Computational Oceanography is Coming of Age

Thomas Haine^{1,1}, Renske Gelderloos^{1,1}, Miguel Jimenez-Urias^{1,1}, Ali Siddiqui^{1,1}, Gerard Lemson^{1,1}, Dmitry Medvedev^{1,1}, Alex Szalay^{1,1}, Ryan Abernathey^{2,2}, Mattia Almansi^{3,3}, and Christopher Hill^{4,4}

¹Johns Hopkins University ²Columbia University ³National Oceanography Centre ⁴MIT

November 30, 2022

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 ultra-resolved 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.

1	Is Computational Oceanography Coming of Age?
2	Thomas W. N. Haine*, Renske Gelderloos, Miguel A. Jimenez-Urias, Ali H. Siddiqui,
3	Earth & Planetary Sciences, Johns Hopkins University, Baltimore, MD
4	Gerard Lemson, Dimitri Medvedev, Alex Szalay,
5	Physics & Astronomy, Johns Hopkins University, Baltimore, MD
6	Ryan P. Abernathey,
7	Earth & Environmental Sciences, Columbia University, New York, NY
8	Mattia Almansi,
9	National Oceanography Centre, Southampton, UK
10	Christopher N. Hill
11	Earth, Atmospheric, & Planetary Sciences, MIT, Cambridge, MA

¹² **Corresponding author*: Thomas Haine, Thomas.Haine@jhu.edu

ABSTRACT

Computational Oceanography is the study of ocean phenomena by numerical simulation, especially 13 dynamical and physical phenomena. Progress in information technology has driven exponential 14 growth in the number of global ocean observations and the fidelity of numerical simulations of the 15 ocean in the past few decades. The growth has been exponentially faster for ocean simulations, 16 however. We argue that this faster growth is shifting the importance of field measurements and 17 numerical simulations for oceanographic research. It is leading to the maturation of Computational 18 Oceanography as a branch of marine science on par with observational oceanography. One 19 implication is that ultra-resolved ocean simulations are only loosely constrained by observations. 20 Another implication is that barriers to analyzing the output of such simulations should be removed. 21 Although some specific limits and challenges exist, many opportunities are identified for the future 22 of Computational Oceanography. Most important is the prospect of hybrid computational and 23 observational approaches to advance understanding of the ocean. 24

Capsule summary. Fast growth in the fidelity of ocean general circulation models is driving the
 maturation of Computational Oceanography as a branch of marine science on par with observational
 oceanography.

1. Introduction

Computational Oceanography is the study of ocean phenomena by numerical simulation, es-29 pecially dynamical and physical phenomena using ocean general circulation models (OGCMs). 30 One early pioneer of this field wrote of the 1960s, 1970s, and 1980s as the "birth", "infancy", 31 and "adolescence" of OGCMs, respectively (Bryan 2006, see also Holland and McWilliams 1987; 32 McWilliams 1996). Similarly, the authors of a comprehensive review of OGCMs wrote at the turn 33 of the century "this field...has entered an era of healthy adolescence" (Griffies et al. 2000). With 34 twenty more years of data, this essay explores the continued growth of OGCMs and speculates 35 on their prospects. We ask: Is Computational Oceanography entering a new era that signifies its 36 coming of age? 37

For motivation, Fig. 1 compares oceanographic measurements and results from a high-resolution 38 OGCM. The region of interest is a topographic constriction called the Denmark Strait, between 39 Greenland and Iceland. The Denmark Strait Overflow (DSO) flows south through this gap and is 40 an important current for the Atlantic Meridional Overturning Circulation and thus for the ocean's 41 role in North Atlantic climate. The two timeseries in Fig. 1a show DSO volume flux (transport). 42 One timeseries is from *in situ* measurements, the other is from a high resolution regional OGCM 43 (and they have been processed similarly with similar smoothing). The question is this: Which is 44 which? Fig. 1b compares *in situ* hydrographic measurements along a section north of Denmark 45 Strait with a synthetic hydrographic section from the OGCM. And Fig. 1c shows the trajectories 46 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 is the real data and which is 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 51 distinguish between, and correctly identify, a person and an intelligent machine. The game involves 52 asking questions through an interface that obscures whether the responses are from the person or 53 the machine (Turing 1950). The difficulty of the OGCM Turing tests in Fig. 1 reflects the small 54 systematic error in the OGCM and therefore its realism. Some OGCM solutions are reaching the 55 point that they are essentially indistinguishable from observations, so they pass Turing tests like 56 those in Fig. 1. In the words of Ed Lorenz, numerical experiments will eventually "duplicate the 57 circulation to any desired degree of accuracy" (Lorenz 1967).¹ 58

⁵⁹ With these themes in mind, this essay explores the growth of OGCMs and compares it to the ⁶⁰ growth of ocean observations. The focus is on the computer science and information technology ⁶¹ improvements that contribute to the growth. We then speculate on limits, opportunities, and ⁶² prospects for OGCMs.

