Is Computational Oceanography Coming of Age?

Thomas W. N. Haine, Renske Gelderloos, Miguel A. Jimenez-Urias, Ali H. Siddiqui, Gerard Lemson, Dimitri Medvedev, Alex Szalay, Ryan P. Abernathey, Mattia Almansì, and Christopher N. Hill

ABSTRACT: Computational oceanography is the study of ocean phenomena by numerical simulation, especially dynamical and physical phenomena. Progress in information technology has driven exponential growth in the number of global ocean observations and the fidelity of numerical simulations of the ocean in the past few decades. The growth has been exponentially faster for ocean simulations, however. We argue that this faster growth is shifting the importance of field measurements and numerical simulations for oceanographic research. It is leading to the maturation of computational oceanography as a branch of marine science on par with observational oceanography. One implication is that ultraresolved ocean simulations are only loosely constrained by observations. Another implication is that barriers to analyzing the output of such simulations should be removed. Although some specific limits and challenges exist, many opportunities are identified for the future of computational oceanography. Most important is the prospect of hybrid computational and observational approaches to advance understanding of the ocean.

KEYWORDS: General circulation models; Numerical analysis/modeling; Ocean circulation; Ocean dynamics; Ocean models

<https://doi.org/10.1175/BAMS-D-20-0258.1>

Corresponding author: Thomas Haine, thomas.haine@jhu.edu

In final form 21 February 2021

©2021 American Meteorological Society

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

AFFILIATIONS: Haine, Gelderloos, Jimenez-Urias, and Siddiqui—Earth and Planetary Sciences, The Johns Hopkins University, Baltimore, Maryland; Lemson, Medvedev, and Szalay—Physics and Astronomy, The Johns Hopkins University, Baltimore, Maryland; Abernathey—Earth and Environmental Sciences, Columbia University, New York, New York; Almansi—National Oceanography Centre, Southampton, United Kingdom; Hill—Earth, Atmospheric, and Planetary Sciences, Massachusetts Institute of Technology, Cambridge, Massachusetts

Computational oceanography is the study of ocean phenomena by numerical simulation, especially dynamical and physical phenomena using ocean general circulation models (OGCMs). One early pioneer of this field wrote of the 1960s, 1970s, and 1980s as the “birth,” “infancy,” and “adolescence” of OGCMs, respectively (Bryan 2006; see also Holland and McWilliams 1987; McWilliams 1996). Similarly, the authors of a comprehensive review of OGCMs wrote at the turn of the century “this field ... has entered an era of healthy adolescence” (Griffies et al. 2000). With 20 more years of data, this essay explores the continued growth of OGCMs and speculates on their prospects. We ask, Is computational oceanography entering a new era that signifies its coming of age?

For motivation, Fig. 1 compares oceanographic measurements and results from a high-resolution OGCM. The region of interest is a topographic constriction called the Denmark Strait, between Greenland and Iceland. The Denmark Strait Overflow (DSO) flows south through this gap and is an important current for the Atlantic meridional overturning circulation and thus for the ocean’s role in North Atlantic climate. The two time series in Fig. 1a show DSO volume flux (transport). One time series is from in situ measurements, the other is from a high-resolution regional OGCM (and they have been processed similarly with similar smoothing). The question is this: Which is which? Figure 1b compares in situ hydrographic measurements along a section north of Denmark Strait with a synthetic hydrographic section from the OGCM. And Fig. 1c shows the trajectories of drifting oceanographic floats approaching Denmark Strait from the north and trajectories of drifting particles in the OGCM released from the same locations. Again, the question is, Which are the real data and which are the synthetic data? In each case, the field measurements and the OGCM results are different, but identifying them is difficult.

These are examples of OGCM Turing tests. They are inspired by Alan Turing’s imitation game to distinguish between, and correctly identify, a person and an intelligent machine. The game involves asking questions through an interface that obscures whether the responses are from the person or the machine (Turing 1950). The difficulty of the OGCM Turing tests in Fig. 1 reflects the small systematic error in the OGCM and therefore its realism. Some OGCM solutions are reaching the point that they are essentially indistinguishable from observations, so they pass Turing tests like those in Fig. 1. In the words of Ed Lorenz, numerical experiments will eventually “duplicate the circulation to any desired degree of accuracy” (Lorenz 1967).¹

With these themes in mind, this essay explores the growth of OGCMs and compares it to the growth of ocean observations. The focus is on the computer science and information technology improvements that contribute to the growth. We then speculate on limits, opportunities, and prospects for OGCMs.

¹ The prescient Lorenz was writing about atmospheric models in the late 1960s, but the message applies to OGCMs today.

Unequal exponential growth

Two examples illustrate the growth of ocean observations. First, consider temperature observations in the global deep ocean over the last half century. Figure 2a shows the cumulative number of temperature observations deeper than 1,000 m. They have grown exponentially (notice the y axis is logarithmic). Averaged over the last century, the exponential growth has

a doubling period of 10.4 years, giving an approximately 60-fold expansion in the deep temperature database since 1960. Technology transitions have maintained this exponential growth, specifically, advances in microelectronics and information technology. In the 1990s conductivity–temperature–depth (CTD) sensors on autonomous profiling floats took over from ship CTD sampling, for example, leading in the 2000s to the transformative Argo global float network (Argo 2020).

Second, consider the history of sea level observations from satellite altimeters. Sea level data have revolutionized physical oceanography by providing information on the surface circulation, mesoscale eddies, tides, and sea level change. Figure 2b shows the sequence of altimeter missions (colored bars) and the cumulative number of observing days (black line). The number of observing days reveals the growth in sea level observations (although there is great variety between missions). The number of sea level observations has grown nearly exponentially since the mid 1980s with a doubling time of about 8.1 years and a ~20-fold expansion in the sea level database since 1985. Again, microelectronic and information technology advances have maintained this growth.

