING
OUTDOOR AREA
SALON A
SALON B
SALON C
CHESAPEAKE BALLROOM
Registration / Check-In
Posters
Main Meeting Room
SALON 1
GENERAL VESSEY BALLROOM
SALON 2
AIR WALL
OUTDOOR
Parallel Sessions 1101/1102
OUTDOOR PATIO
HALL OF DISTINCTION
MEETING ROOM 1101
MEETING ROOM 1102
MEETING ROOM 1103
Staff Office 1103
PATUXENT ROOM
Networking Mixer
LADIES
MEN'S
STORAGE
CHASEN FAMILY ROOM
LEROY MERRET CENTER FOR THE ART OF JOSEPH SHEPPARD
1101-1109
MEETING ROOM 1105
Parallel Sessions

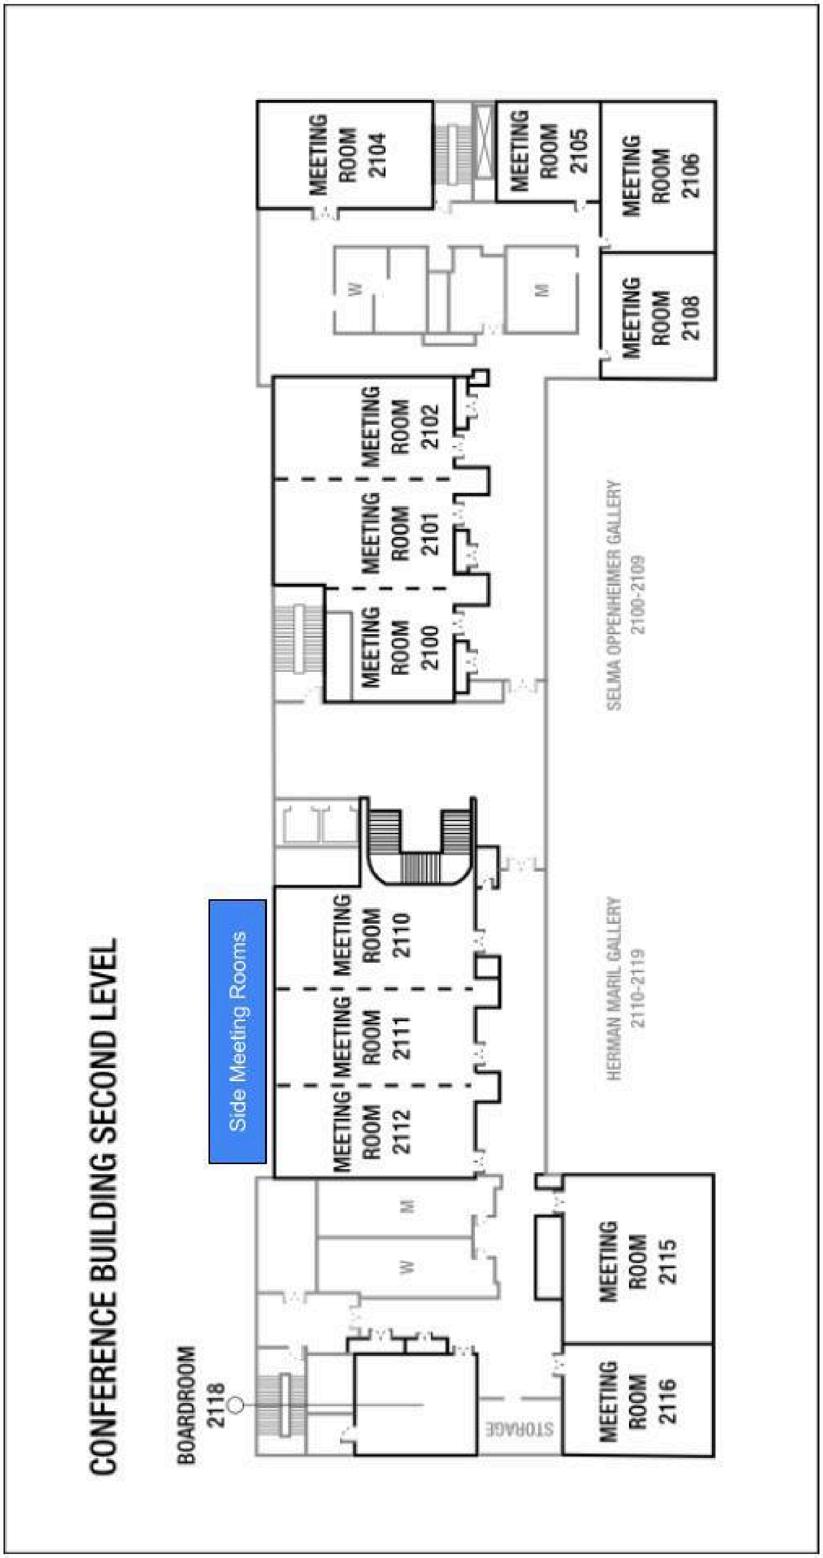
CONFERENCE BUILDING SECOND LEVEL
BOARDROOM 2118
Side Meeting Rooms
MEETING ROOM 2112
MEETING ROOM 2111
MEETING ROOM 2110
MEETING ROOM 2100
MEETING ROOM 2101
MEETING ROOM 2102
MEETING ROOM 2104
MEETING ROOM 2105
MEETING ROOM 2106
MEETING ROOM 2108
MEETING ROOM 2115
MEETING ROOM 2116
STORAGE
W
M
HERMAN MARIL GALLERY 2110-2119
SELMA OPPENHEIMER GALLERY 2100-2109

Virtual

Most presentations will be streamed through Google Meet. The Google Meet Links will be provided to all registered participants. Your online engagement platform will be through Google Meet Q&A and polling functions.

Hybrid

See the above info; come and go as you please.

About EPIC

The Earth Prediction Innovation Center (EPIC) is fundamentally transforming the landscape of Earth system prediction by championing collaborative partnerships across government, academia, and industry. EPIC is dedicated to driving innovative infrastructure, integrating groundbreaking technologies like Artificial Intelligence (AI) into Numerical Weather Prediction (NWP). Beyond technological advancement, EPIC is committed to training the next generation of atmospheric scientists, fostering a skilled workforce prepared to tackle complex environmental challenges. Our core mission includes advancing Earth System Modeling capabilities through the Unified Forecast System (UFS) framework, sustaining an open science, community-driven approach that accelerates research to operations.

About UFS

The Unified Forecast System (UFS) is a community-based, coupled, comprehensive Earth modeling system. UFS numerical applications span local to global domains and predictive time scales from sub-hourly analyses to seasonal predictions. It is designed to support the Weather, Water, and Climate Enterprise and to be the basis for NOAA's operational Earth system modeling applications.

The UFS community includes researchers, developers, and users from NOAA, academia, federal agencies, and the private sector. The UFS supports research and development in the community and accelerates the transition of research to operations.

About UFS-R2O

The UFS-R2O Project was a major collaborative effort focused on transitioning UFS applications and new research innovations into the NWS National Centers for Environmental Prediction

(NCEP) operational modeling suite. The project was supported and managed by both the NWS Office of Science and Technology Integration (OSTI) Modeling Program, and the OAR Weather Program Office (WPO) Joint Technology Transfer Initiative (JTTI), enabling partnerships and opportunities for joint development of next generation, UFS-based operational systems across organizations including the NCEP Environmental Modeling Center (EMC), OAR labs, the Developmental Testbed Center (DTC), and multiple academic partners.

The UFS-R2O Project began in July 2020 and lasted six years (split into two phases, three years each), resulting in significant advancements in developing FV3-based systems such as Air Quality Model (AQMv7), Hurricane Analysis and Forecast System (HAFSv1), regional Rapid Refresh Forecast System (RRFSv1), the Global Forecast System (GFSv17) and Global Ensemble Forecast System (GEFSv13).

Following the conclusion of the UFS R2O project on June 30, 2026, final reports for tasks and deliverables have been submitted by the project teams and are currently undergoing review. Although future plans remain uncertain, a potential path forward involves transitioning to a more targeted, community-oriented framework. This new structure would focus on expanding collaborative development partnerships across NOAA and the broader UFS community, fostering effective co-development to build next-generation Earth system modeling applications that align with NOAA's highest priorities.

Connect With Us

Google Meet Q&A will be your go-to online hub to chat with speakers and connect with other attendees. We will also be sharing updates and information through X when you use **#UIFCW26.**

You can also follow @NOAAEPIC on X, Facebook, @NOAAEPIC on Instagram, or on Earth Prediction Innovation Center | LinkedIn for the latest news from EPIC.

Please email support.epic@noaa.gov with any questions or concerns.

WiFi Information: Connect to the **UMUC-Guest Network** and accept the terms and conditions.

Daily Agenda

*SUBJECT TO CHANGES, please refer to EPIC's website at epic.noaa.gov for important updates

Monday, July 20th

All times Eastern Daylight Time

9:00 am - 12:00 pm **Training 1 | Introduction to Running NWS Global NWP Models in Development Mode**

In-person attendees: Room 1101/1102
Virtual attendees: *See Participant Instruction Email Sent on July 10th*

9:00 am - 12:00 pm **Training 2 | Adding Idealized Test Cases to the UFS Hierarchical System Development Framework**

In-person attendees: Room 1105
Virtual attendees: *See Participant Instruction Email Sent on July 10th*

10:30 am **Virtual Poster Session (Virtual Participants Only)**

Gillian Petro, NOAA/EPIC/Element 84, Inc.

In-Person Attendees: General Vessey Ballroom
Virtual Attendees: https://meet.google.com/jbp-adha-uee
The virtual poster session provides a unique opportunity for the virtual audience to participate and engage in the virtual poster session to hear authors present their posters. If you have any general questions about UIFCW, you can ask them via the Google Meet Q&A.

Abstracts for virtual posters can be found next to the names of the authors below.

VIRTUAL POSTER PRESENTERS

10:32 am - Fariborz Daneshvar - Integration of Neural Operator and Physics-Based Ocean Models for Probabilistic Prediction of Tropical Cyclone Driven Coastal Flooding (abstract)

10:36 am - Yu-Shin Kim - Strongly Coupled Assimilation of Atmospheric Boundary Layer Observations for the Rapid Intensification of Hurricane Ida (2021) using the CADRE Self-cycled HAFS-JEDI System (abstract)

10:40 am - Margarita Mora - An Evaluation of Convective Schemes for Intensity and Track Forecasts for Typhoon Usagi (abstract)

10:44 am - Irena Ivanova - Explicit Effects of Forest Canopy Shading and Turbulence in the Unified Forecast System-Air Quality Model (abstract)

10:48 am - Yaping Wang - JEDI-Based Three-Dimensional Variational Data assimilation for the NOAA Global Chemistry and Aerosol Forecast System (GCAFS) (abstract)

10:52 am - Li (Kate) Zhang - Development of the CCPP-Based GEFS-Aerosols Component in the Unified Forecast System for Subseasonal Prediction (UFS-Chem v1.0) (abstract)

10:56 am - Maria Gehne - MDTF x UFS (abstract)

11:00 am - Muhammad Indra Bendi - Evaluation of Subseasonal Forecast Skill in a Tropical Island Environment: Implications for Forecasting Capability over East Nusa Tenggara, Indonesia (abstract)

11:04 am - Jeff Hamilton - Identifying the Error: Process-Oriented Verification in support of Model Development (abstract)

11:08 am - Chuan-Kai Wang - Impact of High Resolution and Scale-Aware Physics on HAFS: Results from Idealized Tests and Hurricane Milton (2024) (abstract)

11:12 am - Michael Lalime - The NOS Coastal Cloud Sandbox: A Prototype for Operational Modeling in the Cloud (abstract)

11:16 am - Badreddine Alaoui - Leveraging Deep Learning and Satellite Imagery for Enhanced Convective Parameter Diagnosis (abstract)

11:20 am - Linlin Cui - Improving the Intensity of Tropical Cyclones for the AIGFS (abstract)

11:24 am - Ellie Hojeily - Evaluating the Aurora Machine Learning Model Utilizing NCEP Aerosol Initial Conditions for PM2.5 Forecasting (abstract)

11:28 am - Qi Shi - Development of Operational Verification Capability for NOAA Artificial Intelligence Global Forecast System (AIGFS) (abstract)

11:32 am - Feng Hsiao - Physics Sensitivity of All-Sky Infrared Radiance Assimilation in HAFS-JEDI and Its Impact on Tropical Cyclogenesis Forecasts (abstract)

11:36 am -Deepak Kumar - Spatio-Temporal Characterization of Wildfire Dynamics Utilizing High-Resolution WRF-Based Fire Weather (abstract)

11:40 am - Cigdem Akan - Feedback Between Ocean Waves and Vertical Heat Distribution in the High North: A UFS-Based Modeling Approach (abstract)

11:44 am - Patrick Campbell - Dynamic Meteorology-Aware Emissions to Improve UFS Air Quality Model Predictions (abstract)

11:48 am - Joannes Westerink - Integrating Thermohaline Physics into STOFS 2D Global: A Heterogeneous Mode Splitting Approach with Global RTOFS and Regional Model Coupling (abstract)

11:52 am - Haiqin Li - Improving the aerosol representation and microphysics coupling in NOAA's RRFSv2 numerical weather prediction model with MPAS dynamical core (abstract)

11:56 am - Rongqian Yang - Evaluation of Remotely Sensed Leaf Area Index (LAI) Using the Noah-MP Land Surface Model (Noah-MP LSM): Impacts on

Surface Energy and Water Cycles over the Continental United States (CONUS) (abstract)

12:00 pm	**Lunch**
1:00 pm - 4:00pm	**Training 3 | Project EAGLE Training: Machine Learning Weather Forecasting on Microsoft Azure** **In-person attendees:** General Vessey Ballroom **Virtual attendees:** *See Participant Instruction Email Sent on July 10th*
5:00 pm	**Informal Happy Hour/Bowling Gathering** Location: Lucky Strike College Park - 9021 Baltimore Ave, College Park, MD 20740

Tuesday, July 21st

All times Eastern Daylight Time

8:00 am **Registration**
In-Person Attendees: Registration will begin outside of the Chesapeake Ballroom.

Virtual Attendees: If you have any general questions about UIFCW, you can ask them in the Google Meet Q&A space.

9:00 am **Welcome and Kickoff**
In-person attendees: General Vessey Ballroom
Virtual attendees: https://meet.google.com/jbp-adha-uee

Jan Ising, NOAA/OAR/WPO/EPIC

9:10 am **Keynote Presentation from the NOAA Administrator, Dr. Neil Jacobs**
Dr. Neil Jacobs, NOAA Administrator

9:40 am **NOAA Leadership Panel (NOS, NWS, OAR)**
Jennifer Mahoney, NOAA/OAR
Michelle Mainelli, NOAA/NWS
Steve Smith, NOAA/NOS

Moderated by Jan Ising

10:40 am **Morning Break**

11:10 am **Six Years of the UFS-R20 Project**
Vijay Tallapragada, NOAA/NWS/OMD (abstract)

11:30 am **EPIC: Retrospective + Goals Moving Forward**
Maoyi Huang, NOAA/OAR/WPO/EPIC
Keven Blackman, NOAA/EPIC/Raytheon (abstract)

12:00 pm	**Lunch**
1:00 pm	**American Meteorological Society (AMS) Community Modeling Committee Updates** *Fred Carr, University of Oklahoma* *Jim Kinter, George Mason University*
1:30 pm	**NOAA's High Performance Computing and Communications (HPCC) Program Overview** *Kathryn Shontz, NOAA OCIO (Acting) Deputy Director for Strategy*
2:00 pm	**Model for Prediction Across Scales (MPAS) Integration**

2:00 pm - Gretchen Mullendore - The NSF NCAR MPAS Program (abstract)

2:10 pm - Eric Aligo - EPIC's Role in Facilitating the Integration of MPAS into UFS (abstract)

2:20 pm - Ligia Bernardet - Integrating the Model for Prediction Across Scales dynamical core into the Unified Forecast System Weather Model (abstract)

2:30 pm - Ann Tsay - Choose Your Adventure: The Many Ways to Integrate MPAS within ESMs (abstract)

2:42 pm - Nusrat Yussouf - Transitioning to an MPAS-JEDI-based Warn-on-Forecast System (WoFSv2): Next-Generation Storm-Scale Probabilistic Guidance (abstract)

2:54 pm - Clark Evans - NOAA/GSL Model Development and Forecasting Activities using MPAS for Future NOAA Operational Modeling Systems (abstract)

3:06 pm - Adam Clark - MPAS evaluations for severe weather forecasting during the 2026 NOAA/Hazardous Weather Testbed Spring Forecasting Experiment (abstract)

3:18 pm - Discussion

3:30 pm **Group Picture**

3:45 pm **Afternoon Break**

4:15 pm **Poster Session I (In-person Participants)**
If you are present in person, you will walk throughout Chesapeake Ballroom B&C and engage with our in-person presenters as you would at any other poster session.

Abstracts for the in-person poster session I can be found next to the names of the authors below.

IN-PERSON POSTERS

Ayoola Abe - Model Intercomparison of UFS and MPAS Forecasts for a Severe Convection Initiation Event: A Case Study (abstract)

Michael Barlage - Impacts of a Modified Land Surface Representation in UFS/SCM/NoahMP Simulations (abstract)

Kristin Barton - Updates on the UFS-Arctic Application Development (abstract)

Alex Burrows - NOAA's EAGLE Project: Advancing AI/ML Weather Forecasting from Research to Operations (abstract)

Benjamin Cash - Known Unknowns: Impacts of Sampling Variability on ENSO Teleconnections (abstract)

Albert Cerrone - Multi-Agent in Hydrodynamics: An Agentic Approach to Scientific Workflow Management (abstract)

Bowen Chen - OceanAI: A Conversational Agent for Accurate, Transparent, Near-Real-Time Oceanographic Insights (abstract)

Ming Chen - Wave Evaluation for the next Global Forecast System GFSv17 (abstract)

L. Gwen Chen - Systematic Evaluation of NOAA's Artificial Intelligence Modeling Suite at the National Centers for Environmental Prediction (abstract)

Joshua DaRosa - Ensemble Nested EAGLE and Probabilistic Forecasting using Anemoi (abstract)

Jun Du - A weather-regime-dependent verification metric for ensemble forecasts: Ensemble Performance Index for High-Impact Weather (abstract)

Elena Fernandez - Stratospheric Polar Vortex Representation in the Unified Forecast System's Subseasonal Prototypes (abstract)

Cong Gao - PepC-Global: A Basin-Tuned, Environment-Dependent Probabilistic Tropical Cyclone Model (abstract)

Zhichang Guo - Evaluation of Land Surface Variables in the Global Forecast System (GFSv17) Retrospective and Real-Time Forecast Experiments (abstract)

Cenlin He - Subseasonal-to-Seasonal Prediction using MPAS/NoahMP (abstract)

Horeb - Advancing AI Emulations of Physics-Based NWP Models through Investigating Training Data Generation Strategies (abstract)

Wei Huang - Porting and Optimizing the Global-Workflow on AWS (abstract)

Faozia Anzum Itu - Track Evolution and Landfall Dynamics of Cyclone Remal (2024) over the Bay of Bengal using the WRF-ARW Model (abstract)

Young-Joon Kim - Evaluating Global Model Fidelity via Troposphere-Stratosphere Coupling and Teleconnection Indices (abstract)

Mu-Chieh Ko - DA Nudging: Bridging AI Weather Prediction Systems and GFS (abstract)

Jianjun Liu - AIGEFS Improvements via CRPS-based Optimization with the ANEMOI Framework (abstract)

Sarah Lu - Utilizing Water Isotopic Composition in the Unified Forecast System to Constrain Precipitation and Hydrologic Processes: Preliminary Results (abstract)

Somnath Luitel - AirCast-SR: A Foundation Model for Kilometer-Scale Atmospheric Super-Resolution via Latent Consistency Diffusion (abstract)

Xiao Luo - Sensitivity Analysis of Fine-Tuning GraphCast with Adjusted Mean Square Error Using GDAS Data (abstract)

Murali Nageswara Rao Malasala - Watershed-Scale Evaluation of ARAFSv2 for High-Impact Atmospheric River Rainfall over the Western United States (abstract)

Kathryn Newman - From UFS Applications to AIWP: Leveraging the METplus Verification Workflow for Consistent and Scalable Model Evaluation (abstract)

Kaitlin Pereira - Winter Validation of the Operational Great Lakes Wave System Using In-Situ and SWOT Observations (abstract)

Nick Silverman - Harnessing Generative AI and Emulation for High-Resolution Fire Weather Forecasting: A Multi-Scale Modeling Framework (abstract)

Lydia Stefanova - Examining the Shift in Large-Scale Biases and RMSE: From GFSv16 to the Fully Coupled GFSv17 (abstract)

Aman Tejaswi - Leveraging Subgrid Methods for a High-Fidelity Global-Scale Hydrodynamic Flood Inundation Modeling (abstract)

Yu-Cian Tsai - Representation of East Pacific Easterly Waves in Subseasonal-to-Seasonal Forecasts of the Unified Forecast System (abstract)

Sarah Weidman - Reducing state-independent bias in subseasonal weather forecasts using a correction derived from a running mean nudging (abstract)

Brian Weir - Graph-Oriented Workflow Orchestration: A Design Framework for UFS Research-to-Operations Pipelines (abstract)

Keqin Wu - Development of the Atmospheric River Analysis and Forecast System to Enhance Atmospheric River Forecasting (abstract)

Zheng Wu - Assessing subseasonal predictions of stratosphere-troposphere coupling of a data-driven weather forecasting model (abstract)

Tianning Wu - Forecasting the Loop Current System with Physics-Guided Machine Learning (abstract)

Li Xu - Seasonal Forecast System (SFS) Applications to NOAA Early Warning System (abstract)

Abu Ubayeda Ibnal Zarrah - Spatiotemporal Estimation and Forecasting of Particulate Matter in the Pacific Northwest using Machine Learning with Emphasis on Wildfire Smoke (abstract)

Xuejin Zhang - A Prototype AI-Based Quality Control Package for Cloud Physics Observations of Hurricane Field Program at NOAA (abstract)

5:30 pm - 7:00 pm **Networking Mixer at the Marriott Hotel's Patuxent Room & Patio**

Wednesday, July 22nd

All times Eastern Daylight Time

8:00 am **In-Person Attendees:** Registration will begin outside of the Chesapeake Ballroom.

Virtual Attendees: If you have any general questions about UIFCW, you can ask them via the Google Meet Q&A.

9:00 am **Welcome and Kickoff**
In-person attendees: General Vessey Ballroom
Virtual attendees: https://meet.google.com/jbp-adha-uee

Shira Francis, NOAA/NWS/OMD/PPU

9:10 am **Weather Briefing**
Ben Woods, Forensic Meteorology

9:20 am **AIDA: Operational AI Data Assimilation from Level 1 Observations**
Ryan Keisler, Brightband

9:30 am ***New User/Student Parallel Track Begins***
In-person attendees: *Room 1101/1102*
Virtual attendees: *https://meet.google.com/hoz-xbdt-dwh*

9:50 am **Data Assimilation (Part 1)**

Moderated by Daniel Holdaway

9:50 am - Soyoung Ha - NSF NCAR's Next-Generation Online-Coupled Air Quality and Weather Analysis and Forecasting System using MPAS-GOCART2G-JEDI (abstract)

10:02 am - Shun Liu - Toward the Implementation of JEDI in the RRFS Data Assimilation System (abstract)

10:14 am - Jason Sippel - A Summary of HAFS Data Assimilation Activities at the Atlantic Oceanographic and Meteorological Laboratory (abstract)

10:26 am - Discussion

10:30 am **Morning Break**

11:00 am **Data Assimilation (Part 2)**

Moderated by Daniel Holdaway

11:00 am - Yongming Wang - MAPCast: A Convection Allowing Emulator Toward Multi-Scale Data Assimilation (abstract)

11:10 am - Daniel Kubalek - Combining Data-driven and Physics-based Background Ensemble in Data Assimilation (abstract)

11:20 am - Vanderlei Vargas Jr. - Towards a Continuously Cycled AI Weather Prediction Model: An LETKF-Based Data Assimilation System for HRRRCast (abstract)

11:30 am - Henry Santer - Fully non-Parametric Ensemble Data Assimilation: Idealized and Sea Ice Model Experiments Using Data-driven Estimates of Likelihood Functions (abstract)

11:40 am - Man-Yau (Joseph) Chan - Distribution-respecting Model-Space Localization for the Non-Variational Ensemble Data Assimilation of Non-Local Observations: Formulation and Experiments with Clear-sky Infrared Radiances (abstract)

11:50 am - Discussion

12:00 pm **Lunch**

1:00 pm **Parallel Sessions:**

PARALLEL SESSION 1: UFS APPLICATIONS: (HAFS, COASTAL)
In-person attendees: General Vessey Ballroom

Virtual attendees: https://meet.google.com/jbp-adha-uee

Moderated by Tracy Fanara

1:00 pm - Zhan Zhang - UFS Hurricane Application HAFS: 2026 Upgrades and Future Plans (abstract)

1:10 pm - Jing Cheng - Transitioning HAFS Data Assimilation to JEDI: Observation Dataset Readiness and Integration in HAFS (abstract)

1:20 pm - Sundararaman Gopalakrishnan - Hurricane Analysis and Forecast System: R&D and R2O (abstract)

1:30 pm - William Ramstrom - Efficient High Resolution Forecasts of Hurricanes through Basin-Scale and Telescoping Nest HAFS (abstract)

1:40 pm - Rusty Benson - Crossing the Boundaries: Global Moving Nests (abstract)

1:50 pm - Tim Marchok - Recent GFDL development updates: T-SHiELD and Vortex Tracker (abstract)

2:00 pm - Jessica Meixner - NWS Marine Component Synergies with the Unified Forecast System (abstract)

2:10 pm - Saeed Moghimi - UFS-Coastal: Engineering the Enterprise Backbone for NOS's Next-Generation Coastal Earth System Model and Operational Applications (abstract)

2:20 pm - Yi Hong - Developing and Assessing Two-Way Hydrologic–Hydrodynamic Coupling for Coastal Backwater-Affected and Compound Flood Prediction (abstract)

2:25 pm - Scott Durski - Coupling CICE to SCHISM in UFS coastal:Working towards a coupled ice-ocean forecast system for Alaska (abstract)

2:30 pm - Richard Stumpf - Harmful Algal Bloom Forecasts: Examples in Adaptive Unification (abstract)

2:35 pm - Yunfang Sun - The applications of UFS-Coastal: wave-current coupling of SCHISM and WAVEWATCH III (abstract)

2:40 pm - Patrick Burke - National Ocean Service Core Model Evaluation Plan (abstract)

2:45 pm - Discussion

PARALLEL SESSION 2: UFS APPLICATIONS (AIR QUALITY, LAND, SRW/RRFS)
In-person attendees: Room 1105
Virtual attendees: https://meet.google.com/rsg-prex-pso

Moderated by Jong Kim

1:00 pm - Israel Jirak - Evaluation of RRFSv1 during the 2026 NOAA HWT Spring Forecasting Experiment (abstract)

1:12 pm - Curtis Alexander - RRFS Development Towards the NWP Moonshot (abstract)

1:24 pm - Terra Ladwig - New RRFS: MPAS plus JEDI Development and 2026 Convective Season Performance (abstract)

1:36 pm - Wei Li - Progress on the Development of CATChem Version 2 within NOAA's Unified Forecast System (abstract)

1:48 pm - Partha Bhattacharjee - Development and Evaluation of NOAA's global aerosol model: Global Chemistry and Aerosol Forecast System (GCAFS v1) (abstract)

2:00 pm - Ravan Ahmadov - Evaluation of Smoke Transport in NOAA's RRFS-SD Model Using Satellite Observations (abstract)

2:12 pm - Johana Romero Alvarez - Representing Organic Soil Fire Emissions in the RRFS-Smoke Model: Implementation and Performance Assessment (abstract)

2:24 pm - Weizhong Zheng - Performance Assessment of Land Surface Fields in NOAA's Seasonal Forecast System (SFS) (abstract)

2:36 pm - Andrew Newman - Community tools and workflows for hierarchical Noah-MP modeling and optimization to improve UFS hydrometeorological prediction (abstract)

2:48 pm - Ananyo Bhattacharya - Influence of the Laurentian Great Lakes on Summer Time Precipitation Revealed by Isotopic Enrichment Measurements in Michigan (abstract)

3:00 pm **Afternoon Break**

3:30 pm **Parallel Sessions:**

PARALLEL SESSION 3: CROSS CUTTING NWP TOPICS
In-person attendees: General Vessey Ballroom
Virtual attendees: https://meet.google.com/jbp-adha-uee

Moderated by Ligia Bernardet

3:30 pm - Jin Huang - NOAA's Precipitation Prediction Grand Challenge (PPGC) Initiative in Supportive of UFS Improvement and Applications (abstract)

3:42 pm - Jason Levit - An Overview of Verification and Evaluation Activities at the NOAA Office of Modeling and Development (abstract)

3:54 pm - Barry Baker - Unifying Composition in the UFS through CATChem and CECE (abstract)

4:06 pm - Santha Akella - Application targeted configuration of the ocean component in the UFS framework towards seamless predictions (abstract)

4:18 pm - Ali Abdolali - TRITON-C Development Update: A Consistent Implicit Framework with On-the-Fly V&V and Coupling strategy (abstract)

4:30 pm - Andrew Newman - Advancing Land-Atmosphere Evaluation of NOAA UFS Models Through Community Datasets and Software Enhancements (abstract)

4:42 pm - Pranay Reddy Kommera - Portable High-Performance TEMPO Microphysics — A Unified CPU/GPU Codebase via Hybrid Loop Restructuring (abstract)

4:54 pm - Daniel Sarmiento - Addressing Core Modeling Infrastructure Barriers in Next-Generation UFS Weather Model Development (abstract)

5:06 pm - Discussion

PARALLEL SESSION 4: UFS APPLICATIONS (GLOBAL WORKFLOW, SFS, S2S, AR)

In-person attendees: Room 1105
Virtual attendees: https://meet.google.com/rsg-prex-pso

Moderated by Fanglin Yang

3:30 pm - Yan Xue - NOAA's Seasonal Forecast System Project: Progress, Challenges, and Stakeholder Engagement (abstract)

3:42 pm - Phillip Pegion - Practical Strategies for Initialization of the Seasonal Forecast System (abstract)

3:54 pm - Wanqiu Wang - Evaluating SFS Beta: Early Skill Relative to CFSv2 (abstract)

4:06 pm - Catherine Thomas - Overview of the Next Global Forecast System GFSv17 (abstract)

4:18 pm - Ruiyu Sun - Physics Advancements in GFSv17 (abstract)

4:30 pm - Shan Sun - An NSST Alternative in UFS: Skinsst (abstract)

4:42 pm - Maria Aristizabal Vargas - The Ocean Coupling Impact on the forecast of Atmospheric Rivers in the U.S. West Coast (abstract)

4:54 pm - Walter Kolczynski - Public Release of NWS Global Workflow (abstract)

5:15 pm **Adjourn**

Wednesday, July 22nd - *New User/Student Parallel Track*

All times Eastern Daylight Time

9:30 am **Welcome and Kickoff**
In-person attendees: Room 1101/1102
Virtual attendees: https://meet.google.com/hoz-xbdt-dwh

Jenna Holbrook, WPO/EPIC Lapenta Intern, Texas A&M University

9:40 am **Introduction to UFS: A Video Overview**

9:50 am **UFS Applications Overview**
Vijay Tallapragada, NOAA/NWS/OMD
D. Alex Burrows, NOAA/EPIC

10:30 am **Morning Break**

11:00 am **Student Career Panel: Careers in the Weather Enterprise**
Andrew Markowitz, Xcel Energy
Cody Ledbetter, NOAA/NWS WFO Sterling
Kate Fossell, NSF NCAR
Sim James, U.S. Navy

12:00 pm **Lunch**

1:00 pm **Career Building Workshop *(Interactive session)***
Jennifer Vogt, NOAA/OAR/WPO/EPIC
Alison Gregory, UCAR CPAESS
Benjamin Cash, George Mason University
Hendrik L. Tolman, NOAA/NWS

Jan Ising, NOAA/OAR/WPO/EPIC
Linden Wolf, NOAA/OAR/WPO/STC
Segayle Thompson, NOAA/OAR/WPO/EPIC

1:45 pm **HANDS-ON TRAINING SESSION:**
JEDI-Edu Tutorial: Introducing Data Assimilation Methods in JEDI Using a Simple Two-Layer QG Model
Aaron Johnson, University of Oklahoma/CADRE
Nick Gasperoni, University of Oklahoma/CADRE
Jonathan Poterjoy, University of Maryland/CADRE

3:00 pm **Afternoon Break**

3:30 pm **HANDS-ON TRAINING SESSION Cont…**
JEDI-Edu Tutorial: Introducing Data Assimilation Methods in JEDI Using a Simple Two-Layer QG Model

4:15 pm **Optional: (Weather-Themed) Movie Viewing**

5:15 pm **Adjourn**

Thursday, July 23rd

All times Eastern Daylight Time

8:00 am **Registration**
In-Person Attendees: Registration will begin outside of the Chesapeake Ballroom.

Virtual Attendees: If you have any general questions about UIFCW, you can ask them in the Google Meet Q&A.

9:00 am **Welcome and Kickoff**
In-person attendees: General Vessey Ballroom
Virtual attendees: https://meet.google.com/jbp-adha-uee

Aaron Jones, NOAA/EPIC

9:10 am **Parallel Sessions:**

PARALLEL SESSION 1: RESEARCH TO OPERATIONS TO RESEARCH (R2O2R)
In-person attendees: General Vessey Ballroom
Virtual attendees: https://meet.google.com/jbp-adha-uee

Moderated by Aaron Pratt

9:10 am - Fanglin Yang - An Overview of NOAA Unified Forecast System (UFS) Applications for Operational Forecasting Across Scales (abstract)

9:22 am - Jacob Carley - An Efficient and Sustainable Development and Implementation Lifecycle for the NWS Production Suite (abstract)

9:34 am - Youngsun Jung - Enhancing High-Impact Weather Decision-Making through the Warn-on-Forecast System (WoFS) (abstract)

9:46 am - Youhua Tang - Enhancing UFS Air Quality Predictions with Aircraft-Derived Regional Emission Adjustments (abstract)

9:58 am - Samantha Kramer - From Forecast to Fireline: Advancing UFS Wind Prediction for Wildland Fire Operations (abstract)

10:11 am - Soroosh Man - UPSurgeML: Real-Time Framework for Unstructured-Mesh Model Component for Hurricane Surge Simulation (abstract)

10:23 am - Mariah Pope - Near-Real-Time AI Weather Forecasting with Nested-EAGLE: A Framework for Running EAGLE Models in Operations (abstract)

10:35 am - Discussion

9:20 am **Parallel Sessions:**

PARALLEL SESSION 2: DATA ASSIMILATION (PART 3)
In-person attendees: Room 1105
Virtual attendees: https://meet.google.com/rsg-prex-pso

Moderated by Daniel Holdaway

9:20 am - Aaron Johnson - Improving land-atmosphere background error covariance estimation toward RRFS strongly coupled data assimilation (abstract)

9:30 am - Aaron Johnson - Impact of Greenness Vegetation Fraction Perturbations on Convection-Allowing Ensemble-Based Land-Atmosphere Coupled Background Error Covariance. (abstract)

9:40 am - Derek Posselt - A Question of Scale: Using Ensemble Data Assimilation to Understand Convection-Environment Interactions (abstract)

9:50 am - Braedon Stouffer - Assimilating Boundary Layer Depths Observed by WSR-88Ds to Improve the Prediction of a Derecho (abstract)

10:00 am - Chandler Pruett - Improving Forecasts of Hurricane Beryl (2024) Genesis through Assimilation of All-Sky Infrared Radiances (abstract)

10:10 am - Hyun Sook Kim - Preliminary Results of Weakly Coupled Data Assimilation - Hurricane Analysis and Forecast System (WCDA-HAFS) Forecasts (abstract)

10:20 am - Hee-Sung Jung - Constraining Atmospheric Forcing Errors for Sea Ice Data Assimilation in Short-Range Forecasts (abstract)

10:30 am - Discussion

10:45 am **Morning Break**

11:15 am **Community Engagement**
In-person attendees: General Vessey Ballroom
Virtual attendees: https://meet.google.com/jbp-adha-uee

Moderated by Alison Gregory

11:15 am - Tracy Fanara - Introducing the UFS to K12 Students: Strategies and Opportunities (abstract)

11:24 am - Jenna Holbrook - Increasing UFS Student Engagement Through Lesson Plans (abstract)

11:33 am - Kaitlin Pereira - Bridging Academia and Operations: A WINGS Fellow's Experience (abstract)

11:42 am - Brian McDonald - UFS-Chem modeling capability - opportunity for NOAA and external community collaboration and engagement (abstract)

11:51 am - Alison Gregory - Measuring our Community: Social Science Metrics for the Unified Forecast System (abstract)

12:00 pm **Lunch**

1:00 pm **AI4NWP at NOAA: Progress, Partnerships, and the Path Forward**
Isidora Jankov, NOAA/OAR/GSL

1:30 pm **Artificial Intelligence (AI) for Numerical Weather Prediction (NWP) (Part 1)**

Moderated by Keven Blackman

1:30 pm - Jun Wang - EAGLE global weather forecast development (abstract)

1:45 pm - David Harrison - Assessing the Skill and Sensitivity of AIWP Models During the 2026 NOAA HWT Spring Forecasting Experiment (abstract)

2:00 pm - Corey Potvin - WoFSCast Performance in the 2026 HWT Spring Forecasting Experiment (abstract)

2:15 pm - Paul Madden - Designing EAGLE: A Portable and Reproducible Framework for AI/ML Weather Prediction (abstract)

2:30 pm - Isidora Jankov - Next-Gen HRRRCast Ensemble: Enhancements and the Leap into Data Assimilation (abstract)

2:42 pm - Daniel Abdi - Towards incorporating data-assimilation in HRRRCast: a score-based approach (abstract)

3:00 pm **Poster Session II + Afternoon Break (In-person Participants)**
If you are present in person, you will walk throughout Chesapeake Ballroom B&C and engage with our in-person presenters as you would at any other poster session.

Abstracts for the In-Person Poster Session II can be found next to the names of the authors below.

IN-PERSON POSTERS

Barry Baker - The Community Emissions Computing Engine (CECE): A High-Performance, Portable Framework for Atmospheric Emissions in UFS (abstract)

Michael Benneh - UFS/JEDI LIDAR DATA ASSIMILATION: LiDAR Forward Operator Development (abstract)

Kristopher Booker - Accelerating Forecast Innovation: EPIC's Collaborative Framework and HPC-Driven Transparency (abstract)

Maggie Bruckner - Progress update on incorporating the MICM gas-phase chemical solver in CATChem (abstract)

Albert Cerrone - STOFS 2D+ Global: Skill Assessment and Latest Advancements (abstract)

James Correia Jr. - Facilitating the 2 in R2O2R as part of the HydroMeteorology Testbed (HMT) forecast and evaluation activities (abstract)

Felix Deng - Real-Time Ocean State Monitoring and Loop Current Forecasting with CNAPS2 Regional Ocean Data Assimilative Forecast System (abstract)

Leandro Fernandez - Development of a one-way coupled wave forecasting model for the Mouth of the Columbia River with comparisons to data (abstract)

Christina Holt - AIGFS as an Incubator for Modern Software Practices Supporting R2O2R (abstract)

Lucila Houttuijn Bloemendaal - Modernized tools to evaluate and transition operational models to the National Ocean Service (abstract)

Ming Hu - Progress on integrating JEDI with MPAS for the next version of the operational convection-allowing Rapid Refresh Forecast System (abstract)

Bo Huang - Advancing Local Ensemble Solvers in JEDI: Cross-Validation and Multiscale Localization for Improved Data Assimilation (abstract)

Hananeh Jafary - Evaluation of GFDL and EMC Hurricane Tracking Systems During the 2025 Hurricane Season (abstract)

Mansur Ali Jisan - A Modular YAML-Based Unified Workflow Framework for NOAA National Ocean Service's Operational Ocean Forecast Systems and UFS-Coastal Integration (abstract)

Benjamin Johnson - Advancing Radiance Observation Operators for Data Assimilation: Recent Developments in the Community Radiative Transfer Model (CRTM) (abstract)

Jong Kim - Standalone Diagnostic Tools for UFS Data Assimilation Evaluation with Applications in the CADRE Year 2 Data Assimilation Training and Science Workshop (abstract)

Aaron Jones - Empowering Forecasting Innovation Through EPIC Community Engagement and User Support (abstract)

Joshua Kublnick - Empowering and Expanding the UFS Community Through EPIC User Support (abstract)

Bin Liu - 2025 and 2026 HAFS Real-Time Parallel Experiments: Innovations, Performance, and Operational Transition (abstract)

Yonggang Liu - Storm Surge and Coastal Inundation Nowcasts/Forecasts During Hurricanes Helene and Milton (abstract)

Congmeng Lyu - Simultaneous Assimilation of Multi-Wavelength Aerosol Optical Depth Retrievals from NASA's PACE Mission in JEDI for NOAA's Global Aerosol Model (abstract)

Dylan Major - Assessing Hourly Cycled Bayesian Data Assimilation via Localized Particle Filters within a Global Atmospheric UFS Workflow (abstract)

Reza Marsooli - Seasonal Prediction of Coastal Sea-Level Anomalies Using Climate-Driven Hydrodynamic Modeling (abstract)

Cory Martin - Vision Towards a Next-Generation Observation Ingest, Storage, and Access Framework for the National Weather Service (abstract)

Caitlyn McAllister - Let's Track Some Cyclones - Running the GFDL Vortex Tracker Made Easy (abstract)

Chris Miller - ADD-DA: A Machine Learned Model for Rapid Data Assimilation (abstract)

Maria Nikolaitchik - Assessing Structural Differences in Analysis and Forecast Error Covariance between Ensemble Kalman Filters and Local Particle Filters (abstract)

Eric Peino - How Does Land Proximity Reshape Tropical Cyclone Precipitation Asymmetry? Intensity-Stratified Composites from Open Ocean to Landfall (abstract)

Malaquias Pena - Progress on a Wind-Wave-Current Data Assimilation Scheme for the 3D-RTMA: Graph-Based Localization and Coupled Ensemble Analysis (abstract)

Jiayi Peng - HAFS Ensemble Development on Cloud Platforms (abstract)

Joey Rotondo - Advancing Sea Ice Data Assimilation in UFS-Arctic Toward Arctic Reanalysis Applications (abstract)

Rebecca Schwantes - Evaluating UFS and other Models Against Atmospheric Composition and Meteorological Observations Using MELODIES MONET (abstract)

Laura Slivinski - Evaluating the Local Ensemble Transform Kalman Filter for the Ocean in JEDI (abstract)

Sonia Sharma Banjade - CNN-LSTM Based PM2.5 Forecasting and Anomaly Detection in Delhi, India (abstract)

John Steffen - Implementation of a JEDI-based marine data assimilation system for GFSv17 (abstract)

Keston Smith - The Regional Wave Prediction System: Development and Preliminary Studies (abstract)

Chong-Chi Tong - Hydrometeor-Aware All-Sky Microwave Data Assimilation for Atmospheric River Prediction (abstract)

Raelynn Wonnacott - Exploring Generative Methods to Inform Likelihoods for Bayesian Data Assimilation in Earth System Surrogate Models (abstract)

Fei Ye - A Physics-Based Multi-Decadal Compound Flood Reanalysis for the U.S. East and Gulf Coasts (abstract)

Valery Yudin - Adapting Dynamics and Physics of UFSWAM for the Storm-Time Space Weather Predictions (abstract)

4:00 pm **Artificial Intelligence (AI) for Numerical Weather Prediction (NWP) (Part 2)**

In-person attendees: General Vessey Ballroom
Virtual attendees: https://meet.google.com/jbp-adha-uee

Moderated by Keven Blackman

4:00 pm - Ufuk Turuncoglu - A Hybrid Physics–Machine Learning Coupled Model for Subseasonal-to-Seasonal (S2S) Prediction (abstract)

4:15 pm - Niraj Agarwal - Ocean and Coupled Earth System Emulator Development for NOAA-EAGLE (abstract)

4:30 pm - Lin Zhu - Improving Tropical Cyclone Track and Intensity Forecasts via Integrated AI (abstract)

4:45 pm - Hua Leighton - TC-Cast: A Generative AI Emulator for High-Resolution Tropical Cyclone Forecasting (abstract)

5:00 pm - Sergey Frolov: Why bother building variable resolution AI forecast models? (abstract)

5:15 pm - Discussion

5:30 pm **Adjourn**

Friday, July 24th

All times Eastern Daylight Time

8:00 am **Registration**
In-Person Attendees: Registration will begin outside of the Chesapeake Ballroom.

Virtual Attendees: If you have any general questions about UIFCW, you can ask them via the Google Meet Q&A.