2. 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. Fig. 2a shows the cumulative number of temperature observations deeper than 1000m. 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

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

⁷⁰ in micro-electronics and information technology. In the 1990s conductivity-temperature-depth
 ⁷¹ (CTD) sensors on autonomous profiling floats took over from ship CTD sampling, for example,
 ⁷² 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 73 have revolutionized physical oceanography by providing information on the surface circulation, 74 mesoscale eddies, tides, and sea level change. Fig. 2b shows the sequence of altimeter missions 75 (colored bars) and the cumulative number of observing days (black line). The number of observing 76 days reveals the growth in sea level observations (although there is great variety between missions). 77 The number of sea level observations has grown nearly exponentially since the mid 1980s with a 78 doubling time of about 8.1 years and a \approx 20-fold expansion in the sea level database since 1985. 79 Again, micro-electronic and information technology advances have maintained this growth. 80

Technology advances have also fueled growth in the fidelity of OGCMs. For example, Fig. 2c 81 shows the history of global OGCM resolution. The black dots show five pioneering (cutting-edge) 82 models over the last 40 years. The Bryan and Lewis (1979) model had a peak resolution of 2.4° 83 with 12 vertical levels and the Rocha et al. (2016) model had a peak resolution of 0.02° with 90 84 vertical levels. The growth in OGCM resolution (number of grid points) is exponential with a 85 doubling time of 2.2 years and a 10^5 -fold increase since 1980. We also show the global ocean 86 models from the Intergovernmental Panel on Climate Change (IPCC) reports. The peak resolution 87 of the ocean OGCMs in the first IPCC report was 2.7° with 9 vertical levels and the peak resolution 88 in the latest (sixth) IPCC report is 0.067° with 75 vertical levels.² This growth is also exponential 89 with a doubling time of 2.8 years. For the most highly-resolved models in each assessment, the 90 doubling time is close to the cutting-edge OGCM doubling time. 91

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

⁹² Now compare the horizontal resolution of ocean measurements with OGCM resolution. The ⁹³ Argo profiling float network operates about 4000 floats at any one time. Each float makes a vertical ⁹⁴ profile from 2000m depth to the surface every ten days. The global average spacing of profiles ⁹⁵ is therefore 300 km.³ The spacing between altimeter tracks for the TOPEX/Poseidon and Jason ⁹⁶ satellite altimeters is also about 315 km (at the equator), with a repeat period of ten days. The ⁹⁷ present day peak OGCM resolution of $0.02^{\circ} \approx 2$ km is therefore 140 times higher.⁴

3. 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 United Nations 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 ten. Therefore, the prospects for the altimetry record to continue growing exponentially in the 2020s are good.

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

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

For OGCMs, resolution improves as supercomputer technology advances. Historically, that 111 follows Moore's "law," which says that transistor density in microprocessors doubles every two 112 years (Moore 1975). For instance, machines first achieved petaflop speeds (10¹⁵ floating point 113 operations per second) in 2008 and exaflop speeds (10^{18}) in March 2020, a doubling every 1.1 years 114 (see Fig. 2d). Computers available to the oceanographic, atmospheric, and climate communities 115 are less powerful. Still, the machines at NCAR and ECMWF⁵ also show exponential growth over 116 recent decades with a doubling every 1.1 years, albeit lagging the cutting-edge machines by about 117 five years (Fig. 2d). On this basis, the OGCM resolution will probably continue to double every 2.2 118 years, at least for several more years (assuming funding remains at historic levels). It is reasonable 119 to expect cutting-edge exascale global OGCMs with horizontal resolution around one kilometer 120 by the mid 2020s. After that, with widespread anticipation that Moore's law will end (Waldrop 121 2016), future growth is uncertain. 122

4. Maturation of Computational Oceanography

This evidence shows that information technology advances are driving exponential growth in 124 ocean observations and exponential growth in OGCM resolution. But the OGCM growth rate is 125 faster. Therefore, OGCM resolution is also growing exponentially faster than the growth in ocean 126 field data. In 1990, OGCMs were obviously biased compared to measurements, for example, of 127 deep temperatures or sea level. In 2020, OGCMs are achieving resolutions that are substantially 128 greater than the gaps between measurements, at least for some regimes, like deep and abyssal 129 ocean currents. We should expect this trend to continue for the foreseeable future (the next 130 several years). Therefore, the question arises: When, and in what ways, will OGCMs become as 131 important as observations for advancing knowledge in physical oceanography? Historically, most 132

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

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.

¹³⁸ What are the criteria to claim that OGCM solutions should be treated, in some cases, as seriously ¹³⁹ as real measurements? Realizing them would mark the maturation of Computational Oceanography. ¹⁴⁰ These criteria are on our checklist:

1. Confidence in the fidelity of the basic tools and methods. Consider two types of tool: 141 First, consider the theoretical definition of the ocean circulation problem. Computational 142 Oceanography relies on software to compute approximations to the incompressible rotating-143 stratified Navier Stokes equations, with equations for the conservation of dissolved salts and 144 heat (McWilliams 1996; Griffies 2004; Fox-Kemper et al. 2019). There is little doubt that 145 these are the right equations for ocean circulation. The software is mature, stable, and diverse. 146 The issue of unresolved processes, and parametrizing their effects remains an important area 147 of research. For example, it is still unclear how to represent the impacts of unresolved 148 submesoscale processes on the larger scale flow. Although much progress has been made on 149 this problem in the last 30 years (Gent 2011; Le Sommer et al. 2018; Fox-Kemper et al. 2019), 150 resolution improvements have surely played an essential part in refining OGCM accuracy 151 (Griffies et al. 2000). In other words, we believe that the problem of parametrizing unresolved 152

⁶For 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."

