Open Innovation and the Case for Community Model Development

Neil A. Jacobs

ABSTRACT: Despite having the largest associated research community and a rapidly growing private sector, the lack of a well-coordinated national research and development effort for U.S. numerical weather prediction continues to impede our ability to utilize more of the scientific and technical capacity of the nation more efficiently. Over the last few years, considerable progress has been made toward developing a community-friendly Unified Forecast System (UFS) by embracing an open innovation approach that is mutually beneficial to the public, private, and academic sectors. Once fully implemented, the UFS has the potential to catalyze a significant increase in the efficacy of our nation's weather, water, and climate science and prediction.

KEYWORDS: Education; Numerical weather prediction/forecasting; Policy

<https://doi.org/10.1175/BAMS-D-21-0030.1>

Corresponding author: Neil Jacobs, isentropicpv@gmail.com

In final form 22 June 2021

©2021 American Meteorological Society

For information regarding reuse of this content and general copyright information, consult the [AMS Copyright Policy](#).

AFFILIATIONS: Jacobs—Durham, North Carolina

Numerical weather prediction (NWP) has been a cornerstone of weather forecasting for decades, and with advancements in scientific understanding of Earth systems, combined with rapidly expanding access to greater computational resources, our ability to use quantitative methods to predict the future state of Earth's atmosphere and oceans is entering an era of unconstrained innovation. As we become more technologically advanced, and areas prone to hazardous weather events see a rising surge in urbanization, dependency on accurate forecasts will drive a significant increase in the value of NWP, including many socioeconomic aspects of society.

From a scientific perspective, problems such as optimizing the initial state, improving numerical models, and enhancing the postprocessing of model output should have relatively straightforward solutions. Unfortunately, when it comes to management, there are many reasons why various sectors within the NWP community have struggled to work together, and most of them come down to a lack of access, coordination, and resources (Mass 2006). While the level of outreach from the National Weather Service (NWS) is greater now than ever before, the lack of a well-coordinated community-wide research and development effort is restricting the ability to use existing resources more efficiently. This lack of coordination arises from limited and unequal access to a user-friendly global modeling system that is used for both research, development and operations, as well as other resources needed to support that including meteorological observations, a common data assimilation platform, and access to high-performance computing (HPC).

Coordination is essential for success

To address some of these obstacles, the National Oceanic and Atmospheric Administration (NOAA) and the National Center for Atmospheric Research (NCAR) signed a Memorandum of Agreement in 2019 that establishes a new partnership to design a common modeling infrastructure. This common framework, which I discuss in more detail below, will streamline the process and give researchers and forecasters across the global weather enterprise the same tools to accelerate model development. To maximize the benefit, engagement between NOAA and the rest of the NWP community is essential whether managed through relationships with individual participants or through a consortium such as the University Corporation for Atmospheric Research (UCAR).

Additionally, in response to the Weather Research and Forecasting Innovation Act of 2017, the Office of Science and Technology Policy (OSTP) and NOAA created the Interagency Council for Advancing Meteorological Services (ICAMS) in August of 2020.¹ This was the most significant restructuring of coordination efforts among federal agencies since the Federal Committee for Meteorological Services and Supporting Research was established in 1964 (Harrison 1988). While the streamlined interaction among federal agencies facilitated through the four ICAMS committees (Research and Innovation; Observational Systems; Cyber, Facilities and Infrastructure; and Services) will more efficiently leverage shared resources and reduce duplication of effort, there are still more technical challenges to overcome to optimize our collective NWP capabilities.

¹ www.icams-portal.gov/resources/icams/signed_icams_charter.pdf

Barriers impeding broader collaboration

Historically, the NWP community has been an integral part of the development and use of limited-area mesoscale models such as the Weather Research and Forecasting (WRF) Model and its predecessors; however, these models require a global model for lateral boundary conditions, and in some cases, initial conditions. Given that dependency and the fact that global models provide valuable forecast guidance in their own right, it is worth asking why the broader community has not been more engaged in the development and operational use

of NOAA's Global Forecasting System (GFS). This is especially true given that both models share many of the same scientific and technical challenges, such as extracting more value from existing observations, improving model physics, or developing novel approaches to utilize exascale computing. The answer is actually quite simple: the GFS was its own inherent barrier.