9:00 am **Welcome and Kickoff**
In-person attendees: General Vessey Ballroom
Virtual attendees: https://meet.google.com/jbp-adha-uee

Alison Gregory, UCAR CPAESS

9:10 am **Machine Learning**
Moderated by Hendrik Tolman

9:10 am - Wei-Ting Hung - Evaluation of UFS-DeepAQM: an deep learning extension of the Unified Forecast System Air Quality Model (abstract)

9:25 am - Lixin Lu - High Resolution (1-km) simulation of leaf area index at sub-seasonal to seasonal scale for earth system model using deep learning (abstract)

9:40 am - Zhengui Wang - Physics-Based Machine Learning for Compound Flooding Forecast (abstract)

9:55 am - Augustin Vintzileos - Accelerating UFS Model Development Cycles Through a Process-Oriented, Machine Learning–Based Framework (abstract)

10:10 am **Morning Break**

10:20 am **High-Performance Computing (HPC)/Cloud**
Moderated by Jan Ising

10:20 am - Raghu Reddy - Accelerating UFS Development and R2O: NOAA's Research and Development High Performance Computing System (RDHPCS) (abstract)

10:31 am - Kristopher Booker - Accelerating Rapid Innovation: EPIC's Pre-Configured Public Cloud HPC environments (abstract)

10:46 am - Tom Robinson - Portable and Reproducible Earth System Modeling: Containerizing ESM4.5 Case Study (abstract)

10:57 am - Natalie Perlin - Build and run UFS Weather Model tests using containerized spack-stack libraries with GNU or Intel compilers on HPC platforms or cloud environments (abstract)

11:08 am - Lauren Chilutti - Leveraging Cloud-Integrated CI/CD and Containerization to Modernize NOAA's Workflows (abstract)

11:19 am - Hendrik Tolman - Developing WAVEWATCH IV™ to enable exascale wind wave modeling (abstract)

11:24 am - Discussion

11:30 am **"Pair and Share" Conclusion Activity**
Facilitated by
Alison Gregory, UCAR CPAESS
Jennifer Vogt, NOAA/OAR/WPO/EPIC

12:30 pm **Adjourn – Travel Home!**

Additional Training Opportunities

Workshops/Training & Room Location (All times Eastern Daylight Time)

Date	Time	Event	Location
Monday, July 20th	9:00 am - 12:00 pm	Training 1 – Introduction to Running NWS Global NWP Models in Development Model Register here	Room 1101/1102
Monday, July 20th	9:00 am - 12:00 pm	Training 2 -Adding Idealized Test Cases to the UFS Hierarchical System Development Framework Register here	Room 1105
Monday, July 20th	1:00 pm - 4:00 pm	Training 3 – Project EAGLE Training: Machine Learning Weather Forecasting on Microsoft Azure Register here	General Vessey Ballroom
Wednesday, July 22nd	1:45 pm - 4:15 pm	**NEW USER/STUDENT TRACK TRAINING** CADRE-JEDI-Edu Tutorial: Introducing Data Assimilation Methods in JEDI Using a Simple Two-Layer QG Model *This training is scheduled during the new user/student track on Wednesday of the workshop and is gauged towards those new to model code and development.* Register here	Room 1101/1102

Abstracts

Organized by session

Virtual Posters

Integration of Neural Operator and Physics-Based Ocean Models for Probabilistic Prediction of Tropical Cyclone Driven Coastal Flooding

Fariborz Daneshvar*, Atieh Alipour, Soroosh Mani, William J. Pringle, Kamyar Azizzadenesheli, Saeed Moghimi and Edward P. Myers

Tropical Cyclone (TC) driven surge causes widespread losses and damages to the coastal communities. Only in 2024, TC-driven surge in the Atlantics caused more than 400 fatalities and about $130 billion in damage. Early forecasting of the storm's extent and the areas at risk is crucial for timely evacuation and mitigation. The National Oceanic and Atmospheric Administration (NOAA) National Hurricane Center (NHC) has a probabilistic surge (P-Surge) framework for prediction of coastal flooding. It uses a large ensemble (~600) of water elevation simulations from a numerical model corresponding to a wide range of possible storm paths, sizes and strengths for probabilistic analysis. Our research in Machine Learning (ML) based probabilistic surge modeling proved promising in the predictions of TC-driven coastal flooding with a much smaller ensemble of hydrodynamic simulations (~20). This ML framework uses Karhunen–Loève expansions and Polynomial Chaos to develop surrogate models, which is not computationally efficient for continuous spatiotemporal predictions at different times in the storm progression. We propose a state-of-the-art deep learning model with neural operators to develop a fast and resolution invariant prediction of probabilistic coastal flooding based on limited hydrodynamic simulations. It uses coordinates, along with time series of projected wind and pressure fields of an ensemble of perturbed storm tracks, to predict time series of water elevations at any given location. The outcome of this work could benefit the early warning and evacuation ahead of TC-driven storms.

Strongly Coupled Assimilation of Atmospheric Boundary Layer Observations for the Rapid Intensification of Hurricane Ida (2021) using the CADRE Self-cycled HAFS-JEDI System

Yu-Shin Kim* and Xuguang Wang

This study investigates the development and application of air–sea strongly coupled data assimilation (SCDA) in the Hurricane Analysis and Forecast System (HAFS) using the Joint Effort for Data assimilation Integration (JEDI) framework. In a previous study, Kim and Wang (2026) showed that assimilating atmospheric boundary layer (ABL) observations constrains thermodynamic and dynamical conditions within the ABL, enabling forecasts to capture Category 4 intensity, compared with Category 2 when ABL observations were withheld. However, the analysis increments were limited to the atmosphere because traditional data assimilation updates only the single Earth

system component in which observations are taken. SCDA allows observations in one component (e.g., the atmosphere) to update another component (e.g., the ocean) through coupled background error covariances, enabling coherent updates across the air–sea interface. This study examines whether assimilating ABL observations can update both atmospheric and oceanic states and influence the rapid intensification of Hurricane Ida (2021). Results reveal strong correlations between atmospheric and oceanic conditions across the air–sea interface, suggesting the potential to improve tropical cyclone forecasts.

An Evaluation of Convective Schemes for Intensity and Track Forecasts for Typhoon Usagi

Margarita Mora*, Feng Hsiao, Zhan Zhang and Jong-Dao Jou

Super Typhoon Usagi (2024) formed over the North Philippine Sea on 9 November 2024 and underwent rapid intensification before making landfall in Luzon, Philippines. The Hurricane Analysis Forecast System (HAFS), as well as its legacy version, the Hurricane Weather Research and Forecasting (HWRF) model, demonstrated high uncertainty in both intensity and track at lead times of 48 to 72 hours. Neither model accurately captured the magnitude and timing of the rapid intensification, as well as the landfall locations in Luzon and Taiwan. Improving HAFS performance over the northwest Pacific is a critical component of advancing forecasting skill for tropical cyclones. In this study, the convective schemes were changed from the operational GFS Scale-Aware Simplified Arakawa Deep Convection Scheme (sa-SAS) and the GFS SAS-based Mass-Flux Scheme for Shallow Convection (sa-MF) to Tiedtke. Model output with Tiedtke showed improved minimum 10-m wind and track forecasts, along with a more organized structure with stronger inflows and outflows compared to the operational HAFS suite. Further testing of alternative physical and convective schemes will provide more insight on best methods of TC modeling over the northwest Pacific, specifically those that undergo rapid intensification or have environmental flows with high uncertainty.

Explicit Effects of Forest Canopy Shading and Turbulence in the Unified Forecast System-Air Quality Model

Irena Ivanova*, Patrick C. Campbell, Paul Makar, Wei-Ting Hung, Barry Baker, Youhua Tang, Zachary Moon, Rick Saylor, Fanglin Yang, Kai Wang, Raffaele Montuoro, Benjamin Koziol

The chemistry of the Earth's atmosphere close to the surface is known to be strongly influenced by vegetation. We introduce two important aspects of forests' physical, dynamical, and chemical environment in the operational Unified Forecast System-Short Range Weather-Air Quality Model (UFS-AQM) at NOAA: the impacts of forest shading on the reduction of photolysis rates and the modification of vertical transport due to the presence of foliage. We hypothesize that the lack of these processes may be a significant cause for systematic biases in boundary layer weather and air quality (e.g., ozone) forecasts over North America in and around contiguous forested regions within the continent.

Here we report our progress on including such canopy effects in an experimental version of NOAA's online-coupled UFS-AQM. We describe the contiguous forest canopy regions in North America with new global, gridded high-resolution (1 km) satellite-derived vegetative canopy datasets, and explicitly represent the canopy shading and turbulence effects on boundary layer ozone with three canopy sub-layers.

Our results show that the photolysis attenuation and the gas-phase chemistry within the forest canopy produces a realistic canopy shading effect that corrects in part the long-standing systematic positive ozone bias in the UFS-AQM. The splitting of the vertical diffusion processes that affect mass on the new sub-canopy layers is consistent with previously published work by our collaborators, and suggests that the relative importance of the resolved canopy turbulence versus the reduction in photolysis rates on boundary layer ozone in AQM is even greater. Further implications will be presented, including our efforts to transition the explicit sub-canopy parameterizations into a future version of the operational UFS-AQM at NOAA.

JEDI-Based Three-Dimensional Variational Data assimilation for the NOAA Global Chemistry and Aerosol Forecast System (GCAFS)

Yaping Wang*, Andrew Tangborn, Cory Martin, Barry Baker, Bo Huang, Mariusz Pagowski, Guillaume Vernieres, Jérôme Barré, Li Pan, Partha S. Bhattacharjee, Fanglin Yang, Li Zhang, Raffaele Montuoro, Daniel Holdaway

The National Oceanic and Atmospheric Administration (NOAA) is transitioning to its next-generation atmospheric composition framework, the Global Chemistry and Aerosol Forecast System (GCAFS). Integrating the FV3 dynamical core with the GOCART aerosol module, GCAFS provides high-resolution (~25 km) global aerosol predictions. A key advancement is the operational implementation of the Joint Effort for Data Assimilation Integration (JEDI) infrastructure within the system's data assimilation component, the Global Constituent Data Assimilation System (GCDAS). GCDAS utilizes a 3DVar-FGAT algorithm coupled with flow-dependent background error covariances. The aerosol forecast error covariance is modeled using a variance partition method for standard deviations and a diffusion-based approach for error correlations. The system's performance was evaluated over an extensive nine-month retrospective period (December 2023–September 2024), assimilating multi-satellite VIIRS AOD retrievals from the S-NPP, NOAA-20, and NOAA-21 platforms. Validation against independent ground-based AERONET observations and global reanalyses (NASA MERRA-2 and ECMWF CAMS) demonstrates skill improvements. Results indicate a 10–20% reduction in global AOD Root Mean Square Error (RMSE) relative to a non-assimilating free run. Notably, the system exhibits superior performance during extreme aerosol episodes, including significant error reductions during Saharan dust transport events and improved representation of organic carbon plumes during intense wildfire seasons. This work demonstrates the readiness of GCAFS for operational implementation and highlights the advantages of the JEDI framework in advancing atmospheric composition forecasting.

MDTF x UFS

Maria Gehne* (CIRES CU Boulder and NOAA PSL), Kevin Counts (CIRES CU Boulder and NOAA PSL) and Aparna Radhakrishnan (CIMES Princeton University and GFDL)

The Model Diagnostics Task Force (MDTF) Diagnostics package is an open-source platform that applies "process-oriented diagnostics" (PODs) to Earth system models. The MDTF-Diagnostics Framework allows subject-matter experts to easily develop and deploy these PODs for universal application to support model output data. It currently includes nearly 30 community-contributed PODs for atmospheric, oceanic, land, and coupled processes.

A tool like MDTF-Diagnostics would greatly benefit the Unified Forecast System (UFS) model development activities. However, the current framework is focused on climate applications and only supports free-running model data spanning multiple years. This design does not accommodate typical forecast datasets, which are frequently initialized and run only to specific lead times.

This work aims to extend the MDTF-Diagnostics package to support this "forecast paradigm". This extension will provide NOAA and the research community with a crucial platform for process-oriented evaluation of short-range to seasonal forecast models. We show how to read seasonal UFS forecasts from AWS in .zarr format into the MDTF Diagnostics package and apply select MDTF PODs to those forecasts.

Evaluation of Subseasonal Forecast Skill in a Tropical Island Environment: Implications for Forecasting Capability over East Nusa Tenggara, Indonesia

Muhammad Indra Bendi*, Fera Adrianita, Supari and Robi Muharsyah

Subseasonal precipitation forecasting over the Maritime Continent remains a challenge for global prediction systems due to strong local variability, complex land–sea interactions, and the influence of large-scale climate drivers. Regions composed of small islands with strong rainfall gradients provide a useful environment for evaluating the performance of forecasting systems, particularly at subseasonal lead times where predictability is limited. East Nusa Tenggara, Indonesia, represents one of these challenging regions, where reliable precipitation forecasts are important for water management, agriculture, and early warning applications.

This study evaluates the subseasonal forecast skill of global S2S prediction systems over East Nusa Tenggara using ERA5 reanalysis as a reference dataset (Hersbach et al., 2020). Weekly accumulated precipitation forecasts from global subseasonal prediction systems participating in the Subseasonal to Seasonal Prediction Project (Vitart et al., 2017) are verified at multiple lead times to assess how forecast skill changes from early to extended ranges. Anomaly-based verification metrics are used to better represent the model's ability to capture wet and dry conditions, which are more relevant for practical forecasting than absolute rainfall values. The analysis focuses on the peak rainy season, when precipitation variability is high and forecast errors can have significant impacts.

The results show that forecast skill decreases with increasing lead time, with useful skill generally limited to the first forecast weeks. However, forecast performance is not uniform and varies depending on the large-scale background state. Periods influenced by stronger large-scale variability tend to show higher predictability, while weakly forced conditions lead to rapid loss of skill. These findings indicate that subseasonal forecasting capability in tropical island environments is strongly conditional and sensitive to the representation of large-scale circulation in global prediction systems.

This regional evaluation highlights the importance of targeted verification in challenging environments and provides useful feedback for understanding the strengths and limitations of current subseasonal forecasting systems within the broader forecasting community.

Identifying the Error: Process-Oriented Verification in support of Model Development

Jeff Hamilton*, Molly Smith and Michelle Harrold

Verification is essential to the development and operational transition of NOAA-based numerical weather prediction modeling (NWP) systems, notably the High Resolution Rapid Refresh (HRRR), Global Forecast System (GFS), and prototype Rapid Refresh Forecast System (RRFS). To support model development efforts, GSL has developed and maintained a verification system, the Model Analysis Tool Suite (MATS), with three main requirements: responsiveness, flexibility, and consistency. Various verification types are included, with slightly different methodologies, but generally metrics are produced from partial aggregated statistics derived from model-obs pairs, including both real-time and retrospective sources. The derived partially aggregated statistics are stored within a database, which is accessible from an interactive web application.

While this approach to verification has historically been successful in the NOAA Research-to-Operations (R2O) process, changing requirements in NWP development are facilitating the need for more in-depth and targeted analysis than pre-generated partial sums can currently provide. Greater user control of metric generation and assessment of atmospheric processes is becoming crucial to finer-scale model development, as causes of systematic error become harder to determine when limited to large-scale static statistical aggregation. To meet this challenge, the GSL model verification team, in a collaborative effort with the Developmental Testbed Center's (DTC) METplus team, is developing new process-oriented verification capabilities that incorporate dynamic schema design, scalability, data driven processing, and enhanced evaluation methodologies, culminating in a shared platform for metric generation and evaluation. This approach aims to provide a high level of flexibility over the types of evaluations that can be done for various atmospheric situations, guiding NWP developers in their attempts to address deficiencies during the R2O development cycle.

In this presentation, we will highlight the ongoing model verification enhancements at GSL and DTC, specifically addressing the limitations of the legacy partial sum paradigm, as well as the impact of process-oriented verification capabilities in isolating the sources of NWP shortcomings in the development process. Additionally, we will explore future capabilities and uses for the system within numerical model development across the community.

Impact of High Resolution and Scale-Aware Physics on HAFS: Results from Idealized Tests and Hurricane Milton (2024)

Chuan-Kai Wang*, Weiguo Wang, Bin Liu, Zhan Zhang, Ping Zhu, Kwun Yip Fung, Xuejin Zhang, Jongil Han and Fanglin Yang

As numerical weather simulations move toward higher spatial resolutions, models increasingly operate within the "gray zone," where traditional turbulence and convection parameterizations become less reliable. Recently, the Scale-Aware Three-Dimensional Turbulent Kinetic Energy (SA3DTKE) GFS PBL scheme has been incorporated into the ufs-weather-model for the Hurricane Analysis and Forecasting System (HAFS) application (Zhu et al., 2025). In this study, impact of horizontal resolution and its combination with the SA3DTKE PBL and Scale-Aware Simplified Arakawa-Schubert (SASAS) convection schemes were evaluated across various horizontal resolutions (from km to sub-km grid spacings) by both idealized experiments and a real-case analysis of Hurricane Milton (2024).

Idealized results demonstrate that at higher resolutions, the SA3DTKE scheme produces stronger maximum winds and higher TKE near the eyewall compared to traditional 1D scheme, leading to an earlier onset of intensification. Furthermore, the inclusion of the SASAS convection scheme in idealized environments resulted in a stronger eyewall and earlier intensification compared to runs without SASAS at coarser resolutions. For the Hurricane Milton (2024) case, the implementation of SA3DTKE improved forecast skill for both track and intensity. The impact of the convection scheme was found to be resolution-dependent, with a more pronounced effect on track prediction at high resolutions and on intensity at lower resolutions. These findings highlight the critical role of scale-aware physics in enhancing the predictive capabilities of HAFS for both idealized simulation and operational tropical cyclone forecasting.

The NOS Coastal Cloud Sandbox: A Prototype for Operational Modeling in the Cloud

Breanna Vanderplow, Michael Lalime*, Tiffany Vance, Katherine Moore Powell, Patrick Tripp, Zachary Wills, Mykel Alvis, Micah Wengren, Minna Ho and Kelly Knee

Ensuring accurate and timely coastal and ocean forecasts is paramount to meeting NOAA's mission to understand and predict our changing environment, and subsequently sharing this information with the public for use in managing coastal and marine ecosystems and resources. Delivering these forecasts relies on the development and operation of models, which often face challenges including access to computational resources and efficient collaboration across organizations. To address these challenges, NOAA's Integrated Ocean Observing System and partners developed the Coastal Modeling Cloud Sandbox - a scalable, cloud-based high-performance computing platform, hosted by

Amazon Web Services, that facilitates collaborative model development, testing, optimization, as well as post-processing, validation, and data sharing among NOAA and external collaborators, streamlining the transition from research-to-operations.

While the Sandbox's cloud-based infrastructure circumvents limitations associated with traditional on-premises computing, other challenges have emerged, particularly in providing access to non-NOAA users. While navigating these challenges, the Sandbox hosted a variety of projects spanning NOAA operational and academic research models, demonstrating the benefits of cloud computing for operational modeling. Two recent use cases demonstrate the Sandbox's cost savings and efficiency: a 43-year total water level retrospective of NOAA's National Water Model and training of machine learning models to detect sensor measurement drift. The Sandbox currently serves as a prototype for cloud-based operational modeling, with the goal of advancing innovation and community-driven model development to ensure communities impacted by coastal hazards receive reliable, accurate, and timely forecasts.

Leveraging Deep Learning and Satellite Imagery for Enhanced Convective Parameter Diagnosis

Badreddine Alaoui*, Bari Driss and Chakib Bounoun

This study introduces a novel Semi-Observational Analog (AnOb) framework designed to diagnose convective parameters by integrating traditional analog forecasting with modern computer vision. Unlike conventional point-wise analog methods, the AnOb system utilizes deep learning to extract complex spatial features from high-resolution RGB satellite imagery. By employing three fine-tuned architectures—VGG16, Xception, and Inception-ResNetV2—the model generates feature vectors that represent atmospheric states. Using cosine distance, the system identifies historical analogs from a 2020–2023 training period to diagnose current convective conditions, specifically targeting CAPE, TCWV, Vertical Velocity (W600), and the Lifted Index (LI). Evaluated over the 2024 test period across the region 40°N–20°N and 20°E–4°W, the system was tested for lead times of 12, 15, and 18 hours. Results indicate that the AnOb approach consistently outperforms climatology and persistence benchmarks. In a focused sample of 100 events (split between extreme convection and non-convective cases), the model maintained low hourly mean biases, ranging from -200 to 200 J/kg for CAPE and -0.3 to 0.3 Pa/s for vertical velocity. Furthermore, a fuzzy logic-based verification yielded a Probability of Detection (POD) of 0.6 and a Brier Score of 0.4, demonstrating the system's robustness in identifying severe weather. This research presents a computationally efficient, high-performance tool that complements traditional numerical weather prediction for operational meteorology.

Improving the Intensity of Tropical Cyclones for the AIGFS

Linlin Cui*, Jun Wang, Jianjun Liu, Xiao Luo, Wei Li, Jiayi Peng, Zhan Zhang, Bin Liu, Lin Zhu, Qi Shi, Alicia Bentley and Jacob Carley

Recent advances in artificial intelligence and machine learning have significantly accelerated progress in numerical weather prediction (NWP), enabling data-driven models to complement and, in some cases, rival traditional physical-based models. Despite these advances, accurately representing high-impact, small-scale phenomena such as tropical cyclone intensity remains a persistent challenge.

A key limitation in many current AI-based deterministic models is the reliance on mean squared error (MSE) as the primary training objective. While MSE is effective for minimizing global average error, it tends to oversmooth spatial fields and underrepresent extreme events, leading to systematic underestimation of tropical cyclone intensity. This limitation is particularly problematic for hazard prediction, where accurate representation of extremes is critical.
In this study, we present improvements to tropical cyclone intensity forecasts in AIGFS through the introduction of a spectral-based adjusted MSE (AMSE) loss function and targeted fine-tuning using more recent analysis data. The AMSE loss function explicitly emphasizes errors across different spatial scales by directly aligning the power spectral density with the training data, thereby ensuring wave energy, sharp gradients, and high-frequency structures are preserved. Complementing this, fine-tuning with more recent GDAS data helps the model adapt to GDAS dynamics. Results demonstrate that the combined approach reduces tropical cyclone intensity bias compared to AIGFS version 1. Headline scores also show improvement over the previous AIGFS version for most variables and metrics.

Evaluating the Aurora Machine Learning Model Utilizing NCEP Aerosol Initial Conditions for PM2.5 Forecasting

Ellie Hojeily*, Jun Wang, Linlin Cui and Cheng-Hsuan Lu

Background/objectives
Aurora is a global machine learning foundational model created by Microsoft. Aurora learns the governing dynamics of the atmosphere through pre-training and can be fine-tuned to downstream tasks. For air quality prediction, Aurora is fine-tuned on CAMS Reanalysis and Analysis data. It is of interest to NOAA to evaluate Aurora's accuracy relative to its global atmospheric composition forecast, GEFS-Aerosols, and determine the feasibility of driving Aurora with GEFS initial conditions.

Methods
The June 5-10 2023 severe PM2.5 pollution episode was assessed to represent an extreme event. To determine output variables NOAA should consider in future iterations of GEFS to run Aurora in parallel, we evaluate Aurora's sensitivity to various initial condition configurations by omitting trace gases (i.e., CO, NOx). Aurora is then initialized using 3 configurations: (1) CAMS data only, (2) GEFS Aerosols only, and (3) GEFS Aerosols and meteorology.

Results

Preliminary results suggest O3, CO, and NOx are critical predictors for PM2.5 prediction as their omission can destabilize Aurora, especially as lead-time increases. Aurora has a tendency to disperse wildfire smoke by forecast hour 48 while GEFS-Aerosols sustains plumes due to wildfire emissions inventories. Generally, the best performing Aurora forecast uses GEFS Aerosols and meteorology initial conditions rather than CAMS.

Conclusions
We show that Aurora can complement GEFS by reducing forecast uncertainty by providing a secondary forecast outcome from the same set of initial conditions at significantly lower computation costs. Future efforts include fine-tuning Aurora on GEFS with emissions information.

Development of Operational Verification Capability for NOAA Artificial Intelligence Global Forecast System (AIGFS)

Qi Shi*, Alicia Bentley, Mallory Row, L. Gwen Chen, Jason Levit, Linlin Cui, Jun Wang, Jianjun Liu and Jacob Carley

The NOAA Artificial Intelligence Global Forecast System (AIGFS), launched in December 2025, represents a major step toward integrating artificial intelligence into operational forecasting. Built upon Google DeepMind's GraphCast machine learning framework and fine-tuned with NOAA Global Data Assimilation System (GDAS) data, AIGFS aims to complement traditional numerical weather prediction (NWP) systems by providing skillful forecasts with significantly reduced computational cost and latency. However, AI-based models exhibit distinct error characteristics and variable-dependent performance, making robust, standardized verification essential to establish trust and support operational adoption.

Operational AIGFS verification has been developed at the National Centers for Environmental Prediction (NCEP) through the Environmental Modeling Center (EMC) Verification System (EVS) v2.0. This system provides a comprehensive suite of statistical metrics and graphical products generated with the enhanced Model Evaluation Toolkit (METplus). Verification results are displayed on a public website to support real-time evaluation and intercomparison. Results from EVS v2.0 indicate that AIGFS achieved a statistically significant improvement relative to GFSv16 in anomaly correlation coefficient (ACC) for geopotential height during winter 2025–2026. In addition, AIGFS reduced 2-m temperature bias over the Continental United States (CONUS) for the same period.
Future work includes collaboration with the Developmental Testbed Center (DTC) to develop new verification metrics tailored to AI model evaluation. The evaluation of AIGFS v1.1 performance is also ongoing at EMC, with a particular emphasis on improving tropical cyclone intensity prediction.

Physics Sensitivity of All-Sky Infrared Radiance Assimilation in HAFS-JEDI and Its Impact on Tropical Cyclogenesis Forecasts

Feng Hsiao*, Xingchao Chen, Zhan Zhang and Sen Chiao

This study investigates how cumulus and microphysics parameterizations influence all sky infrared radiance assimilation and, in turn, affect forecasts of Hurricane Beryl's cyclogenesis using the Hurricane Analysis and Forecast System with the Joint Effort for Data Assimilation Integration system (HAFS JEDI). Experiments are performed over a basin wide 6 km domain, where both parameterization types strongly influence the simulated moisture distribution, convective organization, hydrometeor fields, and cloud structures.
We first conduct no data assimilation experiments to assess the sensitivity of the simulated storm to different physics configurations. Compared with the scale aware Simplified Arakawa–Schubert (saSAS) cumulus scheme, the Tiedtke scheme tends to over activate convection in the large-scale environment, producing diffuse outer rainbands, a less organized primary circulation, and a weaker developing storm. In contrast, saSAS confines intense updrafts closer to the inner core, yielding a more compact rainband structure and stronger vortex organization. For microphysics, we compare the single moment GFDL scheme with the double moment Thompson scheme. Thompson generally produces broader and colder cloud fields, likely linked to excessive upper-level snow and ice.
These physics dependent cloud and hydrometeor differences substantially affect the background fields used in all sky infrared radiance assimilation. Consequently, assimilating the same GOES all sky infrared radiances can produce different analysis increments and subsequent deterministic forecasts depending on the physics combination. The impact is most pronounced in cloudy regions, where prior hydrometeor, thermodynamic, and dynamical fields differ markedly among configurations. The physical mechanism linking physics dependent background errors, radiance increments, and cyclogenesis forecast impacts will be discussed.

Spatio-Temporal Characterization of Wildfire Dynamics Utilizing High-Resolution WRF-Based Fire Weather

Deepak Kumar* and Brian Ancell

Accurate wildfire risk mitigation depends critically on the high-resolution representation of atmospheric forcing variables; however, the transformation of raw numerical weather prediction (NWP) outputs into a coherent spatio-temporal depiction of fire potential remains methodologically underdeveloped. This study evaluates the performance of the Weather Research and Forecasting (WRF) model in resolving hourly fire-weather dynamics, with particular emphasis on the accuracy of key atmospheric state variables and the characterization of associated aleatoric uncertainty. The WRF model is configured at convection-permitting spatial scales ($\Delta x<=4$km) to simulate the evolution of primary meteorological drivers, including 2-m air temperature, specific humidity, and 10-m wind vectors. Model outputs are systematically compared against high-density in situ observations using established verification metrics, including root mean square error (RMSE), mean bias, and Pearson correlation coefficients. In addition, an error propagation framework is employed to quantify the nonlinear sensitivity of derived fire-weather indices to uncertainties in the underlying meteorological variables. Results indicate that WRF reproduces diurnal variability and broad spatial gradients with a high degree of fidelity. However, increased variability and reduced predictive skill are observed in the wind field, particularly during transient mesoscale phenomena. To ensure reproducibility and computational scalability, the analysis is implemented within a standardized workflow incorporating containerization and automated verification procedures. This work presents a structured framework for integrating high-resolution NWP outputs into multi-dimensional representations of wildfire dynamics, thereby supporting the development of spatio-temporally explicit and operationally relevant risk assessment tools.

Feedback Between Ocean Waves and Vertical Heat Distribution in the High North: A UFS-Based Modeling Approach

Cigdem Akan*, Theodore W. Letcher, Ali Abdolali, Sandra L. LeGrand and Denise Worthen

This work presents the development and implementation of a fully coupled model (atmosphere-ocean-wave-ice) within the Unified Forecast System (UFS)-Weather framework designed to investigate the interaction between ocean waves and vertical mixing processes in the High North. The coupled atmosphere-ocean-wave-ice system enables analysis of feedback mechanisms between the surface waves, ocean stratification, and heat transport, which are currently underrepresented in operational models. We assess the effect of wave-driven mixing on regional heat budgets and its implications for seasonal to sub-seasonal forecasting, especially in the context of declining sea ice and changing surface energy budgets. The model has been applied to selected high-impact events during 2024, with validation performed against available in-situ and remote sensing datasets, including but not limited to stationary NDBC and drifting SWIFT buoys, satellite altimeters, SWOT, and CFOSAT/SWIM data for wave, ice, and meteorological variables. This work advances UFS's capability in high latitude applications and contributes to NWP goals in physics and system architecture. The results will inform next-generation UFS configurations and enhance predictive skill for Arctic climate and weather.

Dynamic Meteorology-Aware Emissions to Improve UFS Air Quality Model Predictions

Patrick C. Campbell*, Bok Haeng Baek, Siqi Ma, Chi-Tsan Wang, Irena Ivanova, Zachary Moon, Youhua Tang, Kai Wang and Barry Baker

Background/Objectives: Emissions representation remains a key limitation in UFS atmospheric composition forecasting, as current approaches rely on static, offline inventories with limited sensitivity to meteorological variability. This mismatch introduces structural inconsistencies that degrade forecast skill, particularly for short-term air quality and high-impact events.

Under the NOAA Joint Technology Transfer Initiative (JTTI), we are advancing the research-to-operations transition of the Meteorology-aware Emissions (MetEmis) dynamic coupler to address this gap. MetEmis enables online, meteorology-aware emissions that evolve consistently with modeled atmospheric conditions within the UFS.

Methods: MetEmis onroad mobile emissions have been integrated into the NOAA Emission and Exchange Unified System (NEXUS) for the UFS-Air Quality Model (UFS-AQM). Implemented as an online extension within Harvard's Harmonized Emissions Component (HEMCO), MetEmis applies meteorology-dependent adjustments using ambient temperature and relative humidity, improving spatiotemporal emissions variability relative to static inventories. Ongoing development expands MetEmis to additional meteorology-sensitive and operationally relevant sectors, including agricultural livestock emissions, residential wood combustion, and fugitive dust.

Results: Preliminary results indicate improved boundary layer ozone predictions, particularly reduced summertime daytime overestimates in humid regions. Additional MetEmis sectors show substantial impacts on fine particulate matter ($PM_{2.5}$) predictions, with evaluations ongoing.

Conclusions: MetEmis advances fully coupled meteorology–chemistry modeling in the UFS and provides a scalable pathway for improving operational air quality forecasting in UFS-AQM Version 8 and next-generation atmospheric composition systems at NOAA.

Integrating Thermohaline Physics into STOFS 2D Global: A Heterogeneous Mode Splitting Approach with Global RTOFS and Regional Model Coupling

Joannes J. Westerink*, Aman Tejaswi, Damrongsak Wirasaet, Albert Cerronea, Shohei Ohata, Yuji Funakoshi, Greg Seroka, Saeed Moghimi, Edward Myers, Chris Massey and Margaret Owensby

Our team has advanced the capabilities of NOAA's Storm and Tide Operational Forecast System (STOFS) 2D Global by efficiently incorporating thermohaline processes. The STOFS 2D+ Global implementation has been designed to operate as a heterogeneous internal/external mode model with temperature and salinity fields derived and downscaled from Global RTOFS and used to force ADCIRC's baroclinic pressure gradient term in the depth integrated external mode. We report on the performance of this approach in achieving excellent values of the Topography of the Sea Surface (TSS) and 30 day mean sea level trends over three year simulations. The system captures the depression of mean sea level in the wake of a landfalling hurricane due to cold water upwelling and the generation of the cyclonic ocean gyre that impinges against the coast. We examine the energetics of both a 24 to 2.5 km and a 12.5 to 1.5 km model in replicating the HYCOM simulations. The system works remarkably efficiently with only about a 10% increase in costs as compared to the barotropic only mode. In order to penetrate inner coastal zones and the nearshore which are not resolved by HYCOM, we are working on adding ADCIRC based internal mode regional insets that operate within the global external model mesh. The boundaries of the regional internal mode models are driven by thermohaline data from HYCOM while the mesh exactly overlies the regional component of the Global external model mesh.

Improving the aerosol representation and microphysics coupling in NOAA's RRFSv2 numerical weather prediction model with MPAS dynamical core

Haiqin Li*, Jordan Schnell, Anders Jensen, Ravan Ahmadov, Joseph Olson, Clark Evans, Man Zhang, Guoqing Ge, Minsu Choi, Sudheer Bhimireddy, and Johana Romero-Alvarez

NOAA's next generation storm-scale (3-km resolution grid) Rapid Refresh Forecast System version 2 (RRFSv2) is currently under development with the MPAS dynamical core and an advanced physics suite, which includes the aerosol-aware Thompson-Eidhammer Microphysics Parameterization for Operations (TEMPO) microphysics scheme, MYNN-EDMF boundary layer scheme, RRTMG radiation scheme, and the RUC and Noah-MP land-surface models. In the current configuration, the TEMPO microphysics uses water-friendly (WFA) and ice-friendly (IFA) aerosol number concentrations from GOCART monthly climatologies and overly simplified empirical emissions (based on the climatological datasets) to activate cloud condensation and ice nuclei number concentrations. The WFA and IFA also affect radiation in the model. This method constrains the model such that it can only be used to study the effects of aerosols on radiation, clouds, and precipitation in terms of the aerosol climatology. In this study, we replaced the GOCART climatology (a 7-year GOCART simulation in 2001-2007) with a contemporary aerosol reanalysis derived from the MERRA2 climatology (2014-2023). Furthermore, we integrate the prognostic wildfire smoke, dust, sea salt and anthropogenic aerosol emissions together with the MERRA2 climatology to represent WFA and IFA for the TEMPO microphysics and radiation processes. Preliminary results demonstrating the impact of the improved aerosol representation on cloud microphysics and radiation processes will be presented.

Evaluation of Remotely Sensed Leaf Area Index (LAI) Using the Noah-MP Land Surface Model (Noah-MP LSM): Impacts on Surface Energy and Water Cycles over the Continental United States (CONUS)

Rongqian Yang*, Michael Barlage and Fanglin Yang

Leaf Area Index (LAI) is a critical dimensionless metric representing the one-sided green leaf area per unit ground surface area. Accurate representation of LAI is required for the modeling of many surface processes, including surface energy and water budgets, land-atmosphere interaction among others as they are important to daily weather forecast and climate studies in numeric models. This study presents an evaluation of the new National Environmental Satellite, Data and Information Service (NESDIS) high resolution LAI satellite product using the Noah MP land Surface Model (Noah-MP LSM), demonstrate its impacts on key surface energy and water budget terms over the Continental United States (CONUS), such as surface net shortwave and longwave radiations, soil moisture, evapotranspiration, snow depth and water storage terms when compared to the dataset currently used in the operational models, suggesting an improved/updated representation of land surface property is vital to numeric models and has higher priority.

Invited Presentations

Six Years of the Unified Forecast System Research to Operations (UFS-R2O) Project: Accomplishments, Challenges, and Insights

Vijay Tallapragada*

The Unified Forecast System Research to Operations (UFS-R2O) Project completed six years (2020-2026) of collaborative effort, significantly advancing UFS applications for operational use. Funded jointly by NOAA's NWS Office of Science and Technology Integration (OSTI) and the Oceanic and Atmospheric Research (OAR) Weather Program Office (WPO), the project fostered crucial partnerships between NCEP/EMC, NOAA research labs, NCAR, NRL, and various academic institutions.

The project was led by three representatives (NOAA operations, NOAA research, and the academic community) and structured around three integrated applications teams: Medium Range Weather/Subseasonal to Seasonal (MRW/S2S), Short Range Weather/Convection Allowing Modeling (SRW/CAM), and Hurricane Modeling. These were supported by cross-cutting teams focused on areas including Atmospheric Composition, Data Assimilation, and Modeling Infrastructure. The leadership structure featured shared responsibility between NWS operations and community representatives for each team.

UFS-R2O project achieved substantial successes, including the operational transition of several key systems along with numerous enhancements to UFS components and infrastructure:

- Hurricane Analysis and Forecast System (HAFS) and its annual upgrades.
- Rapid Refresh Forecast System (RRFSv1) for North America.
- A fully coupled Global Forecast System (GFSv17).
- A beta version of the Seasonal Forecast System (SFSv1).
- The Global Aerosol and Chemistry Forecast System (GCAFS).
- The EMC Verification System (EVS).

While successfully uniting UFS developers with a sharp focus on operational outcomes, challenges included maintaining consistent funding levels and ensuring alignment with the broader UFS team structure and community. Key lessons learned and insights from the project leadership will be shared, highlighting their potential application to guide future similar R2O projects.

Community Earth System Modeling Framework within NOAA's Earth Prediction Innovation Center - the first five years

Keven Blackman*, D. Alex Burrows, Farida Adimi, Cameron Book, and Aaron Jones

Over the past five years, the Earth Prediction Innovation Center (EPIC) has been at the forefront of advancing the Unified Forecast System (UFS), fostering a collaborative ecosystem that accelerates innovation to advance our nation's forecasting and modeling capabilities. Through a steadfast commitment to engaging the broader community, EPIC has not only provided the critical tools and knowledge necessary for innovation but has also worked tirelessly to break down barriers to progress in weather prediction.

The EPIC team has established a comprehensive framework designed to foster and facilitate community contributions to Numerical Weather Prediction (NWP), incorporating both physics-based models and those utilizing Artificial Intelligence/Machine Learning (AI/ML). This framework includes the development of an infrastructure for the expedited and automated verification of all community enhancements and submissions to weather modeling. Moreover, this structure integrates a Continuous Integration/Continuous Delivery (CI/CD) pipeline, refined management processes, Hierarchical System Development (HSD), community outreach and instructional events, advanced user support, community resources and tools, and video tutorials, with the EPIC community Portal serving as a central hub for engagement and support.

EPIC's overarching mission has remained steadfast: to cultivate an inclusive, collaborative environment where diverse perspectives are celebrated, and groundbreaking ideas take root. By harnessing the collective expertise and input of the community, EPIC has driven significant advancements in weather prediction science, improving the accuracy and reliability of forecasts that are essential for public safety and economic resilience. As we reflect on the progress of the last five years, EPIC remains committed to advancing weather prediction innovation for the benefit of all.

This presentation will focus on the pivotal role of community engagement and support in EPIC's success. We will offer a summary of the current status of our various community engagement initiatives, detailing the mechanisms and platforms established for collaboration, contribution, and feedback.

Model for Prediction Across Scales (MPAS) Integration

The NSF NCAR MPAS Program

Gretchen Mullendore* and Kate Fossell

The adoption of the Model for Prediction Across Scales (MPAS) has accelerated over the past few years across academic, private, and public sectors both domestically and internationally. Most recently, NOAA announced intent to implement a global, fully coupled 3-km Earth system prediction system using the MPAS dynamical core. NSF NCAR has recently launched the NCAR MPAS Program—a strategic program designed to strengthen alignment and coordination across MPAS development activities, partnerships, and community engagement to support the growing community of MPAS users and developers, including NOAA. This presentation will provide an overview of the NCAR MPAS Program, highlights of MPAS projects, key development updates and strategic directions of MPAS.

EPIC's Role in Facilitating the Integration of MPAS into UFS

Eric Aligo*, Ben Koziol, Jong Kim, Jennifer Vogt, and Maoyi Huang, Youngsun Jung, Ligia Bernardet, and Jacob Carley

At the annual American Meteorological Society (AMS) meeting in January 2026, the National Weather Service announced an ambitious goal for NOAA to adopt a global, fully coupled and convection allowing 3-km unified system using the Model Prediction Across Scales (MPAS) for both research and operations*. The Earth Prediction and Innovation Center (EPIC) will play a key role in accelerating MPAS into the Unified Forecast System (UFS), a community-based system that supports NOAA's operational modeling systems.

Over the past few years, the EPIC has developed a publicly accessible, community-based weather research modeling system. It has also established testing and development capabilities to foster strong collaboration across the Earth system modeling community. In addition, EPIC has created platforms and protocols that provide training and tutorials for developers and users with varying levels of experience. Leveraging the established community modeling infrastructure and community engagement framework, EPIC can enable continual enhancement of operational models by providing tools, resources, and shared knowledge to the broader community.

To support the desired adoption of MPAS for both research and operations, in partnership with National Center for Atmospheric Research, NOAA modeling laboratories and centers, EPIC's community modeling infrastructure will be enhanced to incorporate MPAS. This will include enhancements to pre-processing, postprocessing, verification and validation packages in the UFS framework augmented with documentation, user support and community events. EPIC's community engagement efforts will be expanded to include the MPAS community. The EPIC program will also provide project management to establish an agile model development team to accelerate model development, and funding support to accelerate infrastructure enhancement toward an MPAS-based numerical weather prediction system pending resource availability. In this presentation, we will provide updates on EPIC's progress on those fronts and solicit input for future directions.

Integrating the Model for Prediction Across Scales dynamical core into the Unified Forecast System Weather Model

Ligia Bernardet*, Dustin Swales, Grant Firl, Soren Rasmussen, and Lulin Xue

In January 2026, NOAA launched the ambitious "Moonshot" initiative: the deployment of an operational, global, fully-coupled, and convection-allowing 3-km Earth system prediction system, based on the Model Prediction Across Scales (MPAS) dynamical core (dycore). This groundbreaking objective, coupled with a prior NOAA intention to use the MPAS dycore for future operational implementations of the Rapid Refresh Forecast System (RRFS), necessitated a major expansion of the Unified Forecast System (UFS) atmospheric model to include the MPAS dycore.

- Recent UFS Weather Model development has focused on a number of key areas:
- Establishing preliminary code management and testing for MPAS within the UFS ecosystem Integrating MPAS into the UFS Weather Model (WM) build system
- Developing the capability to initialize MPAS within the UFS WM
- Utilizing the MPAS native Parallel Input/Output (IO) capability within the UFS
- Enabling the MPAS-enabled UFS WM to run in dycore-only mode
- Reducing system dependencies by adopting common infrastructure

Additionally, work is in progress to connect the MPAS-enabled UFS WM to the Common Community Physics Package (CCPP) to access parameterizations needed for representing physical processes.

In this presentation, we will detail the progress made to date toward creating a prototype suitable for future operational implementations. We also discuss how this effort interfaces with and builds on current NOAA efforts that utilize the MPAS standalone outside of the UFS context.

Choose Your Adventure: The Many Ways to Integrate MPAS within ESMs

Ann Tsay*, Dan Amrhein, Ben Barr, Ligia Bernardet, Jacob Carley, Brian Dobbins, Michael Duda, Aubrey Dugger, Falko Judt, Sujay Kumar, Tim Lahmers, Jesse Nusbaumer, Robert Oehmke, Patricia Parker, Soren Rasmussen, Dan Rosen, Daniel P Sarmiento, Dustin Swales, Ufuk Turuncoglu, Sheri Voelz, and Denise Worthen

The Model for Prediction Across Scales (MPAS) is a collaborative project for developing earth-system simulation components for use in geophysical modeling applications. The atmosphere component of MPAS, MPAS-Atmosphere (MPAS-A), includes a dynamical core and atmospheric physics parameterizations. Growing interest in coupling MPAS with other earth science models (i.e. hydrology, ocean, and fire) and external modeling systems, such as NOAA's Unified Forecast System Weather Model (UFS-WM), NASA's Land Information System (LIS), and NCAR's Community Earth System Model (CESM) including regional applications using CROCODILE*, has led to at least three distinct software architecture designs, each shaped by the requirements and constraints of the respective community.