Technology advances have also fueled growth in the fidelity of OGCMs. For example, Fig. 2c shows the history of global OGCM resolution. The black dots show five pioneering (cutting-edge) models over the last 40 years. The Bryan and Lewis (1979) model had a peak resolution of 2.4° with 12 vertical levels and the Rocha et al. (2016) model had a peak resolution of 0.02° with 90 vertical levels. The growth in OGCM resolution (number of grid points) is exponential with a doubling time of 2.2 years and a 10^5 -fold increase since 1980. We also show the global ocean models from the Intergovernmental Panel on Climate Change (IPCC) reports. The peak

Fig. 1. OGCM Turing tests. In each row oceanographic field measurements are compared with OGCM results, but they are unlabeled (and processed similarly). The Turing test is to identify which is which. (a) Denmark Strait Overflow (DSO) volume flux (Sv; $1 \text{ Sv} = 10^6 \text{ m}^3 \text{ s}^{-1}$; negative means equatorward). Adapted from Haine (2010). (b) Salinity (colors) on a section north of Denmark Strait (annual average; the heavy contour is the $27.80 \sigma_\theta$ density anomaly). (c) Lagrangian trajectories of RAFOS floats and synthetic RAFOS floats. Adapted from Saberi et al. (2020).

Fig. 2. Unequal exponential growth. (a) History of deep (deeper than 1,000 m) ocean temperature measurements. The colored dots show different instruments and observing platforms. (b) History of sea level measurements from satellite altimetry expressed by the cumulative number of days of measurement. The satellite missions and their durations are indicated with the colored bars. (c) History of cutting-edge global OGCM and IPCC ocean model resolution expressed by the length scale of the horizontal grid and the number of model grid points. Each dot represents one ocean model and the OGCMs are from Bryan and Lewis (1979), Semtner and Chervin (1992), Maltrud et al. (1998), Maltrud and McClean (2005), and Rocha et al. (2016). (d) History of top supercomputers using $R_{\max}$ speed (FLOPS = floating point operations per second) for fastest machines (open circles) and ECMWF and NCAR machines (closed circles). The lines show best fit exponential growth in each panel (τ_{2x} is the doubling time).

resolution of the ocean OGCMs in the first IPCC report was 2.7° with 9 vertical levels and the peak resolution in the latest (sixth) IPCC report is 0.067° with 75 vertical levels.² This growth is also exponential with a doubling time of 2.8 years. For the most highly resolved models in each assessment, the doubling time is close to the cutting-edge OGCM doubling time.

Now compare the horizontal resolution of ocean measurements with OGCM resolution. The Argo profiling float network operates about 4,000 floats at any one time. Each float makes a vertical profile from 2,000 m depth to the surface every 10 days. The global average spacing of profiles is therefore 300 km.³ The spacing between altimeter tracks for the TOPEX/Poseidon and Jason satellite altimeters is also about 315 km (at the equator), with a repeat period of 10 days. The present day peak OGCM resolution of $0.02^\circ \approx 2$ km is therefore 140 times higher.⁴

Prospects for future growth

Looking ahead, the future is bright for the Argo network. The reason is that Argo is part of the Global Ocean and Global Climate Observing System, which implements the Paris Agreement on climate change and UN sustainable development goals. New capacities, like deep profiling floats, and new technologies, like biogeochemical sensors, are planned over the next few years (GCOS 2016). It is unclear how the network can double in size in the next decade and maintain long-term exponential growth, but it is plausible.

The future is also bright for sea level measurements. The Surface Water and Ocean Topography (SWOT) mission, scheduled for launch in 2022, will start a new era of sea level observation. SWOT will observe sea level over a swath, rather than over a single patch. It will have 15 km resolution, or better, covering most of the global ocean every 21 days (Morrow et al. 2019). It will improve the spatial resolution of sea level data by a factor of about 10. Therefore, the prospects for the altimetry record to continue growing exponentially in the 2020s are good.

For OGCMs, resolution improves as supercomputer technology advances. Historically, that follows Moore's "law," which says that transistor density in microprocessors doubles every 2 years (Moore 1975). For instance, machines first achieved petaflop speeds (10^{15} floating point operations per second) in 2008 and exaflop speeds (10^{18}) in March 2020, a doubling every 1.1 years (see Fig. 2d). Computers available to the oceanographic, atmospheric, and climate communities are less powerful. Still, the machines at NCAR and ECMWF⁵ also show exponential growth over recent decades with a doubling every 1.1 years, albeit lagging the cutting-edge machines by about 5 years (Fig. 2d). On this basis, the OGCM resolution will probably continue to double every 2.2 years, at least for several more years (assuming funding remains at historic levels). It is reasonable to expect cutting-edge exascale global OGCMs with horizontal resolution around 1 km by the mid-2020s. After that, with widespread anticipation that Moore's law will end (Waldrop 2016), future growth is uncertain.

Maturation of computational oceanography

This evidence shows that information technology advances are driving exponential growth in ocean observations and exponential growth in OGCM resolution. But the OGCM growth rate is faster. Therefore, OGCM resolution is also growing exponentially faster than the growth in ocean field data. In 1990, OGCMs were obviously biased compared to measurements, for example, of deep temperatures or sea level. In 2020, OGCMs are achieving resolutions that are substantially greater than the gaps between measurements, at least for some regimes, like deep and abyssal ocean currents. We should expect this trend to continue for the foreseeable future (the next several years). Therefore, the question

² The AR6 data points on Fig. 2 are from the HighResMIP experiments, which is a subproject on high-resolution models that does not run the full suite of CMIP6 experiments.

³ The vertical resolution of Argo profile data is about 5 m, which is about 7 times higher than the best AR6 OGCMs and about 3 times higher than the Poseidon Project run mentioned below.

⁴ This comparison avoids the issue of time dependence in the circulation. It simply (and conservatively) imagines the Argo and altimetry data from one 10-day period are used to constrain the time-mean OGCM state over that period.

⁵ Meaning the U.S. National Center for Atmospheric Research and the European Centre for Medium-Range Weather Forecasts.

arises: When, and in what ways, will OGCMs become as important as observations for advancing knowledge in physical oceanography? Historically, most knowledge came from observations of the real ocean.⁶ The growth of OGCMs suggests that the field is approaching an era in which numerical circulation models are as important as observations for advancing knowledge. For example, diagnosing and understanding the rectified effects of mesoscale eddy variability on the large-scale, low-frequency circulation will probably rely heavily on high-resolution OGCMs.