Until recently, the GFS was a complex maze of legacy code written for specific operating systems and hardware architectures and never designed to be externally used or supported. When the GFS was originally written in the late 1970s, engaging the broader research community was a foreign concept because nearly all of the work on operational NWP models was being done within government laboratories. As the NWP community expanded outside of government laboratories, both in reach and expertise, interest in working on the GFS grew. However, the code was specific to NOAA-owned HPC, so it would not run on other hardware. That limitation combined with the policy that requires a security clearance to work on NOAA HPC created significant barriers to the broader NWP community's involvement in the GFS development. Other barriers, which likely stemmed from these long-standing technical and policy challenges, included a lack of user support and documentation, no community-wide strategic plan, and limited access to real-time observations. These barriers played a significant role in separating NOAA from the broader academic and research community. The result was the evolution of parallel and uncoordinated development programs that exacerbated the already existing cultural divide because both sides, despite having different motivations and missions, viewed the other as a largely duplicative effort competing for limited federal funding.

The origin of a solution

In a covert development plot worthy of the big screen, it was the banking industry, which was trading in specific commodity markets that fluctuated based on the 1200 UTC GFS forecast, that first successfully ran the entire modeling system, initialized from observations, outside of the National Centers for Environmental Prediction (NCEP) in the early 2000s. This crucial milestone was originally known only to a small group I was working with at the time. The start-up company, AirDat, which was later acquired by Panasonic, was primarily known for the deployment and management of the Tropospheric Airborne Meteorological Data Reporting (TAMDAR) system on commercial airlines (e.g., Daniels et al. 2004; Moninger et al. 2010).

Early versions of the GFS were hardcoded for the IBM AIX operating system, and from 2005 until 2013, my team was using various IBM Power series HPC systems that were identical to NCEP's operational systems. Once NCEP transitioned to Linux in 2013, the financial barrier to entry for industry was greatly lowered because Linux was an open-source operating system that could run on almost any hardware. Unfortunately, whether it was a government laboratory or industry development program, on-premise HPC, which is inherently constrained by a finite number of nodes, would still be the primary bottleneck in the research-to-operations (R2O) pipeline because it is not dynamically scalable to accommodate surges in workload demand.

Over the last five years, HPC in the cloud has not only become a viable option, but its flexibility, costing, and scalability have made it the preferred option for many NWP applications, particularly for the broader community (Ahmad et al. 2020). Starting around 2015, after benchmark runtime tests for the GFS on cloud HPC were showing superior performance, my team transitioned to a cloud-only compute solution. When the research community learned of our capability, we began to receive proposals to test different ideas to improve the model. During the review process, I would always ask if they had discussed their hypothesis with NCEP. The answer was always yes, but the queue to test any new idea was years long and limited by lack of compute access.

Cloud HPC was appealing because the number of parallel tests that could be conducted was not constrained by the limited capacity of traditional on-premises HPC. To expand the

number of parallel tests, we would simply contact the Cloud Service Provider (CSP) and purchase time on additional nodes. Most CSPs charge by node-hour, so whether the tests are run sequentially or in parallel, the total cost is the same. However, the advantage of conducting numerous experiments simultaneously is the work gets completed considerably faster.

Having the GFS running on a scalable cloud platform allowed us to conduct many experiments ranging from testing the impact of new satellite data to accelerating the radiative transfer model to turning off the scheme that relocates a tropical cyclone vortex to the advisory position. Many of the ideas that were tested actually came from NOAA's own scientists. The company I was working for at the time, Panasonic, was supporting faculty, postdocs, and students at multiple universities to help with model development. I would often share the results with Bill Lapenta, the former NCEP director, when I would visit NCEP or see him at a conference, and we would discuss what a game-changing capability this would be to accelerate NOAA R2O.