All three architectures share a common coupling infrastructure based on the Earth System Modeling Framework (ESMF), which provides a standardized superstructure for Earth system applications. Built on ESMF, the NUOPC** layer serves as an interface to specify how components interact and exchange data within a coupled system.
This presentation aims to describe these three ESMF/NUOPC-based coupling architectures for MPAS from a software engineering perspective. It includes performance comparisons of various configurations across different driver systems. Additional technical considerations affecting

coupling performance, such as the choice between single and double floating-point precision, will also be discussed.

Because each community operates under unique infrastructure constraints and has distinct scientific requirements, this presentation is intended to provide examples and help the scientific communities to identify the most suitable approach for integrating MPAS-A or an MPAS subcomponent (such as the dynamical core) into their system.

Transitioning to an MPAS-JEDI-based Warn-on-Forecast System (WoFSv2): Next-Generation Storm-Scale Probabilistic Guidance

Nusrat Yussouf*, Craig Schwartz, Soyoung Ha, Yunheng Wang, Kent Knopfmeier, Thomas Jones, Joshua Martin, Brian Matilla, Louis Wicker, Patrick Skinner, Ted Mansell, Miranda Silcott, Holly Obermeier, Derek Stratman, Christopher Kerr, Chengsi Liu, Adam Clark, Corey Potvin, Patrick Burke, Israel Jirak, and Jim Nelson

This study describes the development and transition of the current Warn-on-Forecast System Version 1 (WoFSv1), based on the Weather Research and Forecasting (WRF) model and a GSI Ensemble Kalman Filter (EnKF), to a framework built on the Model for Prediction Across Scales (MPAS) dynamical core and the Joint Effort for Data Assimilation Integration (JEDI). As NCEP's Environmental Modeling Center (EMC) adopts MPAS and JEDI as the foundation for future operational data assimilation, migrating WoFS is a critical step toward aligning storm-scale research with emerging operational standards.

To support R2O transition, testbed demonstrations are planned during 2026–2028 in NOAA's Hazardous Weather Testbed (HWT) Spring Forecasting Experiment (SFE) and Hydrometeorology Testbed (HMT) Flash Flood and Intense Rainfall (FFaIR) Experiment. This study is the first testbed demonstration of a WoFS framework using the MPAS dynamical core. The goal is to establish a regional, 15-min cycling ensemble system that provides 0–6-h probabilistic guidance for severe thunderstorms and associated hazards, including tornadoes, large hail, damaging wind, lightning, and flash flood producing rainfall. A key component is evaluation of the MPAS-based system against the current WRF-based WoFSv1 in NOAA testbed settings. The initial WoFS-MPAS prototype uses NCAR's DART DA software.
This presentation highlights findings from the testbed evaluations. Evaluation includes both subjective forecaster feedback and objective verification. Ongoing work is transitioning the data assimilation to NCEP's JEDI-based ensemble system. This presentation also provides an update on the status of that transition of the data assimilation workflows into JEDI, helping establish a prototype pathway toward WoFSv2.

NOAA/GSL Model Development and Forecasting Activities using MPAS for Future NOAA Operational Modeling Systems

Clark Evans*, Curtis R. Alexander, Ligia Bernardet, Terra T. Ladwig, Ming Hu, David C. Dowell, Ravan Ahmadov, and Guoqing Ge

NOAA's Global Systems Laboratory (GSL) seeks to improve Earth-system prediction across all weather-related hazards including severe convective storms, intense rainfall, winter storms, tropical systems, and phenomena such as wildfire smoke, dust, and pollen. GSL has a rich legacy in developing modeling systems for NOAA operations, including the WRF-ARW–based Rapid Refresh (RAP) and High-Resolution Rapid Refresh (HRRR), and in contributing to the development of Finite-Volume Cubed-Sphere (FV3)-based Unified Forecast System (UFS) applications such as the Rapid Refresh Forecast System (RRFS) and Global Forecast System (GFS).

There is increasing stakeholder need for high-resolution modeling systems that can explicitly resolve precipitating phenomena, fire-weather behavior, aerosol radiative and cloud feedbacks, and urbanization influences to improve forecast skill and serve end-user predictive needs. These considerations inform NOAA's Modeling Moonshot goal to develop a global, storm-scale, coupled atmosphere-ocean ensemble forecast system using the Model for Prediction Across Scales (MPAS). For over two years, GSL has been using MPAS to develop prototypes for future NOAA operational modeling systems to address these needs.

This presentation will provide an overview of MPAS development and forecasting activities led by GSL. These include high-resolution regional MPAS development toward RRFSv2 and purpose-built MPAS-based high-resolution forecast systems for atmospheric rivers, aerosols, fire weather, and tropical cyclones. Forecast performance from real-time and NOAA Testbeds evaluations will be presented. The presentation will also discuss supporting developments and applications at GSL, including JEDI-based data assimilation toward RRFSv2 and other applications, traditional and stochastic physics development, and integrating the MPAS dynamical core in the UFS.

MPAS evaluations for severe weather forecasting during the 2026 NOAA/Hazardous Weather Testbed Spring Forecasting Experiment

Adam Clark*, Kent Knopfmeier, Yunheng Wang, Nusrat Yussouf, Israel Jirak, Louis Wicker, Clark Evans, Terra Ladwig, Ming Hu, David Dowell, Craig Schwartz, and Soyoung Ha

For the fourth consecutive year, various configurations of the Model for Prediction Across Scales (MPAS) were tested for severe weather forecasting applications during the 2026 NOAA Hazardous Weather Testbed Spring Forecasting Experiment, which was conducted 27 April - 29 May. With the recent announcement that MPAS will be the future dynamical core for global and regional Unified Forecast System (UFS) applications, these evaluations are critical for gauging forecast utility and accelerating development. The experimental MPAS configurations were contributed by the NOAA/National Severe Storms Laboratory (NSSL) and the NOAA/Global Systems Laboratory (GSL). NSSL's contributions included three CONUS-domain, 3-km grid-spacing configurations initialized from the Rapid Refresh Forecast System version 1 (RRFSv1) and High-resolution Rapid Refresh (HRRR). Additionally, for the first time in the HWT, NSSL tested an MPAS-based configuration of the Warn-on-Forecast System (WoFS). GSL provided two MPAS configurations that are being tested to form the foundation for RRFSv2. These configurations contain the most up-to-date physics suites tuned for MPAS with one configuration initialized from RRFSv1, while the other uses its own fully-cycled JEDI-based data assimilation system.

Daily model evaluations assessed performance characteristics alongside the HRRR, RRFS, High-Resolution Ensemble Forecast System (HREF), RRFS Ensemble Forecast System (REFS), and other experimental systems. Ultimately, these tests are helping advance the use of MPAS within NSSL's Warn-on-Forecast System (WoFS) and within NOAA's UFS. This talk will present preliminary results from these evaluations and highlight notable cases of interest.

Poster Session I

Model Intercomparison of UFS and MPAS Forecasts for a Severe Convection Initiation Event: A Case Study

Ayoola Abe*, Nakul N. Karle, and Sen Chiao

Prior studies on convection-permitting weather models (CPMs) such as MPAS and UFS have shown that convective-scale errors can grow and propagate upscale rapidly, leading to a reduction of forecast skill for storm-scale convective features beyond three days. Understanding their performance in the short range is therefore critical for improving high resolution convective-scale forecasts. In this study, we compare the forecast performance of MPAS and the UFS Short-Range Weather (SRW) applications. MPAS employs a high-resolution, nonhydrostatic approximation using unstructured spherical centroidal Voronoi meshes and C-grid discretization, supporting both global and regional high-resolution simulations. UFS uses the FV3 dynamical core on a structured latitude longitude cubed-sphere grid to produce high-resolution regional forecasts. A high-impact severe weather event on 16-17 May 2017 was investigated, when the interaction of a dryline and a cold front triggered severe storms producing large hail and tornadoes Southern great plains. Model forecasts are verified against radar reflectivity using metrics that assess initiation, timing, movement, and placement of convection. Preliminary results without Data Assimilation indicate that both models reproduced similar initiation, timing, and movement in both locations. MPAS initiated convection roughly one hour later than observed, while UFS overestimated storm intensity throughout the simulation. Additional experiments incorporating an ensemble Kalman filter (EnKF) data assimilation system with MPAS have been conducted, and results will be analyzed to evaluate potential forecast improvements.

Impacts of a Modified Land Surface Representation in UFS/SCM/NoahMP Simulations

Michael Barlage*, Zhichang Guo, Weizhong Zheng, and Rongqian Yang

The current NoahMP implementation in UFS and MPAS assumes a two-tile representation of the surface where the two tiles are quasi-independent. Separate energy balances are calculated over a vegetated and non-vegetated surface, then area-averaged to send back to the atmosphere. These tiles overlay a single soil column, where the total soil/snow interface flux is used as a top boundary condition for soil/snow thermodynamic updates. In addition, for the vegetated tile, the energy balance is calculated first for the canopy and then for the sub-canopy. This approach, though conservative, can make the calculated fluxes inconsistent with the diagnostic temperature/moisture calculated in the model. Furthermore, due to the low roughness length of bare ground and decreased turbulent interactions below canopy, this method can result in a decoupling of the soil with the atmosphere.

A new approach is implemented and tested where a single "blended" tile is used to represent distributed vegetation elements producing a single energy and water flux for the model tile. The method also attempts to better represent land-atmosphere interactions through approximating the impact of turbulent mixing between canopy elements resulting in stronger coupling between the atmosphere and the soil. To achieve a better representation of the surface, increased communication and collaboration with NOAA/NESDIS has been established to provide additional datasets to estimate the vegetation cover using satellite imagery.

As part of this work, a hierarchical testing environment in the CCPP single column model (SCM) was developed to assess systematic impacts of model changes on land atmosphere coupling. This presentation will show the impact of incorporating a new representation of the NoahMP vegetation structure in SCM simulations evaluated against surface temperature/moisture measurements and flux tower observations.

Updates on the UFS-Arctic Application Development

Kristin N. Barton*, Amy Solomon, I-Kuan Hu, Lisa Bengtsson, Christopher J. Cox, Dmitry S. Dukhovskoy, Philip Pegion, Evelyn Grell, Jian-Wen Bao, Sara Michelson, Matthew Shupe, and Michael Gallagher

The NOAA Physical Sciences Laboratory (PSL) is continuing to develop a fully coupled Arctic application within the Unified Forecast System (UFS) for short- to medium-range forecasts. We present updates on the model configuration, initialization, and Arctic-focused physics parameterizations. Key configuration upgrades include refining the atmospheric grid resolution from ~50 km to ~10 km and transitioning the shared ocean-ice grid from 10 km to a 0.08° (3–5 km) Arctic cap based on the Real-Time Ocean Forecast System (RTOFS) grid. This will also allow for direct initialization from RTOFS data without interpolation. The system successfully executes 10-day hindcasts for both summer and winter test cases. We evaluate these hindcasts against observational datasets (e.g., MOSAiC) and NOAA/PSL's experimental Coupled Arctic Forecast System (CAFS) model, which currently produces daily Arctic forecasts. While the components are initialized using UFS Replay data, we demonstrate the impacts of supplementing this process with an offline Noah-MP land model spin-up and offline corrections to the initial state of snow on sea ice. Additionally, physics improvements include work towards integrating the Parameterization of Unified Microphysics Across Scales (PUMAS) cloud microphysics scheme into the Common Community Physics Package (CCPP) and developing a new mixed-layer parameterization scheme to better represent top-down mixing in Arctic stratocumulus. Future work aims to enable real-time forecasting, explore data assimilation for initialization, and test additional components (e.g. GOCART aerosol model or Wavewatch III).

NOAA's EAGLE Project: Advancing AI/ML Weather Forecasting from Research to Operations

D. Alex Burrows*, Mariah Pope, Anil Kumar, Jacob Carley, Paul Madden, Emily Carpenter, Sergey Frolov, Tim Smith, and Keven Blackman

With advances in computational power and high performance computing platform availability, AI/ML forecast model development has expanded considerably among government agencies, universities, and private companies. These AI/ML forecast systems, once trained, can be faster, cheaper, and less complex to run than their physical-based counterparts. At NOAA, AI weather forecast model development is coordinated under project EAGLE, or Experimental AI Global and Limited-area Ensemble forecast system. Over the past couple of years, NOAA's Environmental Modeling Center has begun development of AI/ML model emulators for its flagship Global Forecast System (GFS) and Global Ensemble Forecast System (GEFS) for deterministic- and ensemble-based atmospheric forecasts, coined AIGFS and AIGEFS. NOAA's AI forecast system development spans from retrospective to near-real time, and operational initiations, using different AI/ML architecture and workflows. Here, we present a synopsis of NOAA EAGLE's evolution, current capabilities, and future directions, and highlight contributions from the Earth Prediction Innovation Center, including public releases of two global-nested EAGLE applications running in 1) retrospective mode and 2) near-real time mode.

Known Unknowns: Impacts of Sampling Variability on ENSO Teleconnections

Benjamin Cash*, Chul-Su Shin, Erik Swenson, and Jim Kinter

Skillful subseasonal to seasonal (S2S) prediction depends critically on understanding and accurately representing the influence of long-lived boundary anomalies, such as the El Niño-Southern Oscillation (ENSO). A tremendous body of literature has been developed around identifying and quantifying these remote influences (teleconnections) on locations of forecasting interest. However, diagnosing these links from observations involves considerable uncertainty. Among the limitations of the observed record (length, changes in observing systems, etc.) is the fact that it is only a single 'realization' of the climate system. By its nature, every observation will include variability that is unrelated to the phenomena of interest (ENSO, Madden-Julian Oscillation, etc.) but may project onto any diagnosed teleconnections. This in turn can produce inaccurate assessments of the strength and location of the teleconnection centers of action and their utility as S2S predictors.

In this work, we use a series of ensemble hindcasts from prototype versions of the new NOAA Seasonal Forecast System, as well as the GEM5-NEMO model, to assess the impact of uncertainty and unrelated variability on diagnosed ENSO teleconnections. We find that teleconnections can vary substantially in both magnitude and pattern based on the particular realization of atmospheric noise chosen, sufficient to impact both quantitative and qualitative assessments of their utility as predictors.

Multi-Agent in Hydrodynamics: An Agentic Approach to Scientific Workflow Management

Albert Cerrone*, Jinpai Zhao, Clint Dawson, and Joannes J. Westerink

Scientific workflows across the geosciences routinely demand the simultaneous integration of heterogeneous data sources spanning differing formats, resolutions, agency mandates, and physical domains. Single-agent large language model (LLM) systems struggle with this complexity, as accumulating tool specifications and observational traces saturate context windows and degrade reliability. Hydrodynamics is a particularly acute example as reconciling tide gauges, bathymetric surveys, hydrodynamic model outputs, and regulatory flood products into a coherent analytical pipeline strains any monolithic architecture. We present a multi-agent system (MAS) prototype that addresses these limitations through a dynamically generated Layer Execution Graph (LEG) orchestration strategy, demonstrated here for hurricane-induced storm surge but readily extensible across geoscientific subdomains invoking observational, hypothetical, and forecast spatiotemporal data.

Rather than following a fixed pipeline, a planner agent constructs query-specific execution topologies at runtime from natural-language routing heuristics, adapting graph structure (viz., specialist selection, layer ordering, and parallel track decomposition) to the demands of each individual query. Consolidator agents fuse parallel outputs between layers to protect downstream context windows, while a reporter agent synthesizes final responses. A structured audit trail supports data provenance throughout execution.

Evaluated on 37 queries spanning six complexity categories using Claude Sonnet 4.6, the system achieves 93.6% factual precision with a 100% pass rate and zero execution errors. Ablation studies confirm that adaptive graph generation is the single most consequential architectural decision, with a fixed pipeline degrading factual precision by 28.7 percentage points. They further reveal that while consolidation is necessary, overly aggressive compaction erodes source attribution by shedding technical tokens (e.g., station IDs, vertical datums, timestamps) that make outputs auditable. The LEG mitigates this tension: by situating every tool call within a graph node carrying explicit role and dependency information, it endows provenance logs with structural context that transforms isolated retrieval events into purposeful, traceable acts within a defined scientific workflow.

Wave Evaluation for the next Global Forecast System GFSv17

Ming Chen*, Jessica Meixner, Saeideh Banihashemi and Mark Willis

This study focuses on the evaluation of the wave component, WAVEWATCH III (WW3), within the Global Forecast System (GFSv17), which is a UFS-based fully coupled atmosphere-land-ocean-ice-wave system. Compared to the current operational version (GFSv16), which features one-way coupling from the atmospheric model to WW3, GFSv17 introduces two-way interactions. In this new configuration, WW3 receives inputs from the atmosphere, ocean, and ice components, and in turn provides feedback to both the atmosphere and ocean components. Another major upgrade to the WW3 component is the adoption of unstructured grids, enabling improved scalability. Evaluation results of the retrospectives are presented to quantify the impact of these updates, along with a discussion of potential future improvements and development.

OceanAI: A Conversational Agent for Accurate, Transparent, Near-Real-Time Oceanographic Insights

Bowen Chen*, Ruoying He, Paul Liu, DongKuan Xu, Greg Dusek, Rob Redmon, and Patrick Hogen

Background:
Large language models (LLMs) can generate inaccurate or unsupported scientific information, limiting their reliability in oceanographic applications. This study presents OceanAI, a conversational platform designed to combine LLM capabilities with authoritative NOAA oceanographic data sources.

Methods:
OceanAI integrates function-calling frameworks with modular NOAA API pipelines. User queries trigger real-time retrieval of NOAA observations and reports. Retrieved data are automatically parsed, visualized, and annotated with timestamps, units, station metadata, and source references before being incorporated into natural-language responses.

Results:
In blind comparisons with three general-purpose LLMs, OceanAI consistently provided verifiable answers linked to original NOAA data sources and processing workflows, while competing models often declined to answer or generated unsupported values. The platform successfully handled water-level observations, sea surface temperature trends, and other environmental data queries with improved transparency and reproducibility.

Conclusions:
By grounding responses in live NOAA observations, OceanAI improves reliability and trust in AI-assisted oceanographic analysis. The platform demonstrates potential applications in marine hazard forecasting, ecosystem assessment, and water-quality monitoring, while providing a scalable framework for trustworthy scientific AI systems."

Systematic Evaluation of NOAA's Artificial Intelligence Modeling Suite at the National Centers for Environmental Prediction

L. Gwen Chen*, Jason Levit, Jun Du, Qi Shi, Alicia Bentley, Jun Wang, Jacob Carley, Linlin Cui, Bo Cui, Jianjun Liu, Bing Fu, and Hong Guan

On 17 December 2025, NOAA launched a new suite of operational, artificial intelligence (AI)-driven global weather prediction models, including the Artificial Intelligence Global Forecast System (AIGFS), the Artificial Intelligence Global Ensemble Forecast System (AIGEFS), and the Hybrid Global Ensemble Forecast System (HGEFS). The AIGFS is a new AI-based deterministic system and the AIGEFS is an AI-based ensemble system that provides a range of probabilistic forecasts to meteorologists and decision-makers. Both the AIGFS and AIGEFS leverage Google DeepMind's GraphCast model, fine tuned with the Global Data Assimilation System (GDAS) analyses, and use initial conditions from their physics-based counterparts, the Global Forecast System (GFS) and Global Ensemble Forecast System (GEFS). The HGEFS is a hybrid "grand ensemble" that combines the new AI-based AIGEFS with physics-based GEFS.

To evaluate AI model performance, new real-time forecast verification tools were developed at the National Centers for Environmental Prediction (NCEP), as part of the EMC Verification System (EVS). The EVS provides a wide range of statistical metrics and graphics for NCEP's operational modeling suite using the METplus software package. Quantitative and qualitative analyses are also conducted to investigate the differences between the AI models and physics-based models. Additionally, case studies are carried out to perform in-depth analysis on selected high-impact events, such as the United States East coast winter storm of 24-25 January 2026.

In this presentation, we will describe the current framework and system for evaluating the AI modeling suite at NCEP, as well as future development plans. We will also provide a comprehensive skill assessment of AI models using operational forecasts and performance metrics from the EVS, as well as ad-hoc diagnostics. Results show that AIGFS/AIGEFS improve forecast skill over the traditional, physics-based GFS/GEFS for many large-scale features. By combining two different modeling systems (one physics-based, one AI-based), the HGEFS creates a larger, more robust ensemble that more effectively represents forecast uncertainty. As a result, the HGEFS consistently outperforms both the GEFS and the AIGEFS across most major verification metrics. Greater differences between the AI-based and physics-based models are observed near the surface. Additionally, AI model performance during the winter storms of 2026 will be showcased to illustrate their strengths and weaknesses. Lastly, we will summarize knowledge gained from these analyses and areas of improvement for the next version of AI models at NCEP.

Ensemble Nested EAGLE and Probabilistic Forecasting using Anemoi

Tim Smith, Joshua DaRosa*, Sergey Frolov, Chris Miller, Mohammad Alam, Nick Silverman, and Conor Lewellyn

Building off the strong results in the deterministic Nested EAGLE variable resolution AI atmospheric forecast model, MITRE and NOAA PSL are developing the Nested EAGLE Ensemble. This is a probabilistic variable resolution AI atmospheric forecast model built on the Anemoi framework. The model architecture retains a 25 km global resolution with a 6 km refinement over the continental United States (CONUS), trained on a combination of operational GFS and HRRR analyses with a 6 hour temporal resolution. To generate calibrated probabilistic forecasts, multiple noise injection schemes have been explored within a CRPS-based ensemble framework. The forecast quality has been assessed through mean RMSE, spread skill ratio (SSR) and spectral analysis to characterize the distribution of forecast energy across spatial scales.

Extending this work, we investigate the incorporation of the multi-scale CRPS loss function within the Anemoi framework alongside the effects of extended training duration, auto-regressive rollout length, and finer temporal resolution on the model's ability to resolve small-scale atmospheric structures over CONUS — with a particular focus on precipitation fields.

A weather-regime-dependent verification metric for ensemble forecasts: Ensemble Performance Index for High-Impact Weather

Jun Du* and Jason Levit

In ensemble forecasting, the spread-error relationship is a central measure of the quality of an Ensemble Prediction System (EPS). First, we demonstrate that important spatial features are not effectively captured and even mistakenly interpreted by current verification measures. To address this, a new weather-regime-dependent verification metric, the Ensemble Performance Index for High-Impact Weather (EPI4HIW), is introduced. This new measure can quantitatively describe the spatial structure of the spread-error relationship and effectively reveal deficiencies in spread, especially over error-prone, high-impact weather regions. Examples are provided using both the NCEP Global Ensemble Forecast System (GEFS) and its data-driven counterpart, an AI-model-based global ensemble forecast system (AIGEFS).

Stratospheric Polar Vortex Representation in the Unified Forecast System's Subseasonal Prototypes

Elena Maria Fernández*, Zheng Wu, Andrea Lopez Lang and Mark Olsen

A well-resolved stratosphere is a critical component of Earth system models, as wintertime stratospheric dynamics provide key information for understanding atmospheric teleconnections and improving subseasonal-to-seasonal (S2S) predictability. However, S2S models differ substantially in their representation of the stratosphere and in how stratospheric information is leveraged in forecasts. This study evaluates how well the NOAA's Unified Forecast System (UFS) S2S prototypes represent the internal variability of the stratospheric polar vortex and its downward influence on Northern Hemisphere 1000-hPa temperatures. The UFS S2S prototypes vary in vertical extent and

resolution and physical parameterizations, including gravity wave representations, motivating an assessment of how these design choices impact stratospheric fidelity and forecast utility.
A total of 56 reforecasts from four UFS S2S prototypes are analyzed for the 2011/2012 and 2017/2018 winter seasons (November – March) and compared against ERA-5 reanalysis. Stratospheric polar vortex geometry and evolution are quantified using a suite of vortex-centric ellipse metrics detailed in Fernández et al. (2026, Under Review). The ellipse metrics are designed to capture enhanced stratospheric variability such as sudden stratospheric warmings and strong vortex events. Mean forecast errors of the ellipse metrics are evaluated over 36-day forecasts. Results indicate that UFS prototypes 5, 6, and 8 exhibit comparatively strong ability in representing the shape, evolution, and strength of the stratospheric polar vortex. Prototype 7 underperforms relative to the other prototypes, and often tends to simulate a weaker, more shifted vortex when compared to ERA-5 reanalysis.

To confirm potential forecasts of opportunity from the stratosphere and explore the possibility of further improving UFS surface predictions on S2S timescales, the UFS-derived stratospheric vortex metrics during the first 10 forecast days from each reforecast are used as input to a random forest (RF) model. The RF generates probabilistic predictions for the sign of regional 1000-hPa temperature anomalies 14 days later. Predictions produced using UFS-based stratospheric metrics for the 2011/2012 – 2017/2018 period are compared against both ERA-5 verification and the original UFS surface forecasts for the corresponding lead times. Ultimately, this analysis demonstrates that UFS S2S forecasts of the stratosphere capture physically consistent signals linking stratospheric variability to the sign of surface temperature anomalies, revealing forecasts of opportunity on S2S timescales.

PepC-Global: A Basin-Tuned, Environment-Dependent Probabilistic Tropical Cyclone Model

Cong Gao* and Ning Lin

The Princeton environment-dependent probabilistic tropical cyclone (PepC) model, initially developed by Jing & Lin (2020), generates synthetic tropical cyclones (TCs) to support robust TC risk assessment for the North Atlantic basin. PepC integrates three components: a hierarchical Poisson genesis model, an analog-wind track model, and a Markov intensity model. Here, we present PepC-Global, an enhanced and extended version of the original PepC model tailored specifically to various oceanic basins, including the western North Pacific, eastern North Pacific, North Indian Ocean, South Indian, Australia, and the South Pacific.

PepC-Global extends the original PepC. In the hierarchical Poisson genesis model, PepC-Global employs a support vector regression (SVR) model, in place of the linear model, to quantify the Poisson parameter, which significantly improves the representation of temporal variability and spatial patterns. The analog-wind track model is complemented by a pure-wind track model, characterized by greater sensitivity to advection flow and independence from prior track information.

Furthermore, the Markov Environment-Dependent Hurricane Intensity Model (MeHiM) demonstrates robust performance in predicting intensity changes, lifetime maximum intensity, and landfall intensity across multiple basins, extending its proven effectiveness from the North Atlantic to other oceanic regions.

Distinct from other global-scale TC risk models such as the Emanuel (2008) statistical-deterministic approach and Lee et al. (2018) Columbia Hazard model (CHAZ), PepC-Global uniquely combines a global modeling framework with basin-specific tuning. This ensures the model effectively represents basin-specific climatological and environmental variabilities, significantly enhancing predictive accuracy and reliability for each basin. Validation against historical TC records demonstrates significant improvements in representing basin-specific characteristics, including annual TC frequency, spatial distribution of genesis locations, track variability, and intensity lifecycle. Comparative analysis between basin-tuned and non-basin-tuned simulations highlights substantial differences in modeled cyclone characteristics. Specifically, the basin-tuned model exhibits enhanced accuracy in representing regional climatic and oceanographic influences on cyclone behavior, such as variations in sea surface temperature patterns, vertical wind shear distributions, and large-scale atmospheric circulation. This globally applicable yet finely tuned approach positions PepC-Global as a powerful tool for comprehensive probabilistic assessments, facilitating improved understanding and management of TC risks worldwide.

Evaluation of Land Surface Variables in the Global Forecast System (GFSv17) Retrospective and Real-Time Forecast Experiments

Zhichang Guo*, Michael Barlage, and Weizhong Zheng

GFSv17, an upgrade to the operational Global Forecast System (GFSv16), has been successfully implemented at the Environmental Modeling Center (EMC). The completion of retrospective experiments and the transition into the evaluation phase mark a significant milestone for the project. This presentation focuses on the evaluation of land surface variables, including 2-m air temperature and specific humidity, daily temperature extremes, soil temperature, and land surface energy and water budget components. Model outputs are assessed against observational and reanalysis datasets.

We also examine the strengths and limitations of GFSv17 forecast skill relative to GFSv16, with particular attention to the model's representation of water partitioning between evaporation and runoff across major river basins.

Subseasonal-to-Seasonal Prediction using MPAS/NoahMP

Zhe Zhang, Cenlin He*, Abby Jaye, Alex Lojko, Judith Berner, Michael Barlage, Meg Fowler, Yaga Richter, and Zong-Liang Yang

Subseasonal-to-seasonal (S2S) prediction, spanning 2 weeks to 2 months, is both a fundamental scientific frontier and a critical societal need. High-impact weathers, including extreme rainfall-flash flooding, heatwave, cold-air outbreak, droughts, are critical to be predicted with an adequate accuracy at S2S timescale to ensure sufficient preparedness. We present the MPAS-A/NoahMP S2S prediction system, featuring weekly-initialized 11-member ensemble reforecast, coupled with the Noah-MP land surface model, and daily ERA5 SST updates. Three hierarchical variable-resolution configurations are designed: a 60-km uniform global mesh, a 60–15 km North American regional refinement, and a 60–4 km convection-permitting refinement centered over the central United States. Our comprehensive evaluation of global 60-km simulations demonstrates good performance in capturing key near-surface meteorological fields and land-atmosphere interaction metrics. As a case study, we examine the May 2015 Texas–Oklahoma extreme rainfall and flooding event, a record flood-drought whiplash that terminated a multi-year drought. The 4-km configuration substantially outperforms the 60-km baseline in reproducing convective intensity and nocturnal MCS propagation at 1-week lead, and critically, maintains a persistent, spatially organized wet signal at 2–3 week leads where coarse-resolution parameterized convection rapidly loses coherence, suggesting convection-permitting resolution actively extends S2S predictability through mesoscale-to-large-scale upscale feedbacks. Some comparison with GEFS and ECMWF S2S results are also conducted. Looking ahead, the completion of a 20+ year 60-km global reforecasts (1999–2022) and a 10-year warm-season regional refinement (2013–2022) will provide the statistical foundation to robustly quantify the added value of kilometer-scale regional refinement across diverse synoptic regimes and extreme event populations.

Advancing AI Emulations of Physics-Based NWP Models through Investigating Training Data Generation Strategies

Horeb*, Man-Yau Chan

AI emulations of physics-based numerical weather prediction (AIWP) models have made significant progress recently. Improvements in training data generation approaches will further that progress. This study investigates three training data generation approaches: (1) an envelope-based approach that draws i.i.d. samples from the model's climatological distribution, (2) a trajectory-based approach that uses a long model trajectory, and (3) a semi-synthetic approach that generates data by combining resampling with short integrations of the physics-based model. The trajectory-based approach is the current de facto approach. The semi-synthetic approach is designed for situations where little training data is available.

As a starting point, we focused on emulating the chaotic 40-variable Lorenz-96 (L96) model. We also characterized the sensitivity of our findings to training data volumes (10, 20, 50, and 100 thousand samples). All semi-synthetic datasets are generated based on 1,000 samples drawn from the climatological distribution of L96. In other words, the semi-synthetic approach generates 10, 20, 50 and 100 thousand semi-synthetic samples from 1,000 original samples.

Results from 1,000 repeated tests indicate that the envelope-based approach produces more faithful AI emulations of the L96 model than the trajectory-based approach. Furthermore, the semi-synthetic approach produces AI emulations of comparable (albeit slightly worse) quality to the envelope-based approach. The performance of the AI emulator produced via the semi-synthetic approach also scales with the semi-synthetic data volume, indicating that the semi-synthetic data samples contain useful information beyond the original 1,000 samples. Further investigation is underway. Our findings potentially advance NOAA's AIWP efforts.

Porting and Optimizing the Global-Workflow on AWS

Walter Kolczynski, Wei Huang*, and Kristopher Booker

The Global-Workflow (GW) was originally developed as a unified package to execute NOAA-EMC numerical models on NOAA's on-premises high-performance computing (HPC) platforms. Following the successful implementation of GW on Amazon Web Services (AWS) via NOAA-ParallelWorks, EPIC has successfully decoupled the workflow from NOAA-specific infrastructure.

EPIC is currently collaborating with multiple stakeholders to enable GW on diverse architectures, including National Center for Atmospheric Research's Derecho supercomputer and, most recently, native AWS environments. This transition allows researchers and external users to deploy the Global-Workflow independently, without the requirement of a NOAA affiliation.

In this presentation, EPIC will share comprehensive benchmark results and technical insights gained during the porting process. These efforts establish a critical performance baseline, ensuring that GW operates efficiently and reliably across varied local and cloud-based infrastructures.
Keywords: Global-Workflow, AWS, Cloud Computing, Portability, Numerical Modeling

Track Evolution and Landfall Dynamics of Cyclone Remal (2024) over the Bay of Bengal using the WRF-ARW Model

Faozia Anzum Itu*

Accurate prediction of tropical cyclone track and landfall dynamics is essential for timely warnings and effective disaster risk reduction. This study simulates Cyclone Remal, which formed on May 24, 2024, following the monsoon onset over the Andaman and Nicobar Islands. The cyclone made landfall along the Bangladesh coast between Sagar Island and Khepupara on May 26, 2024. The Weather Research and Forecasting (WRF) v4.6.0 model was run from May 23, 0000 UTC to May 27, 0900 UTC, using NCEP FNL initial and lateral boundary conditions at a 9 km resolution. Model outputs were compared with the best track data from the India Meteorological Department (IMD). Sensitivity tests were performed and the best-performing configuration was identified by minimizing RMSE in cyclone track (km), central pressure (hPa), and maximum sustained wind (ms^{-1}). The simulated minimum central sea level pressure was slightly overestimated at 980 hPa compared to the observed 978 hPa (25 May, 12 UTC). Maximum sustained wind speed (MSW) was also overestimated, reaching 126 km h^{-1} against the observed 111 km h^{-1}. Track prediction improved closer to landfall, with a mean positional error of 101.49 ± 25.47 km. Vertical profiles showed strong convective activity and a well-defined warm-core structure extending up to 100 hPa. The model realistically captured 6-hour accumulated rainfall, spiral rainbands, and low-level vorticity fields associated with cyclone evolution. However, minor deviations in wind speed and landfall timing simulation indicate areas for further improvement.

Evaluating Global Model Fidelity via Troposphere-Stratosphere Coupling and Teleconnection Indices

Young-Joon (YJ) Kim*

The extended-range predictability of global numerical weather prediction systems depends heavily on a model's ability to faithfully simulate the complex, multi-scale dynamical "chain of events" governing the Earth's atmosphere. This study provides a comprehensive evaluation of the Global Forecast System version 17 (GFSv17) prototype's performance by diagnosing the full atmospheric continuum : from tropical convective triggers—evaluated via Madden-Julian Oscillation (MJO) phase space diagrams and Wheeler-Kiladis (WK) wavenumber-frequency diagrams—to the Quasi-Biennial Oscillation (QBO), stratospheric wave driving, and the resulting tropospheric teleconnection patterns including the Arctic Oscillation (AO), Northern Annular Mode (NAM), North Atlantic Oscillation (NAO), and Pacific-North American (PNA) pattern. A critical component of this evaluation is the model's ability to simulate major Sudden Stratospheric Warming (SSW) events, which serve as a primary metric for assessing downward dynamical coupling.

By utilizing various experimental retrospective streams and specialized datasets, a systematic approach to identifying model biases and deficiencies in dynamical fidelity is demonstrated. A primary focus of this evaluation is the investigation of troposphere-stratosphere coupling and the subsequent downward propagation of anomalies, which serve as critical indicators of extended-range skill. Preliminary results indicate that while the prototype captures initial state patterns and initiates upward wave flux, it suffers from aggressive amplitude damping and an accelerated return to climatology at extended lead times. These findings suggest that systematic issues, such as excessive numerical diffusion and miscalibrated wave-drag parameterizations, may prematurely wash out the potential vorticity gradients required to sustain major teleconnection phases.

This predictability study underscores the utility of teleconnection indices as diagnostic targets for pinpointing exactly where kinetic energy "bleeds" from the system. Based on these systematic findings, targeted model revisions—specifically the tuning of diffusion schemes, convective parameterizations, and drag physics—may be useful to sustain organized tropical forcing and improve global forecast skills, especially in the extended range. The full scope of this analysis, including results from all evaluated streams and specialized data sets, will be presented at the conference.

DA Nudging: Bridging AI Weather Prediction Systems and GFS

Mu-Chieh Ko*, Russ Treadon, Jeffrey Whitaker, and Daniel Holdaway

This project aims to improve numerical weather prediction (NWP) by incorporating large-scale information from AI weather prediction (AIWP) systems into an operational forecast model. Systems such as GraphCast-GFS (AIGFS) and other AIWP models have shown promising results for representing large-scale atmospheric patterns, and Google DeepMind demonstrated encouraging track and intensity forecasts for Hurricane Melissa in 2025. To leverage these strengths, we developed a nudging framework built around the Global Workflow, targeting the Global Forecast System (GFS), which allows AIGFS outputs to be ingested as observational data and used to guide GFS forecasts.

Initial nudging experiments showed that combining hurricane intensity data from Google DeepMind with AIGFS outputs improved GFS track and intensity forecasts for Hurricane Melissa. A separate experiment found that nudging with AIGFS data can reduce RMSE in wind and temperature predictions, though it also introduced bias degradation. Work is ongoing to evaluate these trade-offs and improve how the model weights AI-generated observations.

An alternative approach uses the “replay” capability in the UFS to constrain the large scales in the GFS forecast (total wavenumbers < 25) to the pre-generated AIWP model forecast, but allows the GFS model to determine the smaller scales and fields that the AIWP model does not predict. Results show that “replaying” the GFS to the ECMWF AIFS forecast can yield precipitation forecasts that are more skillful than either the GFS or the AIFS.

AIGEFS Improvements via CRPS-based Optimization with the ANEMOI Framework

Jianjun Liu*, Jun Wang, Linlin Cui, Xiao Luo, Bing Fu, and Jacob Carley

Ensemble weather forecasting provides a more comprehensive framework for predictability than deterministic approaches by quantifying initial condition uncertainties and accounting for the stochastic nature of error growth. Recently, data-driven ensemble weather prediction models have demonstrated superior accuracy over traditional numerical weather prediction systems while requiring orders of magnitude fewer computational resources.

NOAA operationalized the AI-driven Global Ensemble Forecast System (AIGEFS) version 1 by fine-tuning GraphCast with GDAS as well as additional supporting datasets. The model was initially trained using the mean-squared-error (MSE) loss function. While the MSE loss promotes accurate ensemble-mean forecasts, it implicitly favors smoothing and suppresses small-scale variability as lead time increases. This training strategy encourages ensemble members to collapse toward the conditional mean, resulting in blurred forecast fields and under-dispersive ensembles that inadequately represent atmospheric variability.

In this study, we overcome these limitations by introducing an updated version of the AIGEFS trained on the Continuous Ranked Probability Score (CRPS) objective using the Anemoi framework. Unlike MSE loss, which rewards convergence to a mean state, the CRPS objective explicitly balances an accuracy term with a spread term, rewarding both the precision and the calibration of the ensemble.

The updated AIGEFS employs a GNN-based encoder–decoder architecture with a sliding-window transformer processor and is trained on GDAS data with the Anemoi framework. Results demonstrate that CRPS-based AIGEFS substantially mitigates under-dispersion and blurriness in AIGEFS v1, while maintaining comparable deterministic skill across most variables, providing vastly superior probabilistic information for weather event estimation.

Utilizing Water Isotopic Composition in the Unified Forecast System to Constrain Precipitation and Hydrologic Processes: Preliminary Results

Sarah Lu*, Scott Miller, Mathias Vuille, Shih-Wei Wei, Zhiqiang Lyu, SUNY-Albany, Yi Ming, Boston College, Lulin Xue, Louisa Nance, NCAR and DTC, Ligia Bernardet, Grant Firl, NOAA and DTC, Fanglin Yang, NOAA/NCEP/EMC

The overarching goal of this WPO S2S project is to help guide the development and analysis of the Unified Forecast System (UFS) by adopting the latest stable water isotopes (SWI) modeling and observational capabilities, with special focus on precipitation and hydrological processes. Isotopic signals contain valuable information about atmospheric processes and the water cycle across time scales, from weather to climate. As natural tracers of moisture exchange, the water isotopic composition provides a valuable means to examine the integrated and connected hydrological cycle. Recent development of innovative measurement techniques has produced SWI observations from ground, aircrafts, ships, and space. Enhancing observational capabilities, in turn, enable the scientific community to use these novel observations for process-level study, model evaluation, and data assimilation.

This UAlbany-BC-DTC collaborative research will implement an innovative hydrometeorological diagnostic, i.e., SWI in liquid and vapor phases, in the UFS atmosphere and land model. We propose to investigate hydrological processes and extreme precipitation characteristics through isotope-enabled simulations and model-data comparison. The major tasks in year 1 include (1) implementing SWI into UFS physics suite, and (2) developing interface for observation data access (IODA) converter and unified forward operator (UFO) for SWI satellite observations. This presentation will provide an overview of SWI-enabled UFS project and the status of project progress.

AirCast-SR: A Foundation Model for Kilometer-Scale Atmospheric Super-Resolution via Latent Consistency Diffusion

Somnath Luitel*, Manmeet Singh, Joshua Durkee, Abdullah Al Fahad, Naveen Sudharsan, Prabhjot Singh, Cenlin He, Harsh Kamath, Zong-Liang Yang, Krishnagopal Halder, Sandeep Juneja, Parthasarathi Mukhopadhyay, Saptarishi Dhanuka, and Amit Kumar Srivastava

Kilometer-scale weather prediction is computationally prohibitive for traditional numerical weather prediction (NWP) models, restricting fine-grained forecasts to a small set of operational systems run on supercomputing infrastructure. Here we introduce AirCast-SR, a diffusion-based foundation model that downscales global AI weather forecasts from 0.25°(~28 km) to 1 km resolution at hourly cadence, producing 67-hour forecasts of seven coupled near-surface variables simultaneously. AirCast-SR couples a three-dimensional U-Net to a Latent Consistency Model (LCM) diffusion framework, trained on patch-based samples over the contiguous United States (CONUS) with GraphCast forecasts as input and NOAA's Analysis of Record for Calibration (AORC) as the target. Across two extreme events (Winter Storm Elliott in December 2022 and the June 2022 summer convective episode) and one spring transition case (March 2023), AirCast-SR achieves near-zero systematic bias across all variables and all lead times to 48 h, while its radial power spectral density tracks the AORC reference at wavelengths from 10 km to 100 km where coarser models lose spectral power. The model approaches but does not yet exceed the operational High Resolution Rapid Refresh (HRRR) on pointwise skill for most variables; instead, it delivers broader coupled-variable coverage, structural realism via diffusion (the perception–distortion regime of generative super-resolution), and inference in minutes on a single commodity GPU. Without retraining, AirCast-SR generalises zero-shot to India and Germany when verified against surface station observations from StationBench. Released with open weights, AirCast-SR establishes a foundation for kilometer-scale AI weather services and a clear scaling path—more compute, more training years, ensemble inference—toward eventually exceeding operational NWP systems.

Sensitivity Analysis of Fine-Tuning GraphCast with Adjusted Mean Square Error Using GDAS Data

Xiao Luo*, Jun Wang, Linlin Cui, Jianjun Liu, Wei Li, Bing Fu, Jiayi Peng, Zhan Zhang, Bin Liu, Neil Barton, Fanglin Yang, and Jacob Carley

Loss functions strongly shape the behavior of machine-learning weather models. The commonly used mean squared error (MSE) penalizes spatial displacement twice, suppressing variance and producing overly smooth forecasts of high-impact events. To address this, we utilize the Adjusted Mean Squared Error (AMSE), which mitigates this "double penalty" by separating amplitude and phase contributions in spherical harmonic space. We conduct a sensitivity analysis by fine-tuning the pretrained GraphCast model on less than two years of NOAA Global Data Assimilation System (GDAS) data, exploring the impacts of learning rate schedules. Using Hurricane Ian (2022) as a case study, we demonstrate that AMSE-based fine-tuning improves spectral amplitudes and hurricane structure relative to the MSE-loss baseline. We further explore a "Fully-Decoupled" AMSE (FD-AMSE) formulation and a weighted-AMSE objective function. The FD-AMSE further decouples power spectral density (PSD) and coherence terms; initial results for a case study of Hurricane Ian suggest this variant can mitigate the amplitude degradation observed with standard AMSE, yielding hurricane intensity forecasts that more closely align with GDAS sea-level pressure and 10 m wind speeds. Additionally, the weighted-AMSE objective—utilizing a spectral weighting of 1/PSDy—shows notable improvements in spectral alignment with GDAS reanalysis at high wavenumbers (k>100) across multiple variables, particularly for vertically weighted winds and precipitation. These results demonstrate that AMSE—potentially in its FD-AMSE and weighted forms—provides a physically

interpretable approach for leveraging limited GDAS data to sharpen machine learning weather forecasts and improve the prediction of extreme events.