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 con-163 straints. This criterion concerns the state estimation and data assimilation problems. In 164 essence, it is about whether it is possible (in principle) to adjust an OGCM solution to fit 165 the observations exactly or not. If the OGCM can be adjusted to fit the data exactly, the 166 state estimation problem is under-determined. Otherwise, it is over-determined.⁸ The num-167 ber of OGCM degrees of freedom scales as the number of grid points (for large numbers 168 of grid points). The number of observational constraints scales as the number of distinct 169 measurements. Fig. 2 shows evidence that the number of OGCM degrees of freedom per 170 observational constraint exceeds one because, loosely, the peak OGCM resolution is now 140 171

⁷It 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 towards 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).

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

172	times higher than the Argo and Jason data spacing (see footnote 4). This gap is growing
173	exponentially because OGCM resolution is growing exponentially faster than data density.
174	Therefore, the state estimation problem is moving from (in principle) being over-determined
175	to being under-determined.9
176	Crossing this threshold has interesting implications: First, the systematic errors in OGCMs
177	disappear and they pass Turing or Feigenbaum tests (Turing 1950; Feigenbaum 2003; Harel
178	2005), like those in Fig. 1. That is, OGCM solutions become indistinguishable from obser-
179	vations of the real ocean and a subject-matter expert cannot tell them apart. Regional OGCM
180	simulations of the Denmark Strait Overflow (DSO) at resolutions of 0.5–2 km are approach-
181	ing this point (Magaldi and Haine 2015; Almansi et al. 2020; Saberi et al. 2020). Similarly,
182	regional high-resolution state estimates are nearly under-determined (Lea et al. 2006; Dwivedi
183	et al. 2011). A fair comparison (Turing test) requires that the space-time scales of the ob-
184	servations and the model results are the same, which means the power spectra should match.
185	This comparison is a necessary test to realize Lorenz' vision quoted in the Introduction. It
186	is not a sufficient test, however, as the OGCM results can resemble the measurements for the
187	wrong reasons, but we take it as strong evidence of small OGCM bias.
188	Second, the OGCM solutions make accurate, testable predictions about the real ocean. His-
189	torically, advances from theoretical and numerical research in dynamical oceanography have
190	lagged advances from observational research (see footnote 6). Once OGCMs become under-
191	determined by data, it will be common for them to make predictions that can be tested by field

⁹It is possible to argue that any inverse problem with real observations is formally under-determined 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. 2020).

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.

¹⁹⁷ 5. Limits to Computational Oceanography

Although these opportunities are exciting, there are clear limits to Computational Oceanography. 198 First, direct numerical simulation (DNS) of the global ocean circulation is inconceivable today. 199 DNS in this context means running OGCMs that resolve all scales of motion; from the planetary 200 scale to the dissipation scale (around 1mm), and from centuries to seconds. DNS would avoid 201 the challenge of parametrizing the effects of the unresolved scales, but at vast computational cost. 202 Fig. 3 shows why. It shows the full range of space and time scales relevant to the ocean general 203 circulation, about ten orders of magnitude in both. It also shows the space time volumes accessible 204 to present-day supercomputers, including the best AR6 OGCMs shown in Fig. 2, the Poseidon 205 Project run,¹⁰ and turbulence simulations (DNS and large eddy simulations, LES). To span the 206 entire space time plane, supercomputers would need to resolve about 10^{25} grid points and 10^{10} time 207 steps. That is about 16 orders of magnitude more grid points than is possible today. Extrapolating 208 the doubling time of 2.2 years in Fig. 2c, it would take 120 years to achieve this increase, which is 209 impossible to envision. Clearly, the exponential growth must roll off at some point, and, clearly, 210 OGCM simulations cannot replace observations of the natural ocean. 211

¹⁰The 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.

Another potential limit concerns scalability of OGCM codes. Fig. 2 shows that the historic 212 doubling time for the number of OGCM grid points is about twice the doubling time for supercom-213 puter speed. That value is close to the optimal ratio of 3/2, which assumes that machine speedup is 214 spent on increasing horizontal resolution, that model timestep is inversely proportional to the grid 215 spacing (for numerical stability), and that all other factors are equal. In other words, the historic 216 OGCM growth has nearly maintained pace with the supercomputer acceleration. It is unclear how 217 this trend will continue, however, because of the overhead of communication from processor cores 218 to other cores, to memory, and to disk (Le Sommer et al. 2018). Moreover, exascale supercomput-219 ers will not resemble petascale supercomputers: they will have different architectures and greater 220 diversity (Giles and Reguly 2014). These changes are driven by physical limits on clock speed and 221 power densities in silicon microprocessors, as well as economic forces. To harness exascale ma-222 chines OGCM software must radically change (for discussion of this point for atmospheric general 223 circulation models, see Lawrence et al. 2018 and Gropp and Snir 2013). The developers of next 224 generation OGCMs should adopt collaborative, open community habits (Le Sommer et al. 2018). 225 Promising paths are to define and refine modular sub-components, and to develop domain-specific 226 languages, performance tools, and data models that separate different levels in the software stack 227 for optimization by experts (Lawrence et al. 2018). OGCM computational intensity (the fraction of 228 time spent performing floating point calculations versus memory operations) is low: Le Sommer 229 et al. (2018) estimate OGCMs run at 5% peak speed, for example. So there is potential to accelerate 230 OGCMs by reducing this bottleneck (for example by exploiting time parallelism, Schreiber et al. 231 2017; Hamon et al. 2020). Exploiting new application-specific hardware accelerators and new 232 OGCM solver paradigms, like lower precision (Palmer 2012; Palem 2014), will also be important. 233 These developments will mitigate the saturation of transistor density and the demise of Moore's 234 Law, and they offer hope to continue the refinement of OGCM meshes. 235

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.