What are the criteria to claim that OGCM solutions should be treated, in some cases, as seriously as real measurements? Realizing them would mark the maturation of computational oceanography. These criteria are on our checklist:

- 1) Confidence in the fidelity of the basic tools and methods. Consider two types of tool: First, consider the theoretical definition of the ocean circulation problem. Computational oceanography relies on software to compute approximations to the incompressible rotating-stratified Navier–Stokes equations, with equations for the conservation of dissolved salts and heat (McWilliams 1996; Griffies 2004; Fox-Kemper et al. 2019). There is little doubt that these are the right equations for ocean circulation. The software is mature, stable, and diverse. The issue of unresolved processes, and parameterizing their effects remains an important area of research. For example, it is still unclear how to represent the impacts of unresolved submesoscale processes on the larger-scale flow. Although much progress has been made on this problem in the last 30 years (Gent 2011; Le Sommer et al. 2018; Fox-Kemper et al. 2019), resolution improvements have surely played an essential part in refining OGCM accuracy (Griffies et al. 2000). In other words, we believe that the problem of parameterizing unresolved scales is not so pathological that it contaminates all of the resolved scales.⁷ A corollary is that OGCMs are less complicated than the real ocean, meaning that OGCM variability is a lower bound on the variability in the real system. These are de facto working hypotheses of all theoretical and numerical approaches to understanding the ocean circulation.

Second, we need tools to adjust OGCM solutions to agree with observations; that is, to solve the data assimilation and state estimation problem (Bennett 1992; Wunsch 1996, 2006; Kalnay 2002). For example, state estimation is used to produce retrospective reanalyses (hindcasts) of the time-evolving ocean state and data assimilation is used to initialize prospective forecasts of the future. Although many questions remain open, these methods are also now mature, stable, and diverse.

- 2) The number of OGCM degrees of freedom exceeds the number of observational constraints. This criterion concerns the state estimation and data assimilation problems. In essence, it is about whether it is possible (in principle) to adjust an OGCM solution to fit the observations exactly or not. If the OGCM can be adjusted to fit the data exactly, the state estimation problem is underdetermined. Otherwise, it is overdetermined.⁸ The number of OGCM degrees of freedom scales as the number of grid points (for large numbers of grid points). The number of observational constraints scales as the number of distinct measurements. Figure 2 shows evidence that the number of OGCM degrees of freedom per observational constraint exceeds one because, loosely, the peak OGCM resolution is now 140 times higher than the Argo and Jason

⁶ For example, Stewart (2008) writes: “The theory describing a convecting, wind-forced, turbulent fluid in a rotating coordinate system has never been sufficiently well known that important features of the oceanic circulation could be predicted before they were observed. In almost all cases, oceanographers resort to observations to understand oceanic processes.”

⁷ It is likely that errors in parameterized physics *influence* all resolved scales, not least because of error growth due to deterministic chaos. But the issue is whether the errors in parameterized physics cause systematic errors in the resolved scales, such as biases in statistics of resolved quantities. It is reasonable to suppose that (i) resolution improvements and parameterization improvements reduce these systematic biases toward zero, and (ii) the systematic biases are not so bad as to preclude use of models to understand (and hindcast and predict) the natural system. Of course, these are quantitative (not qualitative) hypotheses that vary from case to case (models, parameterizations, resolved metrics, science questions).

⁸ Ignoring the atypical case of the problem being exactly determined.

data spacing (see footnote 4). This gap is growing exponentially because OGCM resolution is growing exponentially faster than data density. Therefore, the state estimation problem is moving from (in principle) being overdetermined to being underdetermined.⁹

Crossing this threshold has interesting implications: First, the systematic errors in OGCMs disappear and they pass Turing or Feigenbaum tests (Turing 1950; Feigenbaum 2003; Harel 2005), like those in Fig. 1. That is, OGCM solutions become indistinguishable from observations of the real ocean and a subject-matter expert cannot tell them apart. Regional OGCM simulations of the DSO at resolutions of 0.5–2 km are approaching this point (Magaldi and Haine 2015; Almansi et al. 2020; Saberi et al. 2020). Similarly, regional high-resolution state estimates are nearly underdetermined (Lea et al. 2006; Dwivedi et al. 2011). A fair comparison (Turing test) requires that the space–time scales of the observations and the model results are the same, which means the power spectra should match. This comparison is a necessary test to realize Lorenz’s vision quoted in the Introduction. It is not a sufficient test, however, as the OGCM results can resemble the measurements for the wrong reasons, but we take it as strong evidence of small OGCM bias.

Second, the OGCM solutions make accurate, testable predictions about the real ocean. Historically, advances from theoretical and numerical research in dynamical oceanography have lagged advances from observational research (see footnote 6). Once OGCMs become underdetermined by data, it will be common for them to make predictions that can be tested by field programs. For example, DSO simulations show exchange of dense water out of the overflow onto the east Greenland continental shelf, and vice versa (Magaldi et al. 2011). They also show entrainment of near-surface waters south of Iceland into the DSO within a few months, at least during hard winters (Saberi et al. 2020). It remains to be seen if these predictions occur in the real ocean.

Limits to computational oceanography

Although these opportunities are exciting, there are clear limits to computational oceanography. First, direct numerical simulation (DNS) of the global ocean circulation is inconceivable today. DNS in this context means running OGCMs that resolve all scales of motion; from the planetary scale to the dissipation scale (around 1 mm), and from centuries to seconds. DNS would avoid the challenge of parameterizing the effects of the unresolved scales, but at vast computational cost. Figure 3 shows why. It shows the full range of space and time scales relevant to the ocean general circulation, about 10 orders of magnitude in both. It also shows the space–time volumes accessible to present-day supercomputers, including the best AR6 OGCMs shown in Fig. 2, the Poseidon Project run,¹⁰ and turbulence simulations [DNS and large-eddy simulations (LES)]. To span the entire space–time plane, supercomputers would need to resolve about 10^{25} grid points and 10^{10} time steps. That is about 16 orders of magnitude more grid points than is possible today. Extrapolating the doubling time of 2.2 years in Fig. 2c, it would take 120 years to achieve this increase, which is impossible to envision. Clearly, the exponential growth must roll off at some point, and, clearly, OGCM simulations cannot replace observations of the natural ocean.