Watershed-Scale Evaluation of ARAFSv2 for High-Impact Atmospheric River Rainfall over the Western United States

Murali N Malasala*, Keqin Wu, Xingren Wu, Binbin Zhou, Maria Aristizabal Vargas, and Vijay Tallapragada

Atmospheric river (AR) events are a major driver of high-impact winter precipitation across the western United States, often producing intense rainfall and significant hydrological responses at watershed scales. However, forecasting these events remains challenging due to complex terrain interactions and limitations in representing precipitation processes. This study evaluates the performance of the near-real-time Atmospheric River Analysis and Forecast System version 2 (AR-AFSv2) for selected high-impact rainfall events during the Winter 2025–2026 season, with a comparative assessment against GFSv16. An event-based framework is adopted, focusing on major western U.S. watersheds where observed 24-hour precipitation exceeded critical thresholds. Model forecasts are verified against Stage IV quantitative precipitation estimates, used as the observational reference. A comprehensive evaluation strategy is employed, combining spatial diagnostics with statistical verification. Spatial analyses assess the representation of precipitation structure and terrain-influenced patterns, with emphasis on watershed-scale aggregation. Continuous metrics, including mean bias and Root Mean Square Error (RMSE), are used to quantify magnitude errors. In addition, categorical metrics such as Hit Rate, Equitable Threat Score (ETS), Success Rate, and Frequency Bias are applied across multiple precipitation thresholds to evaluate forecast skill in capturing extreme rainfall. Performance diagrams further synthesize key aspects of forecast quality, including hit rate, false alarm ratio, and critical success index, across multiple lead times. This multi-faceted assessment highlights the relative strengths of AR-AFSv2 and GFSv16 in representing high-impact precipitation and provides insights into the added value of high-resolution regional modeling for improving operational forecasts of extreme hydrometeorological events.

From UFS Applications to AIWP: Leveraging the METplus Verification Workflow for Consistent and Scalable Model Evaluation

Kathryn Newman*(NSF NCAR, DTC), Michael Kavulich (NSF NCAR, DTC), Gerard Ketefian (CIRES, NOAA GSL, DTC), Michelle Harrold (NSF NCAR, DTC), Jeff Beck (NOAA GSL, DTC), Vanderlei Vargas (CIRA, NOAA GSL, DTC), Tina Kalb (NSF NCAR, DTC), Will Mayfield (NSF NCAR, DTC), George McCabe (NSF NCAR, DTC), Sam Trahan (CIRES, NOAA GSL, DTC), John Halley Gotway (NSF NCAR, DTC), and Julia Simonson (CIRES, NOAA GSL, DTC)

The Developmental Testbed Center (DTC) has developed the enhanced Model Evaluation Tools (METplus) Verification Workflow, a unified, configurable verification workflow system designed to accelerate model evaluation across a wide range of applications. By consolidating observation data

handling, standard metrics, and application-specific configurations into a shared, centralized system, the framework reduces effort and spin-up costs while enabling consistent, reproducible evaluation and easy sharing of configurations for running verification workflows across projects.

The METplus Verification Workflow is currently deployed across several DTC testing and evaluation projects with an emphasis on UFS Applications. It supports evaluation of MPAS-based Rapid Refresh Forecast System (RRFS) and RRFS Ensemble Forecast System (REFS) prototype comparisons against operational baselines for high-impact events. For the Hurricane Forecast and Analysis System (HAFS), it enables track and intensity verification while also extending ensemble-based verification capabilities to Tropical Cyclones (TCs) using HAFS Ensemble in Realtime on the Cloud (HERC) forecasts. A growing focus of the framework is supporting the intercomparison of Artificial Intelligence (AI), hybrid, and physics-based modeling systems, including NOAA Project Experimental AI Global and Limited-area Ensemble forecast system (EAGLE) models. The breadth of metrics available within MET, including feature-based spatial verification and the newly added power spectra, enables AI-based model evaluation that goes beyond basic scores, providing deeper insight into model behavior relative to traditional Numerical Weather Prediction (NWP) systems. Looking ahead, we will discuss ongoing efforts in AIWP-specific metric development, common data standards, portable workflow capabilities, and visualization enhancements that will further position the framework as a comprehensive model evaluation tool.

Winter Validation of the Operational Great Lakes Wave System Using In-Situ and SWOT Observations

Kaitlin Pereira* and Eric Anderson

The Great Lakes constitute a massive, dynamic, fetch-limited system characterized by rapid atmospheric forcing transitions and significant seasonal ice coverage. These enclosed basins experience their most operationally consequential wave conditions in winter, driving coastal, navigation hazards, and damage to coastal infrastructure. At the same time however, winter is also the most under-observed and under-validated season in the observational record. The operational Great Lakes Wave Unstructured system (GLWU), built on WAVEWATCH III, feeds NWS marine guidance for the region, yet winter validation of WW3 has been limited by persistent observational gaps. Addressing this gap necessitates the integration of two complementary observational anchors: temporally dense point measurements from winter-resilient in-situ buoys, and spatially comprehensive two-dimensional(2D) surface snapshots from the Surface Water and Ocean Topography (SWOT) satellite. Together, they provide a systematic way of errors introduced by surface forcing, static ice representations, and source-term physics. This work has direct implications for UFS-Wave development on the Great Lakes by characterizing winter dynamics and using these observational records to evaluate parameterizations originally calibrated on open-ocean conditions.

Harnessing Generative AI and Emulation for High-Resolution Fire Weather Forecasting: A Multi-Scale Modeling Framework

Nick Silverman*, MITRE, Nick Krall, MITRE, Daniel Abdi, NOAA-GSL, Natalie Wagenbrenner, USFS, Jason Forthofer, USFS, Robert Roth, MITRE, Isidora Jankov, NOAA-GSL, and Conor Lewellyn, MITRE

Accurate fire weather forecasting requires capturing complex atmospheric interactions at extremely high spatial resolutions, particularly for variables like surface winds and relative humidity in complex terrain. While the High-Resolution Rapid Refresh (HRRR) model provides critical operational guidance, its 3km resolution often fails to resolve sub-grid scale features essential for fire behavior modeling. This presentation outlines a novel research framework designed to bridge this gap using a suite of Artificial Intelligence (AI) architectures.

Our methodology explores three primary research directions: (1) Generative Downscaling via CorrDiff: We use NVIDIA's CorrDiff model architecture to downscale HRRR, using a custom 1km WRF-based dataset—Weather1k—to produce a 1km ensemble forecast that captures the probabilistic nature of fire weather variables. (2) HRRRCast Emulation: Using similar architecture to the 3km HRRRCast model, we aim to produce rapid-refresh ensembles of high-resolution atmospheric states that maintain the physical consistency of the underlying Weather1k physics. (3) Hyper-Local Wind Emulation: Recognizing that fire-line winds are driven by micro-topography, we are building a WindNinja emulator. This model is designed to translate 1km forecasts down to 100m resolution, providing the granular wind fields necessary for tactical fire behavior modeling without the overhead of the full WindNinja computational fluid dynamics (CFD) solver.

We will present our model architectures and preliminary results, focusing on a comparative analysis of the 1km ensemble forecasts. By validating these AI-enhanced outputs against traditional WindNinja simulations and surface observations, we aim to identify the most effective pathway for delivering actionable, high-resolution fire weather intelligence in real-time operational environments.

Examining the Shift in Large-Scale Biases and RMSE: From GFSv16 to the Fully Coupled GFSv17

Lydia Stefanova*, Alicia Bentley, Neil Barton, Dmitry Dukhovskoy, Xiao Luo, Jessica Meixner, Ruiyu Sun, and Catherine Thomas

The transition to GFSv17 marks a significant evolution in NOAA's global modeling, moving from the partially coupled GFSv16 toward a more integrated Earth system framework. As GFSv17 nears operational implementation, characterizing how its extensive updates to physics suites and multi-component coupling have altered systematic error profiles is essential. This study compares the large-scale biases and RMSE of both versions, evaluating global forecast performance across multiple seasons.

We present a seasonal analysis of errors in key variables to identify where long-standing biases have been mitigated and how total error magnitudes have evolved. Given the simultaneous introduction of multiple system updates, we provide a diagnostic attribution of these shifts. By using spatial and temporal clues, we associate observed error patterns with specific changes in model physics or the newly integrated atmosphere-ocean-ice feedback loops. These findings offer a practical assessment of the GFSv17 error climate and provide a baseline for future tuning and development within the UFS framework.

Leveraging Subgrid Methods for a High-Fidelity Global-Scale Hydrodynamic Flood Inundation Modeling

Aman Tejaswi*, Shintaro Bunya, Damrongsak Wirasaet, Albert Cerrone, Andrew Kennedy, Rick Luettich, Chris Massey, and Joannes Westerink

Flooding is a highly localized phenomenon characterized by complex interactions between pluvial (rainfall), fluvial (river), and coastal (tide and surge) processes. Accurately modeling these compound flood hazards requires high-resolution topographic data, typically derived from Digital Elevation Models (DEMs). However, traditional numerical modeling faces a fundamental trade-off: while fine-scale grids are necessary to capture critical terrain features, applying them across global scales is computationally prohibitive for operational forecasting. Even state-of-the-art unstructured mesh models must limit high-resolution nodes to specific regions to maintain feasibility. This study addresses a critical gap in current frameworks: the inability to simulate simultaneous coastal and inland flood drivers within a single, scalable system. Existing methods often rely on model coupling, which creates significant performance bottlenecks over large domains. To overcome these limitations, this research investigates the application of subgrid bathymetry theory within a unified 2D hydrologic-hydrodynamic framework. We hypothesize that integrating subgrid physics allows for the representation of fine-scale topographic effects without the computational burden of an explicitly high-resolution mesh. This approach aims to enable the accurate and efficient simulation of compound flooding across multi-regional scales, offering a potential solution for global operational forecasting.

Representation of East Pacific Easterly Waves in Subseasonal-to-Seasonal Forecasts of the Unified Forecast System

Yu-Cian Tsai* and Eric D. Maloney

East Pacific easterly waves (EPEWs) are critical precursors to tropical cyclogenesis and regional rainfall yet are poorly captured in subseasonal-to-seasonal (S2S) forecast models. This study evaluates how EPEWs are simulated in NOAA's Unified Forecast System (UFS) Prototypes 5–8 during May-October relative to ERA5 reanalysis, using wave-tracking algorithm and process-based diagnostics to identify biases that increase with forecast lead time. UFS Prototypes exhibit strong differences in EPEW activity and moisture-vorticity coupling strength, with all prototypes except for Prototype 7 exhibiting excessive EW activity relative to reanalysis that becomes more vigorous with forecast lead from weeks 1-2 to weeks 3-5. Variations in the background thermodynamic environment are a dominant driver of differences in EW activity across prototypes, with the strength of meridional moisture gradients showing the strongest correlation with EPEW track density. In particular, Prototype 7 has weaker mean moisture gradients and EW activity compared to reanalysis. Composite diagnostics reveal EPEW convective structures evolve from relatively bottom-heavy to more top-heavy with increasing lead time, accompanied by accelerated convective adjustment timescales and exaggerated moisture-vortex instability, with overall biases being most exaggerated for Prototypes 5, 6, and 8. The vorticity maintenance budget shows excessive horizontal advection and vortex stretching contributions across all Prototypes, while the MSE maintenance budget reveals an excessively strong column process associated with excessive longwave radiative heating and surface latent heat flux contributions. However, these maintenance budget terms are generally more muted for Prototype 7. These findings highlight the critical need for improved parameterizations in S2S prediction systems.

Reducing state-independent bias in subseasonal weather forecasts using a correction derived from a running mean nudging

Sarah Weidman* (Harvard) and Zhiming Kuang (Harvard)

At seasonal-to-subseasonal (S2S) timescales, forecasts are less strongly tied to their initial conditions, resulting in “climate drift,” where models drift towards their preferred mean state. This mean state is often biased from the observed atmosphere due to pervasive challenges of numerical forecasting, including coarse model resolution and imperfect representations of physics. For weather phenomena that depend on the background state, climate drift is a difficult forecasting challenge.

To address this challenge, we derive a set of correction tendencies using an algorithm that slowly nudges the mean state of a climate model towards a target mean state, taken here to be a reanalysis product. Nudging parameters are chosen to preserve weather-timescale phenomena while forcing longer timescales towards the target mean climatology. The nudging tendencies converge towards a set of “optimal” correction tendencies that can vary with season and time of day, but are approximately constant between years. These correctors are then applied as a state-independent forcing in a separate, freely-running simulation. In preliminary results, we find that global mean bias is reduced by 40-80% compared to the control model.

The optimal correction tendencies derived from the nudging procedure have a spatial and temporal pattern that can reveal how the mean state depends on local, short-term model errors and potentially highlight missing processes in the model. We expect that an improved simulation of the background state can better represent S2S processes like Madden-Julian Oscillation teleconnections, and thus help us better understand the dynamics behind these processes.

Graph-Oriented Workflow Orchestration: A Design Framework for UFS Research-to-Operations Pipelines

Brian Weir*, Jong Kim, Fernando Andrade Maldonado, Michael Lueken, Chan-hoo Jeon, Gillian Petro, and Benjamin Koziol

The Unified Forecast System (UFS) is a community-driven, end-to-end Earth system modeling framework that supports a wide range of configurations, applications, and research pathways. Managing this variability—while maintaining a reliable and efficient Research-to-Operations (R2O) pipeline—presents significant challenges in workflow governance, reproducibility, and user experience. To streamline and enhance this process, we propose a framework for supporting the broad range of workflows and toolkits required throughout the R2O pipeline.

In this presentation, we explore the advantages of introducing a lightweight orchestration layer designed to bridge gaps between research and operations. We highlight how graph-oriented design principles enable a scalable, modular system that efficiently leverages every available tool across all stages of the R2O pipeline. This approach ensures transparent workflow execution, improves maintainability, and supports consistent behavior across heterogeneous environments.

Additionally, we discuss complementary design strategies that provide the flexibility needed by diverse users while unifying their overall workflow experience. By aligning workflow governance with modern orchestration practices, this framework supports a more resilient, extensible, and user-friendly ecosystem for the UFS community, ultimately accelerating the transition of scientific innovation into operational capability.

Development of the Atmospheric River Analysis and Forecast System to Enhance Atmospheric River Forecasting

Keqin Wu*, Xingren Wu, Murali N Malasala. Maria Aristizabal, Binbin Zhou, and Vijay Tallapragada

Atmospheric rivers (ARs) are primary drivers of extreme precipitation and water supply along the U.S. West Coast. Operational numerical weather prediction (NWP) models often struggle to accurately represent AR landfall timing, location, and associated heavy precipitation due to complex terrain and intrinsic system complexity. To address these challenges and advance NOAA's Research-to-Operations (R2O) goals, the Atmospheric River Analysis and Forecast System (AR-AFS)—a standalone, high-resolution (3-km) regional model based on the FV3 dynamical core—was developed.

This study details the transition to AR-AFSv2, which incorporates system-level upgrades aligned with the Hurricane Analysis and Forecast System version 2 (HAFSv2). Key enhancements include increasing vertical resolution from 64 to 82 levels, expanding the domain to Alaska, and utilizing an optimized physics suite. This suite features Thompson microphysics coupled with the turbulent kinetic energy-based moist hybrid eddy-diffusivity mass-flux (TEDMF) planetary boundary layer (PBL) scheme, a configuration previously identified as superior for moisture transport and precipitation processes.

AR-AFSv2 was operated in near-real time from November 2025 to March 2026. A brief skill assessment of the 2025–2026 season and selected retrospective cases from past seasons will also be presented, including forecast performance for 24-hour precipitation and Integrated Vapor Transport (IVT) evaluation against Stage IV observations and GDAS reanalysis. We will show how these upgrades impact 1–5 day forecasts. The work offers valuable insights into the system's operational readiness within NOAA's Unified Forecast System framework for Atmospheric River predictions.

Forecasting the Loop Current System with Physics-Guided Machine Learning

Tianning Wu*, Ruoying (Roy) He, Ashesh Chattopadhyay, Leonard Lupin-Jimenez, Michael Gray, Bowen Chen, and Felix Deng

The Loop Current (LC) and associated eddies are dominant oceanographic features in the Gulf of Mexico that significantly impact regional weather, ecosystem, and marine environment. Skillful subseasonal-to-seasonal prediction of the LC system remains a major challenge for both numerical and machine learning models. We present OceanNet3D, a physics-guided deep learning ocean emulator developed for rapid and accurate subseasonal-to-seasonal prediction of the LC system. The model is based on the Fourier Neural Operator and trained with 32 years of high-resolution ocean reanalysis data. Spectral regularization and physical loss are used to promote physical consistency across scales. Ensemble forecasts are generated through initial condition perturbation and provide sea surface height, upper ocean currents, as well as probabilistic maps of the LC position and high-speed currents. Process-oriented evaluations demonstrated that OceanNet3D can skillfully resolve the Loop Current evolution and its eddy shedding events over several weeks. Compared to numerical models, OceanNet3D achieves comparable forecast skill with a fraction (1/3,000) of the computational cost, making it a promising tool for operational and research applications.

Assessing subseasonal predictions of stratosphere-troposphere coupling of a data-driven weather forecasting model

Zheng Wu*, Department of Atmospheric and Environmental Science, SUNY University at Albany, Albany, NY, USA

While data-driven weather forecasting models outperform state-of-the-art numerical models in surface and lower-tropospheric variables, their skill in predicting and representing the stratosphere and its dynamical coupling with the troposphere remains unclear. More importantly, whether such models can be used to study and advance subseasonal predictability from the stratosphere still needs to be investigated. In this study, we evaluate predictions by GraphCast, one of the best existing data-driven models, developed by Google DeepMind, of the stratospheric variability and the stratosphere-troposphere coupling, focusing on sudden stratospheric warmings (SSWs) - one type of stratospheric extreme circulation event considered an important source of subseasonal to seasonal predictability for surface weather. We generate predictions for the recent SSW events using GraphCast initialized with ERA5 reanalysis data. The results show that GraphCast predicts the stratospheric polar vortex strength up to two weeks, although errors increase rapidly with height. However, GraphCast fails to predict any SSW events even at lead times of one week or less. The primary reason is its inability to accurately simulate wave-mean flow interactions in the stratosphere, as it fails to satisfy the Transformed Eulerian Mean momentum budget. We also found that GraphCast cannot accurately predict the downward propagation of stratospheric anomalies, resulting in larger surface errors after SSW events compared to normal winters. Our study reinforces the link between the stratosphere and surface predictive skill and identifies current limitations in data-driven weather-forecasting models, which could guide the improvements of next-generation AI models for S2S applications.

Seasonal Forecast System (SFS) Applications to NOAA Early Warning System

Li Xu*, Tasnuva Rouf, Aissatou Faye, Zewdu Segele and Wassila Thiaw

NOAA Early Warning System (NEWS) is a NWS initiative led by the Climate Prediction Center and funded by the Department of State to advance the forecasts of extreme weather events on sub-seasonal to seasonal (S2S) time scales. It supports national and regional cross-institutional partnerships to help vulnerable nations in Africa and the Small Islands Developing States in the Pacific develop early warning systems with actionable weather and climate forecasts and action plans to reduce weather-related risks. The NEWS–Africa project prioritizes collaboration between meteorologists and professionals from across economic and social sectors to develop real time forecasts that capture and predict a broad range of impacts.

The new release of the Seasonal Forecast System (SFS) beta 1.0 is anticipated to significantly enhance the CPC's S2S prediction capability, thereby replacing the currently operational model, CFSv2. This paper document details the current status of development efforts aimed at integrating the SFS beta 1.0 into the NOAA Early Warning System (NEWS) project. The primary focus involves transforming SFS outputs to effectively address the diverse requirements of priority regions, which include Africa, Central Asia, Central America, and the Pacific islands. Key development includes Model Evaluation and Validation, to assess subseasonal-to-seasonal (S2S) prediction skill, benchmarked against the operational CFSv2, particularly concerning critical rainfall and temperature anomalies. Early warning Products are currently in the design and prototyping phase, concentrating on the development of post-processing routines essential for converting raw meteorological data into actionable, impact-based forecast products, applicable to sectors such as agriculture, public health, and water management. Furthermore, the project is advancing efforts in Regional Downscaling & Bias Correction to achieve improved resolution and accuracy.

The JCSDA extensively uses the MPAS (NCAR) and UFS core FV-3 (NOAA) models. Initializations can be readily executed by using the NOAA GFS and GEFS outputs available publicly through services such as Amazon Web Services (AWS) simple storage service (S3) buckets. Using the JEDI Skylab system we have developed and demonstrated the various UFO observation filters, variational bias correction and observation error models over wide-ranging observations used by the partner agencies. We have used JEDI Skylab to assist our partners and in particular NOAA NESDIS and the USAF to examine new observations provided from the commercial sector. Using JEDI Skylab we have created end-to-end evaluations ingesting all data and initial conditions, pre-processing the data to the common JEDI format defined collaboratively in the Interface for Observation Data Access (IODA) component. Recently we have updated our evaluation capability to include the WeatherBenchX tools and evaluation framework developed by Google. We will present results from work done using a JEDI Skylab environment, verified against the ERA-5 re-analysis and to other observation types. We will provide examples of how JEDI configuration offers flexibility and ease in both integration and customization, creating opportunities for new users to engage more readily with the JEDI system. JEDI-enabled systems are in use currently at NOAA, NASA, the US Air Force and Navy and the UK Met Office; additionally, JEDI is being used by the private sector such as The Weather Company and Tomorrow.io, providing new opportunities to interact and collaborate.

Spatiotemporal Estimation and Forecasting of Particulate Matter in the Pacific Northwest using Machine Learning with Emphasis on Wildfire Smoke

Abu Ubayeda Ibnal Zarrah*, MD. Jahidul Islam and Jun Meng

Despite decades of effective air pollution control in the United States, fine particulate matter (PM2.5) continues to pose significant health risks, even at relatively low concentrations. This challenge is particularly acute in the Pacific Northwest (PNW), where recurrent wildfires generate highly variable spatiotemporal pollution patterns that are not adequately captured by sparse ground monitoring networks. While chemical transport models (CTMs) provide valuable spatial insights, their high computational cost and reduced reliability during extreme wildfire events highlight the need for complementary approaches.

This study develops a machine learning–based framework for high-resolution PM2.5 estimation and forecasting across the PNW by integrating multi-source datasets, including satellite-derived aerosol optical depth (MODIS, VIIRS), WRF meteorology (3kmx3km), emerging observations (TEMPO), regional air quality model forecasted outputs, Fire inventory data (GFED) and ground-based measurements. Advanced feature engineering incorporates downscaled meteorological fields, physics-informed chemical components, and geographic clustering to capture complex atmospheric processes.

Multiple algorithms, including convolutional neural networks (CNN), Random Forest, and Gradient Boosting, are systematically evaluated within an optimization pipeline using cross-validation to maximize predictive skill. Applied to a decade of operational data, the framework demonstrates improved performance over baseline approaches (cross validated daily R2: 0.83), enabling fine-resolution gridded PM2.5 estimation and short-term forecasting. Importantly, ensemble modeling and real-time bias correction enhance robustness during wildfire episodes.
Overall, this approach provides a scalable and computationally efficient pathway for improving air quality assessment in wildfire-prone and underserved regions, supporting more equitable and informed air quality management across the PNW.

A Prototype AI-Based Quality Control Package for Cloud Physics Observations of Hurricane Field Program at NOAA

Xuejin Zhang*, Xin Yang, Andrew Mercer, Jun Zhang, and Frank Marks Jr.

The Hurricane Field Program has observed cloud physics for more than forty years. The cloud physics instruments and observing strategies are evolving significantly. During 1980-1992, NOAA's P3 flew into some significant hurricanes and obtained a significant amount of cloud physics data in the strong convective regimes. After 1992, NOAA's P3 gradually shifted to fly into the stratiform cloud regimes due to the hazards to the plane and safety concerns in strong convective regimes. Scientifically, as deep convection of tropical cyclones matures, transitions to primarily stratiform clouds were observed and characterized by different microphysical processes. For example, ice processes, such as vapor deposition and aggregation, take on a more dominant role as vertical motion becomes weaker, and evaporation below the melting level also becomes important. These processes are critical for TC intensity prediction, yet are poorly parameterized in current high-resolution microphysics parameterization schemes. The objective of this study is to understand the sophisticated evolving processes of cloud physics and improve the representation of these processes within numerical models based on these historical observations in convective regimes and more recent observations in the stratiform regimes. The AI-based quality control package for cloud physics observations is under development. This is the first step to create a microphysics emulator for microphysics in the numerical models.

Data Assimilation

NSF NCAR's Next-Generation Online-Coupled Air Quality and Weather Analysis and Forecasting System using MPAS-GOCART2G-JEDI

Soyoung Ha*, Shih-Wei Wei, Jun Park, Maryam Abdi-Oskouei, Cheng Dang, Cheng-Hsuan Lu, Ben Johnson, Rajesh Kumar, Mary Barth, Gabriele Pfister, Forrest Lacey, and Michael Duda

Accurate aerosol prediction remains challenging due to uncertainties in atmospheric composition arising from imperfect initial conditions, errors in emission inventories, and limited understanding of aerosol processes and properties interacting with meteorological variables. Because of their short lifetime and strong spatial and temporal variability, global observations of aerosols and clouds rely heavily on satellite remote sensing.

The U.S. National Science Foundation (NSF) National Center for Atmospheric Research (NCAR) has recently been developing the atmospheric Model for Prediction Across Scales (MPAS-A) coupled with the next-generation Goddard Chemistry Aerosol Radiation and Transport model (GOCART-2G) and interfaced with the Joint Effort for Data Assimilation Integration (JEDI) system.
Leveraging various data assimilation algorithms within JEDI, this integrated framework facilitates coupled data assimilation of multi-sensor, hyperspectral satellite aerosol retrievals and all-sky

radiances across a wide spectral range within a unified atmosphere-aerosol analysis and forecasting system.

This talk introduces the new MPAS-GOCART2G-JEDI system, which is now publicly available, with an emphasis on the assimilation of NASA Plankton, Aerosol, Cloud, ocean Ecosystem (PACE) Aerosol Optical Depth (AOD) retrievals and their systematic evaluation against legacy AOD products such as MODIS, VIIRS, and AERONET in the context of atmosphere-aerosol coupled data assimilation for heavy pollution events.

Toward the Implementation of JEDI in the RRFS Data Assimilation System

Shun Liu*, NOAA NWS NCEP EMC, Ming Hu, NOAA GSL, Daniel Holdaway, NOAA NWS NCEP EMC, Samuel Degelia, Lynker@ NOAA NWS NCEP EMC, Donald Lippi, Lynker@ NOAA NWS NCEP EMC, Ting Lei, Lynker@ NOAA NWS NCEP EMC, Xiaoyan Zhang, SAIC@ NOAA NWS NCEP EMC, Hui Liu, Lynker@ NOAA NWS NCEP EMC, Cory R Martin, NOAA NWS NCEP EMC. Masanori Oigawa, JMA/Numerical Prediction Development Center, Jacob R. Carley, NOAA NWS NCEP EMC, Daryl Kleist, NOAA NWS NCEP EMC, Matthew Pyle, NOAA NWS NCEP EMC, Guoqing Ge, CU/CIRES@ NOAA GSL, David Dowell, NOAA GSL, Chunhua Zhou, CU/CIRES@ NOAA GSL, Terra Ladwig, NOAA GSL, Ruifang Li, NOAA GSL and Curtis Alexander, NOAA GSL

NOAA's Rapid Refresh Forecast System (RRFS) is a next-generation, regional convection-allowing ensemble forecast system developed collaboratively by EMC, GSL, and the broader UFS community for the National Weather Service. The RRFS Data Assimilation System (RDAS) utilizes hourly spin-up and product cycles, coupled with an hourly cycled EnKF data assimilation system. As RRFSv1 moves toward operational implementation in 2026, development is focused on transitioning the core DA engine from the legacy Gridpoint Statistical Interpolation (GSI) package to the Joint Effort for Data assimilation Integration (JEDI) system. This presentation details the ongoing efforts to evaluate JEDI's functional performance, and to optimize its computational efficiency to meet operational requirements. We provide a comprehensive comparative evaluation between the JEDI-based RDAS and the GSI-based RRFSv1 baselines, highlighting the benchmarks necessary for a successful transition to operations.

A Summary of HAFS Data Assimilation Activities at the Atlantic Oceanographic and Meteorological Laboratory

Jason Sippel*, Hyun-Sook Kim, Altug Aksoy, Sarah Ditchek, Dan Wu, Kathryn Sellwood, Brittany Dahl, Trey Alvey and Bachir Annane

This presentation reviews data assimilation (DA) activities within the Hurricane Analysis and Forecast System (HAFS) being conducted at NOAA Atlantic Oceanographic and Meteorological Laboratory's (AOML).

One major focus is improving the use of observations with a high readiness level. We have developed the tools to create operational superobservations from tropical cyclone reconnaissance data that improve data representativeness and quality control, leading to improved tropical cyclone (TC) forecasts in HAFS. AOML also showed that assimilating data from small uncrewed aircraft improves HAFS forecasts, and the model will assimilate such observations for the first time in 2026. Assimilation of SailDrone observations is also being tested.

Further focus is on low-readiness-level observations to fill gaps in our current observing capabilities. For example, AOML is currently exploring assimilation of OCONUS ground-based radar data and high-resolution data from additional airborne instruments to augment existing capabilities. In addition, AOML is exploring the use of observations from novel satellite platforms in HAFS, including microwave radiances from the NASA TROPICS CubeSat constellation.

AOML also leads major new DA technology developments. One major development effort is a JEDI-based HAFS-MOM6 coupled DA system to improve use of ocean observations in HAFS. Within this framework, we are also developing observation-space diagnostics and online quality control techniques to optimize assimilation of observations.

Finally, AOML also conducts observing-system experiments to evaluate and improve the efficacy of TC reconnaissance missions. Notable activities include an impact assessment of G-IV data and the exploration of the effectiveness of using reconnaissance for improving post-landfall hazard forecasts.

MAPCast: A Convection Allowing Emulator Toward Multi-Scale Data Assimilation

Yongming Wang*, Multiscale data Assimilation and Predictability (MAP) Lab, University of Oklahoma; Consortium for Advanced Data Assimilation Research and Education (CADRE)
Xuguang Wang, Multiscale data Assimilation and Predictability (MAP) Lab, University of Oklahoma; Consortium for Advanced Data Assimilation Research and Education (CADRE)
Machine learning (ML) methods offer a cost-efficient alternative to numerical weather prediction models for generating convection-allowing background ensembles in ensemble-based data assimilation (DA). However, few studies have explored ML-based surrogates for this application. This study develops MAPCast, a convection-allowing emulator trained on historical convection-allowing forecasts from the Model for Prediction Across Scales (MPAS), and evaluates its ability to estimate background-error covariances (BECs), paving the baseline toward multiscale DA. The evaluation uses 10 retrospective convective cases at 15- and 60-min forecast intervals relevant to subhourly and hourly DA. MAPCast reproduces MPAS forecasts with good fidelity, including realistic storm coverage, temporal evolution, and similar spatial and spectral characteristics of state variables. Discrepancies are primarily confined to small spatial scales near sharp gradients and convective features. For BEC statistics, MAPCast captures ensemble spread magnitude and spatial distribution for most variables, although larger errors occur for storm-related variables, particularly vertical velocity and reflectivity. Correlation structures are best replicated at large scales, followed by small scales; univariate correlations outperform cross-variable correlations,

revealing limitations in multivariate coupling. BEC estimates derived from 15-min forecasts consistently outperform those from 60-min forecasts, suggesting that shorter lead times better preserve flow-dependent error structures. Weaker replication of full-scale versus decomposed large and small-scale covariances underscores the importance of multiscale DA frameworks for effectively incorporating surrogate ensembles. These results suggest that further development of the baseline MAPCast needs to improve representation of near-grid-scale structures and cross-variable correlations.

Combining Data-driven and Physics-based Background Ensemble in Data Assimilation

Dan Kubalek*, Xuguang Wang, and Yongming Wang

The success of machine learning (ML) in providing cost-effective emulations of NWP model forecasts motivates the exploration of exploiting data-driven emulators to generate short-term background ensembles in ensemble-based DA. While an earlier study (Slivinski et al. 2025) has indicated the inadequate replication of the physics-based ensemble covariances by the pure data-driven ensemble covariances, properly combining data-driven and physics-based background ensemble to maximize the potential of data-driven surrogate ensemble in ensemble DA has not been explored. The goal of this study is to investigate whether combining the background ensembles produced by data-driven emulators and the physics-based models can efficiently improve the performance of ensemble-based cycling DA. This study explores this question using the surface quasi-geostrophic (SQG) model, an Attention Augmented U-Net (UNET) that emulates the SQG model and the Local Gain form of the Ensemble Transform Kalman Filter (LGETKF; Wang 2021) DA framework. A set of ensemble experiments was created: 20-member SQG ensemble (SQG20), 100-member UNET (UNET100), 60-member hybrid (10 SQG + 50 UNET; HV1) and another 60-member hybrid (10 SQG + 50 UNET; HV2) with the UNET bias removed by centering the 50-member UNET perturbations about the 10-member SQG ensemble members' mean. All ensembles had the same cost as the 20-member SQG ensemble. It was found that HV2 produced the most accurate analysis and forecast, outperforming both the pure physics and pure data-driven ensembles. A series of diagnostics demonstrated that UNET100 suffered from a small-scale error bias that translated to poor cycling DA performance. Once the physics-based ensemble members were added in HV1, there was improvement upon at least UNET100. However, HV1 performance was still only comparable to and slightly worse than SQG20. HV2 substantially improved the cycling DA experiments and significantly improved free-forecasts relative to both UNET100 and SQG20 with statistical significance over a 6-day forecast period.

Towards a Continuously Cycled AI Weather Prediction Model: An LETKF-Based Data Assimilation System for HRRRCast

Vanderlei Vargas Jr.*, Daniel Abdi, Isidora Jankov, and Paul Madden

Artificial Intelligence Weather Prediction (AIWP) models can now match some aspects of traditional Numerical Weather Prediction (NWP) systems at a fraction of the computational cost, yet they still depend on physics-based NWP for their training data. This work presents the first data assimilation system developed for HRRRCast, a convection-allowing AI model trained on High-Resolution Rapid Refresh (HRRR) analysis data, and aims to build a continuously cycled AI forecast system that runs without external NWP analyses. We couple HRRRCast with the Local Ensemble Transform Kalman Filter (LETKF), an ensemble-based data assimilation method that estimates flow-dependent background error covariances from a forecast ensemble and updates each grid point using only nearby observations within a local domain. This local formulation makes the LETKF computationally efficient and well suited to high-resolution regional models. Each cycle uses HRRRCast's own ensemble forecast as the background, assimilates surface observations from ground stations, and produces an analysis that initializes the next HRRRCast forecast cycle. To prevent ensemble inbreeding and spread collapse, we apply an LETKF cross-validation scheme in which each member is updated with a Kalman gain computed from an independent subset of the ensemble. The study will explore whether HRRRCast produces reasonable forecasts when initialized from these cycled LETKF analyses, or whether the model must be retrained on cycled LETKF analyses in place of the original HRRR analysis training data. The work points toward an observation-driven AI forecasting system that could lower even further the costs associated with weather forecasting.

Fully non-Parametric Ensemble Data Assimilation: Idealized and Sea Ice Model Experiments Using Data-driven Estimates of Likelihood Functions

Henry Santer* and Jonathan Poterjoy

Most data assimilation schemes currently used for research or operational prediction rely on an accurate characterization of observation error statistics. When this assumption is not well met, errors are likely to arise in both the resulting analysis and in the analysis uncertainty itself. Observations of sea ice concentration have errors that are difficult to characterize due to their high degree of non-Gaussianity and state dependence. As a result, the current operational approach of prescribing a fixed Gaussian observation error variance for sea ice concentration is insufficient. Motivated in part by these challenges, we introduce a method for directly estimating the likelihoods of model states given observations based on kernel embeddings of conditional distributions. When combined with data assimilation methods that can directly use likelihood functions to perform their updates, this approach has the dual benefit of automatically characterizing state-dependent uncertainty and enabling the assimilation of observations without the use of a (potentially expensive) measurement operator. Results with idealized models indicate that the method leads to forecast improvements in the presence of complex observation uncertainty. Next, we leverage the method to estimate likelihood functions for passive microwave-sensed sea ice concentration observations, based on data assimilation experiments using the SOCA system. Finally, we show preliminary results using the estimated sea ice concentration likelihood functions in Observing System Simulation Experiments conducted with CICE6.

Distribution-respecting Model-Space Localization for the Non-Variational Ensemble Data Assimilation of Non-Local Observations: Formulation and Experiments with Clear-sky Infrared Radiances

Man-Yau Chan*, Elizabeth Satterfield, Juli Rubin, and Edward Joseph Hyer

Ensemble data assimilation (EnsDA) for Earth systems is often limited to fewer than 100 ensemble members, resulting in substantial sampling errors. This study seeks to ameliorate those sampling errors through a novel efficient high-dimensional resampling algorithm: Local Probit-space Ensemble Size Expansion for Gaussian Copulas (LPGC). LPGC combines the statistics of the prior ensemble with users' expertise to generate many virtual ensemble members. These members are then used in EnsDA. LPGC constructs virtual members from a wide range of distribution families and imposes spatially localized statistics (i.e., B-localization). As such, LPGC has the potential to improve EnsDA.

This study investigates that potential through offline Observing System Simulation Experiments (OSSEs) using archived prior ensembles produced by the US Navy's Global Environmental Model (NAVGEM). These OSSEs assimilate synthetic clear-sky infrared radiance (clrIR) observations from a global ring of geostationary satellite infrared imagers. LPGC's impacts on EnsDA performance are investigated via OSSEs with various EnsDA algorithms (Gaussian or non-Gaussian), original ensemble sizes, and LPGC settings.

Results indicate that applying LPGC generally results in statistically significant improvements to the analyses of thermodynamic fields. Furthermore, OSSEs with LPGC outperform OSSEs that employ a popular model-space localization method (ensemble modulation; EMOD) because LPGC produces virtual ensembles that are faithful to user-chosen marginal distributions whereas that popular method does not. Finally, results (surprisingly) indicate that the flow-dependent prior marginal distributions of specific humidity might be well approximated by rectified Gaussian distributions. Our findings have the potential to improve NOAA’s EnsDA systems.

Improving land-atmosphere background error covariance estimation toward RRFS strongly coupled data assimilation

Aaron Johnson* and Xuguang Wang

Operational convective-scale data assimilation (DA) and forecast systems, such as RRFS, currently use a moderately-coupled DA system to enforce consistency between increments in the low-level atmosphere and soil state variables. Strongly-coupled DA (SCDA), in which both soil and atmosphere variables are simultaneously updated via the ensemble-based covariance, has the potential to further improve both consistency and accuracy of the land-atmosphere coupled analyses, especially during strong land-atmosphere interactions such as dryline convection initiation (CI) events. However, effective ensemble-based land-atmosphere SCDA at convective scales requires ensemble perturbations that sample uncertainty in the land-atmosphere interactions as well as appropriate methods of treating sampling error.

This presentation will provide an update on the status of SCDA in the CADRE JEDI-MPAS DA system followed by a dryline CI case study that is used to explore the structure and sampling error of land-atmosphere coupled correlations. It is found that the coupled correlations are particularly pronounced along the dryline, in the vicinity of the subsequent CI. While the auto-correlation within land and atmosphere variables has different spatial scales, the cross-correlation and its associated sampling error is flow dependent and can be highly anisotropic. Sampling error is not necessarily monotonically increasing with horizontal or vertical distance from a hypothetical observation located near the dryline CI. The results suggest that treatments of sampling error other than distance-based covariance localization may ultimately be necessary for optimizing convective-scale SCDA, especially in dryline CI cases. Experiments with ensemble perturbations related to uncertainty in the green vegetation fraction will also be briefly introduced.

Impact of Greenness Vegetation Fraction Perturbations on Convection-Allowing Ensemble-Based Land-Atmosphere Coupled Background Error Covariance.

Brett Castro, Aaron Johnson* and Xuguang Wang

With the introduction of convection allowing models, the ability to explicitly resolve deep moist convection has allowed for substantial improvements in short range forecasting. However, properly simulating the timing and location of convection requires accurate representation of land-atmosphere interactions which often play a considerable role in convective-scale forecasts, particularly in cases of dryline convection initiation (CI) with weak synoptic scale forcing. Land-atmosphere strongly-coupled data assimilation (SCDA) therefore has the potential to improve convective-scale analyses and forecasts. However, the design of ensemble perturbations used in data assimilation (DA) typically neglects the uncertainty related to land surface parameters such as the greenness vegetation fraction (GVF). We hypothesize that representing this uncertainty in the ensemble design will have meaningful impacts on the ensemble spread and correlation structures, and should therefore be included in ensemble-based background error covariances for convective-scale SCDA. With this in mind, the objective of this study is to compare the ensemble forecasts of a dryline CI case in the southern Plains with GVF perturbations to those with soil state perturbations. It is found that the forecast variance from GVF perturbations is approximately half of that from soil state perturbations for most variables, while the spatial scale of coupled covariance structures associated with GVF perturbations are larger than for soil state perturbations. Based on the encouraging impacts of GVF perturbations on the convective-scale forecasts, GVF perturbations are integrated into the CADRE MPAS-JEDI DA system. Evaluation of the impact on a conventional observation DA experiment that is continuously cycled over a 4-day period will also be presented.

A Question of Scale: Using Ensemble Data Assimilation to Understand Convection-Environment Interactions

Derek J. Posselt*, Longtao Wu, and Masashi Minamide

Severe convective storms are implicated in the majority of the billion-dollar disasters in the United States year after year. The outcomes of these storms are dictated in no small part by the strength of upward motions within convective updrafts. The strength and duration of convective updrafts have been shown to be sensitive to environmental conditions on a range of scales, but determining which environmental variables exert an influence on convection, as well as the time and space scales over which they operate, remains a challenge. High resolution modeling and data assimilation systems can be used to address a set of coupled research questions: 1) how can ensemble data assimilation make optimal use of observations near convection? 2) given an observation-constrained ensemble, what can be learned about the relationships between convective updrafts and their environments?

In this presentation, we show results from experiments that utilize ensemble-based all-sky assimilation of geostationary brightness temperatures for two very different convective cases: a weakly forced tropical oceanic squall line, and a rapidly intensifying hurricane. We demonstrate how assimilation of all-sky geostationary brightness temperatures improves analyses and forecasts of both events. We then show how an analysis that separates the ensemble covariances into different scales (meso-alpha, beta, and gamma) can be used to interrogate convection-environment interactions. We finish by showing how perturbations to moisture at the meso-gamma scales control the organization of convection and in particular hurricane rapid intensification.

Assimilating Boundary Layer Depths Observed by WSR-88Ds to Improve the Prediction of a Derecho

Braedon Stouffer*, David Stensrud and Yunji Zhang

Recent work has shown that dual-polarization WSR-88D radars can be used to accurately estimate the height of the planetary boundary layer (PBLH). While assimilation of these observations is relatively novel, it has been shown to improve model representation of moisture in the lower troposphere which is a key component for severe thunderstorm prediction. We expand on this previous work by developing techniques to assimilate WSR-88D PBLH observations using the Joint Effort for Data assimilation Integration (JEDI) system. Forecasts are generated using a configuration of the Model for Prediction Across Scales (MPAS) that is similar to the Rapid Refresh Forecast System (RRFS) which is currently under development as part of NOAA's Unified Forecast System. We evaluate the impacts of PBLH data assimilation on the predicted evolution of a derecho event over North Dakota on 20-21 June 2025. Observed PBLH from nearby WSR-88Ds are also compared to forecasts from the operational High Resolution Rapid Refresh (HRRR) model which show notable biases during the day of the event.

Improving Forecasts of Hurricane Beryl (2024) Genesis through Assimilation of All-Sky Infrared Radiances

Chandler Pruett* and Dr. Xingchao Chen

Predicting tropical cyclogenesis remains a major challenge because it often occurs over data-sparse oceans and is governed by complex multiscale interactions among tropical convection, precursor disturbances, and the larger-scale environment. This study investigates the impact of assimilating all-sky infrared (IR) brightness temperatures from geostationary satellites on forecasts of Hurricane Beryl (2024) genesis, a particularly difficult case to predict. The experimental design employs two state-of-the-art data assimilation systems: the Hurricane Analysis and Forecasting System with Joint Effort for Data Assimilation Integration (HAFS-JEDI) and the Pennsylvania State University WRF Ensemble Kalman Filter system (WRF-EnKF). Notably, all-sky IR radiances are not currently assimilated in the operational HAFS.

Results show that assimilating all-sky IR radiances, in addition to conventional observations, improves forecasts of Hurricane Beryl's genesis in both HAFS-JEDI and WRF-EnKF. The additional IR radiance data substantially modifies the initial moisture distribution and the circulation structure of the precursor easterly wave. These changes promote more realistic coupling between convection and the parent wave, leading to improved forecasts of tropical cyclogenesis and the subsequent rapid intensification of Beryl. We will further compare and discuss the potential dynamical pathways through which all-sky IR radiance assimilation improves the forecasts in HAFS-JEDI and WRF-EnKF, highlighting both shared and system-dependent mechanisms.