6. Opportunities for Computational Oceanography

²⁴⁰ The opportunities for Computational Oceanography to advance marine science include:

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.

Discovery and characterization of intermittent, time-dependent, three-dimensional phenom ena, 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
 extra-terrestrial 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 section 4.

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 under-determined 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

14

280

281

ensure that the increasing OGCM realism does not outpace understanding of the basic physics (Held 2005; Vallis 2016; Coveney et al. 2016; Emanuel 2020).

7. 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 285 "embrace the null space." Imagine computing an ensemble of high resolution (high degrees of 286 freedom per observation) state estimates that fit the observations (exactly or equally well within 287 instrumental errors). These state estimates would differ, for example, in the characteristics 288 of their eddies, or in their deep circulations, or in their internal wave fields, or in their 289 diapycnal mixing. In such a situation, the different state estimates should all be treated 290 seriously. The ensemble would characterize the null space (indeterminacy) in the inverse 291 problem and therefore quantify the variety of ocean states consistent with observations and 292 ocean circulation physics. This vision for uncertainty quantification echoes the probabilistic 293 practice of ensemble atmospheric model runs to forecast the weather (see also McWilliams 294 2007; Le Sommer et al. 2018). 295

Barriers to dissemination of OGCM simulation output should be lowered—we should "democratize the data." The output should be freely available, including to non-professional users.
 Traditionally, effort has focused on the challenges of calculating OGCM solutions with super-computers. 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-302 performance compute infrastructure (Overpeck et al. 2011). These data portals should be 303 open and have low thresholds to getting started. We should be able to sample the simulations 304 the way that we sample the real ocean. For example, it should be easy for an observational 305 oceanographer to plot a synthetic hydrographic section or mooring timeseries. The data portals 306 should include open software and significant compute resources to process and analyze the 307 simulation data. We should avoid the inefficient practice in which users are forced to download 308 voluminous data to their local machines and then write their own code to analyze them. 309 Technologies and infrastructure to achieve these goals are under development, such as the 310 OceanSpy OGCM data analysis package (Almansi et al. 2019), the Pangeo community in 311 geoscience big data (pangeo.io), and the SciServer and JASMIN big data science platforms 312 (Medvedev et al. 2016, www.jasmin.ac.uk). 313

• "Benchmark" OGCM reference solutions should be computed using the best available com-314 pute resources and served to the public. They are of intrinsic value to all oceanographers, 315 not just ocean modellers, for the reasons stated above. Benchmark solutions for regional 316 ocean circulation problems are valuable for the same reasons, as are idealized simulations 317 of specific ocean dynamical processes. The track record of other fields using this approach 318 is impressive. For instance, the Johns Hopkins Turbulence Database exposes cutting-edge 319 turbulence simulation data to researchers and provides easy-to-use interfaces to retrieve and 320 interact with the data using novel metaphors like immersing virtual sensors into the 4-D data 321 (turbulence.pha.jhu.edu; Perlman et al. 2007; Li et al. 2008). 322

• 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

³²⁵ only about 0.1% of the OGCM solution can be permanently stored for analysis. The problem ³²⁶ arises because of the prohibitive time needed to transfer the massive output volume to long-³²⁷ term storage media, and the prohibitive expense of the media. This loss of OGCM data will ³²⁸ be much worse on exaflop machines.

To mitigate this problem consider the strategy adopted by the Large Hadron Collider (LHC), 329 the world's most sophisticated experimental facility. The LHC provides a single source of 330 data on subatomic particle collisions. Several experiments tap the data stream in so-called 331 "beam-lines." Within each experiment, customized hardware monitors the stream. Only 332 about one event in ten million is retained for storage and detailed analysis. In exascale 333 oceanography the analogous idea (see section 3.3.5 in Asch et al. 2018) is to enable automatic 334 identification of selected circulation events and trigger storage while the OGCM runs. For 335 example, we could target intermittent intense mixing events, plus their antecedents and fates. 336 An implication is that we should build a software interface for community-supplied software 337 plugins to implement the custom triggers. Also, we need to enable posterior re-computation 338 of small space-time chunks of the full solution, with customized diagnostics, and possibly at 339 higher resolutions. 340