Bill was responsible for the formation of a subcommittee of the UCAR Community Advisory Committee for NCEP (UCACN) called the UCACN Model Advisory Committee (UMAC). The UMAC was charged with providing a comprehensive technical review of the NCEP production suite strategy. Several of the findings in the 2015 report,² from streamlining the production suite to better leveraging the external community, ultimately inspired the creation of both the Unified Forecast System (UFS) and the Earth Prediction Innovation Center (EPIC). At an AMS meeting in 2016, Bill, Fuqing Zhang, and I sketched out a way to transition the UFS to a community model, manage the parallel testing of external contributions, and subsequently integrate them back into the NCEP pipeline to operations (Fig. 1).

Two years later, when I found myself at the helm of NOAA, Bill and I realized we had a unique opportunity to make this a reality. NCEP was already following a Strategic Implementation Plan it had formulated with help from the broader community for simplifying and transitioning the existing production suite quagmire to a more streamlined UFS modeling system with various configurations (applications) for different NWP purposes.³ The development of the UFS was the perfect opportunity to refactor the code into a well-documented, user-friendly community model that would run on various external platforms.

Devising a community modeling program that results in a successful public-private-academic partnership, where all stakeholders are contributing and benefitting, is the ultimate goal. Private companies that run models will have greater opportunity for customization, while the back-end value-add companies will have more accurate and reliable products. The CSPs will have a rapidly growing customer base, while platform-agnostic software keeps pricing competitive. Further, the UFS can be used for teaching in the classroom, as well as a tool for basic research and NWP experiments, where subsequent improvements and gained knowledge would result in scientific journal publications.

NOAA can tap into a wealth of community expertise to improve the operational UFS, while working with the community to objectively validate and integrate contributions. Efficiencies gained through eliminating additional work required to translate innovations from one modeling system to another would greatly accelerate scientific progress. The overall collaborative process will ultimately benefit the operational meteorologist by leveraging this expertise to advance numerical guidance skill, thereby improving the accuracy and confidence of the forecaster.

Crowdsourcing model development is a radical paradigm shift for NOAA, and one that would require overcoming technical challenges, procurement hurdles, and cultural resistance; however, open innovation, which has been adopted by a broad swath of private industry in almost every market, is not a new concept (Chesbrough and Bogers 2014).

² www.ncep.noaa.gov/director/ucar_reports/ucacn_20151207/UMAC_Final_Report_20151207-v14.pdf

³ www.weather.gov/media/sti/nggps/UFS%20SIP%20FY19-21_20181129.pdf

Fig. 1. The original sketch of what would ultimately become the UFS community modeling process of EPIC.

Chesbrough et al. (2006) define open innovation as a distributed process that manages the inflow and outflow of knowledge to simultaneously accelerate innovation internally, while expanding the use of it externally. For open innovation to be successful, the GFS code needed to be ported over to community-accessible hardware and made more user friendly. This is exactly what happened, and in 2020, the UFS community development team released the Short-Range, Medium-Range (MR), and Subseasonal-to-Seasonal (S2S) Applications on the code-hosting platform GitHub.⁴

Rollout of the community UFS

The UFS MR Weather Application consists of the entire model workflow, including the build system and Finite-Volume Cubed-Sphere (FV3; Harris and Lin 2013)-based GFSv15, as well as support documentation.⁵ Subsequent releases will be the same versions of code used in production, thereby ensuring that the research community has access to the same code used operationally at NCEP. The amount of traffic on GitHub has ramped up significantly since the initial release, and the co-development across institutions using the common repository is beginning to work as envisioned.