Another potential limit concerns scalability of OGCM codes. Figure 2 shows that the historic doubling time for the number of OGCM grid points is about twice the doubling time for supercomputer speed. That value is close to the optimal ratio of 3/2, which assumes that machine speedup is spent on increasing horizontal resolution, that model time step is inversely proportional to the

⁹ It is possible to argue that any inverse problem with real observations is formally underdetermined because the observational error can be considered as an unknown parameter to be solved for (Wunsch 1996; Stammer et al. 2002). Regardless, no global ocean circulation state estimate has characterized the null space associated with the indeterminacy (to our knowledge), or presented different solutions that fit the observations equally well. Instead, the practice has been to stop the state estimation once an acceptable fit has been achieved (Stammer et al. 2002; Nguyen et al. 2021).

¹⁰ The Poseidon Project intends to run a global OGCM at (nominally) 1 km horizontal resolution (poseidon.idies.jhu.edu). The Poseidon Project is unrelated to the TOPEX/Poseidon altimeter.

grid spacing (for numerical stability), and that all other factors are equal. In other words, the historic OGCM growth has nearly maintained pace with the supercomputer acceleration. It is unclear how this trend will continue, however, because of the overhead of communication from processor cores to other cores, to memory, and to disk (Le Sommer et al. 2018). Moreover, exascale supercomputers will not resemble petascale supercomputers: they will have different architectures and greater diversity (Giles and Reguly 2014). These changes are driven by physical limits on clock speed and power densities in silicon microprocessors, as well as economic forces. To harness exascale machines OGCM software must radically change [for discussion of this point for atmospheric general circulation models, see Lawrence et al. (2018) and Gropp and Snir (2013)]. The developers of next generation OGCMs should adopt collaborative, open community habits (Le Sommer et al. 2018). Promising paths are to define and refine modular subcomponents, and to develop domain-specific languages, performance tools, and data models that separate different levels in the software stack for optimization by experts (Lawrence et al. 2018). OGCM computational intensity (the fraction of time spent performing floating point calculations versus memory operations) is low: Le Sommer et al. (2018) estimate OGCMs run at 5% peak speed, for example. So there is potential to accelerate OGCMs by reducing this bottleneck (e.g., by exploiting time parallelism; Schreiber et al. 2017; Hamon et al. 2020). Exploiting new application-specific hardware accelerators and new OGCM solver paradigms, like lower precision (Palmer 2012; Palem 2014), will also be important. These developments will mitigate the saturation of transistor density and the demise of Moore's law, and they offer hope to continue the refinement of OGCM meshes.

Finally, there are challenging issues to couple OGCMs to other parts of the Earth system at horizontal resolutions around one kilometer. For example, air–sea interaction, sea ice dynamics, and biogeochemistry are all poorly understood and hard to simulate at these scales.

Opportunities for computational oceanography

The opportunities for computational oceanography to advance marine science include the following:

- Migration from the study of specific instances of phenomena to the study of statistics of these phenomena. The DSO is one of many currents that is affected by rotation, stratification, and bathymetry. It is inconceivable to observe all of them, but they can all be simulated in an exascale OGCM. Empirical characterization of these numerical overflows would be an important step forward.

Fig. 3. Characteristic space and time scales of the ocean general circulation. Various geophysical and theoretical scales are shown with gray patches and colored lines (for a discussion, see Klinger and Haine 2019). The colored rectangles show cutting-edge circulation models (direct numerical simulation of turbulence, large-eddy simulation, the Poseidon Project run, AR6 HighResMIP, and TAR OGCMs). The black dot shows the sampling characteristics of the Argo profiling floats, and the TOPEX/Poseidon–Jason altimeters. The diagram is indicative, not definitive, because it suppresses the anisotropies and inhomogeneities present in the general circulation.

- Discovery and characterization of intermittent, time-dependent, three-dimensional phenomena, which are hard to observe. Submesoscale currents are in this class, which occur at horizontal scales shorter than several kilometers (Thomas et al. 2008). Diapycnal mixing is another example, which occurs at scales shorter than meters (MacKinnon et al. 2017).
- Comprehensive and illuminating analyses of ocean mass, heat, salt, momentum, energy, and vorticity budgets in a way that is nearly impossible with direct observations.
- Discovery and characterization of ocean circulation regimes that cannot be observed. Examples include the circulation during the last glacial maximum (paleo-oceanography) or in extraterrestrial oceans (exo-oceanography). For these ocean circulation problems, the data-sparseness challenge is much worse than for the modern ocean (LeGrand and Wunsch 1995; Amrhein et al. 2018; Way et al. 2017). Criterion 2 was achieved with smaller computational resources for these fields, and therefore they have already entered the era of computational oceanography by the rationale in the “Maturation of computational oceanography” section.
- Robust observing system design using OGCM solutions as synthetic data. These observing system simulation experiments (Errico et al. 2012) should become the best-practice standard for fieldwork design. There are implications for making the OGCM output accessible and easy to work with (see below), but the payoff from engaging observational oceanographers is great.
- Insight from OGCM state estimation to support fieldwork, ideally in real time. The community should recognize the fact that the underdetermined state estimates imply an infinite number of OGCM solutions that match the data exactly. This means that techniques are needed to characterize and handle the OGCM null space (indeterminacy). For example, observational oceanographers at sea could make decisions about where, when, and how to observe using OGCM information that captures the range of possible circulation states consistent with data. This practice is common in atmospheric science already.
- More efficient identification of interesting phenomena using automatic methods, like artificial intelligence and data mining (Kutz 2017; Lguensat et al. 2019). In fact, such automatic methods will become essential as the size of OGCM output grows exponentially and overwhelms manual feature identification (see below).
- Increasing transition of dynamical oceanography to an experimental (computational) science. It has long been recognized that idealized models isolate physical mechanisms relevant to the general circulation and thereby build dynamical understanding. We still require idealized models; in particular, we need a hierarchy of models that span the gap between geophysical fluid dynamics problems and realistic simulations of the circulation. This hierarchy will ensure that the increasing OGCM realism does not outpace understanding of the basic physics (Held 2005; Vallis 2016; Coveney et al. 2016; Emanuel 2020).

Prospects for computational oceanography

How can these priorities be achieved and what are the prospects for computational oceanography? We should focus on these issues in the next several years:

- The indeterminacy of OGCM solutions by observations should be recognized—we should “embrace the null space.” Imagine computing an ensemble of high-resolution (high degrees of freedom per observation) state estimates that fit the observations (exactly or equally well within instrumental errors). These state estimates would differ, for example, in the characteristics of their eddies, or in their deep circulations, or in their internal wave fields, or in their diapycnal mixing. In such a situation, the different state estimates should all be treated seriously. The ensemble would characterize the null space (indeterminacy) in the inverse problem and therefore quantify the variety of ocean states consistent with observations and ocean circulation physics. This vision for uncertainty quantification echoes

the probabilistic practice of ensemble atmospheric model runs to forecast the weather (see also McWilliams 2007; Le Sommer et al. 2018).