Preliminary Results of Weakly Coupled Data Assimilation - Hurricane Analysis and Forecast System (WCDA-HAFS) Forecasts

Hyun-Sook Kim*, HeeSook Kang, Biju Thomas, Bin Liu and Zhan Zhang

The ocean component of the Hurricane Analysis and Forecast System (HAFS) has been significantly enhanced through the integration of the Marine JEDI module within the UFS-HAFS framework, hereafter weakly coupled data assimilation -HAFS (WCDA-HAFS). This advancement enables DA for both the atmospheric and oceanic models in HAFS, a critical step towards improving tropical cyclone (TC) forecasts by accounting for essential air-sea interaction processes. Recent case studies highlighted two issues: 1) data latency of remote-sensing Absolute Dynamic Topography (ADT) and various in-situ ocean profile observations within the Global Telecommunication System (GTS) and 2) the limited quantity of in-situ data. These issues present a significant challenge for accurately constraining the subsurface ocean conditions in an operational environment, especially since the assimilation of outdated information reduces analysis accuracy. Retrospective case studies conducted for Hurricane Fiona demonstrated the system's potential, showing an average improvement of up to 20% improvement in intensity forecast skills when all observations are ingested, both mixed degradation and improvement when only ADT are excluded, and a degradation of less than 20% when only glider observations are removed. A detailed analysis of the results from the WCDA-HAFS forecasts will be presented.

Constraining Atmospheric Forcing Errors for Sea Ice Data Assimilation in Short-Range Forecasts

Hee-Sung Jung*, Jonathan Poterjoy and Alek Petty

Benefits of a dynamic sea ice component in global weather prediction systems for enhancing short-range forecast skills have been widely noted. To maintain realistic sea ice and associated atmosphere/ocean states throughout model forecasts, sea ice data assimilation (DA) has become an increasing focus. Ensemble DA has been a popular algorithm choice due to its ability to propagate flow-dependent error statistics through time. A key challenge in applying ensemble DA to the sea ice component within ESMs is correctly accounting for atmospheric/oceanic influence on the sea ice forecast error (commonly referred to as forcing errors), which have been argued to be a prominent error source at shorter timescales.

We construct an observing-system simulation framework for sea ice DA research that aims to faithfully capture atmospheric influence on short-range multivariate sea ice forecast errors by directly modeling/prescribing uncertainty in the atmospheric forcing. Specifically, we leverage a decade-long coupled ESM dataset to train an atmospheric covariance model that dynamically perturbs atmospheric forcings across short-term sea ice forecasts during DA. We show that our model for atmospheric forcing uncertainty can be tuned to match short-range (6-h) error growth rates found in sea ice predictions from a prototype operational UFS-based DA system, while properly quantifying the known source of error in the ensemble spread. We also show that a common approach for tuning sea ice ensembles based on model parameter perturbation has limited impacts in short-range forecasts, further highlighting the importance of constraining the atmospheric forcing errors for sea ice DA at such timescales.

UFS Applications (HAFS, Coastal)

UFS Hurricane Application HAFS: 2026 Upgrades and Future Plans

Zhan Zhang* and Fanglin Yang

The Hurricane Analysis and Forecast System (HAFS) is NOAA's flagship operational tropical cyclone (TC) prediction system, providing TC track and intensity forecasts across all global ocean basins. Since the operational implementation of HAFSv2.1 on July 16, 2024, substantial upgrades and enhancements have been introduced. HAFS will continue to undergo annual improvements, incorporating emerging science and technology, adding new capabilities, and resolving issues identified in earlier versions.

This presentation will review the performance of the operational HAFS during the 2025 hurricane season, present the latest upgrades of HAFSv2.2, which is scheduled in July 2026. It will also highlight current and planned research activities, as well as research-to-operations (R2O) plans associated with ongoing and future HAFS development. Development priorities for HAFS will be discussed, with an emphasis on emerging needs from HFIP and stakeholders following the retirement of the legacy HWRF and HMON models. The current status and future plans for AI based hurricane applications at EMC will also be presented.

Transitioning HAFS Data Assimilation to JEDI: Observation Dataset Readiness and Integration in HAFS

Jing Cheng*, Xu Lu, Yonghui Weng, Bin Liu, Zhan Zhang, Shun Liu, Daniel Holdaway, Daryl Kleist, Fanglin Yang, and Vijay Tallapragada

To support future operational upgrades of the Hurricane Analysis and Forecast System (HAFS), a data assimilation (DA) capability is being actively developed to transition the system to the next-generation Joint Effort for Data assimilation Integration (JEDI) framework. A critical component of this transition is ensuring the complete readiness and robust integration of the complex observation datasets required for tropical cyclone initialization.

This presentation will highlight the migration of a comprehensive suite of observation types from the legacy GSI system to the new JEDI framework. The modernized data assimilation system has been engineered to natively process both the standard observation data utilized by Global Forecast System (GFS) and the specialized, high-impact observation unique to HAFS. Key datasets successfully transitioned include Tail Doppler Radar (TDR) data, reconnaissance dropsondes, High-Density Observations (HDOB), and high-resolution Atmospheric Motion Vectors (AMVs). Results demonstrate that by seamlessly integrating this extensive observational suite, the JEDI-based system yields inner-core and environmental analyses that are comparable to the current GSI baseline. Finally, the presentation will outline our strategic roadmap for transitioning the JEDI-based HAFS into operational use, which will be one of the main objectives for the future HAFSv3 upgrade, detailing the rigorous evaluation and parallel testing scheduled for the upcoming 2026 summer hurricane season to fully validate the system's readiness.

Hurricane Analysis and Forecast System: R&D and R2O

Sundararaman Gopalakrishnan*, William Ramstrom, Lew Gramer, Andrew Hazelton, Ghassan Alaka, Zhan Zhang, Aaron Poyer, Fanglin Yang, Biin Liu, Jason Sippel, William Komaromi and Vijay Tallapragada

NOAA's Hurricane Forecast Improvement Program (HFIP) has reduced hurricane track and intensity errors by more than 50% since 2009. As we approach the 2026–2027 seasons, the focus is expanding beyond the "three numbers" (latitude, longitude and maximum winds) to actionable, high-fidelity guidance for hazards—wind gusts, flash flooding, and tornadoes—often embedded in tropical cyclone rainbands.

At the center of research and development (R&D) and its transition to operations (R2O) is the Hurricane Analysis and Forecast System (HAFS), NOAA's flagship model for hurricane prediction. Since its operational implementation in 2023, HAFS has evolved toward a basin-scale parent domain with multiple telescopic moving nests, enabling multi-storm simulation, improving storm–storm and land–storm interactions, and supporting future JEDI-DA development. This framework is demonstrating measurable forecast skill gains while providing a more realistic environmental context for tropical cyclone evolution. HAFS is further grounded in observations through boundary layer physics informed by NOAA P-3 Hurricane Hunter data, improving representation of inflow, turbulence, and air–sea coupling.

HFIP Real-time experiments (HREx) during the 2024–2025 seasons demonstrated the performance of this multi-nest system, with improvements in track, intensity, and structure, and increasing skill in representing rainbands and associated hazards. They also highlight a key limitation: the "gray zone" of resolution. While current 1–2 km nests improve capability, capturing rapid intensification and rainband-scale hazards in small or highly sheared systems requires finer resolution.

We will discuss the transformative potential of these advances through a forward-looking R&D roadmap toward ~1 km and sub-km HAFS. By integrating advanced data assimilation (JEDI-DA) with scale-aware and sub-grid physics, and leveraging next-generation computing such as URSA, HAFS is progressing toward a fully convection-permitting, multiple-nest system—critical for next-generation probabilistic, hazard-centric guidance for landfalling tropical cyclones.

Efficient High Resolution Forecasts of Hurricanes through Basin-Scale and Telescoping Nest HAFS

William Ramstrom*,Lewis Gramer, Andrew Hazelton, Sundararaman Gopalakrishnan

Model resolution is critical for capturing the fine-scale features of tropical cyclones that drive forecast accuracy. Observations gathered from NOAA hurricane reconnaissance flights have documented many features in intense hurricanes that are at 1 kilometer or smaller scales. In the watch/warning timeframe, hurricane specialists need accurate guidance on probability of eyewall replacement cycles, extent of hurricane-force and tropical storm-force wind fields, and hazards such as 10m wind and precipitation along coastlines and inland from landfall.

For many years into the foreseeable future, it is clear that there will remain a demand for ever-finer resolution for the inner core of hurricanes – both for operational and research purposes. To run on available computing resources, these resolutions will require the use of nesting.

Basin-scale HAFS, which has been run in real-time parallel starting from July 2025, has demonstrated accuracy improvements and computational efficiency – 4 storms can be simulated for the same number of core hours as 2 storm-centered HAFS runs, while also better capturing storm-storm interactions across scales.

Moving telescoping nests are currently under development at AOML for future releases of HAFS. Telescoping nests allow us to run 6km resolution for the large-scale environment, 2km resolution for an intermediate nest, and 660m resolution for the storm-centered nest. We will demonstrate results from sub-kilometer scale simulations of Hurricane Melissa that run on current operational-sized clusters.

Crossing the Boundaries: Global Moving Nests

Rusty Benson* (NOAA/GFDL), Lucas Harris (NOAA/GFDL), and Jesse Lentz (SAIC)

The powerful nesting capabilities of the FV3 dynamical core used by NOAA and NWS requires the nest to be co-incident with the parent grid and contained within a single tile of the cubed-sphere grid. NWS and WPO have funded a project at GFDL to overcome this limitation and implement a capability for nests to span across tile boundaries and move freely throughout the cubed-sphere grid. This new feature will enable tracking of multiple atmospheric features like tropical cyclones, atmospheric rivers, and mesoscale convective systems in a global domain. This talk/poster will demonstrate two approaches and progress towards implementation of global moving nests.

Recent GFDL development updates: T-SHiELD and Vortex Tracker

Tim Marchok*, Kun Gao and Caitlyn McAllister

This presentation will cover two different areas of tropical cyclone research & development at NOAA/GFDL: the T-SHiELD tropical cyclone model and the GFDL Vortex Tracker.

T-SHiELD is the configuration of the NOAA/GFDL System for High-resolution modeling for Earth-to-Local domains (SHiELD) that is designed for application to tropical cyclone prediction, and it features a high-resolution nest over the Atlantic Basin that is embedded within a global model. Details will be provided on forecast improvements from an increase in resolution from 3.25 km to 2.6 km. Details will also be provided on the implementation this season of an AI-NWP version of T-SHiELD that uses a selective nudging approach to relax the high-resolution T-SHiELD fields toward AI-NWP forecast data during model integration.

A project continues to develop the GFDL Vortex Tracker as a community-supported vortex tracker for the Unified Forecast System (UFS). The genesis detection algorithm in the tracker has recently been updated to utilize methods that more closely mimic operational methods for diagnosing observed storm genesis, and this new algorithm also reduces the number of false alarms for genesis detection as compared to the previous version of the tracker. Also, new functionality has been added to the tracker to diagnose the degree of vertical tilt in a TC vortex, allowing the user to analyze the tilt using metrics of relative vorticity, wind circulation, temperature or geopotential height.

NWS Marine Component Synergies with the Unified Forecast System

Jessica Meixner*, Arthur Taylor, Michael Churma, Jung-Sun Im, Huiqing Lliu, and Mark Willis

As part of the National Weather Service transformation and the merging of the Meteorological Development Laboratory (MDL) and the Environmental Modeling Center (EMC), the marine components of storm surge, rip currents and waves are joining forces to improve future marine and coastal hazard forecasts. Currently, this team supports many existing operational models such as: Probabilistic tropical cyclone storm Surge (P-Surge), Probabilistic Extra Tropical Storm Surge (P-ETSS), Nearshore Wave Prediction System (NWPS), and Great Lakes Wave Unstructured (GLWU). In addition the team supports the wave components of the Global Forecast System (GFS), the Global Ensemble Forecast System (GEFS), and the Hurricane and Analysis Forecast System (HAFS) models. Furthermore, the team supports the Rip Current Model Output Statistics (RCMOS) and the Regional Wave Prediction System (RWPS) which are in development for future operations.

This talk aims to introduce this new team to the Unified Forecast System (UFS) community, sharing overviews of currently supported operational systems and examples of future opportunities for synergy between these systems and the UFS community. Currently the wave components of GFS, GEFS and HAFS all use WAVEWATCH III and the UFS weather model. Rip current models and coastal impact predictions, while not traditionally seen as a UFS model, rely on outputs of the wave model and provide critical information for dangerous phenomena. Among the four operational models currently maintained by the team, P-ETSS exhibits the highest degree of UFS compatibility due to its reliance on global model wind and pressure guidance rather than forecaster-generated or parametric inputs. Consequently, investigating its potential synergies with the UFS framework represents a significant opportunity for further development. Our team's hope is to collaborate with the UFS community to improve model guidance.

UFS-Coastal: Engineering the Enterprise Backbone for NOS's Next-Generation Coastal Earth System Model and Operational Applications

Saeed Moghimi* (NOAA/NOS/OCS), Maoyi Huang (NOAA/OAR/WPO/EPIC), Ufuk Turuncoglu (UCAR/NCAR), Yunfang Sun (NOAA/NOS/OCS), Mansur Jisan (NOAA/NOS/CO-OPS), Felicio Cassalho (NOAA/NOS/OCS), Alexander Kurapov (NOAA/NOS/OCS), Soroosh Mani (NOAA/NOS/OCS), Fariborz Daneshvar (NOAA/NOS/OCS), Liujuan Tang (NOAA/NOS/OCS), Ed Myers (NOAA/NOS/OCS), Ann Tsay (UCAR/NCAR), Greg Seroka (NOAA/NOS/OCS), Ayumi Fujisaki-Manome (CIGLR), Lucila J. Houttuijn Bloemendaal (NOAA/NOS/CO-OPS), Ali Abdolali (USACE/ERDC/CHL), Carsten Lemmen (Hereon), Joseph Zhang (VIMS), Damrongsak Wirasaet (UND), Joannes Westerink (UND), Hernan G. Arango (Rutgers), Scott Durski (OSU), Jihun Jung (OSU), Joseph Smith (OSU), Kristopher Booker (EPIC Project), Shachak Peeri (NOAA/NOS/NGS). Corey Allen (NOAA/NOS/OCS), Patrick Burke (NOAA/NOS/CO-OPS), Stephan Smith (NOAA/NOS)

NOAA's National Ocean Service (NOS) is leading the transition toward a unified, high-fidelity modeling enterprise. In partnership with OAR, NWS, USACE/ERDC and the broader academic community, NOS is developing UFS-Coastal, a specialized coupling infrastructure integrated into the NOAA Unified Forecast System (UFS) portfolio. This framework represents a paradigm shift from siloed regional models to a standardized, inter-operable Earth System approach.

This presentation outlines the plans and current development status of the ufs-coastal-model infrastructure and its mission-critical downstream applications (ufs-coastal-app). By leveraging the

ESMF/NUOPC coupling architecture, UFS-Coastal integrates community models (including ROMS, ADCIRC, SCHISM, and FVCOM) with other physical domains such as sea-ice, inland hydrology and surface waves (WW3). We will highlight the development of a set of applications and tools towards operational implementations designed to support the Blue Economy and Coastal Resilience, including:

- SECOFS & STOFS-3D-Alaska: Advanced 3D unstructured hydrodynamic (SCHISM), surface wave (WW3) and sea-ice (CICE) coupling for high-latitude and regional navigation safety.
- Hurricane-Surge (PAHM+SCHISM+WW3): Next-generation compound flooding prediction for life-saving emergency response in collaboration with NWS NHC.
- NOS Enterprise Workflow: The modernization of operational pipelines to ensure 99.9% up-time and standardized delivery of navigational products.

Developing and Assessing Two-Way Hydrologic–Hydrodynamic Coupling for Coastal Backwater-Affected and Compound Flood Prediction

Yi Hong*, Joseph Zhang, Jason Ducker, David Canon, Dan Titze and Jesse Feyen

Accurate prediction of coastal flooding depends on representing dynamic interactions between watershed runoff and downstream water levels. In low-gradient, lake- and ocean-connected systems, backwater effects from coastal water levels, seiches, and surges can substantially alter river stage and routing, challenging traditional hydrologic models that assume discharge-driven flow. This study develops and assesses a two-way coupling framework between hydrologic routing and hydrodynamic models to enhance compound flood prediction in coastal environments.
A prototype system couples T-Route (a vector-based channel routing model) and SCHISM (Semi-implicit Cross-scale Hydroscience Integrated System Model) within a test domain near Rochester, New York. T-Route is driven by hourly National Water Model retrospective runoff inputs and executed with subhourly internal timesteps, while SCHISM resolves nearshore hydrodynamics and lake-level variability. Two-way coupling is achieved through the exchange of discharge and water surface elevation at defined interface locations, allowing downstream water levels simulated by SCHISM to dynamically influence upstream routing.

USGS gauge observations (15-minute stage and discharge) are utilized to select representative backwater events. The coupled model is evaluated against uncoupled routing to assess improvements in stage prediction during backwater conditions.

This study offers an initial demonstration of hydrodynamic feedback in coastal river systems and underscores the potential of an inland and coastal two-way coupling application to support the Next Generation Water Resources Modeling Framework (NextGen) development and NOAA coastal compound flooding applications. Future work will expand the analysis to additional coupling configurations, including extended-domain and simplified nearshore coupling scenarios.

Coupling CICE to SCHISM in UFS coastal:Working towards a coupled ice-ocean forecast system for Alaska

Scott Durski*, Joseph Smith, Alexander Kurapov, Joseph Zhang, Felicio Cassalho, Dan Yu, Edward Myers, and Saeed Moghimi

Efforts continue on the development of a three-dimensional (3D) prediction system for the Alaska region that two-way couples the Los-Alamos National Laboratory ice model CICE to the unstructured grid ocean circulation model SCHISM (the Semi-implicit Cross-scale Hydroscience Integrated System Model) as parts of the Unified Forecast System (UFS). The objective is to develop a system that will provide predictions of the total sea level along the Alaska coasts, tidal and non-tidal currents over the shelf and around the Aleutian Islands, and ice parameters of interest to the navigational community (ice extent, concentration, thickness, motion, deformation, etc.). Additionally, the 3D unstructured mesh approach of SCHISM provides the ability to capture navigable inlets and passages of particular importance. A two-way coupled version for a domain that encompasses the Bering Sea has been successfully developed and implemented in UFS-coastal. Currently we are validating oceanic and ice fields against satellite and in situ estimates for the 2019-20 winter seasons. This presentation will focus on demonstrating the strengths and limitations of standard settings of the CICE model for replicating observed event- to seasonal- scale sea ice variability in the Bering Sea. Current work to both extend the model domain, into the Chukchi and Beaufort Seas to the north and the Gulf of Alaska to the south, and to introduce wave coupling will also be presented.

Harmful Algal Bloom Forecasts: Examples in Adaptive Unification

Richard P. Stumpf*, NOAA National Centers for Coastal Ocean Science, Yizhen Li, Xin Yu, CSS, Michelle Tomlinson, Kari St.Laurent, NOAA National Centers for Coastal Ocean Science

Harmful algal blooms (HABs) are significant hazards for our marine and Great Lakes coasts. HABs pose differing risks to human health, protected species, fisheries, aquaculture, coastal economies, and public water supplies, depending on the HAB species and region. The NOAA National Centers for Coastal Ocean Science produces operational HAB forecasts to reduce the detrimental effects of these blooms, and these forecasts are supported by a combination of models (e.g. ROMS, FVCOM, and SCHISM) and modeling strategies. The forecasts use output either from standard operational hydrodynamic models or from custom models if we need a specific ecological solution. Ecological models and other data types (including satellite-derived distributions and in situ observations) are combined with the hydrodynamic models in order to provide solutions to each specific HAB risk. Examples presented here include Lake Erie (which also includes a hypoxia model), Gulf of Maine, Chesapeake Bay, and Florida. We will also present our ‘future' goals to produce HAB forecasts on a truly national scale.

The applications of UFS-Coastal: wave-current coupling of SCHISM and WAVEWATCH III

Yunfang Sun*, Ufuk Turuncoglu, Mansur Jisan, Hao-Cheng Yu, Joseph Zhang, Carsten Lemmen, Ali Abdolali, Denise Worthen, Ishan A Tsay, Jana Haddad, Panagiotis Velissariou, Felicio Cassalho, Soroosh Mani, Fariborz Daneshvar, Saeideh Banihashemi, Ali Salimi-Tarazouj, Joseph Smith, Edward Myers, and Saeed Moghimi

Accurate prediction of coastal flooding and nearshore processes requires a coupled modeling framework that resolves interactions among ocean circulation, surface waves, ice, and atmospheric forcing across complex coastal geometries. To address this need, we present Unified Forecast System Coastal model (UFS-Coastal), a coupled coastal modeling framework extending the Unified Forecast System Weather Model (UFS-WM). The UFS-Coastal was developed to address this need by disentangling extending the prior infrastructure (called CoastalApp), and combining CoastalApp and UFS-WM, and by integrating four different coastal hydrodynamic models: SCHISM, ADCIRC, FVCOM, and ROMS, along with the spectral wave model WAVEWATCH III (WW3) and CICE within a unified coupled framework. Here we introduce UFS-Coastal to the community, and focus on one particular aspect, the wave-current coupling between SCHISM and WW3 with two methods: the Longuet-Higgins radiation stress formulation and a three-dimensional (3D) vortex approach. We provide a comprehensive description of UFS-Coastal's infrastructure, coupling methodologies, and its automated regression testing framework. The system's performance is demonstrated through experiments in the Duck, North Carolina (NC) region against Field Research Facility observations. Both wave-current coupling schemes between the SCHISM and WW3 could successfully reproduce observed coastal hydrodynamics and wave conditions; notably, the 3D vortex coupling improved simulations of wave-induced mixing in shallow and stratified environments. Overall, UFS-Coastal offers a robust framework for accurate prediction of coastal flooding, storm surge events, and wave-current interactions, supporting both research and operational forecasting applications.

National Ocean Service Core Model Evaluation Plan

Patrick Burke*, NOS/CO-OPS, Aijun Zhang, NOS/CO-OPS, Edward Myers, NOS/OCS, Saeed Moghimi, NOS/OCS, Kathleen Bailey, NOS/IOOS, Becky Baltes, NOS/IOOS, Stephan Smith, NWS, Hendrik Tolman, NWS/OMD, Hite Spencer, NOS, Will Tyburczy, NOS, Joanna Peth, NOS

The National Ocean Service (NOS) currently employs a diverse suite of community-supported physical oceanographic models, including ADCIRC, SCHISM, FVCOM, and ROMS. While this diversity enables NOS to meet a wide range of mission requirements, it also fragments resources by requiring support for multiple modeling frameworks. To address this, NOS will conduct a formal evaluation to identify an optimal set of core model codes capable of meeting its operational and statutory needs while positioning the program to adopt emerging advances in computing, data integration, and predictive methodologies.

Physical oceanographic models are foundational to NOS operations, particularly for marine navigation and integration with NOAA's Unified Forecasting System (UFS). They also underpin

numerous downstream applications across NOS and other NOAA line offices. Accordingly, this evaluation will focus on these models and their underlying core codes.

NOS model and statutory mission requirements will guide the assessment, which will include standardized “apples-to-apples” comparisons using objective performance metrics such as accuracy, forecast skill, resolution, and computational efficiency. In parallel, qualitative criteria—such as projected community support and sustainability over the next 5–10 years—will be considered.

The evaluation will also examine each model’s capacity to leverage advances in high-performance and cloud computing, GPU and emerging hardware architectures, machine learning and hybrid modeling approaches, automated data assimilation, and ensemble forecasting. The outcome will guide strategic investment in a streamlined, future-ready modeling portfolio for NOS. Transparent communication and feedback from the modeling community is critical to ensure the success of this evaluation.

UFS Applications (Air Quality, Land, SRW/RRFS)

Evaluation of RRFSv1 during the 2026 NOAA HWT Spring Forecasting Experiment

Israel Jirak*, Adam J. Clark, David Harrison, Jake Vancil , and Tim Supinie

The 2026 NOAA Hazardous Weather Testbed Spring Forecasting Experiment (SFE2026) was conducted from 27 April – 29 May with participation from forecasters, researchers, and model developers from around the world. One of the primary activities of SFE2026 was to evaluate the FV3-based Rapid Refresh Forecast System version 1 (RRFSv1), including RRFS deterministic and RRFS Ensemble Forecast System (REFS) components. The RRFSv1 is scheduled for operational implementation in the National Weather Service later this year to enable the retirement of several deterministic CAMs and the High Resolution Ensemble Forecast (HREF) system. Several deterministic and ensemble evaluations were conducted to compare the performance of the RRFS and REFS to operational baselines and the next version of RRFS (i.e., RRFSv2) based on the Model for Prediction Across Scales (MPAS). The subjective evaluation results of these RRFS and REFS forecasts from the SFE2026 will be discussed alongside objective verification statistics, offering evidence regarding the overall performance of the RRFS and REFS for severe weather forecasting.

RRFS Development Towards the NWP Moonshot

Curtis Alexander* (NOAA/OAR/GSL), Ligia Bernardet (NOAA/OAR/GSL), Terra Ladwig (NOAA/OAR/GSL), Matt Pyle (NOAA/NWS/OMD), and Jacob Carley (NOAA/NWS/OMD)

At the annual American Meteorological Society (AMS) meeting in January 2026, the National Weather Service announced an ambitious goal for NOAA to adopt a global, fully coupled and convection allowing 3-km unified system using the Model Prediction Across Scales (MPAS) for both research and operations*. The Earth Prediction and Innovation Center (EPIC) will play a key role in accelerating MPAS into the Unified Forecast System (UFS), a community-based system that supports NOAA's operational modeling systems.

Over the past few years, the EPIC has developed a publicly accessible, community-based weather research modeling system. It has also established testing and development capabilities to foster strong collaboration across the Earth system modeling community. In addition, EPIC has created platforms and protocols that provide training and tutorials for developers and users with varying levels of experience. Leveraging the established community modeling infrastructure and community engagement framework, EPIC can enable continual enhancement of operational models by providing tools, resources, and shared knowledge to the broader community.

To support the desired adoption of MPAS for both research and operations, in partnership with National Center for Atmospheric Research, NOAA modeling laboratories and centers, EPIC's community modeling infrastructure will be enhanced to incorporate MPAS. This will include enhancements to pre-processing, postprocessing, verification and validation packages in the UFS framework augmented with documentation, user support and community events. EPIC's community engagement efforts will be expanded to include the MPAS community. The EPIC program will also provide project management to establish an agile model development team to accelerate model development, and funding support to accelerate infrastructure enhancement toward an MPAS-based numerical weather prediction system pending resource availability. In this presentation, we will provide updates on EPIC's progress on those fronts and solicit input for future directions.

New RRFS: MPAS plus JEDI Development and 2026 Convective Season Performance

Terra Ladwig*, Ming Hu, Clark Evans, David Dowell, Jeff Beck, Chunhua Zhou, Adam Clark, Kent Knopfmeier, Israel Jirak, and GSL RRFS development team

Hourly-updating numerical weather guidance at convective-allowing scales is the foundation of severe storms forecasting. We are on the precipice of a new era of community modeling at convection-allowing scales with NOAA regional modeling research and development focused on the Rapid Refresh Forecast System (RRFS). The implementation of the first version of RRFS is scheduled for summer 2026 and it leverages the UFS with the FV3 dynamical core, GSI for assimilation, and UPP for post processing. A significant pivot for RRFS has taken place and future versions will leverage the Model for Prediction Across Scales (MPAS) and the Joint Effort for Data Assimilation Integration (JEDI), along with UPP for post processing.

This presentation will share research results from the first real-time RRFS prototype built with MPAS and JEDI, named RRFSv2X. The RRFSv2X began running in real-time in early 2026 and will be evaluated by the Spring Forecast Experiment (SFE) and the Flash Flood and Intense Rainfall Experiment (FFaIR). It includes hourly data assimilation cycling using both JEDI 3DEnVar and GETKF with 1-hour 30 ensemble member MPAS forecasts and subsequent deterministic forecasts out to 48 hours. Current development includes efforts to increase the observation types being assimilated toward matching all observations assimilated by RRFSv1. Improving observation processing including thinning, quality control, and error settings are being investigated in the MPAS + JEDI framework. The cycling strategy is evolving from cold start and partial cycling towards more continuous cycling with blending. The assimilation methodology is being advanced and optimized to improve accuracy and computational performance. The physics schemes have also undergone significant development. Quantitative and qualitative evaluations of RRFSv2X in the 2026 convective season will be shared in this presentation. MPAS forecasts initialized from RRFSv1, called “GSL MPAS RT” and MPAS forecasts initialized from HRRR, called “NSSL MPAS HT”, will be compared to evaluate the initial conditions provided by the JEDI cycling. The operational HRRR and RRFSv1 will also be used as benchmarks for comparisons with RRFSv2X forecasts including reflectivity statistics and case study evaluations for several 2026 severe weather events.

Progress on the Development of CATChem Version 2 within NOAA's Unified Forecast System

Wei Li*, Barry Baker, Zachary Moon, Youhua Tang, Patrick Campbell, Jian He, Maggie Bruckner, Quazi Rasool, Rebecca Schwantes, Li (Kate) Zhang, and Brian McDonald

The Configurable Atmospheric Chemistry (CATChem) component is being developed to unify all chemistry-related code within a single, flexible framework in NOAA's Unified Forecast System (UFS). CATChem v2 features a modular core architecture in which a Process Manager orchestrates chemistry, aerosol, emission, and deposition components through centralized state management. A column virtualization approach transforms 3D atmospheric fields into independent 1D columns, reducing memory pressure and enabling efficient OpenMP/MPI parallelization with improved cache performance. To accelerate the addition of new scientific processes, a Python-based Process Generator uses YAML configuration and Jinja2 templating to automatically produce standardized Fortran code, test suites, and documentation, reducing the typical development cycle from weeks to days while eliminating common errors in field access and memory management.

CATChem v2 couples to the UFS through a NUOPC cap that handles multi-phase initialization, field advertisement, and bidirectional state transformation between NUOPC and CATChem data structures. Major GOCART aerosol processes — including dust, sea salt, organic carbon and black carbon, and sulfate — have been implemented and tested within the UFS framework. All model behavior is controlled through a hierarchy of YAML configuration files covering simulation settings, species properties, emission mappings, and NUOPC field mappings. Preliminary comparisons between CATChem v2 and the Global Chemistry and Aerosol Forecast System (GCAFS) will be presented to evaluate the fidelity of the newly implemented processes. CATChem v2 is designed to serve as the unified atmospheric chemistry component for NOAA's next-generation forecasting systems, providing a configurable and extensible foundation for applications ranging from global aerosol forecasting to regional air quality prediction.

Development and Evaluation of NOAA's global aerosol model: Global Chemistry and Aerosol Forecast System (GCAFS v1)

Partha S. Bhattacharjee*, Li Pan, B. Baker, Yaping Wang, Andrew Tangborn, Cory Martin, Ho-Chun Huang, Li Zhang and F. Yang

The National Oceanic and Atmospheric Administration (NOAA) has developed a new global aerosol forecast model named Global Chemistry and Aerosol Forecast System (GCAFS) version 1.0 with the purpose of improving predictions of dust, sea salt, smoke from wildfires, and other particles in the atmosphere that can impact visibility or human health replacing the current capability of the Global Ensemble Forecast System (GEFS). Alongside GCAFS version 1.0 comes the Global Constituent Data Assimilation System (GCDAS) which provides analyses of aerosol species to initialize the GCAFS forecasts using the next-generation assimilation software, the Joint Effort for Data assimilation Integration (JEDI). The aerosol component in the GCAFSv1 is based on NASA's Goddard Chemistry Aerosol Radiation and Transport model (GOCART). GCAFS includes bulk modules from the GOCART as well as an updated FENGSHA dust emission scheme from the previous GEFS-Aerosol model. The Blended Global Biomass Burning Emissions Product (GBBEPx version 5) provides biomass burning emission. The global anthropogenic emission inventories are derived from the Community Emissions Data System (CEDS) and all sub-grid-scale transport and convective wet deposition are handled inside the atmospheric physics routines while the large-scale wet deposition and dry deposition is handled within the GOCART component. GCAFS v1.0 will provide a 5-day GCAFS forecast with 3-hourly output for the 00 and 12 UTC cycles. The model is scheduled to be operational in the later part of 2026.

This work will describe the details of GCAFS model configuration and model performance evaluation of real-time and retrospective runs. The model's performance will be assessed against a variety of observational data, including in situ measurements, satellite observation and reanalysis data (MERRA2). GCAFS v1 predictions demonstrate substantial improvements for both composition and variability of aerosol distributions compared to the former operational GEFS-Aerosol system with the inclusion of Aerosol DA and model updates.

Evaluation of Smoke Transport in NOAA's RRFS-SD Model Using Satellite Observations

R. Ahmadov, J. Romero-Alvarez*, H. Li, J. Schnell, S. Bhimireddy, E. James, M. Hu, S. Kondragunta, C. Xu, F. Li, and H. Zhang, A. Huff

NOAA's Global Systems Laboratory, in collaboration with other laboratories across NOAA, has developed a next-generation coupled weather forecast model: the Rapid Refresh Forecasting System with Smoke and Dust (RRFS-SD). Currently, the model is running in experimental mode at NOAA/NCEP (https://rapidrefresh.noaa.gov/RRFS-SD/). RRFS-SD is scheduled to be implemented operationally in summer 2026.

The RRFS-SD domain covers all of North America with 3 km spatial resolution. The forecasting system ingests the RAVE satellite fire radiative power product from NESDIS to estimate biomass-burning emissions and fire heat fluxes in real time. A smoke tracer (primary particulate matter) from biomass burning is advected and mixed within the model during the simulation. One of the key challenges in smoke forecasting is representing the vertical distribution of smoke, which can reach high altitudes due to fire plume buoyancy. The RRFS-SD model parameterizes plume injection based on fire heat flux estimates and simulated meteorological conditions. In this presentation, I will discuss how the model simulates smoke transport driven by processes such as advection, turbulent mixing, and plume rise.

I will also present results from an evaluation of the horizontal and vertical distribution of smoke using observations from the JPSS, TEMPO, and EarthCARE satellites. Aerosol optical depth, aerosol layer height, and aerosol backscatter profiles derived from satellite retrievals are used to evaluate smoke transport in the model across North America. An analysis of several cases of long-range smoke transport from Canadian wildfires into the United States during summer 2025 will also be presented using both model output and satellite data.

Representing Organic Soil Fire Emissions in the RRFS-Smoke Model: Implementation and Performance Assessment

J Romero-Alvarez*, R Ahmadov, H Li, J Schnell, E James, S Bhimireddy, and G Grell

During the record-breaking 2023 Canadian wildfire season, extreme drought exposed normally saturated peatlands and deep organic soils to high-severity combustion, revealing underestimations in the top-down emissions inventories used by flagship operational systems, such as HRRR-Smoke and the upcoming RRFS-Smoke. While these systems captured smoke transport pathways well, they underestimated smoke PM2.5 concentrations across North America.
Using high-resolution geospatial masks, we isolate regions dominated by organic soils and permanent wetlands and then update both the emission factors (EFs) and the biomass combustion coefficient (β), which is used to convert Fire Radiative Energy (FRE) into total biomass consumed. We test two β formulations: one based on land-use-dependent values from published literature and another ecosystem-dependent β derived by calibrating NOAA's Regional ABI and VIIRS fire Emission product (RAVE) FRP against total biomass consumption estimates from the bottom-up wildfire emissions inventory Global Fire Emissions Database (GFED5.1).

Initial RRFS-Smoke simulations using these adjusted coefficients show a substantial increase in smoke emission estimates for fires influenced by organic-soil combustion. This update significantly reduces the model biases observed during the 2023 season and brings model output into closer agreement with surface observations. These improvements address a low bias likely resulting from limitations in satellite-derived Fire Radiative Power (FRP) to detect weak smoldering signals from vulnerable carbon reservoirs under canopy, as well as retrieval uncertainty under the dense smoke conditions in 2023.

Performance Assessment of Land Surface Fields in NOAA's Seasonal Forecast System (SFS)

Weizhong Zheng*, Michael Barlage, Zhichang Guo and Fanglin Yang

As a major Earth system component of the NOAA Unified Forecast System (UFS), the community Noah land surface model with multi-parameterization options (Noah-MP) has been coupled to the Seasonal Forecast System version 1 (SFSv1) to simulate exchanges of energy, water, and momentum between the land surface and the atmosphere. Improvements in land surface physics and land–atmosphere interactions are expected to reduce systematic biases in SFSv1 and enhance forecast skill.

In this study, eight one-year simulations were conducted using SFSv1 at C192 resolution, initialized at 00Z on May 1 and November 1 for each year from 2021 to 2024. Land surface variables, including 2-m air temperature and humidity, snow cover fraction, snow depth, and soil moisture, as well as surface water and energy budget, were evaluated against the reanalysis datasets and the ground measurements. Forecast biases and errors in these surface and near-surface fields are analyzed and discussed.

Community tools and workflows for hierarchical Noah-MP modeling and optimization to improve UFS hydrometeorological prediction

Andrew Newman*, Yifan Cheng, Andrew Bennett, Ronnie Abolafia-Rosenzweig, Kathryn Newman, Arezoo RafieeiNasab, Tzu-Shun Lin, Cenlin He, Ethan Gutmann, Andrew W Wood and Michael Barlage

As Noah-MP will be the land surface model in next generation UFS applications, it is essential to ensure it can simulate high-fidelity land states and fluxes in the coupled modeling system. To support that effort we have been developing community tools and workflows to support Noah-MP sensitivity experiments and optimization within the UFS Hierarchical System Development (HSD) framework. HSD is useful in coupled model optimization because uncoupled land model optimization is performed using data atmospheres, which ignores the coupled land-atmosphere feedback and may compromise the usability of optimized land parameters in coupled systems. Specifically, we are focused on optimizing the parameters in Noah-MP to better simulate hydrometeorology as well as terrestrial hydrology in the coupled system, with an emphasis on precipitation, snow, soil moisture, and streamflow.

We have developed offline and SCM-based optimization workflows that will be released to the Noah-MP and UFS communities. We have also developed a new ERA5 based SCM forcing generator that is also nearing community release. Finally, we have performed extensive testing of the SCM for seasonal and longer simulations and are developing documentation of our experience to support further SCM use and development. We present an overview of these capabilities as well as an initial testbed where we have configured Noah-MP at nearly all CONUS AmeriFlux sites. We have benchmarked Noah-MP with default parameters, identified sensitive parameters, and optimized Noah-MP for both the offline and SCM versions. We will present initial results quantifying biases in energy and hydrologic fluxes and process oriented metrics.

Influence of the Laurentian Great Lakes on Summer Time Precipitation Revealed by Isotopic Enrichment Measurements in Michigan

Ananyo Bhattacharya*, Adriana Bailey and Christiane Jablonowski

The Laurentian Great Lakes account for one of the largest freshwater lake sources in the world, and serve as a critical transboundary water resource for the USA and Canada. Representation of lake water hydrodynamics and thermodynamics is known to influence the regional climatology and enhance mesoscale processes. Isotope ratios of hydrogen and oxygen in water vapor measured at the University of Michigan Biological Station (UMBS) are analyzed to understand the roles of Great Lakes evaporation and regional convective processes in summertime humidity patterns. We used the NOAA Unified Forecast System (UFS) Short Range Weather Application (SRW) to evaluate precipitation outcomes over a limited-area domain at 3-km horizontal resolution. We investigated four synoptic-scale events associated with strong isotopic anomalies at UMBS. NOAA UFS is configured to near-operational Rapid Refresh Forecast System (RRFSv1) settings and coupled with the Great Lakes FVCOM lake model in unidirectional lake-to-atmosphere coupling. Short-term forecast over a two to three-day period provides a systematic assessment of lake enhancement in summer-time precipitation.

Cross Cutting NWP Topics

NOAA's Precipitation Prediction Grand Challenge (PPGC) Initiative in Supportive of UFS Improvement and Applications

David Novak (NOAA/NCEP Weather Prediction Center), John Ten Hoeve (NOAA/OAR Weather Program Office) and Jin Huang (NOAA/OAR Climate Program Office)

Despite advances in numerical modeling and observations, precipitation forecast skill—particularly for high-impact events—has improved only modestly in recent decades, with persistent errors in the timing, location, and intensity of extremes. Addressing these gaps requires coordinated innovation across model physics, data assimilation, observations, verification, and research-to-operations pathways. In response, NOAA established the Precipitation Prediction Grand Challenge (PPGC) Initiative to accelerate improvements in precipitation prediction across weather and S2S timescales.

PPGC is a cross-line office and end-to-end effort from observations, research, modeling, to products and services. PPGC aligns strongly with the collaborative goals of the Unified Forecast System (UFS) and the Earth Prediction Innovation Center (EPIC). Current activities are supported through NOAA research grants and supplemental funds. The recent $21M FY25 supplemental fund is focused on improving probabilistic warm season precipitation on weather timescales through enhanced high-resolution of rainfall estimates from observations, better understanding of precipitation processes through field studies, advances in data assimilation and modeling, development of tools and products from observations and model guidance, improved forecast communications.

This presentation will highlight how PPGC advances precipitation forecasting capabilities relevant to the UFS, including model improvement and validation within the UFS framework, data assimilation (DA) with AI/ML approaches; improvement of physics parameterizations across multiple UFS applications; verification for precipitation extremes. Opportunities for stronger collaboration with the broader UFS/EPIC community will also be discussed, including shared model development, data assimilation, verification, community engagement, and accelerated research-to-operations-to-research (R2O2R) pathways.

An Overview of Verification and Evaluation Activities at the NOAA Office of Modeling and Development

Jason Levit*

To meet the growing demand for reliable environmental forecasting, the NOAA Office of Modeling and Development (OMD) is modernizing its verification and evaluation (V&E) frameworks. This presentation outlines OMD's strategic V&E roadmap, highlighting the recent merger of the Environmental Modeling Center (EMC), Meteorological Development Laboratory (MDL), and the Analyze, Forecast, and Support (AFS)-11 verification teams to unify product assessment across the NWS enterprise.

A key technical modernization milestone is the March 2026 release of the EMC Verification System version 2.0 (EVS 2.0), which expands comprehensive verification capabilities for OMD models, utilizing the METplus (Model Evaluation Toolkit) software ecosystem. The unified V&E team is also advancing National Blend of Models (NBM) verification and National Digital Forecast Database (NDFD) comparisons. Furthermore, to accelerate model transitions to operations, OMD is adopting continuous DevOps and Agile life cycles for rapid, iterative evaluation. As part of the DevOps strategy, OMD is developing the use of Summary Assessment Metrics (SAMs) that aggregate spatially and temporally weighted Primary Assessment Metrics (PAMs) such as RMSE and Bias into a single summary score, to understand differences between current and future model performance. Finally, the presentation will highlight thoughts on using novel verification techniques for the exploration of the development of AI/ML "agentic" systems. These AI agents aim to automate verification tasks. Ultimately, unifying the EMC, MDL and AFS-11 teams, modernizing architectures, and leveraging some emerging AI tools will significantly accelerate the delivery of rigorously validated guidance to forecasters and the public.

Unifying Composition in the UFS through CATChem and CECE

Barry Baker*, Patrick Campbell, Rebecca Schwantes, Wei Li, Jian He, Maggie Bruckner, Zachary Moon, Quazi Rasool, Ben Koziol, Colin Harkins, Li Zhang, Ravan Ahmadov, Youhua Tang, Fanglin Yang and Brian McDonald

The Unified Forecast System (UFS) is evolving into a comprehensive modeling system spanning sub-seasonal to seasonal (S2S), global, regional, and air quality applications. A critical challenge in this transition is the historically fragmented nature of atmospheric chemistry representations, which often vary significantly across application domains, leading to inconsistencies, increased maintenance overhead, and barriers to scientific innovation. To address this, we propose and present the integration of two key community-developed components: the Configurable Atmospheric Chemistry (CATChem) package and the Community Emissions Computing Engine (CECE), into a unified chemistry framework within the UFS.

CATChem serves as the core chemical solver, designed to provide a single, consistent interface for managing atmospheric transformations with gas and aerosol species. Its modular design allows users to select from a library of chemical mechanisms—ranging from simple, computationally efficient schemes necessary for long-range global forecasts, to complex, detailed mechanisms required for regional air quality predictions—without altering the underlying transport model structure. This configurability ensures scientific rigor while maintaining the necessary computational efficiency across the UFS's diverse applications.