The backbone of the UFS is the NOAA Environmental Modeling System, which employs a community-developed and community-governed software process for building and coupling model components called the Earth System Modeling Framework (ESMF; Hill et al. 2004). The National Unified Operational Prediction Capability (NUOPC) Layer is a set of extensions to

⁴ <https://github.com/ufs-community>

⁵ <https://github.com/ufs-community/ufs-mrweather-app/wiki>

ESMF that increases component interoperability and provides enhanced architectural options to simplify model coupling (Theurich et al. 2016). The Common Community Physics Package contains the Noah land surface model and interoperable physics.⁶ Various mediators used for coupling are now available on GitHub, including the Community Mediator for Earth Prediction Systems (CMEPS), which is a NUOPC-compliant coupler developed through the NOAA–NCAR collaboration. These interface with models such as the Modular Ocean Model (MOM6; Adcroft et al. 2019), WAVEWATCH III (Tolman et al. 2011), and the Community Sea Ice Code (CICE), which are further along in the S2S application development, but will ultimately work with the full UFS (Stajner et al. 2020). The CMEPS framework is being designed in such a way that even different dynamic cores, such as the Model for Prediction Across Scales (MPAS; Skamarock et al. 2012) and FV3, can be interchangeable. The workflow also includes the Unified Post Processor, which interpolates native output to a variety of grids and file formats for use in visualization and forecast product generation.

All of the software has extensive documentation on both the UFS website and the GitHub wiki.⁷ There is also a community discussion forum for posting and answering questions, accessing the latest releases, and reporting bugs. The wiki includes a Getting Started section that leads users through a quick example of how to download, compile, and run the UFS MR Weather Application.

A fun exercise called the Graduate Student Test is included in this tutorial as a way for students, postdocs, and others to provide feedback on the usability of the code by filling out a questionnaire about their experience.⁸ As the UFS wiki says, you do not have to be a graduate student to take the test. I took the test and shared my experience on various weather forums and listservs. Admittedly, my coding is a little rusty these days, but I was able to download, compile, and run the entire UFS workflow on a spare MacBook (not even a MacBook Pro) with only 8 GB of memory. The install and build took less than 2 h, including downloading the code. I then ran the full FV3-GFS at C96 (100 km) resolution out to 12 h, which took about 1 h. Not long after posting my experience, I received dozens of emails from people ranging from tenured faculty to high school students who successfully ran the UFS on platforms ranging from commercial cloud HPC to video game consoles. This is the essence of the open innovation philosophy that will accelerate development of a true community-based modeling system.

Many in the NWP field, myself included, believe that the largest gains in forecast skill will arise from advancements in data assimilation techniques, so ensuring this critical component of the GFS is part of the UFS is essential. The Joint Center for Satellite Data Assimilation (JCSDA), which has been supporting the Community Radiative Transfer Model (CRTM), recently issued the first public release of the Joint Effort for Data assimilation Integration (JEDI) system.⁹ The JEDI-FV3 release, which is linked to the UFS in GitHub, includes generic observation operators provided by the Unified Forward Operator (UFO), an Interface for Observation Data Access (IODA), the System-Agnostic Background Error Representation (SABER), and the Object-Oriented Prediction System (OOPS), which includes 3D-Var, 4D-EnVar, and 4D-Var data assimilation options. Not only will JEDI allow users to test and optimize various data assimilation options, it also serves as the basis to objectively quantify impacts from current and future observing systems through data denial, Forecast Sensitivity to Observation Impact (FSOI), and Observing System Simulation Experiments (OSSE). The implications of having this capability in the public domain are significant, ranging from the ability to utilize crowdsourced weather observations to performing cost–benefit analyses on large-scale commercial satellite programs. To initialize the FV3-GFS using JEDI, users will need access to global observations that come in a variety of preprocessed levels and formats (NetCDF, BUFR, etc.). NOAA is currently working with CSPs to design and build

⁶ <https://dtcenter.org/community-code/common-community-physics-package-ccpp>

⁷ <https://ufsccommunity.org>

⁸ <https://ufsccommunity.org/science/gst/>

⁹ www.jcsda.org/news-blog/2020/10/29/jedi-fv3-100-and-crtm-240-releases

a public-facing Data Lake that will serve as a centralized repository, where observations can be converted to a common IODA format used by JEDI applications in real time.