- Barriers to dissemination of OGCM simulation output should be lowered—we should “democratize the data.” The output should be freely available, including to nonprofessional users. Traditionally, effort has focused on the challenges of calculating OGCM solutions with supercomputers. The OGCM output has become increasingly hard to use, because of the massive data volume, and the technical complexities that attend the high-performance computation. Access to high-resolution OGCM output is restricted to a few experts in practice.

The remedy is to build high-performance data science infrastructure to match the high-performance compute infrastructure (Overpeck et al. 2011). These data portals should be open and have low thresholds to getting started. We should be able to sample the simulations the way that we sample the real ocean. For example, it should be easy for an observational oceanographer to plot a synthetic hydrographic section or mooring time series. The data portals should include open software and significant compute resources to process and analyze the simulation data. We should avoid the inefficient practice in which users are forced to download voluminous data to their local machines and then write their own code to analyze them. Technologies and infrastructure to achieve these goals are under development, such as the OceanSpy OGCM data analysis package (Almansi et al. 2019), the Pangeo community in geoscience big data (pangeo.io), and the SciServer and JASMIN big data science platforms (Medvedev et al. 2016; www.jasmin.ac.uk).

- “Benchmark” OGCM reference solutions should be computed using the best available compute resources and served to the public. They are of intrinsic value to all oceanographers, not just ocean modelers, for the reasons stated above. Benchmark solutions for regional ocean circulation problems are valuable for the same reasons, as are idealized simulations of specific ocean dynamical processes. The track record of other fields using this approach is impressive. For instance, the Johns Hopkins Turbulence Database exposes cutting-edge turbulence simulation data to researchers and provides easy-to-use interfaces to retrieve and interact with the data using novel metaphors like immersing virtual sensors into the 4D data (turbulence.pha.jhu.edu; Perlman et al. 2007; Li et al. 2008).
- OGCMs will migrate to exascale compute resources in the next few years. This migration will involve new paradigms to access the data. For example, with today’s petaflop supercomputers only about 0.1% of the OGCM solution can be permanently stored for analysis. The problem arises because of the prohibitive time needed to transfer the massive output volume to long-term storage media, and the prohibitive expense of the media. This loss of OGCM data will be much worse on exaflop machines.

To mitigate this problem consider the strategy adopted by the Large Hadron Collider (LHC), the world’s most sophisticated experimental facility. The LHC provides a single source of data on subatomic particle collisions. Several experiments tap the data stream in so-called beam lines. Within each experiment, customized hardware monitors the stream. Only about 1 event in 10 million is retained for storage and detailed analysis. In exascale oceanography the analogous idea (see section 3.3.5 in Asch et al. 2018) is to enable automatic identification of selected circulation events and trigger storage while the OGCM runs. For example, we could target intermittent intense mixing events, plus their antecedents and fates. An implication is that we should build a software interface for community-supplied software plugins to implement the custom triggers. Also, we need to enable posterior recomputation of small space–time chunks of the full solution, with customized diagnostics, and possibly at higher resolutions.

It is instructive to compare computational oceanography with computational meteorology, which is the analogous field in atmospheric sciences. Computational meteorology has

somewhat different science objectives. Numerical weather prediction (NWP) is an important task, for example. The main advances in NWP attributable to growth in computer power are (i) improved model resolution (now also approaching global 1 km horizontal resolution; Fuhrer et al. 2018) and (ii) improved forecast uncertainty quantification through larger ensembles of forecast runs. Computational meteorology also concerns reanalysis products to hindcast the historical atmospheric state. The reanalysis state estimation tools tolerate unphysical adjustments (increments), however, which give more accurate fits to observations at lower computational cost. This practice is different than the ocean state estimation tools discussed here, which firmly constrain the model solutions to satisfy the model equations.

Nevertheless, there are several useful lessons from computational meteorology: First, NWP has steadily improved since the 1980s (Bauer et al. 2015). The rate is an improvement in forecast skill of about 1 day decade⁻¹ (meaning a 3-day forecast in 2015 is about as skillful as a 2-day forecast in 2005). The improvement derives mainly from better forecast initialization and better atmospheric general circulation models (AGCMs; Magnusson and Källén 2013; see also Simmons and Hollingsworth 2002). In this context, better AGCMs means models that have higher resolution, have more accurate parameterizations and/or complexity, and have larger forecast ensembles that better estimate forecast uncertainty. Computing advances have played an enormous role in these improvements (Bauer et al. 2015). Second, as AGCM resolution increases, new phenomena begin to be resolved. For example, with AGCM grid spacing of a few kilometers, convective scales are partly resolved (convective systems) but partly unresolved (convective cells). This partial resolution of convection is called the “gray zone,” akin to eddy-permitting resolution in OGCMs. The best approach to set up convection parameterization schemes in the AGCM gray zone is unclear and forecast skill does not always improve at all lead times as resolution increases (Hong and Dudhia 2012). Moreover, at cloud-resolving resolution, data density is mismatched with AGCM resolution (the number of degrees of freedom exceeds the number of observations) and the model solution is not well constrained (Hong and Dudhia 2012).

Conclusions

Global OGCMs have a rich history that stretches back to the 1970s and regional OGCMs stretch back to the 1960s (models of the tides stretch back even further; see Cartwright 2012). OGCMs have been valuable to elucidate the ocean circulation since their inception. More broadly, numerical solution of rotating, stratified flow has roots in numerical weather prediction (NWP) from the early twentieth century [Abbe 1901; Bjerknes 1904; Richardson 1922; see also Lynch (2008) and Benjamin et al. (2019) for historical perspectives on NWP and climate models]. Since 2000, global OGCMs have continued their exponential improvement in resolution. They are now becoming unconstrained by observations. Benchmark OGCM solutions have increasing value to a growing community and should be permanently archived and freely available. Clear limits, opportunities, and prospects for computational oceanography are in sight. For these reasons, our answer to the question posed in the title of this essay is yes: computational oceanography is entering a new era and is coming of age.