Complementing CATChem, CECE standardizes the crucial process of emissions handling. CECE acts as a centralized engine for ingesting and processing all relevant emission types (anthropogenic, biogenic, and fire). By decoupling emissions pre-processing from the model physics, CECE ensures that emissions data are consistently handled, interpolated, and coupled to the UFS, facilitating the rapid incorporation of new emission inventories (e.g., global fire products or evolving national emissions data) and improving reproducibility across different UFS runs.

The unification of chemistry within the UFS through the CATChem and CECE framework represents a significant step towards a truly cohesive and sustainable forecasting system. This talk will detail the technical implementation strategy, demonstrate the flexibility of CATChem's mechanism switching, and present initial validation results from UFS testbed simulations. We will specifically highlight how this unified structure improves consistency between UFS components (e.g., seamless transition of chemical states between global and nested regional domains) and significantly reduces the Research-to-Operations (R2O) timeline for novel atmospheric chemistry parameterizations. This work provides a foundation for the next generation of UFS air quality and coupled chemistry-climate applications.

Application targeted configuration of the ocean component in the UFS framework towards seamless predictions

Santha Akella*, Dmitry Dukhovskoy, Neil Barton, Fanglin Yang and Vijay Tallapragada

The Unified Forecasting System (UFS) framework, using the Community Mediator for Earth Prediction Systems (CMEPS) as a coupler, supports a wide range of applications. These range from high-resolution, short-term forecasts (e.g., Hurricane, Atmospheric Rivers (AR) at <= 10 km) to lower-resolution, long-term Seasonal Forecast System (SFS) (at >= 50 km). This application spectrum necessitates targeted component-specific tuning and configuration.

The focus of this presentation is on addressing the following issues for the coupled model and ocean-sea-ice components: (i) Air-Sea-Ice Coupling: Enhancing air-sea coupling to reduce model biases across various time- and length-scales (from 5-day to 1-year forecasts); (ii) Coupled Processes: Identifying and better representing coupled processes to improve the assimilation of observations; (iii) Computational Performance: Exploring options to improve performance (both throughput and I/O) for application-targeted resource allocation and optimization.

Currently, only two configurations exist for the MOM6 ocean component: one for the Global Forecasting System (GFS)/SFS (0.25-deg/75-layer) and one for HAFS/AR (0.08-deg/55-levels). We will present alternatives for parameterizations and resolutions to address the issues above and ensure processes relevant to these applications are better represented. Time permitting, we will also cover configuration changes/tuning for high-resolution applications like the Real-Time Ocean Forecasting System (RTOFS; 0.08-deg/41-layer). This presentation is of relevance to all UFS stakeholders, developers, users, and model developers- specifically coupled/earth system modelers.

TRITON-C Development Update: A Consistent Implicit Framework with On-the-Fly V&V and Coupling strategy

Ali Abdolali*, Aron Roland, David Honegger, Tyler Hesser and Jane M. Smith

To overcome the inherent computational slowness of spectral wave models within coupled frameworks, we introduce TRITON-C, a next-generation model built on a mathematically consistent implicit foundation. Developed natively in modern C++, TRITON-C is engineered for the modern HPC era and future GPU acceleration. By utilizing advanced coding practices, node-major structure-of-arrays memory layouts, and minimal library dependencies, requiring only MPI and a standard compiler, the model achieves a transformative leap in efficiency. Benchmarks demonstrate that TRITON-C is two orders of magnitude faster than conventional curvilinear WAVEWATCH III (WW3) and one order of magnitude faster than our team's previous unstructured implementations. This architecture defines a scalable roadmap for coupled Earth system modeling, enabling high-fidelity integration with atmospheric, sea ice, and hydrodynamic models such as, CICE, AdH and SCHISM. The coupling to hydrodynamic models employs classical 2D Longuet-Higgins radiation stress and 3D vortex-force representations for rigorous wave-averaged momentum exchange.

A critical advancement in this development update is the implementation of on-the-fly validation and verification (V&V) capabilities. This system serves as a powerful diagnostic tool and a major leap forward for data assimilation, capable of ingesting millions of in-situ and remote sensing data points with less than 1% computational overhead. Validated against 2023–2024 hindcasts using SWOT mission, CFOSAT/SWIM, and high-density drifter data, TRITON-C accurately resolves processes across wind-sea, swell, and infragravity bands. Collectively, this work establishes the conditions for the next generation of seamless, scalable global-to-coastal simulations.

Advancing Land-Atmosphere Evaluation of NOAA UFS Models Through Community Datasets and Software Enhancements

Scott Miller, Andrew Newman*, Anna Glodzik, Dan Adriaansen, Cheng-Hsuan (Sarah) Lu, John Halley Gotway, Kathryn Newman, George McCabe, Jason Covert and Michael Ek

Here we present key outcomes from a NOAA Joint Technology Transfer Initiative (JTTI) collaborative project between SUNY Albany and NSF NCAR that aims to advance NOAA UFS land and land-atmosphere coupling evaluation. We have focused on 1) enhancing the availability of high-quality community land and near-surface data for model evaluation from the New York State Mesonet (NYSM), and 2) increasing community software capabilities to ingest land related data and provide relevant forecast evaluations in the enhanced Model Evaluation Tools (METplus) software.

First, for observations, we achieved a milestone success, uploading three years of 30-minute eddy covariance flux data from 18 NYSM Flux Network eddy covariance flux sites to the AmeriFlux Data Repository. This effort has included extensive assessment of the NYSM flux data and quality control methodological developments including energy balance closure assessment and mitigation, as well as analysis of footprint heterogeneity impacts on energy budget closure. Second, we have included new land related capabilities in the community METplus software, which is also the pathway to NOAA EMC operational verification. These METplus additions address multiple verification needs identified in the 2021 NOAA and NOAA Developmental Testbed Center community metrics workshop. For example, we have added the ability to ingest soil temperature and moisture data from multiple networks, NYSM flux and standard data, as well as satellite soil moisture data. We have also included process diagnostics such as the convective triggering potential (CTP) and humidity index (HI), the Terrestrial Coupling Index (TCI), and information theory metrics. Finally, we have developed METplus use cases to provide users with example configurations.

Portable High-Performance TEMPO Microphysics — A Unified CPU/GPU Codebase via Hybrid Loop Restructuring

Pranay Reddy Kommera* and David Appelhans, NVIDIA Corporation

TEMPO microphysics, like many physics parameterizations, uses a column-based design in which the outer loops iterate over the horizontal grid points while inner loops iterate through individual vertical columns. This structure is efficient for CPU implementation due to effective cache utilization but limits the GPU parallelism, resulting in lower occupancy and underutilized hardware. The alternative plane-based implementation that collapses all the iterations through the horizontal grid points and vertical column into a grouped nested loop— maximizes the GPU throughput but could have a negative impact on the CPU performance. This limitation would require maintaining two separate codebases for efficient performance on respective hardware.

In this talk, we will demonstrate a hybrid-strided loop structure applied to TEMPO microphysics within a single, unified codebase. The approach introduces a configurable stride parameter that switches between the different loop structures. When stride = 1, the implementation reflects column-based loop structure, ideal for CPU implementation. When stride = grid points, it shifts to a plane-based implementation, targeting GPU parallelism. The implementation is parallelized using OpenACC directives and tested on both the CPU and GPU hardware. We will present performance results comparing the hybrid approach against both the column-based and plane-based implementations, and discuss the readability and performance portability trade-offs of this strategy in the column-based physics schemes.

Addressing Core Modeling Infrastructure Barriers in Next-Generation UFS Weather Model Development

Daniel P Sarmiento*, Denise Worthen, Barry Baker, Ligia Bernardet, Robert L Jacob, Rahul Mahajan, Dustin Swales, Vijay Tallapragada, Ann Tsay and Jacob Carley

The Unified Forecast System (UFS) is foundational to developing a community-based, coupled Earth system modeling framework, with the UFS Weather Model (UFS-WM) at its core. Its successful adoption and advancement across the diverse U.S. weather enterprise are crucial for realizing improved prediction capabilities. However, a significant obstacle to widespread UFS-WM community engagement and accelerated development is the heterogeneity and inherent limitations within the core modeling infrastructure; largely due to a complex software stack that relies on multiple, overlapping frameworks, resulting in a "heavy" system that complicates portability and maintenance.

This presentation addresses key infrastructural barriers impeding seamless UFS-WM development, execution, and integration for a broad user base. Specific challenges include: managing diverse and evolving software dependencies, difficulties in efficiently porting the UFS-WM across different High-Performance Computing (HPC) architectures, and insufficient internal tools and standards for modularity and interface definition.

We propose a path forward focusing on the modernization and standardization of the UFS-WM's internal infrastructure through a Lean UFS architecture. Our strategy is to consolidate from the current three frameworks to a single, unified layer, which includes developing robust abstraction layers and standardized APIs, such as Generic Write-Grid Components to simplify I/O across the UFS and improve code portability across diverse HPC architectures. This approach reduces the technical debt and software footprint, facilitating advances in other areas, such as adoption of GPU architectures. Overcoming these core infrastructural barriers is paramount to unlocking the full collaborative potential of the UFS-WM community and rapidly transitioning fundamental research advances into operational prediction systems.

UFS Applications (Global Workflow, SFS, S2S, AR)

NOAA's Seasonal Forecast System Project: Progress, Challenges, and Stakeholder Engagement

Yan Xue*, Steven Simon, Neil Barton, Phil Pegion, Xiaqiong (Kate) Zhou, Wanqiu Wang, Fanglin Yang, Kevin Garrett, Daryl Kleist, Mark Olsen and John Ten Hoeve

Jointly supported by the NWS Modeling Program and the OAR Weather Program Office (WPO), the Seasonal Forecast System (SFS) Project was established in October 2023. The primary goal is to replace the legacy Climate Forecast System version 2 with a modern, fully coupled system based on the Unified Forecast System (UFS). SFS version 1 (SFSv1) will advance NOAA's seasonal forecasting capability and enhance decision-makers across the entire Weather Enterprise.

The project features a refined four-way coupled Earth system model incorporating FV3 (atmosphere), MOM6 (ocean), CICE6 (sea ice), and Noah-MP (land surface). Updated atmospheric physics and dynamics from the Global Forecast System version 17 (GFSv17) improve computational efficiency and simulate large-scale climate variability across subseasonal and seasonal timescales. A key milestone is the March 2026 launch of SFS Beta v1.0 (SFSbeta) in Near-Real-Time (NRT). Generating 12-month predictions and on-the-fly historical reforecasts (1991-2025), SFSbeta consistently outperforms CFSv2 in forecasting U.S. surface air temperature, precipitation, and sea surface temperatures (SST). SFSbeta currently provides experimental guidance for Climate Prediction Center (CPC) operational seasonal outlooks.

Next steps include developing consistent initialization strategies for SFSv1 operational implementation. Agile development will refine model configuration, physics, and initialization (e.g., ERA5/GLORe replay, and SFS 3DVar coupled reanalysis). Continuously updated SFSbeta NRT forecast and reforecast data are publicly available on AWS via NOAA Open Data Dissemination (NODD) to gather crucial stakeholder feedback. This presentation highlights SFS Project milestones since 2023, ongoing developmental and operational challenges, explores community collaboration opportunities, and seeks stakeholder feedback on the project.

Practical Strategies for Initialization of the Seasonal Forecast System

Philip Pegion*, Shan Sun, Neil Barton, Wanqiu Wang, and Fanglin Yang

Due to model errors and small signals at timescales beyond weather, a set of retrospective forecasts (re-forecasts) are necessary to bias-correct and calibrate the raw model output in order to make forecasts useful. Ideally, an operational system would be initialized with a multi-decadal reanalysis created using the identical system. However, producing a 30-to-40-year reanalysis is time-consuming (at least a year) and often requires reduced resolution compared to the operational system. Some existing systems, like CFSv2 and SEAS5, address this by extending the reanalysis in real time, although CFSv2's realtime initialization is not identical to the re-forecast period.

NOAA, however, is transitioning to a single, high-resolution operational analysis for initializing all global applications, from medium-range to seasonal. This change, coupled with an accelerated implementation schedule for a new Seasonal Forecast System (SFS) that needs to become operational before the required reanalysis is complete, necessitates alternative initialization strategies.

We are currently evaluating options such as using existing reanalyses and employing "replay" where the forecast model is nudged toward an existing reanalysis to emulate it, which improves consistency between the initial state and the forecast model. Our work will compare these methods' impact on skill assessment and systematic bias estimation for forecasts initialized with the high-resolution, operational GFS analysis. Additionally, we will investigate replaying the SFS model to the high-resolution analysis as a means to reduce initialization shock and minimize the discrepancy between re-forecasts and real-time forecasts.

Evaluating SFS Beta: Early Skill Relative to CFSv2

Wanqiu Wang*, Yutong Pan, Robert Long and Yanyun Liu

Development of the UFS-based Seasonal Forecast System (SFS) is progressing toward an operational implementation planned for late 2027, replacing the current Climate Forecast System version 2 (CFSv2). As part of this transition, an interim system—SFS Beta—began producing monthly forecasts in March 2026 using an on-the-fly hindcast configuration. SFS Beta output is now routinely incorporated into CPC's seasonal outlook discussions. This presentation provides an initial evaluation of SFS Beta performance based on available 1991–2020 hindcasts for March, April, and May initial conditions. The assessment focuses on key variables relevant to CPC operations, including sea-surface temperatures (with emphasis on the Niño-3.4 index), surface air temperature, precipitation, soil moisture, and related fields. Results are compared against CFSv2 and the ECMWF seasonal forecast systems to highlight areas of improvement, remaining challenges, and implications for future SFS development.

Overview of the Next Global Forecast System GFSv17

Catherine Thomas*, Jessica Meixner and Ruiyu Sun

The Environmental Modeling Center is working towards the next operational implementation of the Global Forecast System, GFSv17, which leverages the Unified Forecast System (UFS) and the UFS community at large. A major goal for GFSv17 is to employ a UFS-based fully coupled atmosphere-land-ocean-ice-wave model. The coupled model will be part of a weakly coupled data assimilation system, with new non-atmospheric data assimilation components for snow, ocean and sea ice based on the Joint Effort for Data assimilation Integration (JEDI) software. In addition, the existing atmospheric model and data assimilation algorithms will have major upgrades, including updates to atmospheric physics, land modeling, an increase of the horizontal atmospheric resolution to 9 km, and the inclusion of multiscale atmospheric data assimilation.

This presentation will provide an overview of the GFSv17 system along with details of the existing component updates and introduction of the new modeling and assimilation components. Evaluation of the performance of pre-implementation retrospective and real-time parallels compared to the operational configuration will also be shared.

Physics Advancements in GFSv17

Ruiyu Sun* and Fanglin Yang

The Global Forecast System version 17 (GFSv17) is a fully coupled Earth system model integrating atmosphere, land, ocean, sea ice, and wave components through the CMEPS mediator. The horizontal resolution is increased to 9 km (C1152) across 127 vertical levels. Data assimilation is modernized by adopting the JEDI framework for non-atmospheric components and introducing new land moisture and temperature assimilation to optimize the model's initial state.

The updates in the physics suite in the GFSv17 is a comprehensive overhaul to better resolve complex sub-grid interactions:
Microphysics & Radiation: The Thompson hybrid two-moment microphysics scheme replaces the legacy GFDL scheme. Radiation is further refined using a new ten-year mean MERRA2 aerosol dataset.

Land Surface Modeling: The system officially migrates to the Noah-MP Land Surface Model. .
Convection & Boundary Layer: Updates are designed to improve CAPE and inversion forecasts and improve the performance of the schemes. Additionally, updates are made to the surface layer parameterizations.

Gravity Wave Drag (GWD): The GWD suite is modernized with updated orographic drag, mountain blocking, and the introduction of turbulent orographic form drag.

Validation has demonstrated significant improvements across numerous areas, such as the 500 hPa height skill scores and precipitation Equitable Threat Scores (ETS). Further details will be presented at the conference.

An NSST Alternative in UFS: Skinsst

Shan Sun* and Rainer Bleck

Using NOAA's coupled Unified Forecast System (UFS; SFSbeta1.1), we evaluated seasonal forecasts by comparing the default Near-surface Sea Surface Temperature (NSST) algorithm against a simpler skin temperature scheme (Skinsst) developed at GSL, as well as a baseline configuration lacking skin temperature parameterization. While the NSST was originally developed for uncoupled atmospheric applications, our alternative scheme directly leverages ocean model feedback. Results indicate that although seasonal SST forecast skill is comparable across all three configurations, the Skinsst scheme yields a diurnal SST cycle more consistent with satellite and ship observations. These findings underscore that accurate skin temperature representation is essential for improving the physical coupling between surface temperature and energy fluxes within coupled atmosphere-ocean models.

The Ocean Coupling Impact on the forecast of Atmospheric Rivers in the U.S. West Coast

Maria Aristizabal Vargas*, Keqin Wu, Murali N. Malasala, Binbin Zhou , Xingren Wu and Vijay Talapragada

The Atmospheric Rivers Analysis and Forecast System (AR-AFS) is an experimental high-resolution regional model developed at the NOAA/Environmental Modeling Center, under the framework of the Unified Forecasting System (UFS). AR-AFS is based on the Finite Volume Cubed-Sphere (FV3) Dynamical Core. As a next step on the development of AR-AFS, the system has been coupled to the Modular Ocean Model version 6 (MOM6) using the existing infrastructure of the Hurricane Analysis and Forecast System (HAFS).

This study's main objective is to evaluate the impact of ocean coupling on the forecast of land falling atmospheric rivers (ARs) in the West Coast of the U.S. For this we ran the coupled and the uncoupled AR-AFS for six major ARs from 2023 to 2026. We found that the ocean coupling produces substantial changes on the sensible and latent heat fluxes, with the mean heat fluxes being larger in the coupled simulations. The higher heat fluxes in the coupled case are reflected as higher values of the integrated water vapor (IWV). The areas with the highest IWV also exhibit higher values of 24 hours integrated precipitation, but these areas are located mostly in the open ocean. As a result, the ocean coupling has a modest impact on improving the precipitation on land. We will continue investigating if the changes in the IWV in the coupled AR-AFS can ultimately bring substantial forecast improvements for the precipitation during AR events in the U.S West coast.

Public Release of NWS Global Workflow

Walter Kolczynski*, Wei Huang, Kris Booker and Priya Pillai

This spring the Earth Prediction Innovation Center (EPIC), in coordination with the National Weather Service Environmental Modeling Center (NWS/EMC), made a public release of NWS's global workflow. Global workflow is the overarching system responsible for orchestrating many of the NWS's global modelling systems, including the flagship global system in the Global Forecast System

(GFS), as well as the Global Ensemble Forecast System (GEFS) and the upcoming Sub-seasonal Forecast System (SFS), a Climate Forecast System replacement, and Global Chemistry and Aerosol Forecast System (GCAFS), which will replace the current GEFS "chem member".

The goal of the public release is to make the workflow more accessible to those outside the National Oceanic and Atmospheric Administration (NOAA) who do not have prior experience using it. To extend accessibility, EPIC added the capability to run the workflow on National Science Foundation National Center for Atmospheric Research's Derecho and in the Amazon Web Services (AWS) cloud environment. EPIC also restructured and augmented existing system documentation to be more approachable for new users. After release, we will work with EMC to provide support and ensure the new resources maintain functionality as development on the forecast systems continues. We hope this release will facilitate future contributions to the operational systems from the community and a move towards a common platform for global model development.

Here we will discuss the improvements made specifically for this release, the current capabilities and limitations to running outside of NOAA, and ongoing community support as global workflow continues to evolve.

Research to Operations to Research (R2O2R)

An Overview of NOAA Unified Forecast System (UFS) Applications for Operational Forecasting Across Scales

Fanglin Yang*

NOAA, in close collaboration with the broader U.S. weather and climate science community, is advancing a fully coupled Earth system modeling capability through the Unified Forecast System (UFS). The UFS is a community-based modeling framework designed to unify research and operations across a wide range of temporal and spatial scales. It supports applications from sub-hourly, convection-allowing forecasts to subseasonal and seasonal predictions and is central to NOAA's strategy for improving environmental prediction. This presentation will provide an update on several UFS applications that are in advanced stages of development and are being prepared for implementation in support of National Weather Service (NWS) operations. These include the Global Forecast System (GFS) version 17, Global Ensemble Forecast System (GEFS) version 13, Hurricane Analysis and Forecast System (HAFS) version 2.2, Global Chemistry and Aerosol Forecast System (GCAFS) version 1, Air Quality Model (AQM) version 8, Seasonal Forecast System (SFS) version 1, and Real-Time Ocean Forecast System (RTOFS) version 3. The presentation will highlight major development activities, achievements, improvements in forecast skills, remaining challenges, and future directions. This work directly supports the NWS transformation initiative by advancing a unified modeling framework to deliver more accurate, efficient, and societally impactful environmental predictions.

An Efficient and Sustainable Development and Implementation Lifecycle for the NWS Production Suite

Jacob R. Carley*, Steven Earle, Rahul Mahajan, Daniel Holdaway, Jason J. Levit, Vijay Tallapragada, Geoffrey S. Manikin, Scott Scallion, and Arthur Taylor

Together the Environmental Modeling Center, the Meteorological Development Laboratory, and NWS Central Operations are responsible for developing, maintaining, and implementing the complex suite of analysis and prediction systems that form much of the foundation of our nation's weather forecasts. However, the increasing complexity of these systems, coupled with a growing reliance on community-developed scientific software, has exposed the limitations of our traditional development and implementation methods. Our current "waterfall" approach, characterized by infrequent, high-stakes "big bang" implementations, is becoming unsustainable.

We propose a fundamental modernization of our development-to-operations lifecycle, inspired by the proven principles of DevOps and Agile methodologies. Our vision is to transition from our current high-risk, multi-year release schedule to a predictable and sustainable rhythm of smaller, more frequent upgrades. This will be achieved by restructuring our processes, investing in modern tools, and fostering a culture of collaboration and continuous improvement. In this talk we will present the basic premise for this endeavor.

Enhancing High-Impact Weather Decision-Making through the Warn-on-Forecast System (WoFS)

Youngsun Jung*, Kevin Garrett, Albert Rabassa, Chad Gravelle, Kenneth Sperow, Jianping Huang (NWS), Patrick Burke, Pam Heinselman, Adam Clark, Louis Wicker (NSSL), Joshua Martin, Kent Knopfmeier, Brian Matilla, Thomas Jones, Nustrat Yussouf and Miranda Silcott (CIWRO)

The National Weather Service's (NWS) warning and pre-warning decision-making has relied heavily on remote-sensing observations, which constrain achieving longer lead times, higher confidence, and better communication of uncertainty. To improve severe weather warning lead time and enhance community-centric warning services, a critical gap was identified: the need for frequently updated probabilistic hazard information and associated uncertainties available from a convection-allowing ensemble prediction system. The National Severe Storms Laboratory (NSSL) developed the Warn-on-Forecast System (WoFS) to bridge this gap on watch to warning time scales.

WoFS leverages cutting-edge technology on the Microsoft Azure cloud platform to provide graphical probabilistic forecast guidance for up to 6 hours, delivered within 15-40 minutes. These forecasts are provided at 3-km grid spacing over 900×900 km^2 domains. Since its demonstration began in the NWS in October 2024, WoFS has received strong positive feedback across numerous events spanning a spectrum of hazardous weather. During this demonstration period, comprehensive requirements for transitioning to operations have been developed and documented, including the initial demonstration protocol (or concept of operations), cost model and financial analysis, system description document, data footprints, and catalog of challenges.

Running in targeted fashion on 91 days as of February 2026, WoFS provided forecasters with reliable probabilistic guidance for high-impact weather, including severe weather, landfalling hurricanes, excessive rainfall, and fire weather. In early 2026, the protocol was expanded to support Department of Homeland Security designated Special Event Assessment Rating 1 (SEAR 1) events, such as the Super Bowl and the 2026 FIFA World Cup.

In this presentation, we will provide the current status and key achievements of the demonstration project, share lessons learned that will inform the transition to operations, and demonstrate how WoFS will support the warning process to ultimately enhance community preparedness and safety.

Enhancing UFS Air Quality Predictions with Aircraft-Derived Regional Emission Adjustments

Youhua Tang*, Patrick Campbell, Wei Li, Barry Baker, Kai Wang, Jianping Huang, Daniel Tong, Xinrong Ren, Eleanor Waxman, Ken Aikin, Andrew Rollins, John Crounse, Paul Wennberg, Ryan Bennett, Ilana Pollack, Emily Lill, Matthew Coggon, Kelvin Bates, Amy Sullivan, Mike Robinson, Carsten Warneke, Lu Xu, Kristen Zuraski, Jeff Peischl, Chris Jernigan, Wyndom Chace, Steven S. Brown, Caroline Womack, Gordon Novak, Patrick Veres, Chelsea Stockwell, Ann Middlebrook, Rebecca Schwantes, Brian McDonald, Charles Gatebe and Youngsun Jung

In this study, we evaluate the NOAA Unified Forecast System–Air Quality Model (UFS-AQM) using aircraft observations from two field campaigns: the Atmospheric Emissions and Reactions Observed from Megacities to Marine Areas (AEROMMA) conducted in 2023, and the Baltimore Air Quality and Marcellus Survey (BAQMS) conducted in 2025. The UFS-AQM simulations are driven by the National Emission Inventory (NEI) for base years 2016, 2019, and 2022.

The model shows some systematic biases, including underprediction of CO and ethane across the CONUS domain, and overprediction of ethene and aromatics in regions east of the Rocky Mountains. Similar biases are also evident when comparing model results with AQS surface observations for these chemical species. To address these discrepancies, we derive specified emission adjustment factors across six regions: Northern California, Southern California, the San Joaquin Valley, the Rocky Mountain region, the U.S. North Central region, and the Northeast, assuming a linear response.

Applying these regionally derived adjustments based on aircraft observations leads to significant improvements not only in aircraft-based evaluations but also in surface predictions compared with AQS observations. In addition, adjustments to primary emitted species result in improved simulation of secondary pollutants, such as ozone.

From Forecast to Fireline: Advancing UFS Wind Prediction for Wildland Fire Operations

Samantha J. Kramer*, Benjamin Kirtman and Natalie Perlin

Wildfires have become increasingly destructive in recent decades, and the forecasts that guide response, resource allocation, and risk communication must keep pace. Subseasonal forecast products from the Unified Forecast System (UFS) Mid-Range offer extended fire weather guidance through 35–45 days, but their operational utility for wildland fire applications remained largely untested — a gap our cross-sector team set out to address through a research-to-operations-to-research framework.

Evaluation of UFS Mid-Range forecasts revealed promising skill for key fire weather indicators, particularly vapor pressure deficit (VPD) and temperature. Wind, however, presented a persistent challenge: forecasts across all lead weeks exhibited low-magnitude bias and limited variability compared to observations across the Continental U.S. — a critical shortcoming given wind's central role in fire behavior and firefighter safety. To address this gap, we systematically tested multiple statistical and machine learning downscaling approaches, none of which produced sufficient improvement to the wind forecast distribution to meet operational needs.
This finding redirected our focus to dynamical downscaling using the Earth Prediction Innovation Center (EPIC) Short Range Weather Application (SRW App), where physics-based approaches offer a more promising path forward. Preliminary results from SRW App downscaling of wind speed show meaningful improvement when compared to observations, reanalysis, and native grid UFS data. Critically, these results have generated genuine enthusiasm from the operational fire weather community — a signal that this research is on the right track to deliver tools forecasters will trust and use. Next steps focus on expanding case study testing and addressing remaining deficiencies in the UFS Mid-Range wind distribution to support a full transition to operations.

UPSurgeML: Real-Time Framework for Unstructured-Mesh Model Component for Hurricane Surge Simulation

Soroosh Mani*, Fariborz Daneshvar, William J. Pringle, Andrew B. Penny, Allison Brannan, Panagiotis Velissariou, Laura Alaka, Saeed Moghimi and Cody Fritz

Office of Coast Survey (OCS) in collaboration with National Hurricane Center (NHC) has been implementing a real-time framework for testing SCHISM-based deterministic and probabilistic surge simulations during tropical cyclones. The probabilistic workflow, now called Unstructured Probabilistic Surge with Machine Learning (UPSurgeML), will serve as experimental guidance designed to be evaluated in real-time during active events. UPSurgeML is UFS-enabled, which keeps the door open for coupling additional physics such as wave and hydrology. The workflow is developed based on two currently operational workflows; the structure of the code closely mimics STOFS models to ensure NOAA Center for Operations (NCO) compliance and the on-demand auto-trigger mechanism is derived from the NHC's P-Surge workflow. One of the largest obstacles to operationalizing the workflow has been package dependencies. This issue was resolved through the following three approaches:

- Using the existing STOFS workflow components
- The multi-platform nature of P-Surge workflow
- Rewriting segments of the code (dependency and logic)

Since the main objective was augmenting an existing model, the outputs had to be standardized; the experimental system's outputs were either images or largely unused array dump files.

The operational time restrictions are also important to adhere to given the tight scheduling of NHC simulations during the hurricane season. The real-time workflow was first developed with relaxed timing and further optimizations are being worked on.

The UPSurgeML system was tested in deterministic mode in 2025 and after an evaluation using historical storms later this year will be evaluated in real-time during the 2026 hurricane season.

Near-Real-Time AI Weather Forecasting with Nested-EAGLE: A Framework for Running EAGLE Models in Operations

Mariah Pope*, D. Alex Burrows, Anil Kumar, Jacob Carley, Paul Madden, Emily Carpenter, Sergey Frolov, Tim Smith and Keven Blackman

The transition of artificial intelligence (AI)-based weather forecast systems from research to operational environments requires robust, automated, and scalable workflows capable of delivering timely and reliable forecasts. To address this need, the Earth Prediction Innovation Center (EPIC) is developing a near-real-time forecasting capability within the Experimental AI Global and Limited-area Ensemble (EAGLE) forecast system framework, implemented on the Nested-EAGLE model within the EAGLE repository. Near-real-time EAGLE enables continuous execution of AI-based models approximating operational capabilities by integrating data acquisition, model inference, post-processing, and verification within a unified workflow. Designed for scalability and reliability, it will support deterministic and ensemble configurations across global and regional domains using EPIC computing and data infrastructure. We demonstrate initial capabilities with the Nested-EAGLE model on both Ursa and AzureML.

Near-real-time EAGLE also provides integrated diagnostic capabilities that generate standardized verification outputs, enabling continuous monitoring of forecast behavior across lead times and regions. These products can be integrated into EPIC-compatible dashboards for real-time awareness. This presentation will illustrate how near-real-time EAGLE provides a practical pathway for evaluating and advancing AI/ML forecasting systems toward operational deployment.

Community Engagement

Introducing the UFS to K12 Students: Strategies and Opportunities

Dr. Tracy Fanara, Inspector Planet / Project ENKI and Alison Gregory, UCAR CPAESS

The future of Earth system prediction depends not only on advances in modeling systems such as the Unified Forecast System (UFS), but also on cultivating a diverse, technically skilled workforce prepared to contribute to these systems. However, integrating Earth system modeling and coding into K–12 education remains a persistent challenge due to limited teacher capacity, curriculum constraints, and the perceived complexity of numerical models.

This project presents a collaborative, educator-centered approach to introducing grades 5–12 students to Earth system modeling concepts through accessible coding frameworks. Rather than imposing additional burden on educators, we co-develop curriculum with teachers through workshops, pilot programs, and iterative feedback cycles. This ensures alignment with classroom realities while building teacher confidence in computational instruction.

Our approach emphasizes conceptual understanding before technical complexity. Students engage in scaffolded learning modules that combine hands-on experiments, simplified coding exercises, and gamified simulations that mirror the logic of Earth system models—without requiring full operational model execution. Supporting materials include tutorial videos, workbooks, and modular lesson plans adaptable for classrooms, afterschool programs, and STEM camps.

To further bridge education and operations, we propose a student competition integrated into the workshop environment, where participants present model-inspired solutions evaluated by practicing scientists. This creates a direct pathway between learners and the research community while reinforcing real-world relevance.

By leveraging existing infrastructure and embedding educator feedback into every stage of development, this initiative lowers barriers to entry for both teachers and students. The result is a scalable, community-driven framework that builds computational literacy, fosters early engagement with Earth system science, and strengthens the pipeline into operational modeling ecosystems.

Increasing UFS Student Engagement Through Lesson Plans

Jenna Holbrook*

As machine learning based weather models continue to improve and become more widely implemented, collaboration between federal and academic sectors is key. Particularly, working to increase awareness and participation from students will be vital to guiding a new generation of innovators. The role of the Unified Forecast System (UFS) student ambassador is to provide a student's perspective on improving the engagement of students and researchers in the UFS modeling community and expanding current outreach initiatives, such as the UIFCW, social media, and lesson plans.

This project involves creating new lesson plans to support online tutorials for project EAGLE (Experimental AI Global and Limited-area Ensemble forecast system) and the UFS. These lesson plans will enhance student's understanding of the models while also allowing them to gain hands-on experience with running them. Utilizing past ambassador recommendations, an updated Student Engagement Plan is also being developed to better reflect current student understanding and engagement activities, including these new lesson plans. While running existing tutorials, there have been some challenges due to a difference in level of coding skills and familiarity with different programming languages. These initial experiences may be consistent with new student users, and the lesson plans will be designed to address some of these challenges. Some of the code within these tutorials also needs to be brought up to date to ensure the latest version of these tutorials are aligned with current model releases.

Bridging Academia and Operations: A WINGS Fellow's Experience

Kaitlin Pereira*

The research-to-operations pipeline depends on sustained connections between academic researchers and operational centers, yet these connections can be difficult to build for early-career students. The WINGS fellowship is structured to close this gap by providing graduate students with mentors who offer feedback and insights that shape both the scientific framing and the operational relevance of their work. In my case, this took the form of distributed mentorship across a primary advisor at CSM, a NOAA/NCEP mentor, and a NOAA/EMC mentor. What began as validating the winter skill of an operational wave model has matured into a contribution that directly informs UFS-Wave development enabled by bidirectional feedback between research and operations that an academic-only path would not provide.

UFS-Chem modeling capability - opportunity for NOAA and external community collaboration and engagement

Monika Kopacz (NOAA OAR), Brian McDonald (NOAA OAR)*, Barry Baker (NOAA NWS), Ligia Bernardet (NOAA OAR), Lori Bruhwiler (NOAA OAR), Maoyi Huang (NOAA OAR), Fiona Horsfall (NOAA OAR), Shobha Kondragunta (NOAA NESDIS) and Vaishali Naik (NOAA OAR)

NOAA colleagues from across OAR developed the first version of a UFS-compatible atmospheric composition modeling system. In May 2026, NOAA's atmospheric composition modeling community as a whole met to discuss and plan further development of UFS-Chem, including applications to air quality forecasting for short range weather (fire, dust, pollen, ozone, particulate matter), hindcast applications for air quality management (support for air quality management plans and exceptional event analysis), S2S applications (composition-climate interactions), and applications at decadal time scales - from regional to Earth System scales. In addition to discussion of further model development, the May summit focused on engaging the broader community in developing and using the model. This presentation would summarize the discussions from the summit and next steps for community engagement.

Measuring our Community: Social Science Metrics for the Unified Forecast System

Alison Gregory*

As previous assessments of the UFS community have shown, the group of UFS stakeholders and users make up a complex community of practice and therefore require thoughtful strategic assessments to measure success within this community. This paper suggests UFS Community success metrics be interpreted in four distinct ways: Network, Engagement, Strategy, Technical (NEST): (1)Network–Social Network theory is a valuable way to interpret team relationships and community collaboration. Network metrics such as centrality, weighted ties, and network size can help determine if a group of actors is collaborating efficiently. (2) Engagement–Success can be measured by public awareness of the UFS and its associated applications. This can include an assessment of webpage content, social media content, public-facing interviews, webinars, and other mainstream avenues of engagement. (3) Strategic–Success can be measured by the varied representation of scientists, industry partners, and stakeholders involved in UFS committees, plans, and meetings. This includes analysis of UIFCW activity, leadership participation, and university partnerships. (4) Technical— Success can be measured by quantifying and evaluating the people contributing to the UFS codebase, the total number of models used in earth system prediction, and the number of applications running UFS compliant code. This assessment can also include the number of peer-reviewed articles and journals on UFS related contributions and Github commits.

Inherent in all of these metrics is a robust social science methodology and application of social theory. These success metrics are intended to be a guiding framework that grows and develops as the UFS continues to improve and evolve over time. Key tools for understanding this community will continue to be attendance and surveys related to the UFS Annual Meeting (UIFCW), work completed by the UFS Working Groups, and ongoing survey assessments.

Artificial Intelligence (AI) for Numerical Weather Prediction (NWP)

EAGLE global weather forecast development

Jun Wang*, L. Cui, J. Liu, B. Cui, X. Luo, W. Li, B. Fu, Z. Zhang, L. Zhu, J. Peng, B. Liu, Sergey Frolov, M. Pope, K. Blackman and J. Carley

Over the past several years, machine learning has fundamentally transformed weather and climate forecasting. Data-driven Machine Learning Weather Prediction (MLWP) models now demonstrate a level of forecast skill and computational efficiency that traditional Numerical Weather Prediction (NWP) models struggle to match. In December 2025, NCEP operationalized its first AI/ML models: AIGFS (deterministic), AIGEFS (ensemble), and the hybrid global ensemble system (HGEFS). Both AIGFS and AIGEFS systems leverage Google DeepMind's GraphCast architecture and were trained on NCEP's Global Data Assimilation System (GDAS) and supplementary datasets, utilizing operational GFS and GEFS initial conditions for input. The AIGFS and AIGEFS models serve as complements to the operational GFS and GEFS forecasts. Development of these applications was enabled through the EAGLE framework.

In this presentation, we provide an overview of the continuous development for the global EAGLE forecast systems, which support development of the AIGFS, AIGEFS and HGEFS. This includes 1) recent EAGLE solo (i.e., AIGFS) updates with a spherical harmonic loss function and wind component decoupling to improve the hurricane intensity forecasts and to resolve the blurriness issue in long lead times; 2) the EAGLE ensemble (i.e, AIGEFS) application transitioned to the Anemoi framework, and model training was conducted using NCEP GDAS data using the CRPS-based loss function in Anemoi to resolve the under-dispersion and blurriness issue in AIGEFS v1; 3) new quantile probability products developed for AIGEFS and HGEFS to provide more weather forecast guidance to the community; 4) Transition Global EAGLE solo to the Anemoi framework. Future work to further improve the forecast performance for the EAGLE global forecast systems will also be discussed.

Assessing the Skill and Sensitivity of AIWP Models During the 2026 NOAA HWT Spring Forecasting Experiment

David Harrison*, Tim Supinie, Israel Jirak and Adam Clark

AI-generated weather prediction (AIWP) models are rapidly seeing adoption within operational settings across the weather enterprise. These fully AI-based emulators provide global and regional weather forecasts that have been shown to statistically perform equal to or better than that of traditional numerical weather prediction (NWP) while requiring significantly less computational time and resources. To better understand how these models perform in operational environments, a number of global and regional AIWP models and ensembles were evaluated as part of the 2026 Hazardous Weather Testbed Spring Forecasting Experiment (SFE). AIWP models and ensembles assessed during the 5-week SFE include the AIFGS, Pangu-Weather, GraphCast, AIFS, Aurora, EAGLE Nest, AIGEFS, ECENS, WeatherNext2, and HRRRCast. These AIWP forecasts were directly compared to those of traditional NWP models and evaluated based on observations and model analyses. The SFE additionally tested AIWP models' sensitivity to initial conditions by subjectively rating and comparing global models initialized using GFS and ECMWF initial conditions. This presentation will detail the design of each AIWP experiment and explore the subjective evaluation results to provide a first look at the current state of AIWP for operational meteorological applications.

WoFSCast Performance in the 2026 HWT Spring Forecasting Experiment

Corey Potvin*, Joshua Martin, Miranda Silcott, Jorge Duarte, Adam Clark, Brian Matilla and Patrick Burke

NSSL's WoFSCast model was demonstrated in the 2025 Hazardous Weather Testbed Spring Forecasting Experiment (SFE), where it was rated nearly as highly as the cloud-based WoFS (cb-WoFS) in next-day evaluations of 0–3-h reflectivity forecasts. WoFSCast will reappear in the 2026 SFE (27 Apr – 29 May) as a key component of two activities. In one activity, participants will create mesoscale discussions using real-time forecasts from either cb-WoFS, WoFSCast, or an experimental MPAS-based WoFS. In the second activity, participants will rate the previous day's forecasts from each of the three systems. We hypothesize that recent WoFSCast improvements will shrink its performance gap with (WRF-based) cb-WoFS, especially at 3–6-h lead times when forecast smoothing and domain corner artifacts often became prominent during the 2025 SFE. We will summarize our findings from the 2026 SFE and outline future directions for WoFSCast development.

Designing EAGLE: A Portable and Reproducible Framework for AI/ML Weather Prediction

Paul Madden*, Emily Carpenter, Mariah Pope, Anil Kumar, Kristopher Booker and D. Alex Burrows

For the initial release of NOAA's EAGLE (Experimental AI Global and Limited-area Ensemble) forecast system, we implemented a design that leverages the Unified Workflow Tools (uwtools) system to organize the complex configuration and execution of the many components powering the data prep, model training, inference, verification, and visualization steps in a forecast pipeline.

Using uwtools allows us to configure a full EAGLE configuration by composing fragments describing a base configuration, a specific use case (e.g., nested or global), and platform-specific settings; and to reference common values tying together different pipeline steps without repetition or needing to keep values manually in sync. We wrapped underlying component tools in uwtools drivers to provide a consistent execution interface and allow for rigorous schema checking of component configurations.

We supplemented these benefits with a build system to create the runtime software environments needed to execute pipeline steps, and robust code-quality tooling (formatting, linting, typechecking, unit testing – both for developers and in our continuous-integration system) for Python code and shell scripts.

Overall, this design supports reusing configuration, code, build logic, etc. for many potential use cases and execution platforms, promoting both community engagement and R2O2R innovation sharing.

Next-Gen HRRRCast Ensemble: Enhancements and the Leap into Data Assimilation

Daniel Abdi, CIRES affiliated with NOAA/GSL and Isidora Jankov*, NOAA/GSL

Recent advancements in HRRRCast have led to the transition from Version 2 (V2) to Version 3 (V3), introducing key improvements in both model training and predictive performance. V3 is trained at full 3 km resolution, enabling better representation of fine-scale atmospheric processes. In addition, a sigmoid-based weighting function has been implemented during training to better emphasize extreme events, particularly for composite reflectivity and precipitation. These enhancements have resulted in substantial improvements in forecast skill, especially in capturing high-impact weather features.

Most recently, a score-based data assimilation (DA) approach has been incorporated into the HRRRCast framework. This development represents a promising pathway toward generating dynamically consistent and observation-informed initial conditions for future model iterations. Ongoing work is focused on evaluating the impact of this DA methodology on forecast accuracy and system performance. Recent results will be presented at the workshop.

Towards incorporating data-assimilation in HRRRCast: a score-based approach

Daniel Abdi*, Isidora Jankov, Conor Lewellyn, Nick Silverman, Nick Krall and Chris Miller

HRRRCast v3 is a convection-permitting, AI-based weather prediction system trained on multi-year HRRR and GFS datasets, with improvements in resolution (3 km), variable coverage, and treatment of extreme events. However, with the operational phase-out of HRRR anticipated, HRRRCast's current reliance on HRRR analyses as training targets and implicit background states presents a structural limitation. This transition motivates the need for a self-consistent data assimilation (DA) capability that can ingest observations directly and reduce dependence on external analysis products. While prior developments in HRRRCast have improved forecast fidelity–particularly for convective-scale features–the absence of an internal DA framework now represents a primary bottleneck for operational viability.

To address this, we investigate score-based data assimilation (SDA), formulating DA as a Bayesian inference problem in which a learned generative prior (from HRRR analysis) is combined with observation-driven likelihoods. Our implementation reuses the diffusion-based architecture and employs analytic, primarily linear observation operators (e.g., interpolation and neighborhood-based mappings), enabling efficient evaluation of both operators and their adjoints. Results indicate that SDA can recover physically consistent multivariate relationships, even when only a subset of variables is directly assimilated, and can flexibly incorporate varying observation densities. At the same time, performance depends sensitively on operator design and observation coverage. We will share our experience applying SDA within HRRRCast, including practical challenges, mixed outcomes, and comparisons with traditional approaches.

A Hybrid Physics–Machine Learning Coupled Model for Subseasonal-to-Seasonal (S2S) Prediction

Ufuk Utku Turuncoglu* and Wessel P. Bruinsma

The subseasonal-to-seasonal (S2S) applications aim to fill the gap between weather and climate predictions to provide an extended range of forecasts of the Earth system for a relatively longer time scale compared to short-term weather forecasts (10 days to two weeks in advance). Providing accurate predictions at this time scale is especially challenging due to the involved coupled processes such as air-sea interaction, the effect of soil moisture, sea ice, and their two-way and complex interactions (White et al., 2017). To that end, S2S applications require representation of different earth system model components such as atmosphere, ocean, wave, and sea ice and their coupled interactions in a realistic and accurate way.