Open innovation through EPIC

With the UFS community modeling program gaining momentum, how do we distill the potential Wild West of open-source software development down to a refined production-ready forecast system that is stable, efficient, and meets the up-time requirements that NWS stakeholders depend on? Without some type of framework, it is unlikely that community innovation will lead to efficient advancements in the operational UFS, so a rigorous process for objectively vetting and integrating enhancements needed to be established. EPIC should serve as the R2O funnel that aggregates development from the broader community, including NOAA, and manages layers of testing and refinement with the objective of improving NCEP's operational UFS. This will be done through an open and transparent process, where the community plays a role in the development and validation of new ideas. The National Integrated Drought Information System Reauthorization Act (NIDIS), which was signed into law in 2019, codifies EPIC and directs targeted investment in weather research. EPIC is the framework and process to manage, validate, and catalyze innovation; EPIC is not the innovation itself, but the sandbox in which innovators can develop, test, and exchange ideas. Resources to support these new ideas will require separate funding mechanisms, which are discussed below.

A conceptual representation of the EPIC process is shown in Fig. 2. This is a version of the well-known R2O “funnel” originally discussed in MacDonald et al. (2006). The community can propose and contribute ideas ranging from advancements in data assimilation and physics to software engineering and even new observing systems. These contributions will be run through a series of test-driven gateways, where advancement will be determined by transparent evidence-based evaluation. Candidate innovations originating from any corner of the community can enter and exit the process at any stage, as some may not need upstream testing, while others that are not aligned with the operational mission of NOAA may still have significant scientific value. The gateway evaluation criteria in the early part of the development process (green) will be determined by the community.

As a candidate for operational implementation progresses through the EPIC funnel, real-time parallel evaluation will take on a more NCEP-centric focus (blue) because of specific NCEP Central Operations (NCO) requirements; however, these tests will still be performed through an open and collaborative process with the community using either publicly accessible cloud or near-identical architecture. It will not be uncommon for various modifications to be sent back through previous steps if predetermined and agreed-to criteria are not met. This is the essence of operations-to-research (O2R), which is sometimes overlooked by the research community.

Once EPIC community governance, cloud resource management, and parallel testing environments are established, the focus will shift to incentivizing innovation. Contributions from across the global scientific community can be submitted for evaluation, and if deemed beneficial to the NOAA mission, ultimately integrated into NCEP's operational system and subsequent public UFS releases. Certain incentives are inherent in EPIC, such as opportunities for publications and recognition within the community UFS documentation. None of these contributions would be possible without the extensive upfront work performed by the entire UFS team. Refactoring the global model code to run on community-accessible hardware is the great equalizer that enables community participation. Innovative ideas can come from universities, government laboratories, international corporations, a retired software engineer, or a team of 8th graders. The source does not matter, and we, as objective scientists, will evaluate the contributions based on their ability to impart progress.

Fig. 2. The open innovation funnel of research-to-operations and operations-to-research, where innovative ideas coming from the research community (purple) are passed through a series of evidence-based test-driven gateways that are initially determined by the community (green) and become more NWS focused (blue) as a chosen candidate transitions closer to operational implementation.

Since the NWS is one of the principal stakeholders that will benefit from successful contributions to the UFS through EPIC, it is imperative that grants and other financial and educational opportunities be provided by NOAA and the National Science Foundation to further support innovative research. A great example of this is the Collaborative Science, Technology, and Applied Research (CSTAR) Program, which is a NWS initiative to transition research from academic institutions into the operational meteorology community.¹⁰ There is also an opportunity to incentivize friendly competition through NOAA's Office of Education with awards similar to those offered by the XPRIZE foundation.¹¹ As we have seen during the evolution of WRF, workshops, code sprints, and online forums are essential to furthering the collaborative process and broadening community engagement.

The grand challenge

To maximize the potential value of EPIC, the NWP community must embrace risk-taking. This can easily happen if stakeholders in the broader weather enterprise are willing to accept that failure is an essential step in rapid innovation. Overwhelming risk aversion caused by the fear of criticism has resulted in paralysis that propagates across the broader NWP community. Complaints of unstable research-grade code often cause scientists to spend more time fussing over efficiency and stability at the expense of making advancements to improve model skill. We need the software engineers and scientists to be successful, but it will require both ends of the research and operations chain to understand and respect the mission and constraints of the other. These constraints should be continuously reexamined, as rapid advancements in technology are routinely eliminating barriers to reveal new paths forward.