This field promises powerful new tools to address previously intractable problems. It does not aim to supplant observational oceanography. Indeed, observing the natural ocean must never cease. Instead, the greatest opportunity lies in merging these hitherto disparate branches of marine science. Lasting progress will require that we trust computational insights, verify them with real world observations, and understand them with fundamental theory.

Acknowledgments. This material is based upon work supported by the National Science Foundation under Grant OAC-1835640 and by the Institute for Data Intensive Engineering and Science at Johns Hopkins University. Comments by Baylor Fox-Kemper and two other reviewers refined the paper.

Data availability statement. Codes to make the figures are available at github.com/hainegroup/Computational-Oceanography-Commentary. For Fig. 2, the temperature data are from the National Centers for Environmental Information World Ocean Database, the altimeter mission data are from www.altimetry.info, the IPCC data are from the IPCC reports and pcmdi.llnl.gov/CMIP6, and the supercomputer data are from en.wikipedia.org/wiki/List_of_fastest_computers. The data for AR6 are from the HighResMIP project in July 2020, which was incomplete then. The ECMWF and NCAR machine speed data are from www.top500.org.

References

- Abbe, C., 1901: The physical basis of long range weather forecasts. *Mon. Wea. Rev.*, **29**, 551–561, [https://doi.org/10.1175/1520-0493\(1901\)29\[551c:TPBOLW\]2.0.CO;2](https://doi.org/10.1175/1520-0493(1901)29[551c:TPBOLW]2.0.CO;2).
- Almansi, M., R. Gelderloos, T. W. N. Haine, A. Saberi, and A. H. Siddiqui, 2019: OceanSpy: A Python package to facilitate ocean model data analysis and visualization. *J. Open Source Software*, **4**, 1506, <https://doi.org/10.21105/joss.01506>.
- , T. W. N. Haine, R. Gelderloos, and R. S. Pickart, 2020: Evolution of Denmark Strait Overflow cyclones and their relationship to overflow surges. *Geophys. Res. Lett.*, **47**, e2019GL086759, <https://doi.org/10.1029/2019GL086759>.
- Amrhein, D. E., C. Wunsch, O. Marchal, and G. Forget, 2018: A global glacial ocean state estimate constrained by upper-ocean temperature proxies. *J. Climate*, **31**, 8059–8079, <https://doi.org/10.1175/JCLI-D-17-0769.1>.
- Argo, 2020: Argo float data and metadata from Global Data Assembly Centre (Argo GDAC). SEANOE, accessed 16 July 2020, <https://doi.org/10.17882/42182>.
- Asch, M., and Coauthors, 2018: Big data and extreme-scale computing. *Int. J. High Perform. Comput. Appl.*, **32**, 435–479, <https://doi.org/10.1177/1094342018778123>.
- Bauer, P., A. Thorpe, and G. Brunet, 2015: The quiet revolution of numerical weather prediction. *Nature*, **525**, 47–55, <https://doi.org/10.1038/nature14956>.
- Benjamin, S. G., J. M. Brown, G. Brunet, P. Lynch, K. Saito, and T. W. Schlatter, 2019: 100 years of progress in forecasting and NWP applications. *A Century of Progress in Atmospheric and Related Sciences: Celebrating the American Meteorological Society Centennial*, Meteor. Monogr., No. 59, Amer. Meteor. Soc., <https://doi.org/10.1175/AMSMONOGRAPHIS-D-18-0020.1>.
- Bennett, A. F., 1992: *Inverse Methods in Physical Oceanography*. Cambridge University Press, 346 pp., <https://doi.org/10.1017/CBO9780511600807>.
- Bjerknes, V., 1904: Das problem der wettvorhersage, betrachtet vom standpunkte der mechanik und der physik (The problem of weather prediction, considered from the viewpoints of mechanics and physics). *Meteor. Z.*, **21**, 1–7, <https://doi.org/10.1127/0941-2948/2009/416>.
- Bryan, K., 2006: Modeling ocean circulation. *Physical Oceanography: Developments since 1950*, Springer, 29–44, https://doi.org/10.1007/0-387-33152-2_3.
- , and L. J. Lewis, 1979: A watermass model of the world ocean. *J. Geophys. Res.*, **84**, 2503–2517, <https://doi.org/10.1029/JC084iC05p02503>.
- Cartwright, D. E., 2012: *Tides*. Cambridge University Press, 292 pp.
- Coveney, P. V., E. R. Dougherty, and R. R. Highfield, 2016: Big data need big theory too. *Philos. Trans. Roy. Soc. London*, **374A**, 20160153, <https://doi.org/10.1098/rsta.2016.0153>.
- Dwivedi, S., T. W. N. Haine, and C. E. Del Castillo, 2011: Upper ocean state estimation in the Southern Ocean Gas Exchange Experiment region using the four-dimensional variational technique. *J. Geophys. Res.*, **116**, C00F02, <https://doi.org/10.1029/2009JC005615>.
- Emanuel, K., 2020: The relevance of theory for contemporary research in atmospheres, oceans, and climate. *AGU Adv.*, **1**, e2019AV000129, <https://doi.org/10.1029/2019AV000129>.
- Errico, R. M., R. Yang, N. C. Privé, K.-S. Tai, R. Todling, M. E. Sienkiewicz, and J. Guo, 2012: Development and validation of observing-system simulation experiments at NASA's Global Modeling and Assimilation Office. *Quart. J. Roy. Meteor. Soc.*, **139**, 1162–1178, <https://doi.org/10.1002/qj.2027>.
- Feigenbaum, E. A., 2003: Some challenges and grand challenges for computational intelligence. *J. Assoc. Comput. Mach.*, **50**, 32–40, <https://doi.org/10.1145/602382.602400>.
- Fox-Kemper, B., and Coauthors, 2019: Challenges and prospects in ocean circulation models. *Front. Mar. Sci.*, **6**, 65, <https://doi.org/10.3389/fmars.2019.00065>.
- Fuhrer, O., and Coauthors, 2018: Near-global climate simulation at 1 km resolution: Establishing a performance baseline on 4888 GPUs with COSMO 5.0. *Geosci. Model Dev.*, **11**, 1665–1681, <https://doi.org/10.5194/gmd-11-1665-2018>.
- GCOS, 2016: The global observing system for climate. WMO Tech. Rep. 200, 341 pp., https://library.wmo.int/index.php?lvl=notice_display&id=19838#.XxWxPJNKhBw.
- Gent, P. R., 2011: The Gent–McWilliams parameterization: 20/20 hindsight. *Ocean Modell.*, **39**, 2–9, <https://doi.org/10.1016/j.ocemod.2010.08.002>.
- Giles, M. B., and I. Reguly, 2014: Trends in high-performance computing for engineering calculations. *Philos. Trans. Roy. Soc. London*, **372A**, 20130319, <https://doi.org/10.1098/rsta.2013.0319>.
- Griffies, S. M., 2004: *Fundamentals of Ocean Climate Models*. Princeton University Press, 518 pp., <https://doi.org/10.2307/j.ctv301gzg>.
- , and Coauthors, 2000: Developments in ocean climate modelling. *Ocean Modell.*, **2**, 123–192, [https://doi.org/10.1016/S1463-5003\(00\)00014-7](https://doi.org/10.1016/S1463-5003(00)00014-7).
- Gropp, W., and M. Snir, 2013: Programming for exascale computers. *Comput. Sci. Eng.*, **15**, 27–35, <https://doi.org/10.1109/MCSE.2013.96>.
- Haine, T. W. N., 2010: High-frequency fluctuations in Denmark Strait Overflow transport. *Geophys. Res. Lett.*, **37**, L14601, <https://doi.org/10.1029/2010GL043272>.
- Hamon, F. P., M. Schreiber, and M. L. Minion, 2020: Parallel-in-time multi-level integration of the shallow-water equations on the rotating sphere. *J. Comput. Phys.*, **407**, 109210, <https://doi.org/10.1016/j.jcp.2019.109210>.
- Harel, D., 2005: A Turing-like test for biological modeling. *Nat. Biotechnol.*, **23**, 495–496, <https://doi.org/10.1038/nbt0405-495>.
- Held, I. M., 2005: The gap between simulation and understanding in climate modeling. *Bull. Amer. Meteor. Soc.*, **86**, 1609–1614, <https://doi.org/10.1175/BAMS-86-11-1609>.
- Holland, W. R., and J. C. McWilliams, 1987: Computer modeling in physical oceanography from the global circulation to turbulence. *Phys. Today*, **40**, 51–57, <https://doi.org/10.1063/1.881115>.