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 to 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 352 steadily improved since the 1980s (Bauer et al. 2015). The rate is an improvement in forecast 353 skill of about one day per decade (meaning a 2015 three-day forecast is about as skillful as a 2005 354 two-day forecast). The improvement derives mainly from better forecast initialization and better 355 atmospheric general circulation models (AGCMs; Magnusson and Källén 2013; see also Simmons 356 and Hollingsworth 2002). In this context, better AGCMs means models that have higher resolution, 357 have more accurate parameterizations and/or complexity, and have larger forecast ensembles that 358 better estimate forecast uncertainty. Computing advances have played an enormous role in these 359 improvements (Bauer et al. 2015). Second, as AGCM resolution increases, new phenomena begin 360 to be resolved. For example, with AGCM grid spacing of a few kilometers convective scales are 361 partly resolved (convective systems) but partly unresolved (convective cells). This partial resolution 362 of convection is called the "gray zone", akin to eddy-permitting resolution in OGCMs. The best 363 approach to set up convection parameterization schemes in the AGCM gray zone is unclear and 364 forecast skill does not always improve at all lead times as resolution increases (Hong and Dudhia 365 2012). Moreover, at cloud-resolving resolution, data density is mismatched with AGCM resolution 366 (the number of degrees of freedom exceeds the number of observations) and the model solution is 367 not well constrained (Hong and Dudhia 2012). 368

369 8. Conclusion

Global OGCMs have a rich history that stretches back to the 1970s and regional OGCMs stretch 370 back to the 1960s (models of the tides stretch back even further; see Cartwright 2012). OGCMs 371 have been valuable to elucidate the ocean circulation since their inception. More broadly, numerical 372 solution of rotating, stratified flow has roots in numerical weather prediction (NWP) from the early 373 twentieth century (Abbe 1901; Bjerknes 1904; Richardson 1922 see also Lynch 2008 and Benjamin 374 et al. 2019 for historical perspectives on NWP and climate models). Since 2000, global OGCMs 375 have continued their exponential improvement in resolution. They are now becoming unconstrained 376 by observations. Benchmark OGCM solutions have increasing value to a growing community and 377 should be permanently archived and freely available. Clear limits, opportunities, and prospects for 378 Computational Oceanography are in sight. For these reasons, our answer to the question posed in 379 the title of this essay is yes: Computational Oceanography is entering a new era and is coming of 380 age. 381

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 Founda tion under Grant Number 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 391 github.com/hainegroup/Computational-Oceanography-Commentary. For Fig. 2, 392 the temperature data are from the National Centers for Environmental Information World 393 Ocean Database, the altimeter mission data are from www.altimetry.info, the IPCC data 394 are from the IPCC reports and pcmdi.llnl.gov/CMIP6, and the supercomputer data are from 395 en.wikipedia.org/wiki/List_of_fastest_computers. The data for AR6 are from the 396 HighResMIP project in July 2020, which was incomplete then. The ECMWF and NCAR machine 397 speed data are from www.top500.org. 398

References

- Abbe, C., 1901: The physical basis of long range weather forecasts. *Mon. Weather. Rev.*, 29 (12),
 551–561, doi:10.1175/1520-0493(1901)29[551c:tpbolw]2.0.co;2.
- ⁴⁰² Almansi, M., R. Gelderloos, T. W. N. Haine, A. Saberi, and A. H. Siddiqui, 2019: OceanSpy:
 ⁴⁰³ A Python package to facilitate ocean model data analysis and visualization. *Journal of Open* ⁴⁰⁴ Source Software, 4 (39), 1506, doi:10.21105/joss.01506.
- ⁴⁰⁵ Almansi, M., T. W. N. Haine, R. Gelderloos, and R. S. Pickart, 2020: Evolution of Denmark
 ⁴⁰⁶ Strait Overflow cyclones and their relationship to overflow surges. *Geophys. Res. Lett.*, doi:
 ⁴⁰⁷ 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** (**19**), 8059–8079, doi:
 10.1175/jcli-d-17-0769.1.
- Argo, 2020: Argo float data and metadata from global data assembly centre (Argo GDAC).
 SEANOE, doi:10.17882/42182.

Asch, M., and Coauthors, 2018: Big data and extreme-scale computing. *The International* Journal of High Performance Computing Applications, 32 (4), 435–479, doi:10.1177/
 1094342018778123.

Bauer, P., A. Thorpe, and G. Brunet, 2015: The quiet revolution of numerical weather prediction.
 Nature, **525** (**7567**), 47–55, doi: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. *Meteorological Monographs*, 59, 131–1367,
 doi:10.1175/amsmonographs-d-18-0020.1.
- Bennett, A. F., 1992: *Inverse methods in physical oceanography*. Cambridge University
 Press, Cambridge, United Kingdom and New York, NY, USA, 346 pp., doi:10.1017/
 cbo9780511600807.
- ⁴²⁴ Bjerknes, V., 1904: Das problem der wettervorhersage, betrachtet vom standpunkte der mechanik
 ⁴²⁵ und der physik (the problem of weather prediction, considered from the viewpoints of mechanics
 ⁴²⁶ and physics). *Meteorologische Zeitschrift*, **21**, 1–7, doi:10.1127/0941-2948/2009/416, translated
 ⁴²⁷ and edited by E. Volken and S. Brönnimann ^{..}. Meteorol. Z. 18 (2009), 663–667.
- Bryan, K., 2006: Modeling ocean circulation. *Physical Oceanography: Developments since 1950*,
 Springer New York, 29–44, doi:10.1007/0-387-33152-2_3.
- Bryan, K., and L. J. Lewis, 1979: A watermass model of the world ocean. J. Geophys. Res., 84,
 2503–2517, doi:10.1029/jc084ic05p02503.
- Cartwright, D. E., 2012: *Tides*. Cambridge University Press, Cambridge, United King dom and New York, NY, USA, URL https://www.ebook.de/de/product/2966034/david_edgar_
 cartwright_tides.html.