Christensen (1997) explains that creative solutions are far more likely to occur when requirements are based on achieving an end goal versus defining the path to get there. For example, would a node failure during a forecast cycle be as concerning if these runs had multiple real-time parallel mirrors in the cloud? The objective is not to successfully complete a run, but to consistently get model guidance into the hands of stakeholders on time, and there are an increasing number of alternate ways we can achieve that goal that reduce or eliminate

¹⁰ www.weather.gov/sti/r2o_cstar

¹¹ XPRIZE is a nonprofit organization that hosts public competitions to encourage radical breakthroughs that benefit humanity.

the need for elaborate failover and rerun procedures. The rationale for exploring alternative hardware architectures can also be applied to software. It should not matter if actionable information, whether produced by industry for niche markets or NWS's life-saving Impact-Based Decision Support Services, is derived from the classic set of Navier–Stokes equations or artificial intelligence using a Fourier neural operator (Li et al. 2020). When working backward from the end goal to map a solution path, we should ensure present-day constraints do not result in requirements that inadvertently complicate or prevent the adoption of future innovation. This is why it is critical for the EPIC O2R governance to propose challenges or highlight weaknesses versus requesting specific types of solutions to known problems.

Moving forward

Readers can review the EPIC community workshop summary and presentations,¹² as well as the UMAC final report and briefing to the Environmental Information Services Working Group (EISWG) of the NOAA Science Advisory Board,¹³ and decide if the course we have charted aligns with this vision. I believe it does, and I am incredibly grateful to all those who contributed. The amount of progress made over the last few years has been astounding, but we are only at the beginning. No one sector can do this alone. A constructive public–private–academic partnership can only be achieved if we respect our complementary strengths and effectively share the risk. Our success will be determined by our ability to foster a collaborative culture that embraces failure as a valuable contribution to innovation. It is time to set sail on this EPIC journey as partners in the community for the benefit of the Nation.

¹² <https://wpo.noaa.gov/Programs/EPIC>

¹³ ftp://cola.gmu.edu/pub/kinter/UCACN/Carr_Kinter_briefing_EISWG_27Jul2018.pdf

Acknowledgments. I am exceptionally grateful for the thoughtful comments and suggestions provided by Gary Lackmann, Cliff Mass, Shali Mohleji, Peter Neilley, Bill Hooke, Ricky Rood, Paul Higgins, Thomas Auligné, Fred Carr, and three anonymous reviewers. I would also like to thank Dominikus Heinzeller for helping me get the FV3 GFS running on a Mac. The initial version of the diagram used for Fig. 2 was designed by Chris Franks. The UFS and EPIC teams deserve a tremendous amount of credit for all of the work to get the program to where it is today. None of this would have ever happened if it was not for the visionary leadership and contagious enthusiasm of Bill Lapenta and Fuqing Zhang. I am incredibly thankful for the time I had with them as both colleagues and friends.