- Hong, S.-Y., and J. Dudhia, 2012: Next-generation numerical weather prediction: Bridging parameterization, explicit clouds, and large eddies. *Bull. Amer. Meteor. Soc.*, **93**, E56–E59, <https://doi.org/10.1175/2011BAMS3224.1>.
- Kalnay, E., 2002: *Atmospheric Modeling, Data Assimilation and Predictability*. Cambridge University Press, 341 pp., <https://doi.org/10.1017/CBO9780511802270>.
- Klinger, B. A., and T. W. N. Haine, 2019: *Ocean Circulation in Three Dimensions*. 1st ed. Cambridge University Press, 484 pp., www.cambridge.org/9780521768436.
- Kutz, J. N., 2017: Deep learning in fluid dynamics. *J. Fluid Mech.*, **814**, 1–4, <https://doi.org/10.1017/jfm.2016.803>.
- Lawrence, B. N., and Coauthors, 2018: Crossing the chasm: How to develop weather and climate models for next generation computers? *Geosci. Model Dev.*, **11**, 1799–1821, <https://doi.org/10.5194/gmd-11-1799-2018>.
- Lea, D. J., T. W. N. Haine, and R. F. Gasparovic, 2006: Observability of the Irminger Sea circulation using variational data assimilation. *Quart. J. Roy. Meteor. Soc.*, **132**, 1545–1576, <https://doi.org/10.1256/qj.05.77>.
- LeGrand, P., and C. Wunsch, 1995: Constraints from paleotracer data on the North Atlantic circulation during the Last Glacial Maximum. *Paleoceanography*, **10**, 1011–1045, <https://doi.org/10.1029/95PA01455>.
- Le Sommer, J., E. P. Chassignet, and A. J. Wallcraft, 2018: Ocean circulation modeling for operational oceanography: Current status and future challenges. *New Frontiers in Operational Oceanography*, GODAE OceanView, 289–305.
- Lguensat, R., J. Le Sommer, S. Metref, E. Cosme, and R. Fablet, 2019: Learning generalized quasi-geostrophic models using deep neural numerical models. arXiv, <https://arxiv.org/abs/1911.08856>.
- Li, Y., and Coauthors, 2008: A public turbulence database cluster and applications to study Lagrangian evolution of velocity increments in turbulence. *J. Turbul.*, **9**, N31, <https://doi.org/10.1080/14685240802376389>.
- Lorenz, E. N., 1967: The nature and theory of the general circulation of the atmosphere. WMO Tech. Doc. 218, 161 pp.
- Lynch, P., 2008: The origins of computer weather prediction and climate modeling. *J. Comput. Phys.*, **227**, 3431–3444, <https://doi.org/10.1016/j.jcp.2007.02.034>.
- MacKinnon, J. A., and Coauthors, 2017: Climate Process Team on internal wave-driven ocean mixing. *Bull. Amer. Meteor. Soc.*, **98**, 2429–2454, <https://doi.org/10.1175/BAMS-D-16-0030.1>.
- Magaldi, M. G., and T. W. N. Haine, 2015: Hydrostatic and non-hydrostatic simulations of dense waters cascading off a shelf: The East Greenland case. *Deep-Sea Res. I*, **96**, 89–104, <https://doi.org/10.1016/j.dsr.2014.10.008>.
- , —, and R. S. Pickart, 2011: On the nature and variability of the East Greenland Spill Jet: A case study in summer 2003. *J. Phys. Oceanogr.*, **41**, 2307–2327, <https://doi.org/10.1175/JPO-D-10-05004.1>.
- Magnusson, L., and E. Källén, 2013: Factors influencing skill improvements in the ECMWF forecasting system. *Mon. Wea. Rev.*, **141**, 3142–3153, <https://doi.org/10.1175/MWR-D-12-00318.1>.
- Maltrud, M. E., and J. L. McClean, 2005: An eddy resolving global 1/10° ocean simulation. *Ocean Modell.*, **8**, 31–54, <https://doi.org/10.1016/j.ocemod.2003.12.001>.
- , R. D. Smith, A. J. Semtner, and R. C. Malone, 1998: Global eddy-resolving ocean simulations driven by 1985–1995 atmospheric winds. *J. Geophys. Res.*, **103**, 30825–30853, <https://doi.org/10.1029/1998JC900013>.
- McWilliams, J. C., 1996: Modeling the oceanic general circulation. *Annu. Rev. Fluid Mech.*, **28**, 215–248, <https://doi.org/10.1146/annurev.fl.28.010196.001243>.
- , 2007: Irreducible imprecision in atmospheric and oceanic simulations. *Proc. Natl. Acad. Sci. USA*, **104**, 8709–8713, <https://doi.org/10.1073/pnas.0702971104>.
- Medvedev, D., G. Lemson, and M. Rippin, 2016: SciServer Compute: Bringing analysis close to the data. *Proc. 28th Int. Conf. on Scientific and Statistical Database Management*, Budapest, Hungary, ACM, 27, <https://doi.org/10.1145/2949689.2949700>.
- Moore, G. E., 1975: Progress in digital integrated electronics. *Int. Electronic Devices Meeting*, Washington, DC, IEEE, 11–13.