Here we present a generic approach that enables the replacement of any component of the Earth system model with its AI/ML-based equivalent in a seamless and efficient manner. The proposed approach is agnostic to any specific modeling application that uses ESMF/NUOPC and promotes an interoperable modeling environment. To demonstrate the flexibility of the newly introduced coupling approach and reveal its capability to perform numerical simulation in the S2S time scale, a new fully coupled multi-component global S2S application is introduced (HESM-S2S; Hybrid Earth System Model for S2S prediction). Then, we conducted a set of numerical experiments by using UFS WM based coupled model configuration as well as the configurations forced with the fine-tuned version of Microsoft's Aurora foundational model. Then, presents the results of the fully coupled S2S simulations that are initialized from four different seasons.

Ocean and Coupled Earth System Emulator Development for NOAA-EAGLE

"Niraj Agarwa*l, NOAA Physical Sciences Laboratory, Boulder, CO, Joshua DaRosa, The MITRE Corporation, Bedford, MA, Timothy A. Smith, NOAA Physical Sciences Laboratory, Boulder, CO and Sergey Frolov, NOAA Physical Sciences Laboratory, Boulder, CO"

Machine Learning emulators have shown extraordinary skill in replicating atmospheric states and their application to emulate ocean and coupled earth systems offers similar promise. We have adapted the GraphCast architecture into a global ocean-only emulator for medium-range predictions, trained on NOAA's UFS-Replay dataset. Using a 24 hour time step, single initial condition, and without using autoregressive training, we have produced a robust emulator that provides skillful forecasts for 10–15 day lead times. We have further demonstrated the use of Mahalanobis distance as loss that improves the forecast skill compared to the Mean Squared Error loss by explicitly accounting for the correlations between the target variables. This loss acts as a statistical-dynamical regularizer for the slow, covarying dynamics of the global oceans, offering a better background forecast for downstream tasks such as data assimilation.

We have extended this work further by coupling independent ocean-only and atmosphere-only emulators (both trained on UFS-Replay) by exchanging the states of the coupled variables during the inference. This development uses ECMWF's Anemoi ecosystem for model development and NOAA's EAGLE tools for postprocessing and analysis. Using a comprehensive suite of baselines to isolate the impact of coupling, we have shown that the two-way coupling significantly impacts the forecast of key surface ocean quantities such as SST and SSH on 10-15 days lead times. The atmospheric variables however show a relatively lower impact on standard error metrics at these time scales, but this is under active investigation using longer lead times and process-level diagnostics. The framework is also set for further advancement using a more optimized set of coupled variables, increased spatial resolution, and refined ocean-only and atmosphere-only components using iterative fine tuning. The most up-to-date results will be presented from both ocean-only and coupled emulators, along with the lessons learned and future directions.

Improving Tropical Cyclone Track and Intensity Forecasts via Integrated AI

Lin Zhu*, Xu Lu, Yonghui Weng, Biju Thomas, Bin Liu, Zhan Zhang, Fanglin Yang and Vijay Tallapragada

Accurate prediction of tropical cyclone track and intensity remains a challenging problem due to complex nonlinear storm–environment interactions. Recent advances in artificial intelligence and machine learning have accelerated progress in numerical weather prediction, offering new opportunities to improve hurricane forecasts. However, their success depends critically on the availability of high-quality training data. In particular, there is a growing need for high-resolution analysis and reanalysis datasets that better represent storm structure and environmental conditions. To address this need, we are developing high-resolution analysis and reanalysis products, data integration and refinement efforts to better capture storm-scale features and environmental variability.

Here, we present integrated artificial intelligence (AI) methods that enhance tropical cyclone prediction through forecast nudging, statistical post-processing, and deep learning. We apply forecast nudging to improve track prediction by adjusting HAFS model variables toward the AI-based global forecast system (AIGFS) with superior track performance, thereby improving the representation of the large-scale environment. For intensity prediction, XGBoost models are used for post-processing intensity correction/forecasts, leading to improved intensity accuracy. In addition, we are developing a regional graph transformer based model to capture storm-centered structural features and environmental interactions from high-resolution atmospheric fields.
Together, these approaches aim to improve hurricane track and intensity prediction skills relative to baseline dynamical model guidance. This work demonstrates the potential of combining

physics-aware modeling with data-driven methods, supported by improved high-resolution datasets, for advancing tropical cyclone prediction.

TC-Cast: A Generative AI Emulator for High-Resolution Tropical Cyclone Forecasting

Hua Leighton*, Xuejin Zhang, Lin Zhu and Zhan Zhang

Tropical cyclones remain among the most destructive natural hazards, and accurate forecasts of their track, intensity, and structure depend on high-resolution, convection-allowing numerical models. NOAA's Hurricane Analysis and Forecast System (HAFS) provides state-of-the-art tropical cyclone guidance at 1.8 km resolution, explicitly resolving eyewall convection, rainbands, and other storm-scale processes. However, its high computational cost limits the use of large ensembles and reduces flexibility for rapid probabilistic forecasting.

This work introduces TC-Cast, a generative artificial intelligence emulator being developed to reproduce high-resolution HAFS forecasts at greatly reduced computational cost. TC-Cast builds on the StormCast framework and is designed to generate fast, ensemble-based tropical cyclone forecasts using the HAFS storm-following 1.8-km nest, rather than a fixed regional domain. The model is conditioned on large-scale environmental forecast fields together with the latest high-resolution HAFS analysis, with the goal of preserving both storm-scale structure and the surrounding synoptic environment.

The presentation will describe the TC-Cast model design, training strategy, input and output configuration, and real-time inference framework. It will also present initial results comparing TC-Cast forecasts with HAFS for selected tropical cyclone cases, including evaluations of track, intensity, wind structure, and ensemble-based probabilistic guidance. This work aims to demonstrate the potential of generative AI to provide rapid, high-resolution tropical cyclone forecast ensembles for future operational applications.

Why bother building variable resolution AI forecast models?

T. Smith, S. Frolov*, M. Pope, J. DaRosa, B. Bassarab and B. Huang

NOAA PSL in collaboration with EPIC and MITRE utilized the Anemoi framework to develop a variable resolution AI atmospheric forecast model. The model is trained on the combination of the operational GFS and HRRR analysis, has resolution of 25 km globally and 6 km refinement over the contiguous United States (CONUS). The model runs four times daily in the near real time environment on Microsoft Azure initialized with the operational GFS and HRRR analysis. Initial model evaluation shows promising forecast skill over CONUS, with improved forecast skill over GFS and HRRR for surface wind speed out to 10 days, 2-meter temperatures out to 5 days, and promise for positive precipitation evaluation.

This presentation examines the drivers and constraints behind the development of the variable resolution forecast models, including the availability of the training data, availability of the operational initial conditions, and the computational cost of developing global training datasets and operational inference for global convective resolving models.

Poster Session II

The Community Emissions Computing Engine (CECE): A High-Performance, Portable Framework for Atmospheric Emissions in UFS

Barry Baker*, Patrick Campbell, Zachary Moon, Quazi Rasool, Colin Harkins, Meng Li, Rebecca Schwantes, Brian McDonald and Fanglin Yang

Accurate and timely emissions data is a critical input for Earth system models, particularly for air quality, weather, and climate forecasting applications within the Unified Forecast System (UFS) and the wider community. The complexity of managing diverse global and regional emissions inventories and efficiently processing them for model applications presents a significant challenge. The Community Emissions Computing Engine (CECE) addresses this by providing a high-performance C++ component, designed to standardize, automate, and accelerate the preparation of atmospheric emissions data.

CECE is built on the Kokkos programming model, enabling performance portability across heterogeneous architectures, including NVIDIA GPUs and multi-core CPUs (OpenMP), without requiring source code changes. It operates as an ESMF Grid Component to enable seamless integration into the Earth System Modeling Framework (ESMF) and the UFS infrastructure. The core of CECE is the Stacking Engine, which efficiently combines multiple emission layers using hierarchical processing, kernel fusion optimization, and advanced temporal/spatial scaling. It employs a hybrid data ingestion model, pulling live meteorological data from ESMF while simultaneously reading static inventories from disk using TIDE.

Key features and architectural highlights include:

- Performance Portability: Leverages Kokkos for optimal performance across diverse hardware (CPUs, GPUs).
- Modular Physics Engine: Supports easy extension with new physics schemes (e.g., MEGAN, sea salt, dust) via both native C++ (Kokkos) and legacy Fortran plugins.
- Flexible Stacking Engine: Combines and processes emission layers using prioritized categories, geographical masks, and multiple scale factors.
- Python Bindings: Provides a high-level Python API via pybind11 for integration into data analysis and pre-processing workflows.

Standardized Lifecycle: Follows the standard NUOPC/ESMF lifecycle for initialization, high-throughput execution, and finalization, ensuring smooth operation within coupled models.

This presentation will detail the architecture of CECE, demonstrate its performance advantages and application in preparing emissions for UFS-coupled chemistry-aerosol systems, and highlight its foundational role in ensuring the integrity and scalability of community modeling efforts by providing a reliable, efficient source of essential atmospheric inputs.

UFS/JEDI LIDAR DATA ASSIMILATION: LiDAR Forward Operator Development

Michael K. Benneh*, Cheng-Hsuan Lu, Shih-Wei Wei, Cory Martin, Aara'L Yarber, Cheng Dang and Sen Chiao

The joint collaborative effort between Howard University, SUNY Albany, NOAA/EMC and JCSDA focuses on advancing LiDAR data assimilation capabilities within the Joint Effort for Data Assimilation Integration (JEDI). The primary goal of this effort is to better constrain the vertical structure of atmospheric aerosols, addressing a longstanding limitation in current data assimilation systems that rely on column-integrated observations.

The LiDAR Forward Operator transforms model-resolved atmospheric aerosol properties into observation space by simulating the measurement process of active remote sensing instruments. It is an essential capability for the consistent comparison and computation of innovations between lidar observations and model state. The lidar forward operator is part of the core step towards assimilating high-resolution vertical profile observations.

The current version has been used in verifying our baseline NOAA GCAFS model experiment. In the baseline experiment, VIIRS AOD retrievals were assimilated to constrain aerosol fields. However, the aerosol vertical layers were further evaluated to observe the limitations of AOD assimilation. The application implementation involved coupling the LiDAR forward operator to JEDI components and METplus for verification purposes. This involved extracting model states via VIND by interpolating them onto gridded states at observation locations as GEOVALs and passing them into VADER via OOPS. Prior to that, the CALIOP extinction profiles data are already written and passed into IODA. CRTM is used in our forward operator. VADER prepares CRTM-required input variables using the model "ingredients," and they are then passed into CRTM. Aerosol mixing ratios are then converted into mass concentrations, and the extinction coefficients are computed based on Mie theory via precomputed optical tables. The extinction profile is then obtained from the extinction coefficient based on each vertical layer. While the results showed improvements in the horizontal representation of AOD, the evaluation against CALIOP extinction profiles showed vertical profile mismatch in positioning and concentrations in their vertical profiles.

Ongoing work focuses on extending the current JEDI-based data assimilation framework to assimilate lidar-derived extinction profiles to improve the vertical representation of aerosols.

Accelerating Forecast Innovation: EPIC's Collaborative Framework and HPC-Driven Transparency

Kristopher Booker* and Anna Kimball

The NOAA Earth Prediction Innovation Center (EPIC) program has enabled an accelerated pace of numerical weather prediction innovation through their collaboration with government, academic, and private sectors. Utilizing a common modeling framework, the UFS (Unified Forecast System), has allowed once siloed atmospheric modeling groups to effectively collaborate. However, the greater challenge has come from constraints that kept organizations from accessing high performance computing (HPC) platforms capable of running new modeling innovations at scale. EPIC has solved this dilemma by providing UFS community contributors public access to detailed insights from these HPC platforms. Every contribution considered for integration is extensively tested through a continuous integration/continuous deployment (CI/CD) pipeline on various NOAA RDHPCS (research and development high performance computing systems) platforms with performance capabilities similar to the Weather and Climate Operational Supercomputing System (WCOSS-2). Systematically capturing data such as walltime and core hours for each model run. This provides contributors with detailed insights into computational efficiency and resource usage. This enables the community to track the evolution of model updates over time, offering quick insight into how performance and efficiency improve successively. The design supports future integration of scientific forecast skill, allowing users to correlate computational costs with forecast accuracy. The data collected is processed and displayed to contributors through the EPIC Community Portal creating more transparency between EPIC and the broader community. This openness empowers contributors, researchers, and developers to identify trends, compare model versions, and contribute to continuous model refinement. Ultimately, strengthening the feedback loop between operational teams and the user community, accelerating improvements in forecasting quality and operational efficiency. Our presentation will demonstrate how UFS community members without direct access to a NOAA RDHPCS can still contribute and follow the progress of their pull requests through the various insight metrics and public dashboards.

Progress update on incorporating the MICM gas-phase chemical solver in CATChem

Maggie Bruckner*, Rebecca Schwantes, Jian He, Barry Baker, Kyle Shores, Matthew Dawson, David Fillmore, Zachary Moon, Wei Li, Larry Horowitz, Vaishali Naik, Li Zhang, Congmeng Lyu, Quazi Rasool, Ravan Ahmadov, Jordan Schnell, Alan Gorchov Negron, R. Bradley Pierce, Mary Barth, Gabriele Pfister, Louisa Emmons and Brian McDonald

NOAA's Unified Forecast System (UFS) is a community-based Earth system model that provides a framework to efficiently incorporate research advances into NOAA's operational forecasts. Chemistry related code is currently incorporated into the UFS for several component models of varying complexity including UFS-AQM (UFS Air Quality Model), UFS-Aerosols, and UFS-RAQMS (UFS - Realtime Air Quality Modeling System). The Configurable ATmospheric Chemistry (CATChem) library and modeling component is being developed to unify all chemical and aerosol processes needed to perform atmospheric chemistry and composition simulations within a single, flexible framework. CATChem will be integrated into the UFS to create UFS-Chem version 2 or the Unified Forecast System with Chemistry.

NSF NCAR's Model Independent Chemistry Module (MICM) will be used in CATChem to represent and solve the gas phase chemical mechanisms. MICM uses YAML configuration to define chemical mechanisms, which will enable CATChem users to flexibly switch between mechanisms of varying complexity at runtime. We will demonstrate this capability within CATChem, showing initial comparison of box model and UFS testbed simulations utilizing the AM4.1 and RAQMS gas phase mechanisms. Preliminary evaluation will show consistency between CATChem simulations and the standalone box model implementation of MICM. In the future, fidelity of UFS-Chem version 2 will be evaluated against ozone, aerosols and their precursors from UFS-Chem version 1 and satellite, surface and aircraft observations.

STOFS 2D+ Global: Skill Assessment and Latest Advancements

Albert Cerrone*, Aman Tejaswi, Damrongsak Wirasaet, Shohei Ohata, Yuji Funakoshi, Greg Seroka, Mojan Rostaminia, Liujuan Tang, Saeed Moghimi, Edward Myers, Zach Cobell and Joannes J. Westerink

STOFS 2D+ Global is NOAA's global tide and storm surge model. It is operated four-times daily at NOAA and once-daily at the University of Notre Dame where new advancements are staged for production. These advancements include implementation of the direct total tidal potential, inline self-attraction and loading, recalibrating the internal tide dissipation across ocean basins, incorporation of the Earth's thermohaline engine, wave coupling with SWAN, bias correction strategies, streamlined execution (from pre- to post-processing), and subgrid integration. These advances are detailed alongside an exploration of their impact on computational efficiency and predictive accuracy. Additionally, STOFS 2D+ Global skill is assessed over a three-year hindcast period, facilitating a myriad of studies including a sea level analysis incorporating the new assessments of the Topography of the Sea Surface (TSS). This comprehensive evaluation highlights the model's robustness across diverse coastal and oceanic environments. Finally, its forecast skill over 2025 and Q1-Q2 2026, which is accessible on https://chl.crc.nd.edu/, is presented and critical analysis is given for areas of improvement, offering a roadmap for future version iterations.

Facilitating the 2 in R2O2R as part of the HydroMeteorology Testbed (HMT) forecast and evaluation activities

James Correia Jr*, Sarah Trojniak, Massey Bartolini and James Nelson

The HMT has many forecasting and evaluation activities centered around our years of Rapid Refresh Forecast System (RRFS) research to operations to research (R2O2R). In this abstract we discuss the experiments and our community effort to contribute toward problem detection and identification.

The Flash Flood experiment detected that“popcorn storms” had returned. These storms were easily recognized but how to interpret them and assign a cause was difficult. We borrowed an existing parameter from the NAM, hourly maximum precipitation rate, and requested its use in helping diagnose the magnitude of the problem. Precipitation rates may have been mitigated but remain high relative to other models.

The Winter Weather Experiment has been examining this years' RRFS ensemble for its depictions of precipitation types from its 2 microphysical parameterization schemes: NSSL and Thompson-Eidhammer. We tested modeled sleet against a case study from January that resulted in crippling winter weather. Results show that the amounts of sleet and freezing rain are very different across the ensemble, on either extreme of the HRRR. How this problem feeds back on member p-type diagnosis and ensemble p-type depictions will be discussed.

We will conclude with how we create activities that help us identify model differences in which we have no direct model code development role. Despite this, scouring change logs, bug fixes and developer requirements can lead to better understanding of model character. Working to bridge forecaster feedback with objective, statistical assessments can point in the direction of model problems early rather than post implementation.

Real-Time Ocean State Monitoring and Loop Current Forecasting with CNAPS2 Regional Ocean Data Assimilative Forecast System

Yongfei Deng*, Tianning Wu and Ruoying He

The Coupled Northwest Atlantic Prediction System Version 2 (CNAPS2) is a real-time ocean forecasting system combining the Regional Ocean Modeling System (ROMS) with Ensemble Optimal Interpolation (EnOI) data assimilation, designed to monitor and predict mesoscale ocean dynamics in the Gulf of America (GoA) — with a focus on Loop Current and Loop Current eddy (LC/LCE) variability. Open boundary and surface forcing are sourced from the Copernicus Marine Service (CMEMS) Mercator Ocean product and the NOAA Global Forecast System (GFS), respectively. The assimilation framework ingests satellite sea-surface temperature, altimetry, and near-real-time glider profiles, constraining model dynamics both upstream and downstream of the Loop Current. CNAPS2 generates daily hindcast/forecast cycles spanning five days prior through seven days ahead of the analysis time.

System performance was evaluated against independent satellite observations, moored ADCP current records, and autonomous glider data. CNAPS2 accurately reproduces mesoscale variability and western boundary current structure, including recent Loop Current eddy shedding events such as Eddy Feldman in June and August 2025. The integration of high-frequency glider profiles into the assimilation cycle notably improves subsurface stratification and eddy evolution, translating into more reliable short-term forecasts. These capabilities are directly relevant to navigation safety, offshore operations, and ecosystem management in the GoA.

Ongoing work focuses on Observing System Experiments (OSEs) during an observation-intensive field period, alongside the development of adaptive sampling strategies to further optimize assimilation impact and forecast skill.

Development of a one-way coupled wave forecasting model for the Mouth of the Columbia River with comparisons to data

Leandro Fernández*, Charles Seaton, and Merrick C. Haller

Background: The Mouth of the Columbia River (MCR) is considered the world's most dangerous major commercial navigational entrance and is a critical lifeline for shipping and maritime transportation. Accurate wave forecasts are critical for the safe passage of ships across the Columbia River Bar, and since the local currents can reach 7 knots, including wave-current interaction is critical for forecasting of the most dangerous conditions.

Methods: Here we report on the development and testing of a powerful high-resolution wave forecasting system with the critical additional capability of including the effect of currents on the wave forecasts through one-way coupling between the spectral wave model Wavewatch III (WW3) with the hydrodynamic forecasting model SCHISM.

The system is built on a regional domain centered on the MCR. Both models share the same unstructured grid, enabling high-resolution representation of the spatially complex conditions at the estuary mouth and extending into the open ocean. For the hindcast covering the period 11/2025 – 04/2026, the wave model is forced from ERA5 winds and offshore buoy NDBC 46029. WW3 ingests the current fields from SCHISM and we have run tests using 1) surface currents, 2) depth-averaged currents, and 3) no currents to investigate the detailed effects of wave-current interaction at the MCR.

Results: Model testing is conducted via comparisons of 1) model waves to nearshore buoy NDBC 46243, 2) model currents to surface currents measured at CDIP Clatsop Buoy, and 3) model waves to SWOT satellite derived wave heights (Ka-band Radar Interferometer). The results are very good and demonstrate overall errors of less than 30cm (RMSE). In addition, some interesting specific conditions that affect the forecast at instances away from the strongest currents (non-intuitive) are discussed.

AIGFS as an Incubator for Modern Software Practices Supporting R2O2R

Christina Holt*, Cooperative Institute for Research in Environmental Sciences at NOAA GSL
Paul Madden, Cooperative Institute for Research in Environmental Sciences at NOAA GSL
Emily Carpenter, Cooperative Institute for Research in Environmental Sciences at NOAA GSL
Anna Smoot, SAIC for Lynker Supporting NOAA EMC, Jacob Carley, NOAA EMC, Isidora Jankov, NOAA GSL, Steven Earle, NOAA NCO, Jun Wang, NOAA EMC and Rahul Mahajan, NOAA EMC

In December 2025, the National Weather Service launched the Artificial Intelligence Global Forecast System (AIGFS) as part of its first AI-based operational modeling suite. AIGFS is a global deterministic model based on Google DeepMind's GraphCast and boasts the most streamlined R2O process of any current operational system – an excellent example of DevOps ethos and cross-functional, stream-aligned teams. This accelerated process produced a system tailored to an operational application, but it lacked the support important for the ongoing development cycle, including portability across different platforms and configuration adjustment.

Our team is using AIGFS to help pave the way for modern application development that meets the needs of both research and operations. We'll build an application that is flexible enough to run on many HPC platforms, use ecFlow or Rocoto for workflow orchestration (or run without a workflow manager), and support a variety of real-time and retrospective studies. The application will be uwtools-based, code-quality- and schema-checked, and fully unit-tested, providing a new standard of reliability and software quality.

In addition to enhancing the flexibility and portability of AIGFS built for operations, the team is working closely with NWS Central Operations (NCO) to prototype ways to broaden the utility of operational software architecture as the whole of NOAA aims to become more cloud-centric. The tenets of portability, consistency, and thorough testing will be key to modernizing NOAA's modeling suite and ensuring a complete R2O2R loop for all future operational systems.

Modernized tools to evaluate and transition operational models to the National Ocean Service

Rachel Ackerman (Ocean Associates, Inc. in support of Office of NOAA Coast Survey), Atieh Alipour (Ocean Associates, Inc. in support of NOAA Office of Coast Survey), Lucila Houttuijn Bloemendaal* (NOAA Center for Operational Oceanographic Products and Services), Andrew Kammerer (Ocean Associates, Inc. in support NOAA Center for Operational Oceanographic Products and Services), Min Kong (IBSS Corp. in support of NOAA Center for Operational Oceanographic Products and Services), Patrick Limber (Ocean Associates, Inc. in support NOAA Center for Operational Oceanographic Products and Services), Armin Pruessner (NOAA Center for Operational Oceanographic Products and Services), Jack Reeves Eyre (Ocean Associates, Inc. in support of Office of NOAA Coast Survey), Jiangtao Xu (NOAA Center for Operational Oceanographic Products and Services) and Aijun Zhang (NOAA Center for Operational Oceanographic Products and Services)

NOAA's National Ocean Service (NOS) develops coastal and ocean models through a community modeling framework, to provide products and services supporting precision marine navigation, safe and efficient port operations, emergency response and ecosystem management. Under UFS, NOS is modernizing and expanding its ocean model evaluation capabilities to better serve the modeling community and its model-derived application users. NOS is developing a set of tools, including a

Python-based, customizable ocean model skill assessment package and a web-based GIS application that facilitates efficient, routine, and near real-time evaluation of model results against a suite of observations using documented NOS standards and criteria. As part of this process, the team has adopted an agile approach to develop and migrate the Python package to a public GitHub repository to allow for community tracking of releases and input. This package is available to the external modeling community. It is intended to evaluate ocean models within and beyond NOS, and to ensure models developed by NOS partners meet specific criteria before being transitioned to operations within NOAA. We will demonstrate the features of these skill assessment tools and highlight the open development process, encouraging community feedback, contributions, and engagement.

Progress on integrating JEDI with MPAS for the next version of the operational convection-allowing Rapid Refresh Forecast System

Ming Hu*, Terra Ladwig, Chunhua Zhou, Guoqing Ge, Haidao Lin, David Dowell, Junjun Hu, Sijie Pan, Hongli Wang, Ruifang Li, Shawn Murdzek, Liao-Fan Lin, Huiying Luo, Keenan Eure, Eric James, Amanda Back, Jeff Beck and Angela Burke

Hourly-updating numerical weather guidance at convective-allowing scales is the foundation of severe storm forecasting. In the past two decades, NOAA has successfully developed and operated the High Resolution Rapid Refresh (HRRR), which began an era of community-based modeling components in operations. The HRRR includes the Gridpoint Statistical Interpolation (GSI) for observational data assimilation, a specially configured version of the Advanced Research WRF (ARW) model, and unified post-processing (UPP). In the past 6 years, NOAA has made significant efforts to develop the new Rapid Refresh Forecast System (RRFS). While the first version of RRFS (scheduled for summer 2026) leverages the FV3 dynamical core and GSI, a significant pivot has taken place for future versions to leverage the Model for Prediction Across Scales (MPAS) and the Joint Effort for Data Assimilation Integration (JEDI).

Both MPAS and JEDI are community-based systems representing leading development from the research community, with contributions from NOAA and many other agencies. While they include many advanced features that could significantly improve operational system implementations, such systems also require extensive enhancement and a complete assessment before they can be fully integrated into our operational forecast systems.

The RRFS data assimilation team in GSL has made significant effort to develop and tune JEDI through both real-time and retrospective experiments. Recent progress includes the setup of the first real-time RRFS prototype built with MPAS and JEDI, named RRFSv2X. This system began running in early 2026 and will be evaluated by the Spring Forecast Experiment (SFE) and the Flash Flood and Intense Rainfall Experiment (FFaIR). The RRFSv2X runs on CONUS 3km mesh grid and includes hourly deterministic cycles constructed with JEDI hybrid 3DEnVar to provide initialization conditions for long forecasts at 00z and 12z, along with ensemble data assimilation constructed with JEDI GETKF with 30 ensemble members on an hourly cycle to provide ensemble covariances for the EnVar analysis. Current development is focused on increasing observation types, improving observation processing (thinning, quality control, and error settings), and evolving cycling strategies. Quantitative and qualitative evaluations of RRFSv2X from the 2026 convective season will be shared in this presentation.

Advancing Local Ensemble Solvers in JEDI: Cross-Validation and Multiscale Localization for Improved Data Assimilation

Bo Huang*, Sergey Frolov, Christopher Thomas, Joanne Waller, David Simonin, Anna Shlyaeva, Jeffrey S. Whitaker, Craig Bishop, Meelis Zidikheri, Nathan Eizenberg, Cory R. Martin, Daniel Holdaway, and Yannick Tremolet

The computational efficiency of the local ensemble solvers in the Joint Effort for Data assimilation Integration (JEDI) has been significantly improved, enabling testing in real applications including the atmosphere, ocean, sea ice and atmospheric compositions. We are currently focusing on implementing advanced techniques and science to further improve the background error covariance estimates in the JEDI local ensemble solvers. Recently, a cross-validation capability was implemented to utilize independent subsets of ensemble members for member updates and thus mitigate filter inbreeding. We are now implementing an efficient muti-scale localization method for local ensemble solvers in JEDI to better represent physical correlations across varying spatial scales. These new developments will be evaluated within the NOAA Unified Forecast System (UFS) global atmospheric model using near-operational observation datasets. Detailed implementation strategies and experiment results will be presented.

Evaluation of GFDL and EMC Hurricane Tracking Systems During the 2025 Hurricane Season

Hananeh Jafary* and Timothy Marchok

The ensemble tropical cyclone (TC) tracking system developed at the NOAA Environmental Modeling Center (EMC) within the National Centers for Environmental Prediction determines storm position by using diagnostic fields derived from nearby maxima and minima in the vicinity of an initial estimate. This framework produces forecasts of track, intensity, wind radii and genesis for multiple global modeling systems, including the NOAA Global Forecast System (GFS), Global Ensemble Forecast System, Canadian Meteorological Center deterministic and ensemble configurations, European Centre for Medium-Range Weather Forecasts (ECMWF) deterministic and ensemble systems, and Fleet Numerical Meteorology and Oceanography Center products.
Originally introduced in 1998, the tracker has been continually improved through successive updates. Over time, independent development efforts at the NOAA Geophysical Fluid Dynamics Laboratory (GFDL) and EMC have led to divergence in the underlying algorithms due to debugging

and performance enhancements. Within the Unified Forecast System community modeling effort, ongoing maintenance and systematic evaluation of the tracker is being conducted to better diagnose and reduce forecast errors. A unified implementation is being developed to ensure consistent performance across computing environments, platforms, and compilers.
The resulting unified tracker is intended to deliver more reliable tropical cyclone guidance for forecasters in the National Weather Service, as well as stakeholders in both public and private sectors. To demonstrate the impact of these updates, we compare TC genesis forecasts from the ECMWF and the GFS during the 2025 hurricane season, evaluating differences between the updated GFDL-based tracker and the legacy EMC implementation.

A Modular YAML-Based Unified Workflow Framework for NOAA National Ocean Service's Operational Ocean Forecast Systems and UFS-Coastal Integration

Mansur Ali Jisan*, Zachary Cobell, Ufuk Turuncoglu, Yunfang Sun, Soroosh Mani, Yuji Funakoshi, Zizang Yang, Aijun Zhang, Machuan Peng and Saeed Moghimi

NOAA's National Ocean Service (NOS) operates a suite of Operational Forecast Systems (OFSs), including the Surge and Tide OFS (STOFS) suite and Coastal Ocean Modeling Framework (COMF)-based regional OFSs. These systems provide critical water level, circulation, and inundation guidance for navigation, coastal resilience, and ecosystem management. However, system-specific workflows can limit consistency, reproducibility, and adoption of modern software practices. To advance NOS's vision for community-driven development and integration of coastal modeling into the Unified Forecast System (UFS), we are developing NOS-workflow, a modular, YAML-driven Python workflow framework.

NOS-workflow provides a common orchestration layer for defining model settings, forcing inputs, preprocessing options, runtime parameters, and execution steps while retaining flexibility for system-specific configurations. Reusable preprocessing capabilities are encapsulated in NOS-utils, a standalone Python utilities package. Current support spans SCHISM, ROMS, FVCOM, and ADCIRC-based systems, and UFS-Coastal configurations.

As an initial UFS-Coastal pathway, we adopted the pre-operational Southeast Coastal Operational Forecast System (SECOFS) as a testbed and implemented a one-way coupled atmosphere-ocean configuration using the Community Data Models for Earth Predictive Systems (CDEPS) component of UFS and ESMF/NUOPC on WCOSS2. The framework also supports ensemble forecasting, demonstrated by driving SECOFS and STOFS-3D-Atlantic with GEFS members, and by ensemble cycling under the UFS-based SECOFS configuration. For portable development and testing, NOS-workflow is paired with a WCOSS2-like container and GitHub Actions Continuous Integration pipeline targeting NOAA RDHPCS and academic collaboration.

Together, these developments establish NOS-workflow as a unified, extensible, CI-tested foundation for modernizing NOS coastal forecast workflows and supporting the UFS-Coastal transition.

Advancing Radiance Observation Operators for Data Assimilation: Recent Developments in the Community Radiative Transfer Model (CRTM)

Benjamin T. Johnson*, Cheng Dang, Quanhua Liu, Yingtao Ma, Patrick Stegmann, Isaac Moradi and Yong Chen

The Community Radiative Transfer Model (CRTM) serves as a core observation operator for satellite radiance assimilation across multiple numerical weather prediction systems. Continued improvements in forward modeling accuracy, physical consistency, and computational performance are critical for maximizing the impact of satellite observations in data assimilation (DA) frameworks.

This presentation highlights recent developments in CRTM that directly affect DA performance, including updates to spectroscopy and optical property parameterizations, enhancements to scattering and surface emissivity treatments, and improvements in numerical stability and computational efficiency. Particular attention is given to how these changes influence simulated brightness temperatures, Jacobians, and bias characteristics within operational and research DA systems.

We also discuss ongoing efforts to improve consistency between forward and tangent-linear/adjoint components, with implications for variational and ensemble-based DA applications. Results from controlled experiments demonstrate the sensitivity of analysis increments and forecast skill to observation operator fidelity, emphasizing the role of CRTM in reducing representativeness and forward-model errors.

Finally, we outline current integration efforts within community DA frameworks (e.g., JEDI/UFO) and future directions aimed at supporting next-generation satellite observations and coupled Earth system applications.

Standalone Diagnostic Tools for UFS Data Assimilation Evaluation with Applications in the CADRE Year 2 Data Assimilation Training and Science Workshop

Jong Kim*, Chan-Hoo Jeon, Wei Huang and Gillian Petro

This presentation outlines a reproducible, onboarding-friendly environment for learning Unified Forecast System (UFS) Data Assimilation (DA) concepts through hands-on global atmospheric FV3-JEDI applications, where FV3-JEDI is the Joint Effort for Data assimilation Integration (JEDI) interface for the UFS FV3 atmospheric component. This environment was applied in the June 2026 Consortium for Advanced Data Assimilation Research and Education (CADRE) training and science workshop to include a baseline C96/C48 hybrid Three-Dimensional Variational (3D-Var) control case, background-error covariance parameter tuning, observation-error sensitivity tests, and multi-sensor innovation analysis. All experiments are executed using standalone run scripts to ensure transparency and accessibility for new users.

A central component of this environment is the development of a standalone Python diagnostic tool designed to evaluate DA performance independently of the workflow to lower the barrier to entry for new users. This tool computes chi-squared consistency diagnostics, innovation statistics, background-model and analysis-increment power spectra, and vertical-increment structure. These diagnostics allow users to interpret DA behavior without requiring workflow integration. Using this tool, we demonstrate clear diagnostic signals such as the impacts of Advanced Technology Microwave Sounder (ATMS) Observation-Error inflation on the cost-function ratio Jo/p convergence and cross-sensor balance across conventional observations, Global Navigation Satellite System Radio Occultation (GNSS-RO), Advanced Scatterometer (ASCAT), and Geostationary Operational Environmental Satellite (GOES) Atmospheric Motion Vectors (AMVs). Spectral diagnostics can further confirm expected hybrid-3DVar characteristics and provide intuitive visualizations suitable for poster-style communication, for example.

By combining simple standalone experiment scripts with a modular diagnostic toolkit, the barrier is lowered for community members to explore, understand, and tune UFS DA systems. This approach supports training, research, and future workflow integration while maintaining reproducibility and clarity for learners entering the UFS community DA framework and ecosystem.

Empowering Forecasting Innovation Through EPIC Community Engagement and User Support

Aaron Jones, Charlene Barone, Zach Shrader*, Laura Generosa, Maryia Davis, Jef Dodson, Jessica Wheeler, Casey Seyb, Priya Pillai, Joshua Kublnick and Keven Blackman

The Earth Prediction Innovation Center (EPIC) is dedicated to accelerating contributions to the Unified Forecast System (UFS) and engaging the community to ensure community members have the necessary tools and knowledge to contribute innovations to our Nation's Earth System models. The EPIC Community Engagement (ECE) team supports EPIC's innovation through several initiatives, including community training events and the annual Unifying Innovations in Forecasting Capabilities Workshop (UIFCW). It also publishes dynamic content on social media, the EPIC Community Portal (ECP), GitHub, and other platforms to meet the community's evolving needs and to remove barriers to innovation.

This presentation will highlight several ECE-led outreach efforts which resulted in increased awareness of new Unified Forecast System capabilities. Some of these efforts included: publishing Issues 9 and 10 of UFS Insights and supporting major announcements including Project EAGLE, the UFS Air Quality Model (AQM) Community Version, Hierarchical System Development (HSD) Framework expansion with the Aquaplanet Test, the UFS Land Data Assimilation (Land-DA) v3.0.0 Release, and Nested EAGLE v1.0.0 Release. ECE also expanded EPIC's visibility through communication of involvement in the American Geophysical Union Fall Meeting 2025, the 106th American Meteorological Society Annual Meeting, and supporting the AI Versus AI Hackathon.

The User Support team complements these efforts by maintaining documentation, compiling FAQs, monitoring support requests, and providing feedback that helps ECE refine outreach strategies and ensure the community has the resources needed to innovate within national forecasting systems.

Empowering and Expanding the UFS Community Through EPIC User Support

Joshua Kublnick*, Priya Pillai and Brandon Selbig

The Earth Prediction Innovation Center (EPIC) User Support (US) team plays a key role in enabling the Unified Forecast System (UFS) community by developing and delivering a comprehensive suite of user-focused resources. This presentation will provide an overview of the core resources offered by the US team, including support channels such as the EPIC support email and GitHub Discussions, continuously maintained technical documentation, text and video-based tutorials, onboarding materials (Get Code), and Frequently Asked Questions (FAQs) driven by user needs. It will also highlight how the team supports experimentation and training through curated, staged datasets that align with both user and developer workflows. As EPIC continues to expand support for additional repositories such as Global Workflow (GW) and the Experimental AI Global and Limited-area Ensemble forecast system (EAGLE), these resources will continue to evolve to improve accessibility, streamline onboarding, and enhance overall usability. By focusing on these resources, the presentation will demonstrate how the US team will empower users with the tools and knowledge needed to effectively engage with and contribute to the UFS ecosystem.

2025 and 2026 HAFS Real-Time Parallel Experiments: Innovations, Performance, and Operational Transition

Bin Liu*, Zhan Zhang, Xu Lu, Yonghui Weng, Jing Cheng, JungHoon Shin, Biju Thomas, Weiguo Wang, Lin Zhu, Chuan-Kai Wang, Maria Aristizabal, Jiayi Peng, Fanglin Yang and Vijay Tallapragada

Hurricane Analysis and Forecasting System (HAFS), as the Unified Forecasting System (UFS) hurricane application, is NOAA's flagship regional operational hurricane model to provide reliable and skillful dynamical model guidance on hurricane track, intensity, and storm structure. With multiple years of collaborative developments and real-time parallel experiments, HAFS initial operational capability (HAFSv1) was first implemented in 2023. HAFS continued its roadmap with annual improvements through real-time parallel experiments (testing and evaluating new developments and innovations), followed by operational transitions and upgrades of HAFSv2 (2024), HAFSv2.1 (2025) and upcoming HAFSv2.2 (2026). The 2025 HAFSv2.1.1A (HFXA) real-time parallel experiment, built upon the operational HAFSv2.1A configuration, introduced significant innovations including the scale-aware 3DTKE EDMF GFS PBL scheme, the NoahMP LSM, and enhanced 4DEnVar data assimilation (DA) using 40 HAFS coldstart ensembles. Evaluation of HAFSv2.1.1A for 2025 North Atlantic (NATL) storms indicated overall comparable track and intensity skills compared to operational HFSA. These results informed the upcoming 2026 HAFSv2.2 operational upgrade, which incorporates further optimized storm-following 3DIAU, improved moving nest configuration for

NoahMP LSM, and an upgrade to RRTMGP radiation. The 2026 HAFS real-time parallel experiments will continue exploring JEDI-based atmospheric DA, which will be one of the main objectives for the future HAFSv3 upgrade. Besides, the HAFS basin-wide real-time analysis experiment with a self-cycled JEDI-based hybrid 4DEnVar-EnKF DA system is planned to prototype and demonstrate its capability of producing high resolution storm analysis and generating data for developing combined dynamical and AI/ML-based HAFS deterministic and ensemble forecasting.

Storm Surge and Coastal Inundation Nowcasts/Forecasts During Hurricanes Helene and Milton

Yonggang Liu*, Haibo Xu, Kaili Qiao, Sebin John, Sieu-Cuong San, Robert H. Weisberg, Jing Chen, Lianyuan Zheng, Sherryl Gilbert, Steven A. Murawski, Gary T. Mitchum and Thomas K. Frazer

A daily automated coastal water level (storm surge) nowcast/forecast guidance system has been developed by the USF Ocean Circulation Lab based on the West Florida Coastal Ocean Model (WFCOM) and the very high-resolution Tampa Bay Coastal Ocean Model (TBCOM). Both models are configured to perform realistic simulations of ocean circulation and water levels which are then combined with tide gauge observations to provide 3-day hindcasts and 3.5-day forecasts of coastal water level along the West Florida coast (http://ocl.marine.usf.edu/Models/SeaLevel/). The experimental product was maintained during the approach and passage of Hurricanes Helene and Milton, and provided critical storm surge forecasts to a broad suite of stakeholders including the public. The system successfully predicted the water level set-up and set-down along the west Florida coast three days in advance of each hurricane, with improved forecasts realized each day prior to landfall. The TBCOM-inundation forecast system was also activated during Hurricane Helene. This modeling system extends its dense grid onto the land, facilitating simulation of inundation and flooding associated with storm surge in coastal areas. During Hurricane Helene, areas of severe inundation were identified along the coastal periphery of Tampa Bay and forecasts were accessible two days in advance of landfall.

Simultaneous Assimilation of Multi-Wavelength Aerosol Optical Depth Retrievals from NASA's PACE Mission in JEDI for NOAA's Global Aerosol Model

Congmeng Lyu*, Bo Huang, Mariusz Pagowski, Lorraine A. Remer, Andrew M. Sayer, Xiaoguang (Richard) Xu, Li (Kate) Zhang and Brian McDonald

The three-dimensional ensemble-variational aerosol data assimilation system was developed within the Joint Effort for Data assimilation Integration (JEDI) for NOAA's global chemistry and aerosol forecast system (GCAFS v1). This system has been demonstrated in the near-real-time aerosol assimilation and used to produce global aerosol reanalysis by assimilating the aerosol optical depth (AOD) retrievals at 550 nm. While this approach improves the model's representation of total AOD, assimilating only 550 nm AOD has limited impact on aerosol composition and vertical profiles, which remain largely model-driven. To improve detailed aerosol representation in GCAFS v1, we are extending JEDI to assimilate advanced aerosol retrievals from NASA's Plankton, Aerosol, Cloud, ocean Ecosystem (PACE) satellite mission. We first implemented in JEDI simultaneous assimilation of multi-wavelength AOD retrievals from PACE Ocean Color Instrument (OCI). Preliminary results indicated that the simultaneous assimilation of AOD retrievals at 354 nm, 550 nm, and 670 nm showed clearly different aerosol analysis profiles compared to assimilation of AOD at 550 nm alone. Ongoing efforts include estimation and integration of cross-wavelength AOD correlations and systematic evaluation of multi-wavelength PACE OCI AOD assimilation on global aerosol forecasts against in-situ and satellite retrievals through cycled experiments. We will present our results and discuss plans of assimilating other PACE aerosol products within the JEDI framework.

Assessing Hourly Cycled Bayesian Data Assimilation via Localized Particle Filters within a Global Atmospheric UFS Workflow

Dylan Major*, Jonathan Poterjoy and Laura Slivinski

The UFS serves as the foundational modeling framework for global weather prediction in the United States. While current operational centers rely on ensemble Kalman filters and variational methods for data assimilation, these approaches do not account for non-Gaussian forecast errors and limit our ability to assimilate observations that are characterized by complex measurement processes. Localized particle filters provide an alternative methodology that relaxes many of the restrictive assumptions that constrain traditional methods for highly non-Gaussian and non-linear problems. This makes them exceptionally well-suited for complex challenges, such as assimilating all-sky satellite radiances, which require accurately representing non-Gaussian error structures and state-dependent measurement uncertainty. To further motivate design choices for this research, NOAA explicitly prioritizes rapidly updating ensemble prediction systems for the UFS workflow in their 10-Year Strategy for Data Assimilation. This presentation evaluates the local particle filter for rapidly updating (hourly) cycled global weather prediction and identifies potential advantages of particle filters for representing large-scale flows that dominate synoptic patterns.

Seasonal Prediction of Coastal Sea-Level Anomalies Using Climate-Driven Hydrodynamic Modeling

Reza Marsooli* and Serine Ziadi

Long-range predictions of high-tide flooding (HTF) support proactive, risk-informed decision making. The frequency and intensity of HTF are controlled by coastal sea level anomalies, which are strongly influenced by large-scale climate variability and ocean dynamics. Given Earth system models' ability to predict these large-scale processes, dynamic coastal ocean models forced with climate model outputs offer a promising approach for long-range prediction of coastal sea-level anomalies and, thus, HTF. This study investigates the potential of a three-dimensional hydrodynamic model, forced by atmospheric forcing from a climate model, to provide skillful predictions of coastal sea level anomalies. We employ the SCHISM hydrodynamic model within the Unified Forecast System-Coastal application, forced with atmospheric data from NOAA's SPEAR climate model to simulate sea levels along the U.S. East and Gulf Coasts at a seasonal scale. The modeling framework is implemented on the STOFS-3D Atlantic domain. The simulation is conducted by forcing SCHISM with daily-averaged wind and atmospheric pressure fields derived from a SPEAR historical dataset. Model performance for sea level anomaly prediction is evaluated against observations from NOAA tide gauge stations. Predictive skill is quantified using standard statistical metrics, such as anomaly correlation coefficient and root-mean-square error.