References

- Adcroft, A., and Coauthors, 2019: The GFDL global ocean and sea ice model OM4.0: Model description and simulation features. *J. Adv. Model. Earth Syst.*, **11**, 3167–3211, <https://doi.org/10.1029/2019MS001726>.
- Ahmad, N., S. Qamar, N. Khan, A. Naim, M. R. Hussain, Q. N. Naveed, and M. R. Mahmood, 2020: Cloud computing trends and cloud migration tuple. *Innovations in Electronics and Communication Engineering*, H. Saini et al., Eds., Springer, 737–745.
- Chesbrough, H., and M. Bogers, 2014: Explicating open innovation: Clarifying an emerging paradigm for understanding innovation. *New Frontiers in Open Innovation*, Oxford University Press, 3–28, <https://doi.org/10.1093/acprof:oso/9780199682461.003.0001>.
- , W. Vanhaverbeke, and J. West, 2006: *Open Innovation: Researching a New Paradigm*. Oxford University Press, 336 pp.
- Christensen, C. M., 1997: *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*. Harvard Business School Press, 252 pp.
- Daniels, T., G. Tsoucalas, M. Anderson, D. Mulally, W. Moninger, and R. Mamrosh, 2004: Tropospheric airborne meteorological data reporting (TAMDAR) sensor development. Preprints, *11th Conf. on Aviation, Range, and Aerospace Meteorology*, Hyannis, MA, Amer. Meteor. Soc., 7.6, <http://ams.confex.com/ams/pdfpapers/81841.pdf>.
- Harris, L., and S. Lin, 2013: A two-way nested global-regional dynamical core on the cubed-sphere grid. *Mon. Wea. Rev.*, **141**, 283–306, <https://doi.org/10.1175/MWR-D-11-00201.1>.
- Harrison, J. B., 1988: Coordination of meteorological services and supporting research in the Federal Government. *Bull. Amer. Meteor. Soc.*, **69**, 362–367, [https://doi.org/10.1175/1520-0477\(1988\)069<0362:COMSAS>2.0.CO;2](https://doi.org/10.1175/1520-0477(1988)069<0362:COMSAS>2.0.CO;2).
- Hill, C., C. DeLuca, V. Balaji, M. Suarez, and A. da Silva, 2004: The architecture of the Earth system modeling framework. *IEEE Comput. Sci. Eng.*, **6**, 18–28, <https://doi.org/10.1109/MCISE.2004.1255817>.
- Li, Z., N. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart, and A. Anandkumar, 2020: Fourier neural operator for parametric partial differential equations. arXiv, 16 pp., <https://arxiv.org/abs/2010.08895>.
- MacDonald, A. E., R. Fulton, M. Kenny, S. Murawski, P. Ortner, A. M. Powell, A. Sen, and L. Uccellini, 2006: Research location in NOAA: Physical and social sciences. NOAA, 72 pp, ftp://ftp.oar.noaa.gov/SAB/sab/members/2006/07_meeting/PSTT_Final_Report.pdf.
- Mass, C., 2006: The uncoordinated giant: Why U.S. weather research and prediction are not achieving their potential. *Bull. Amer. Meteor. Soc.*, **87**, 573–584, <https://doi.org/10.1175/BAMS-87-5-573>.
- Moninger, W. R., S. G. Benjamin, B. D. Jamison, T. W. Schlatter, T. L. Smith, and E. J. Szoke, 2010: Evaluation of regional aircraft observations using TAMDAR. *Wea. Forecasting*, **25**, 627–645, <https://doi.org/10.1175/2009WAF2222321.1>.
- Skamarock, W. C., J. B. Klemp, M. G. Duda, L. D. Fowler, S. Park, and T. D. Ringler, 2012: A multiscale nonhydrostatic atmospheric model using centroidal Voronoi tessellations and C-grid staggering. *Mon. Wea. Rev.*, **140**, 3090–3105, <https://doi.org/10.1175/MWR-D-11-00215.1>.
- Stajner, I., and Coauthors, 2020: NOAA's unified forecast System for sub-seasonal predictions: Development and operational implementation plans of Global Ensemble Forecast System v12 (GEFSv12) at NCEP. *EGU General Assembly Conf.*, Online, EGU, EGU2020-6212, <https://doi.org/10.5194/egusphere-egu2020-6212>.
- Theurich, G., and Coauthors, 2016: The Earth system prediction suite: Toward a coordinated U.S. modeling capability. *Bull. Amer. Meteor. Soc.*, **97**, 1229–1247, <https://doi.org/10.1175/BAMS-D-14-00164.1>.
- Tolman, H., M. Banner, and J. Kaihatu, 2011: The NOPP operational wave model improvement project. *Ocean Modell.*, **70**, 2–10, <https://doi.org/10.1016/j.ocemod.2012.11.011>.