- Morrow, R., and Coauthors, 2019: Global observations of fine-scale ocean surface topography with the Surface Water and Ocean Topography (SWOT) mission. *Front. Mar. Sci.*, **6**, 232, <https://doi.org/10.3389/fmars.2019.00232>.
- Nguyen, A. T., H. Pillar, V. Ocaña, A. Bigdeli, T. A. Smith, and P. Heimbach, 2021: The Arctic Subpolar Gyre State Estimate (ASTE): Description and assessment of a data-constrained, dynamically consistent ocean-sea ice estimate for 2002–2017. *J. Adv. Model. Earth Syst.*, **13**, e2020MS002398, <https://doi.org/10.1029/2020MS002398>.
- Overpeck, J. T., G. A. Meehl, S. Bony, and D. R. Easterling, 2011: Climate data challenges in the 21st century. *Science*, **331**, 700–702, <https://doi.org/10.1126/science.1197869>.
- Palem, K. V., 2014: Inexactness and a future of computing. *Philos. Trans. Roy. Soc. London*, **372A**, 20130281, <https://doi.org/10.1098/rsta.2013.0281>.
- Palmer, T. N., 2012: Towards the probabilistic Earth-system simulator: A vision for the future of climate and weather prediction. *Quart. J. Roy. Meteor. Soc.*, **138**, 841–861, <https://doi.org/10.1002/qj.1923>.
- Perlman, E., R. Burns, Y. Li, and C. Meneveau, 2007: Data exploration of turbulence simulations using a database cluster. *Proc. 2007 ACM/IEEE Conf. Supercomputing*, Reno, NV, ACM–IEEE, 23, <https://doi.org/10.1145/1362622.1362654>.
- Richardson, L. F., 1922: *Weather Prediction by Numerical Process*. Cambridge University Press, 236 pp., <https://doi.org/10.1017/CBO9780511618291>.
- Rocha, C. B., T. K. Chereskin, S. T. Gille, and D. Menemenlis, 2016: Mesoscale to submesoscale wavenumber spectra in Drake Passage. *J. Phys. Oceanogr.*, **46**, 601–620, <https://doi.org/10.1175/JPO-D-15-0087.1>.
- Saberi, A., T. W. N. Haine, R. Gelderloos, M. F. de Jong, H. Fury, and A. Bower, 2020: Lagrangian perspective on the origins of Denmark Strait Overflow. *J. Phys. Oceanogr.*, **50**, 2393–2414, <https://doi.org/10.1175/JPO-D-19-0210.1>.
- Schreiber, M., P. S. Peixoto, T. Haut, and B. Wingate, 2017: Beyond spatial scalability limitations with a massively parallel method for linear oscillatory problems. *Int. J. High Perform. Comput. Appl.*, **32**, 913–933, <https://doi.org/10.1177/1094342016687625>.
- Semtner, A. J., and R. M. Chervin, 1992: Ocean general circulation from a global eddy-resolving model. *J. Geophys. Res.*, **97**, 5493–5550, <https://doi.org/10.1029/92JC00095>.
- Simmons, A. J., and A. Hollingsworth, 2002: Some aspects of the improvement in skill of numerical weather prediction. *Quart. J. Roy. Meteor. Soc.*, **128**, 647–677, <https://doi.org/10.1256/003590002321042135>.
- Stammer, D., and Coauthors, 2002: Global ocean circulation during 1992–1997, estimated from ocean observations and a general circulation model. *J. Geophys. Res.*, **107**, 3118, <https://doi.org/10.1029/2001JC000888>.
- Stewart, R. H., 2008: *Introduction to Physical Oceanography*. University Press of Florida, 345 pp.
- Thomas, L. N., A. Tandon, and A. Mahadevan, 2008: Submesoscale processes and dynamics. *Ocean Modeling in an Eddying Regime*, *Geophys. Monogr.*, Vol. 177, Amer. Geophys. Union, 17–38, <https://doi.org/10.1029/177GM04>.
- Turing, A. M., 1950: I. Computing machinery and intelligence. *Mind*, **59**, 433–460, <https://doi.org/10.1093/mind/LIX.236.433>.
- Vallis, G. K., 2016: Geophysical fluid dynamics: Whence, whither and why? *Proc. Roy. Soc. London*, **472A**, 20160140, <https://doi.org/10.1098/rspa.2016.0140>.
- Waldrop, M. M., 2016: The chips are down for Moore’s law. *Nature*, **530**, 144–147, <https://doi.org/10.1038/530144a>.
- Way, M. J., and Coauthors, 2017: Resolving Orbital and Climate Keys of Earth and Extraterrestrial Environments with Dynamics (ROCKE-3D) 1.0: A general circulation model for simulating the climates of rocky planets. *Astrophys. J.*, **231**, 12, <https://doi.org/10.3847/1538-4365/aa7a06>.
- Wunsch, C., 1996: *The Ocean Circulation Inverse Problem*. 1st ed. Cambridge University Press, 442 pp., <https://doi.org/10.1017/CBO9780511629570>.
- , 2006: *Discrete Inverse and State Estimation Problems*. 1st ed. Cambridge University Press, 371 pp., <https://doi.org/10.1017/CBO9780511535949>.