- ⁴³⁵ Coveney, P. V., E. R. Dougherty, and R. R. Highfield, 2016: Big data need big theory too.
 ⁴³⁶ *Phil. Trans. R. Soc. Lond. A*, **374** (**2080**), 20160153, doi:10.1098/rsta.2016.0153.
- ⁴³⁷ Dwivedi, S., T. W. N. Haine, and C. E. Del Castillo, 2011: Upper ocean state estimation in ⁴³⁸ the Southern Ocean Gas Exchange Experiment region using the four–dimensional variational
- technique. J. Geophys. Res., **116**, **C00F02**, doi:10.1029/2009JC005615.
- Emanuel, K., 2020: The relevance of theory for contemporary research in atmospheres, oceans,
 and climate. *AGU Advances*, 1 (2), doi: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. *Q. J. R. Meteorol. Soc.*, **139** (**674**), 1162–1178, doi:
 10.1002/qj.2027.
- Feigenbaum, E. A., 2003: Some challenges and grand challenges for computational intelligence.
 Journal of the ACM, 50 (1), 32–40, doi:10.1145/602382.602400.
- ⁴⁴⁸ Fox-Kemper, B., and Coauthors, 2019: Challenges and prospects in ocean circulation models.
 ⁴⁴⁹ *Frontiers in Marine Science*, **6**, doi: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. Mod. Dev.*, **11** (**4**), 1665–1681,
 doi:10.5194/gmd-11-1665-2018.
- GCOS, 2016: The global observing system for climate. Tech. Rep. GCOS No. 200, World
 Meteorological Organization. https://library.wmo.int/index.php?lvl=notice_display&id=19838#
 .XxWxPJNKhBw.

- ⁴⁵⁶ Gent, P. R., 2011: The Gent-McWilliams parameterization: 20/20 hindsight. *Ocean Modelling*,
 ⁴⁵⁷ **39**, 2–9, doi:10.1016/j.ocemod.2010.08.002.
- Giles, M. B., and I. Reguly, 2014: Trends in high-performance computing for engineering calculations. *Phil. Trans. R. Soc. Lond. A*, **372 (2022)**, 20130 319, doi:10.1098/rsta.2013.0319.
- 460 Griffies, S. M., 2004: Fundamentals of Ocean Climate models. Princeton University Press, Prince-
- ton, NJ, USA and Oxford, United Kingdom, 518 pp., doi:10.2307/j.ctv301gzg.
- Griffies, S. M., and Coauthors, 2000: Developments in ocean climate modelling. *Ocean Modelling*,
 2 (3-4), 123–192, doi:10.1016/s1463-5003(00)00014-7.
- Gropp, W., and M. Snir, 2013: Programming for exascale computers. *Computing in Science & Engineering*, **15** (6), 27–35, doi:10.1109/mcse.2013.96.
- Haine, T. W. N., 2010: High-frequency fluctuations in Denmark Strait Overflow transport. *Geo- phys. Res. Lett.*, **37**, L14 601, doi: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, 109 210, doi:
 10.1016/j.jcp.2019.109210.
- Harel, D., 2005: A Turing-like test for biological modeling. *Nature Biotechnology*, 23 (4), 495–496,
 doi:10.1038/nbt0405-495.
- ⁴⁷³ Held, I. M., 2005: The gap between simulation and understanding in climate modeling.
 ⁴⁷⁴ Bull. Amer. Meteor. Soc., 86 (11), 1609–1614, doi:10.1175/bams-86-11-1609.
- 475 Holland, W. R., and J. C. McWilliams, 1987: Computer modeling in physical oceanography from
- the global circulation to turbulence. *Physics Today*, **40** (**10**), 51–57, doi:10.1063/1.881115.

⁴⁷⁷ Hong, S.-Y., and J. Dudhia, 2012: Next-generation numerical weather prediction: Bridging pa⁴⁷⁸ rameterization, explicit clouds, and large eddies. *Bull. Amer. Meteor. Soc.*, **93** (1), ES6–ES9,
⁴⁷⁹ doi:10.1175/2011bams3224.1.

Kalnay, E., 2002: Atmospheric Modeling, Data Assimilation and Predictability. Cambridge
 University Press, Cambridge, United Kingdom and New York, NY, USA, doi:10.1017/
 cbo9780511802270.

Klinger, B. A., and T. W. N. Haine, 2019: Ocean Circulation in Three Dimensions. 1st ed.,
 Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, URL
 http://www.cambridge.org/9780521768436.

Kutz, J. N., 2017: Deep learning in fluid dynamics. J. Fluid Mech., 814, 1–4, doi: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. Mod. Dev.*, **11** (**5**), 1799–1821, doi:10.5194/ gmd-11-1799-2018.

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, doi:10.17125/gov2018.ch12.

Lea, D. J., T. W. N. Haine, and R. F. Gasparovic, 2006: Observability of the Irminger Sea circulation using variational data assimilation. *Q. J. R. Meteorol. Soc.*, **132** (**618**), 1545–1576, doi:10.1256/qj.05.77.