Vision Towards a Next-Generation Observation Ingest, Storage, and Access Framework for the National Weather Service

Cory R Martin* - NOAA/NWS/NCEP/EMC, Daniel Holdaway - NOAA/NWS/NCEP/EMC, Jeff Ator - NOAA/NWS/NCEP/EMC, Andrew Collard - NOAA/NWS/NCEP/EMC and Rahul Mahajan - NOAA/NWS/NCEP/EMC

Currently, NOAA's National Weather Service (NWS) uses the Gridpoint Statistical Interpolation (GSI) software for generating operational analyses to initialize many modeling applications in the production suite. As part of the transition to the Unified Forecast System (UFS), GSI is slated to be replaced by a Unified Data Assimilation (UDA) framework based on the Joint Effort for Data assimilation Integration (JEDI). One of the main motivations for transitioning to JEDI is to more rapidly assimilate new and novel observations by leveraging shared infrastructure and generic code rather than needing to develop individual functions for each observation type. However, the foundational infrastructure for the ingest, pre-processing and storage of observational data still relies on legacy software. While the existing system has proven to be robust, limitations in its architecture are presenting challenges in our ability to adopt modern metadata conventions and process emerging datasets in a timely manner. Here we present a high level overview of the current state of our observation ingest and storage infrastructure, a summary of requirements for the next-generation system (with care taken to ensure the needs of both operational numerical weather prediction and community stakeholders are met), and a proposed roadmap of how those requirements may be met through modern software design practices.

Let's Track Some Cyclones - Running the GFDL Vortex Tracker Made Easy

Caitlyn McAllister* - SAIC/NOAA/GFDL and Timothy Marchok - NOAA/GFDL

The GFDL Vortex Tracker is a flexible tool used throughout the operational and research weather communities to analyze model output and diagnose various forecast metrics that are critical to providing model guidance for tropical and extratropical cyclones. Despite its importance, usability and extensibility have historically been constrained by complex build processes, scarcity of testing options, and a lack of standardized development infrastructure.

This work presents ongoing efforts to modernize the tracker through the application of contemporary software engineering practices. An overview of the most recent upgrades to the

GitHub repository is provided, including: a reproducible build system with multi-compiler support across all NOAA RDHPCS platforms; automated testing for verification and validation; and integrated CI/CD workflows. Fully developed containerized environments are now available. Details will be provided on how these enable consistent and portable deployment of the tracker across high-performance computing (HPC) and cloud platforms.

A forthcoming release introduces user-friendly run scripts supporting NetCDF, GRIB1 and GRIB2 data formats, along with a comprehensive user guide outlining required inputs and workflow execution. These enhancements, delivered through a new versioned release on GitHub, support continuous improvement of the codebase while lowering barriers to use such functions. Containerized solutions extend accessibility to both RDHPCS and non-RDHPCS users, improving portability, reproducibility, and ease of adoption. This work positions the tracker for broader community engagement and more efficient operational and research applications in tropical cyclone analysis.

ADD-DA: A Machine Learned Model for Rapid Data Assimilation

C. Miller*, J. DaRosa, M. Alam, S. Frolov, L. Slivinski, C. Tong, N. Silverman and C. Lewellyn

Traditional data assimilation (DA) systems rely on variational or ensemble-based methods that make restrictive assumptions, such as Gaussian error statistics and linear observation operators, and require substantial computational resources. Machine learning offers a path toward faster, more flexible DA that can learn complex relationships between observations and atmospheric states directly from data. In this work we describe ADD-DA, a deep learning model that produces an analysis by computing an additive increment to a background forecast state (such as the operational GFS analysis). In this early stage of the work, ADD-DA is pretrained to replicate the analysis increments from a fully cycling GSI 3dVar reanalysis, using background forecasts and observation residuals from ATMS, AMSU-A, and conventional sensors as inputs. We present results from this pretraining phase and discuss ongoing work to couple ADD-DA into an end-to-end learnable forecast system alongside learned observation and forecast operators.

Assessing Structural Differences in Analysis and Forecast Error Covariance between Ensemble Kalman Filters and Local Particle Filters

Nikolaitchik, Maria*, Slivinski, Laura and Poterjoy, Jonathan

Ensemble data assimilation techniques leverage a sample of model states to approximate a posterior distribution conditioned on observations. In practice, Earth System Models rely heavily on methods such as Ensemble Kalman Filters (EnKFs), which are optimal when all underlying errors are strictly Gaussian. These techniques only require estimating means and error covariances, but are suboptimal in geophysical applications characterized by nonlinearity in the model dynamics or observation processes. Fully nonlinear ensemble filters, such as particle filters, can be regularized to mitigate sampling deficiencies in a manner comparable to EnKFs. These methods can achieve results that are nearly as skillful for high-dimensional applications while using ensemble sizes

affordable for operational applications. Yet, the impact of moving beyond Gaussian approximations and our reliance on heuristic mechanisms for treating sampling deficiencies has theoretical implications on the time evolution of multivariate uncertainty. As such, there is a clear need for quantifying and comparing the spatiotemporal structure of error covariances across linear and nonlinear filters. Doing so will provide deeper insights into the value of nonlinear filters for multiscale applications, while facilitating a better understanding of influencing factors such as localization and inflation.

The present study utilizes Analysis Error Covariance Singular Vectors (AEC SVs) to quantify and compare the evolution of error covariances between an EnKF and a nonlinear filter, the Local Particle Filter (LPF), within both an idealized framework and a global atmospheric application of the UFS with a reduced observation network assimilating only surface pressure observations. AEC SVs quantify the fastest growing perturbations within an optimization interval with respect to the analysis error covariance, giving insight into the “upper bound” of perturbation evolution. Preliminary results using AEC SVs suggest the LPF demonstrates comparable estimates of covariance structure to the ENKF, despite not explicitly estimating the background covariance. Furthermore, results suggest that inflation plays a non-trivial role in dampening the growth of perturbations across scales in moderately to strongly coupled systems.

How Does Land Proximity Reshape Tropical Cyclone Precipitation Asymmetry? Intensity-Stratified Composites from Open Ocean to Landfall

Eric Peino*, Gregory Foltz and Jun Zhang

Tropical cyclone (TC) rainfall asymmetry remains a major source of uncertainty for coastal flood forecasting, where land–atmosphere interactions strongly modulate the location and intensity of hazardous precipitation. Understanding how precipitation evolves as storms approach and cross the coastline is critical for improving quantitative precipitation forecasts and evaluating operational TC models. We present a global climatology of TC precipitation asymmetry, composited in a wind shear-relative framework across four coastal phases: open ocean, offshore, pre-landfall, and landfall. Composites are stratified by storm intensity and analyzed separately for mean rainfall and extremes. Coastal effects are quantified relative to the open-ocean baseline, and wave-1 asymmetry is characterized using Fourier decomposition. Across all intensity classes, the precipitation maximum is consistently located in the downshear-left quadrant. Environmental shear associated with extreme rainfall increases substantially near land, rising from ~4.5 to ~8.0 m s^{-1} in major hurricanes between open ocean and landfall. The dominant signal is systematic precipitation suppression near and after landfall, concentrated in the downshear-left inner-core region where open-ocean convection is strongest. TS/TD systems show the earliest and most spatially extensive suppression already offshore, while Category 1–2 storms exhibit localized offshore enhancement consistent with coastal convergence. Major hurricanes show the strongest overland reduction but also a statistically significant enhancement exceeding 60% in the downshear-right quadrant during pre-landfall and landfall. Extreme precipitation responds earlier and more broadly to coastal proximity than the mean. These results provide a physically grounded benchmark for evaluating operational TC forecast systems, including HAFS, and improving coastal flood hazard assessment.

Progress on a Wind-Wave-Current Data Assimilation Scheme for the 3D-RTMA: Graph-Based Localization and Coupled Ensemble Analysis

Malaquias Pena*

Building on three years of development under NOAA-WPO JTTI, we report progress toward a physically consistent wind-wave-current data assimilation (DA) system for the 3D Real-Time Mesoscale Analysis (3D-RTMA). The project targets a coupled analysis of 10-m wind speed (U_{10}), significant wave height (Hs), wind-sea energy (W), swell energy (Q), and surface ocean currents (U^c) by expanding LETKF-based ensemble DA to marine variables within a UFS/JEDI framework.

A working alpha-version of the U_{10}–Hs LETKF scheme has been developed, demonstrating ensemble spread–RMSE consistency for both scatterometer-based wind assimilation (ANL RMSE reduced from 1.09 to 0.46 m/s) and altimeter-based wave height assimilation. Observation statistics derived from six months of RTMA-CONUS buoy, altimetry, and scatterometry data reveal that largest analysis errors coincide systematically with satellite sampling gaps, motivating improved observation network characterization and gap-filling strategies. Bias correction diagnostics confirm that the scheme effectively removes state-dependent systematic errors in background Hs and U_{10} fields.

A key theoretical contribution is the development of water-graph localization: replacing Euclidean distance with a barrier-respecting shortest-path metric on a water-only graph. Standard Euclidean localization treats land as transparent, causing spurious analysis increments to leak across peninsulas and islands, a metric failure that cannot be fixed by radius tuning alone. Applied within LETKF, water-graph localization eliminates cross-land contamination by construction, while fat-tailed localization kernels (Cauchy/Laplacian) accommodate the longer along-water pathways of swell energy. Proof-of-concept experiments in a coupled Lorenz-96 coastal testbed show clear improvements in swell RMSE and spread–skill reliability relative to Euclidean baselines. Looking ahead, we are extending the framework from scalar Hs assimilation to partitioned wind-sea and swell DA, and planning a full 3D-RTMA demonstration on the operational CONUS grid.

HAFS Ensemble Development on Cloud Platforms

Jiayi Peng*, Zhan Zhang, Rajendra Panda, Weiguo Wang, Bin Liu, Vijay Tallapragada, Fanglin Yang, Sundararaman Gopalakrishnan, Xuejin Zhang, Aaron Poyer and William Komaromi

The absence of a skillful and reliable high-resolution ensemble presents a significant challenge for National Hurricane Center (NHC) forecasters in issuing probability forecasts for tropical cyclone track, intensity, structural changes, and associated hazards. To address this deficiency, the Hurricane Analysis and Forecast System (HAFS)-based Ensemble Prediction System (EPS) was developed and has been executed in real-time on the Amazon Web Service and Microsoft Azure cloud platforms since the 2023 hurricane season.

The HAFS-EPS incorporates a novel Stochastically Perturbed Parameterizations (SPP) scheme for ocean surface roughness to address hurricane intensity forecast uncertainty, in addition to three conventional stochastic physics schemes: Stochastically Perturbed Physics Tendencies (SPPT), Stochastically Kinetic Energy Backscatter (SKEB), and Stochastically Perturbed PBL Humidity (SHUM). Initial and boundary conditions are derived from forecast data generated by the operational Global Ensemble Forecast System (GEFS) of the National Centers for Environmental Prediction (NCEP). To account for initial uncertainty in hurricane position, intensity, and structure, the HAFS-A analysis vortex and the 6- or 12-hour forecast vortex from preceding cycles were merged into the initial single uniform Atlantic basin domain fields for the various HAFS-EPS members.

The performance of the HAFS-EPS forecasts during the 2024 and 2025 Atlantic seasons was rigorously evaluated against the operational HAFS-A and HAFS-B forecasts. This comparative analysis clearly demonstrates that the HAFS-EPS consistently produced superior track forecasts relative to the deterministic HAFS-A/B models throughout the 2024 and 2025 hurricane seasons. Moreover, the hurricane intensity forecast error of the HAFS ensemble proved to be comparable to that of the deterministic HAFS-A/B models. The 31-member (21-member) HAFS ensemble forecasts, implemented during the 2024 (2025) hurricane season, provided substantial advantages by offering probabilistic guidance for hurricane track, intensity, structure change, and associated hazards (including wind, wave, and storm surge) through the deployment of this reliable and skillful higher-resolution regional HAFS ensemble prediction system on AWS and Azure cloud.

Advancing Sea Ice Data Assimilation in UFS-Arctic Toward Arctic Reanalysis Applications

Joey Rotondo*, Cecilia Bitz and Christopher Cox

Accurate representation of sea ice is critical for Arctic prediction and reanalysis, yet current data assimilation (DA) capabilities within the Unified Forecast System (UFS)-Arctic remain limited. In particular, existing 3D-Var approaches do not fully leverage available observations or capture key coupled processes. This work aims to evaluate and extend UFS-Arctic DA capabilities for sea ice, with the long-term objective of enabling a consistent, large-scale Arctic reanalysis.
We implement and test sea ice DA within UFS-Arctic using the CICE model framework, incorporating satellite-derived sea ice concentration and thickness observations. Experiments focus on configuring and assessing the existing 3D-Var system, including observation handling, background error specification, and coupling with atmospheric and ocean components. Sensitivity studies are conducted to evaluate DA performance under different configurations and observational constraints.

Preliminary results highlight key limitations in the current UFS 3D-Var framework for sea ice, including challenges in error specification and limited multivariate coupling. However, targeted configuration adjustments improve assimilation stability and regional performance. These findings provide a baseline for future development of more advanced DA approaches.

This work demonstrates both the potential and current limitations of sea ice DA within UFS-Arctic. Improving DA methodologies is essential for advancing Arctic prediction and enabling reanalysis efforts. The results inform ongoing development of UFS DA systems and support future community efforts toward more robust and scalable Arctic forecasting capabilities.

Evaluating UFS and other Models Against Atmospheric Composition and Meteorological Observations Using MELODIES MONET

Rebecca Schwantes*, Zachary Moon, Colin Harkins, Maggie Bruckner, Quazi Rasool, Pablo Lichtig, Liam Thompson, Meng Li, Beiming Tang, Youhua Tang, Jordan Schnell, Rebecca Buchholz, David Fillmore, Andrew Rollins, Eleanor Waxman, Kristen Zuraski, Jeff Peischl, Nell Schafer, Charles Gatebe, Gabriele Pfister, Louisa Emmons, Barry Baker and Brian McDonald

MELODIES MONET is a joint project supported by NOAA and NSF NCAR to develop a community, Python-based model evaluation package that evaluates model results against atmospheric composition and meteorological observational datasets (https://melodies-monet.readthedocs.io/). Here we will demonstrate how to use the existing capabilities in MELODIES MONET to evaluate UFS-Chem (Unified Forecast System with Chemistry), UFS-AQM (UFS - Air Quality Model), and other models like WRF-Chem (Weather Research and Forecasting model coupled with Chemistry) against surface, sonde, aircraft, and satellite observations. For aircraft observations, we will evaluate the model results against the NOAA CSL led Atmospheric Emissions and Reactions Observed from Megacities to Marine Areas (AEROMMA) campaign. AEROMMA sampled several large North American cities to investigate urban air quality and long range transport of wildfire smoke in the summer of 2023. We will also demonstrate how MELODIES MONET can be used to evaluate model simulations against surface observations, including US EPA AQS (Air Quality System) and ISH (Integrated Surface Hourly Database), and satellite data, including TROPOMI (Tropospheric Monitoring Instrument). Our goal for MELODIES MONET development is to increase the usability and accessibility of field campaign datasets like AEROMMA by the modeling and UFS communities and to improve understanding of model biases by enabling a more comprehensive evaluation against a variety of surface, sonde, aircraft, and satellite observational datasets, which all have different temporal, vertical, horizontal, and chemical strengths.

Evaluating the Local Ensemble Transform Kalman Filter for the Ocean in JEDI

Laura Slivinski*, Anna Shlyaeva, Travis Sluka and Guillaume Vernieres

Currently, NOAA's operational forecast system uses an atmospheric model that is uncoupled to the ocean. The next operational version of NOAA's Global Data Assimilation System will be upgraded to use the coupled ocean-atmosphere model of the Unified Forecast System (UFS). The data assimilation will be weakly coupled: ocean observations will only update the ocean, and atmospheric observations will only update the atmosphere. The deterministic member will be updated using a hybrid 4DEnVar in the atmosphere and 3DVar-FGAT in the ocean. Ensembles are updated with the gain-form ensemble transform Kalman filter (GETKF) in the atmosphere. In the ocean, the ensemble is only recentered around the deterministic member (without an ensemble update), but with an option of controlling ensemble spread. A further upgrade to this coupled DA system would involve updating the ocean ensemble. This would also allow for a future strongly-coupled ensemble DA system. To that end, the JEDI local ensemble transform Kalman filter (LETKF) is being developed and tested for the marine component of the UFS. Results evaluating the ocean LETKF in the context of coupled ocean-atmosphere DA will be presented, including discussions of ensemble dispersion. Progress and challenges will be discussed.

CNN-LSTM Based PM2.5 Forecasting and Anomaly Detection in Delhi, India

Sonia Sharma Banjade* and Sandip Rijal

Air quality forecasting is essential for protecting community health, yet predicting fine particulate matter (PM2.5) remains challenging due to its complex temporal patterns. In this study, we apply deep learning to daily PM2.5 observations from Delhi, India, to tackle two problems: short-range forecasting and anomaly detection.

We build a hybrid CNN-LSTM model that uses cyclical seasonal encodings, month, and day-of-year as sin/cos features to capture Delhi's strong winter pollution cycles, where PM2.5 levels are nearly three times higher than in summer. We train the model for both 1-step and 7-day ahead forecasting, and separately train an LSTM autoencoder to identify unusual pollution episodes without labeled data.

For 1-step forecasting, the model achieves R^2=0.75, on par with a persistence baseline, showing that purely data-driven models hit a ceiling without meteorological inputs like wind speed or boundary layer height. For 7-day forecasting, the model achieves a mean R^2=0.57, a meaningful result given that no weather data is used. The LSTM autoencoder detects pollution episodes averaging 188 μg/m³, which is 62% above the dataset mean of 116 μg/m³, and flags events that a standard statistical method misses entirely, suggesting the two approaches capture different types of anomalies.

These findings highlight both the potential and current limits of data-driven air quality modeling and motivate future work coupling deep learning with physics-based community modeling frameworks.

Implementation of a JEDI-based marine data assimilation system for GFSv17

John Steffen*, Guillaume Vernieres, Anna Shlyaeva, Andrew Eichmann and Mindo Choi

The major upgrade to the Global Forecast System Version 17 (GFSv17) includes a suite of fully-coupled, earth system modeling components for the atmosphere (FV3, 9 km/C1152), ocean (MOM6, 1/4°), sea-ice (CICE6, 1/4°), wave (WW3) and land (NOAH-MP) models at increased grid resolutions. Additionally, a JEDI/SOCA (Joint Effort for Data Assimilation Integration/Sea-Ice Ocean and Coupled Assimilation) 3DVAR-FGAT marine data assimilation (DA) scheme was developed to initialize the marine components to improve ocean and sea-ice initial states and forecasts. A 4D h(x) and 3D parametric background error with a steric height balance operator are utilized over a 6 hour DA window. The ocean is initialized by a 3D IAU in MOM6 and by direct insertion for CICE6. The DA system processes near real-time observations from VIIRS and AVHRRF L3 products for SST, AMSR2 for sea ice concentration, altimeters from Sentinel, Jason, SWOT, Cryosat-2, and SARAL for absolute dynamic topography (ADT), and in situ data from Argo profilers and surface drifters. Results show a marked improvement in the model SSTs and sea ice concentration in both RMSE and bias when compared to the operational GFSv16. A conservative quality control (QC) of Argo and ADT constrains the ocean subsurface. Future marine DA development at NOAA will focus on a hybrid EnVar approach, coupled DA, AI-based B-matrix and leveraging innovative observation platforms.

The Regional Wave Prediction System: Development and Preliminary Studies

Keston Smith*, Jessica Meixner, Saeideh Banihashemi and Ming Chen, Ali Salimi-Tarazouj

The Regional Wave Prediction System (RWPS) intends to advance operational wave products within the NWS. Built on WAVEWATCH III, the wave component used in the Unified Forecast System (UFS), RWPS will expand the coverage of the Nearshore Wave Prediction System (NWPS) to become a primary source of official wave forecasts for the NWS. While RWPS and NWPS are not traditional UFS models, RWPS will utilize the UFS weather model infrastructure, improvements to the wave component itself are applicable for other UFS applications and future versions will likely explore active coupling of currents and water levels.

This presentation will share the improvements for wave propagation as well as greater exploitation of unstructured grids to improve nearshore resolution, while balancing computational cost. These improvements in numerics and dynamics are illustrated in hindcast case studies for time periods with prominent swell events in the Pacific as well as recent Atlantic hurricanes.

Hydrometeor-Aware All-Sky Microwave Data Assimilation for Atmospheric River Prediction

Chong-Chi Tong*, Laura Slivinski and Sergey Frolov

Atmospheric Rivers (ARs) drive extreme precipitation along the U.S. West Coast, supplying critical water resources while posing major flood risks. Accurate AR prediction remains challenging, especially over oceanic regions with limited in situ observations. All-sky microwave radiances from instruments such as ATMS and AMSU-A provide important constraints on key atmospheric features (e.g., temperature and moisture) linked to AR development. However, hydrometeor variables, which strongly influence microwave radiative signals and AR microphysics, are not consistently updated within current operational data assimilation (DA) control vectors.

This study examines the role of hydrometeor-aware all-sky DA in AR prediction using a one-month cycling experiment within the NOAA GDAS hybrid 4DEnVar/EnKF system. Experiments for January 2023 are performed in a dropsonde-denial configuration with 6-hour cycling. Multiple configurations systematically varied the inclusion and updating of cloud and precipitating hydrometeors in the control vector to distinguish their respective impacts on analysis increments and short-range forecast evolution.

Evaluation considers AR-region statistics over oceanic domains, independent dropsonde verification, and global conventional observations. The findings highlight sensitivities related to hydrometeor representation in DA. Additional diagnostics and discussion will be presented at the conference.

Exploring Generative Methods to Inform Likelihoods for Bayesian Data Assimilation in Earth System Surrogate Models

Raelynn Wonnacott*, Jonathan Poterjoy and Michael Evans

Data assimilation combines numerical models with observations to solve the Bayesian inverse problem of inferring plausible states of a system evolving over time. Particle filters are ensemble methods that target the Bayesian posterior directly, without Gaussian approximations. Therefore, they contain practical advantages for Earth System Models (ESMs), which are often characterized by non-Gaussian errors in the model state and observation process. These advantages of particle filters are expected to become increasingly realizable as ML-based ESMs promise to enable large ensembles.

Nevertheless, particle filters remain difficult to apply in high-dimensional systems without regularization mechanisms analogous to covariance localization and inflation in Kalman filters. Likewise, replacing physics-based models with ML-based surrogate models introduces new challenges for traditional data assimilation methods that require a forward model to compare model state to observations. Motivated by these challenges, we explore generative modeling to address the central component of Bayesian filtering where both issues converge—the likelihood. In many problems, the scientifically relevant state is of a large, well-resolved scale, but observations are affected by fast varying small scale states. We consider running the particle filter in a reduced space using a cheaper forecast model and learning a marginal likelihood distribution that integrates out small-scale variability. The marginal likelihood allows us to assimilate observations of unresolved scales and, by a Rao–Blackwellization effect, can reduce variance in the importance weights. We investigate deterministic and generative models to incorporate data-driven likelihood functions into particle filters, applying this framework to a low-dimensional dynamical model that mimics multi-scale atmospheric interactions.

A Physics-Based Multi-Decadal Compound Flood Reanalysis for the U.S. East and Gulf Coasts

Fei Ye*, H.C. Yu, Y.J. Zhang, Hyungju Yoo, Saeed Moghimi, Gregory Seroka, Zizang Yang, and Edward Myers

Coastal flood hazard assessment is constrained by the lack of spatially continuous datasets that resolve interactions among storm surge, river discharge, rainfall, and tides. Existing approaches often rely on statistical or event-based methods and provide limited capability to represent the coupled processes underlying compound flooding across the coastal transition zone. This study develops a physics-based reanalysis framework to enable process-resolving characterization of flooding across inland–coastal systems.

A 26-year (2000–2025) reanalysis was conducted using the STOFS-3D system, which integrates coastal ocean dynamics, river networks, and floodplains within a unified unstructured mesh. Hydrological inputs from the National Water Model are incorporated with gauge-based corrections. Water levels are further improved through Ensemble Optimal Interpolation (EnOI) data assimilation of the NOAA CO-OPS observations, enabling efficient and dynamically consistent simulations. The EnOI is implemented via ESMF.

The reanalysis produces a spatially continuous water level dataset with strong agreement to coastal observations across the full period, capturing tidal and non-tidal variability as well as major extreme events. Ongoing analyses explore physically derived metrics to characterize long-term mean conditions, variability, and upper-envelope water levels across both persistently wet and intermittently inundated regions.

This work establishes a new paradigm for physics-based characterization of compound flooding at regional scales, bridging operational forecasting and long-term hazard assessment. The resulting dataset supports return-period estimation, infrastructure planning, and future climate applications, while providing a foundation for hybrid and data-driven modeling approaches.

Adapting Dynamics and Physics of UFSWAM for the Storm-Time Space Weather Predictions

Valery A. Yudin*, Irfan S. Azeem, Svetlana. I. Karol, and Timothy J. Fuller-Rowell

During recent (2024-2026) extreme and severe geomagnetic storms the striking enhancements of the thermospheric air density, temperature and winds demonstrated the need to adapt the UFSWAM formulation when the storm-driven heating raises the model top height from ~600-750 km to ~1300-1700 km under the fixed top lid pressure. The key modifications of the UFSWAM dynamics and physics to reproduce the observed storm-time anomalies of dynamics, air density, and composition will be discussed. Among them are: (1) the deep atmosphere formulation of dynamical core, and conservation of energy under the external forcing; (2) extension of the UFSWAM tracer list to simulate the storm-time heating/cooling and air density in the thermosphere; and (3) data-driven dissipation of the electromagnetic energy during storms. The evaluation of UFSWAM predictions will be highlighted for three storms of 2024.

Development of the CCPP-Based GEFS-Aerosols Component in the Unified Forecast System for Subseasonal Prediction (UFS-Chem v1.0)

Li Zhang*, Haiqin Li, Georg A. Grell, Partha S. Bhattacharjee, Gonzalo A. Ferrada, Benjamin W. Green, Shan Sun, Ligia Bernardet, Anders Jensen, Barry Baker, Li Pan, Jian He, Jordan Schnell, Ravan Ahmadov, Samuel Trahan, Dustin Swales, Anning Cheng, Fanglin Yang, Rebecca H. Schwantes, Brian C. McDonald, Dominikus Heinzeller and Shobha Kondragunta

The Global Ensemble Forecast System (GEFS) version 12 has been operational at the National Centers for Environmental Prediction since September 2020, with GEFS-Aerosols serving as its global aerosol forecasting member. In 2023, GEFS-Aerosols was upgraded to version 12.3 in operations, incorporating improvements to wet deposition, anthropogenic and biomass burning emissions, aerosol optical depth calculations, and the FENGSHA dust emission scheme. While GEFS-Aerosols provides valuable operational aerosol forecasts on medium-range timescales, its one-way coupling strategy restricts aerosol–atmosphere interactions to prescribed climatological fields, thereby precluding fully interactive aerosol feedbacks for extended-range prediction. To overcome this limitation, we develop the Unified Forecast System coupled with Chemistry (UFS-Chem), which incorporates the Configurable Atmospheric Chemistry (CATChem) library and modeling component to include both aerosol and gas-phase chemistry schemes. In the current UFS-Chem configuration, the aerosol component is based on the operational GEFS-Aerosols v12.3 and implemented using the Common Community Physics Package (CCPP) framework. Relative to GEFS-Aerosols v12.3, UFS-Chem incorporates several enhancements, including inline aerosol radiative feedback, inline large-scale wet deposition, and updated FENGSHA dust scheme. In particular, the integration of inline large-scale wet deposition and aerosol indirect effects through the Thompson aerosol-aware microphysics scheme improves the representation of aerosol–cloud interactions and supports weather and subseasonal-to-seasonal forecasting when fully coupled with ocean, sea ice, land, and wave components. The performance of UFS-Chem in weather and subseasonal prediction is evaluated against reanalysis data, ground-based measurements, and satellite observations. The relative impacts of aerosol radiative and indirect feedback on weather and subseasonal predictions are investigated and quantified, advancing our understanding of how aerosol–radiation and aerosol–cloud interactions influence forecast skill across timescales.

Machine Learning

Evaluation of UFS-DeepAQM: an deep learning extension of the Unified Forecast System Air Quality Model

Wei-Ting Hung*, Jia Xing, Youhua Tong, Barry Baker, Patrick C. Campbell and Zachary Moon

Air quality forecasting (AQF) is important for foreseeing and protecting the public from potential hazardous air pollutants from nature and/or anthropogenic sources such as wildfires and industries. To accurately predict air quality, AQF models simulate various atmospheric physical and chemical

processes, including but not limited to the transport and dispersion of air pollutants as well as the chemical reactions among air compositions. Explicit simulations often require substantial time and computational resources, which limit the capabilities of AQF models for near-real-time applications. Hence a deep learning AQF model, UFS-DeepAQM, is developed to emulate the atmospheric physical and chemical processes described in the NOAA Unified Forecast System Air Quality Model (UFS-AQM). Current UFS-DeepAQM can provide hourly surface PM2.5 and ozone (O3) predictions. The UFS-DeepAQM uses the similar inputs as UFS-AQM, including its meteorological fields, geographic features, emissions, and initial and boundary conditions. We compare the performance of UFS-DeepAQM with the targeted UFS-AQM outputs for a full year of 2024. Preliminary results show that UFS-DeepAQM can reproduce UFS-AQM O3 results with R2 of 0.7 - 0.9 and RMSE of 2 - 5 ppb for 72 h forecast. As for PM2.5, although UFS-DeepAQM reproduces UFS-AQM results in general with R2 of 0.6 - 0.9 and RMSE of 3 - 6 µg m-3, UFS-DeepAQM could miss the elevated PM2.5 concentrations due to severe events such as wildfires. Considering its high computational efficiency (72 h forecast takes ~ 8 min with 2 cores on 1 node on NOAA Gaea C6), UFS-DeepAQM can serve as an ensemble member of the existing UFS-AQM operational framework. It can also be expanded to high- resolution simulation or global model emulator (under development). Coupling with other techniques, such as data fusion and emission inversions, this ML tool can be made to improve accuracy of air quality prediction.

High Resolution (1-Km) Simulation Of Leaf Area Index At Sub-Seasonal To Seasonal Scale For Earth System Model Using Deep Learning

Lixin Lu*, Zhixing Ruan and Yu Zhang

Modeling vegetation phenology within Earth system models remains challenging due to its interannual variability, spatial heterogeneity, and complex ecosystem dynamics. Current dynamic vegetation (DV) models represent key processes, including vegetation growth and senescence, carbon and nutrient cycling, disturbance-climate interactions, but their complexity limits integration into operational prediction systems. Consequently, many models prescribe Leaf Area Index (LAI) distributions using a 20-year averaged climatology derived from the Normalized Difference Vegetation Index, restricting responsiveness to sub-seasonal to seasonal and interannual variability.

We develop ConvLSTM-LAI, a deep learning model that integrates Convolutional Neural Network (CNN) and Long Short-Term Memory (LSTM) architectures to predict LAI at regional scale. The model is trained on long-term daily precipitation and temperatures from the NCDC Summary of the Day station observations, and MODIS (MCD15A3H), interpolated to daily resolution. Inputs and targets are provided at 1-km spatial resolution over the south-central United States.

Trained on 2001 to 2020 data and evaluated for 2021 to 2023, ConvLSTM-LAI achieves strong predictive skill. At a 30-day lead time, the domain-averaged LAI error is less than 0.11 with a correlation coefficient of 0.86 relative to MODIS observations, capturing both seasonal and interannual vegetation dynamics. Performance is robust across different land cover types, including grasslands, savannas, shrublands, croplands, and mixed forests. Skill degrades with increasing lead time, as indicated by rising RMSE and decreasing correlation, reflecting the intrinsic limits of extended-range vegetation prediction.

Physics-Based Machine Learning for Compound Flooding Forecast

Zhengui Wang*, Joseph Zhang, Fei Ye, Dan Yu and Saeed Moghimi

STOFS-3D-Atlantic (Surge and Tide Operational Forecast System – 3D Atlantic) is an advanced coastal ocean forecasting system jointly developed by the Virginia Institute of Marine Science (VIMS) and NOAA. Built on SCHISM, it predicts coastal water levels, flooding, and ocean state across the northwestern Atlantic with very high spatial resolution in the coastal transition zone. It incorporates the major physical drivers of coastal flooding, including river discharge, precipitation, tides, and storm surge. It has been in operation since 2023 on NOAA's WCOSS2 supercomputers. The resource usage is intensive for this large forecast, with over 3 million nodes in the horizontal dimension and a vertically varying discretization ranging from 2 to 49 layers (mean ≈ 10).

To improve efficiency and get ready for ensemble forecasts, here we develop a machine learning–based compound flooding model (STOFS3D_ML), trained on 26 years of STOFS-3D-Atlantic re-analysis. STOFS3D_ML is computationally much more efficient, but importantly, unlike many purely data-driven approaches, it is built on a physics-informed foundation by retaining the dominant modes of water-level variability across the full model domain. We evaluate the performance of STOFS3D_ML through direct comparison with STOFS-3D-Atlantic outputs to assess its predictive skill. In addition, we investigate the model's sensitivity to variations in external forcing to objectively assess its response against SCHISM's.

Accelerating UFS Model Development Cycles Through a Process-Oriented, Machine Learning–Based Framework

Augustin Vintzileos*, Benjamin Cash and Jim Kinter

The calibration of physical parameterizations in Earth system models remains a challenging and inefficient process, as improvements in one performance metric often come at the expense of others, reflecting fundamental limitations of current global tuning approaches. In particular, interactions among parameterizations necessitate coordinated adjustments, rendering isolated tuning ineffective.

Here, we introduce PROMETHEUS (PROcess-oriented Model dEvelopment Through dEcision tRee HeUristic System), a process-consistent calibration framework that uses decision trees to recursively partition the space of atmospheric columns into Column States, i.e., subsets of similar columnar environments in which a target variable or forecast error exhibits distinct statistical behavior associated with specific physical regimes (e.g., shallow versus deep convection). This enables targeted, physical regime-dependent optimization of physical parameterizations, for example, using Single Column Models.

We demonstrate the physical process analysis capabilities of PROMETHEUS by applying it to tropical convection using hourly ERA5 data for 2012 over 30°S–30°N. The target variable is the tendency of outgoing longwave radiation (OLR), a proxy for convective activity. The analysis uses 48 features, including vertical profiles of thermodynamic and dynamic variables, surface conditions, turbulent fluxes, and vertically integrated diagnostics such as humidity divergence and moist static energy. Two decision trees are trained separately for atmospheric columns over land and ocean. Their comparison reveals the physical mechanisms underlying stronger convective activity over land relative to the ocean. We finally discuss applications to understanding forecast errors and to conditional post-processing.

High-Performance Computing (HPC)/Cloud

Accelerating UFS Development and R2O: NOAA's Research and Development High Performance Computing System (RDHPCS)

Tracey Wilson and Raghu Reddy*

NOAA's Research and Development High Performance Computing System (RDHPCS) program provides powerful, responsive, easily available High Performance Computing to users throughout the NOAA R&D community – across all line offices and including NOAA's partners in academia and industry - from early-stage research through the entire Research to Operations (R2O) lifecycle. RDHPCS maintains a hybrid HPC architecture comprising both on-premises systems and cloud services from multiple CSPs. The systems include both CPUs and GPUs as well as large-scale storage. RDHPCS provides access to a wide range of NOAA data, supported by tools for data movement. This talk will address 1) an overview of the current RDHPCS systems, services, and user experience; 2) the evolving user needs that RDHPCS must address; 3) challenges to providing HPC services in today's computing and fiscal environment; and 4) show a roadmap for future RDHPCS systems and services. We will also address how user needs change over the R2O lifecycle for a complex modeling system like the UFS from an idea to an implemented operational capability and the impact on the HPC services NOAA R&D requires. We will also present thoughts on where HPC technology is heading over the next 3 to 5 years and identify emerging computing technologies that may be beneficial to NOAA in the future.

Accelerating Rapid Innovation: EPIC's Pre-Configured Public Cloud HPC environments

Kristopher Booker*, Cameron Book, Natalie Perlin and Keven Blackman

The NOAA Earth Prediction Innovation Center (EPIC) program has enabled an accelerated pace of numerical weather prediction innovation through their collaboration with government, academic, and private sectors. Utilizing a common modeling framework, the UFS (Unified Forecast System), has allowed once siloed atmospheric modeling groups to effectively collaborate. However, the greater challenge has come from constraints that kept organizations from accessing high performance computing (HPC) platforms capable of running new modeling innovations at scale. Although the EPIC program has established insights and dashboards there is still an underserved need to allow contributors direct access to HPC platforms with all of the required software dependencies. Perhaps the largest barrier to entry for the UFS community is the ability to easily build/run the UFS weather model. Containerization through the use of Apptainer/Singularity solves some of these problems, but the abstraction can still pose unique technical and performance challenges. EPIC has attempted to solve this dilemma by utilizing infrastructure as code tools/processes to allow consistent installation of the basic UFS software environment. These processes coupled with popular public cloud HPC platforms (AWS Parallelcluster) enable UFS community users to rapidly deploy a turnkey public cloud HPC environment. EPIC has demonstrated the success of this approach to provide UFS community workshop users a ready made HPC environment at various training events. Participants have frequently requested access to these software tools and configurations. This approach should prove valuable to those UFS community members without access to traditional academic or government HPC resources. Our presentation will discuss how we used infrastructure as code tools to build an HPC environment on the public cloud capable of running the UFS weather model (UFS-WM). This presentation should serve as an example to those without access to a traditional government or academic HPC environment, how to utilize the public cloud for working with the UFS-WM.

Portable and Reproducible Earth System Modeling: Containerizing ESM4.5 Case Study

Thomas Robinson*, Dana Singh, Subil Abraham, Elijah Maccarthy, Eric Stofferahn and Uriel Ramirez

Scientific journals increasingly mandate that model description papers include publicly shared code. Historically, sharing source code via repositories without standardized build or runtime support has hindered reproducibility, making it difficult for external researchers to initialize complex simulations. This work addresses these barriers by containerizing the Geophysical Fluid Dynamics Laboratory's (GFDL) Earth System Model version 4.5 (ESM4.5), ensuring that models used for long-term production simulations are accessible and executable across diverse computing environments. We developed a workflow that builds on the High-Performance Computing Portable Containers Model Environments (HPC-ME) containers. Dependencies within the container are managed via Spack. The workflow employs YAML-defined configurations for code and platforms, leveraging Podman on NOAA Research and Development High Performance Computing Systems (RDHPCS) for two-stage builds to minimize runtime overhead. The workflow converts OCI images to SIF format for running with apptainer. A critical component includes a specialized module file for the Gaea supercomputer that configures MPI-bindings to ensure communication performance parity with native executables during heterogeneous Slurm executions. We successfully deployed the first model container benchmark for ESM4.5 on Gaea. The workflow captures exact code commits and environment states, facilitating rapid development and precise version tagging. By utilizing MPI-bind, we

overcame container-native communication latencies, achieving high-performance execution within a portable framework. Transitioning to container-based modeling makes high-performance simulations more accessible. This approach bridges the gap between internal production and public transparency, enabling the broader scientific community to reproduce results and accelerate discovery in earth system dynamics.

Build and run UFS Weather Model tests using containerized spack-stack libraries with GNU or Intel compilers on HPC platforms or cloud environments

Natalie Perlin*, Eddie Snyder and Cameron Book

The development of the Unified Forecast System (UFS) is a community-driven effort aimed at expanding participation in model development, evaluation, and feedback. Broad community engagement is facilitated by the use of open-source software stacks that can be deployed across diverse computing environments, including compilers and Message-Passing Interface (MPI) libraries.

To support this effort, the NOAA Earth Prediction Innovation Center (NOAA-EPIC) has developed container images based on GNU compilers and the Open MPI library, including pre-built spack-stack dependencies required to run the UFS Weather Model and its applications. In parallel, container environments using Intel compilers and MPI libraries have also been created. EPIC-supported NOAA RDHPC systems and NOAA Cloud platforms have these containers prestaged in a common location, and include Intel compilers and software. For broader community use, an alternative approach is offered based on container images with no preinstalled licensed software, which is further added to the singularity container using a few extra steps involving building container sandboxes.

This presentation will demonstrate the execution of UFS Weather Model regression tests using both GNU-based and Intel-based container workflows. The GNU-based containers are highly portable and can be deployed across a wide range of high-performance computing (HPC) systems and cloud environments. In addition, portable strategies for using Intel-based containers are being suggested to enable broader adoption by the community.

Leveraging Cloud-Integrated CI/CD and Containerization to Modernize NOAA's Workflows

Lauren Chilutti*, Thomas Robinson, Ryan Mulhall, Ciheim Brown, Dana Singh, Uriel Ramirez and Vithika Shah

Modernizing NOAA's modeling infrastructure requires transitioning toward automated practices and streamlined workflows. Over the past 6 years, the Geophysical Fluid Dynamics Laboratory (GFDL) has successfully transitioned from a completely manual testing and deployment procedure to automated software testing and deployment workflows.

We developed a Continuous Integration and Continuous Deployment (CI/CD) infrastructure utilizing GitHub Actions, GitLab Pipelines, and automated scheduling. We standardized containers for supporting CI/CD pipelines in order to create parity between the different platforms. The testing workflows of publicly developed codes such as FV3 and FMS use GitHub actions, and self-hosted runners are employed on the NOAA RDHPCS cloud to address security concerns related to community contributions. Automated testing of internally developed models such as AM5 uses GFDL's GitLab runners on the NOAA RDHPCS Gaea supercomputer. Deployments of these models' development releases are triggered from GitLab pipelines and production runs are triggered by automated scheduling on RDHPCS. We deploy public facing packages such as FRE-CLI with GitHub actions.

These DevOps pipelines have successfully automated the build-test-release-deploy cycle across GFDL's modeling suite. Results include the automated deployment of AM5 production runs and the continuous verification of FV3 pull requests in cloud environments. By using containers, we achieved consistent performance across local and cloud platforms, significantly reducing the time from code commit to validated release. By leveraging cloud resources for public code vetting and automating deployments, we have enhanced the reliability and efficiency of GFDL's modeling suite.

Developing WAVEWATCH IV™ to enable exascale wind wave modeling

Hendrik L. Tolman*, NOAA/NWS/EMC (OMD)

For the record, this abstract is a minor edit of an abstract of a paper presented at the WISE meeting in Venice Italy, on May 12, 2026. This paper will present the more recent result, and unlike the WISE meeting, will produce an online text or product.

WAVEWATCH has been around for nearly 40 years, with a first full publication in Tolman (1989). Many of the core architecture elements of the present WAVEWATCH III (WW3) model date back two decades with respect to the parallel architecture, the data structures, and general software architecture. A recent BAMS paper on lessons learned from WW3 (Tolman 2025a), a position paper on software modernization for the Unified Forecast System (Tolman 2025b), and a documentation and assessment of wave modeling at the National Weather Service (Tolman and Meixner 2025) all point toward the need for software modernization for WAVEWATCH. Based on these assessments, NOAA / NWS has committed to a systematic, phased modernization of WAVEWATCH. In the first phase of this project, documented in Tolman (2025c), a team of WAVEWATCH developers led by NOAA has decided that this project will take the form of the development of WAVEWATCH IV, moving away from Fortran, and defining tentative subsequent phases of development in this project. At the time of submitting this abstract, Phase II is completed, and its progress report Tolman (2026) is expected to be published in May. Coding of the WW4 core capabilities has started as part of Phase III of the project. Up to date code can be found at the GitHub repository (https://github.com/NOAA-EMC/WW4) with embedded wiki pages where community governance of the code is defined and managed.

The presentation will highlight up-to-date results of the project, including but not limited to governance, IP protection, model features, and code status.

Planning Committee

A huge thank you and congratulations to the UIFCW26 Planning Committee!

Aaron Jones, Aaron Pratt, Alison Gregory, Ben Cash, Ben Green, Ben Johnson, Ben Woods, David Huber, Eric Aligo, Farida Adimi, Gillian Petro, Hendrik Tolman, Iman Gohari, Jan Ising, Jeff Beck, Jenna Holbrook, Jennifer “Jen” Vogt, Jessica Wheeler, Jimmy Dudhia, Jorge Bravo, Linden Wolf, Michael, Benneh, Mike Ek, Natalie Perlin, Shira Francis, Tracy Fanara, Segayle Thompson, Zach Shrader, and Keven Blackman