24

- LeGrand, P., and C. Wunsch, 1995: Constraints from paleotracer data on the North Atlantic circulation during the Last Glacial Maximum. *Paleoceanography*, **10** (**6**), 1011–1045, doi: 10.1029/95pa01455.
- Lguensat, R., J. Le Sommer, S. Metref, E. Cosme, and R. Fablet, 2019: Learning generalized quasi-geostrophic models using deep neural numerical models, 1911.08856v1.
- Li, Y., and Coauthors, 2008: A public turbulence database cluster and applications to study lagrangian evolution of velocity increments in turbulence. *Journal of Turbulence*, **9**, N31, doi: 10.1080/14685240802376389.
- Lorenz, E. N., 1967: *The Nature and Theory of the General Circulation of the Atmosphere*. WMO:
 TP, World Meteorological Organization.
- ⁵⁰⁷ Lynch, P., 2008: The origins of computer weather prediction and climate modeling. *J. Com-*⁵⁰⁸ *put. Phys.*, **227** (7), 3431–3444, doi:10.1016/j.jcp.2007.02.034.
- ⁵⁰⁹ MacKinnon, J. A., and Coauthors, 2017: Climate Process Team on internal-wave driven ocean ⁵¹⁰ mixing. *Bull. Amer. Meteor. Soc.*, **98** (**11**), 2429–2454, doi:10.1175/bams-d-16-0030.1.
- ⁵¹¹ Magaldi, M. G., and T. W. N. Haine, 2015: Hydrostatic and non-hydrostatic simulations of dense ⁵¹² waters cascading off a shelf: the East Greenland case. *Deep Sea Res.*, *Part I*, **96**, 89–104, ⁵¹³ doi:10.1016/j.dsr.2014.10.008.
- ⁵¹⁴ Magaldi, M. G., T. W. N. Haine, 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, doi:10.1175/JPO-D-10-05004.1.
- ⁵¹⁷ Magnusson, L., and E. Källén, 2013: Factors influencing skill improvements in the ECMWF ⁵¹⁸ forecasting system. *Mon. Weather. Rev.*, **141** (**9**), 3142–3153, doi:10.1175/mwr-d-12-00318.1.

25

- ⁵¹⁹ Maltrud, M. E., and J. L. McClean, 2005: An eddy resolving global 1/10° ocean simulation. *Ocean* ⁵²⁰ *Modelling*, **8** (1-2), 31–54, doi:10.1016/j.ocemod.2003.12.001.
- Maltrud, M. E., 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 (C13), 30825–30853, doi:10.1029/1998jc900013.
- ⁵²⁴ McWilliams, J. C., 1996: Modeling the oceanic general circulation. *Ann. Rev. Fluid Mech.*, **28** (1), ⁵²⁵ 215–248, doi:10.1146/annurev.fl.28.010196.001243.
- McWilliams, J. C., 2007: Irreducible imprecision in atmospheric and oceanic simulations. *Proc. Nat. Acad. Sci.*, **104 (21)**, 8709–8713, doi:10.1073/pnas.0702971104.
- Medvedev, D., G. Lemson, and M. Rippin, 2016: SciServer compute. *Proceedings of the 28th International Conference on Scientific and Statistical Database Management SSDBM '16*,
 ACM Press, doi:10.1145/2949689.2949700.
- Moore, G. E., 1975: Progress in digital integrated electronics IEDM Tech Digest. IEEE.
- Morrow, R., and Coauthors, 2019: Global observations of fine-scale ocean surface topography with the surface water and ocean topography (SWOT) mission. *Frontiers in Marine Science*, **6**, doi:10.3389/fmars.2019.00232.
- Nguyen, A. T., H. Pillar, V. Ocaña, A. Bigdeli, T. A. Smith, and P. Heimbach, 2020: The
 arctic subpolar gyre sTate estimate (ASTE): Description and assessment of a data-constrained,
 dynamically consistent ocean-sea ice estimate for 2002-2017. doi:10.1002/essoar.10504669.1.
- ⁵³⁸ Overpeck, J. T., G. A. Meehl, S. Bony, and D. R. Easterling, 2011: Climate data chal-⁵³⁹ lenges in the 21st century. *Science*, **331** (**6018**), 700–702, doi:10.1126/science.1197869, URL

- https://science.sciencemag.org/content/331/6018/700, https://science.sciencemag.org/content/
 331/6018/700.full.pdf.
- Palem, K. V., 2014: Inexactness and a future of computing. *Phil. Trans. R. Soc. Lond. A*, **372 (2018)**,
 20130 281, doi: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. *Q. J. R. Meteorol. Soc.*, **138** (665), 841–861, doi:10.1002/qj.
 1923.
- ⁵⁴⁷ Perlman, E., R. Burns, Y. Li, and C. Meneveau, 2007: Data exploration of turbulence simulations
- using a database cluster. Proceedings of the 2007 ACM/IEEE conference on Supercomputing -

⁵⁴⁹ *SC* '07, ACM Press, doi:10.1145/1362622.1362654.

- Richardson, L. F., 1922: Weather Prediction by Numerical Process. Cambridge University
 Press, Cambridge, United Kingdom and New York, NY, USA, doi:10.1017/cbo9780511618291,
 reprinted 2007.
- ⁵⁵³ Rocha, C. B., T. K. Chereskin, S. T. Gille, and D. Menemenlis, 2016: Mesoscale to submesoscale
 ⁵⁵⁴ wavenumber spectra in Drake Passage. J. Phys. Oceanogr., 46 (2), 601–620, doi:10.1175/
 ⁵⁵⁵ 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** (8), 2393–2414,
 doi:10.1175/JPO-D-19-0210.1.
- ⁵⁵⁹ Schreiber, M., P. S. Peixoto, T. Haut, and B. Wingate, 2017: Beyond spatial scalability limitations ⁵⁶⁰ with a massively parallel method for linear oscillatory problems. *The International Journal of*
- ⁵⁶¹ *High Performance Computing Applications*, **32** (**6**), 913–933, doi:10.1177/1094342016687625.

- Semtner, A. J., and R. M. Chervin, 1992: Ocean general circulation from a global eddy-resolving
 model. J. Geophys. Res., 97 (C4), 5493, doi:10.1029/92jc00095.
- Simmons, A. J., and A. Hollingsworth, 2002: Some aspects of the improvement in skill of
 numerical weather prediction. *Q. J. R. Meteorol. Soc.*, **128** (580), 647–677, doi: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** (C9), doi:10.1029/
 2001jc000888.
- ⁵⁷⁰ Stewart, R. H., 2008: *Introduction to physical oceanography*. September 2008 edition ed., pub-⁵⁷¹ lished online by the author at: http://oceanworld.tamu.edu/home/course_book.htm.
- Thomas, L. N., A. Tandon, and A. Mahadevan, 2008: Submesoscale processes and dynamics.
 Ocean Modeling in an Eddying Regime, American Geophysical Union, 17–38, doi:10.1029/
 177gm04.
- ⁵⁷⁵ Turing, A. M., 1950: I. Computing machinery and intelligence. *Mind*, **LIX (236)**, 433–460, doi:10.1093/mind/lix.236.433.
- ⁵⁷⁷ Vallis, G. K., 2016: Geophysical fluid dynamics: whence, whither and why? *Proc. R. Soc. Lond. A*,
 ⁵⁷⁸ 472 (2192), 20160 140, doi:10.1098/rspa.2016.0140.
- ⁵⁷⁹ Waldrop, M. M., 2016: The chips are down for Moore's law. *Nature*, **530** (**7589**), 144–147, doi:10.1038/530144a.
- ⁵⁸¹ Way, M. J., and Coauthors, 2017: Resolving orbital and climate keys of earth and extraterrestrial ⁵⁸² environments with dynamics (ROCKE-3d) 1.0: A general circulation model for simulating

- the climates of rocky planets. *The Astrophysical Journal Supplement Series*, **231** (1), 12, doi: 10.3847/1538-4365/aa7a06.
- ⁵⁸⁵ Wunsch, C., 1996: *The ocean circulation inverse problem*. 1st ed., Cambridge University Press,
 ⁵⁸⁶ Cambridge, United Kingdom and New York, NY, USA, doi:10.1017/cbo9780511629570, 442
 ⁵⁸⁷ pp.
- Wunsch, C., 2006: Discrete Inverse and State Estimation Problems. 1st ed., Cambridge University
- Press, Cambridge, United Kingdom and New York, NY, USA, doi:10.1017/cbo9780511535949,

⁵⁹⁰ 371 pp.

591 LIST OF FIGURES

592 593 594 595 596 597 598	Fig. 1.	OGCM Turing tests. In each of (a)–(c) 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, $1Sv = 10^6 \text{m}^3 \text{s}^{-1}$, negative means equatorwards). Adapted from Haine (2010). (b) Salinity (colors) on a section north of Denmark Strait (annual average; the heavy contour is the 27.80 σ_0 density anomaly). (c) Lagrangian trajectories of RAFOS floats and synthetic RAFOS floats. Adapted from Saberi et al. (2020).	. 31
599	Fig. 2.	Unequal exponential growth. (a) History of deep (deeper than 1000 m) ocean temperature	
600		measurements. The colored dots show different instruments and observing platforms. (b)	
601		History of sea level measurements from satellite altimetry expressed by the cumulative	
602		number of days of measurement. The satellite missions and their durations are indicated	
603		with the colored bars. (c) History of cutting-edge global OGCM and IPCC ocean model	
604		resolution expressed by the lengthscale of the horizontal grid and the number of model grid	
605		points. Each dot represents one ocean model and the OGCMs are from Bryan and Lewis	
606		(1979); Semtner and Chervin (1992); Maltrud et al. (1998); Maltrud and McClean (2005),	
607		and Rocha et al. (2016). (d) History of top supercomputers using Rmax speed (FLOPS =	
608		noating point operations per second) for fastest machines (open circles) and ECMWF and	
609		NCAR machines (closed circles). The lines show best fit exponential growth in each panel	22
610		$(\tau_{2\times}$ is the doubling time).	. 32
611	Fig. 3.	Characteristic space and time scales of the ocean general circulation. Various geophysical	
612		and theoretical scales are shown with gray patches and colored lines (for a discussion, see	
613		Klinger and Haine 2019). The colored rectangles show cutting-edge circulation models	
614		(direct numerical simulation of turbulence, large eddy simulation, the Poseidon Project run,	
615		AR6 HighResMIP, and TAR OGCMs). The black dot shows the sampling characteristics	
616		of the Argo profiling floats, and the TOPEX/Poseidon-Jason altimeters. The diagram is	
617		indicative, not definitive, because it suppresses the anisotropies and inhomogeneities present	
618		in the general circulation.	. 33



FIG. 1. OGCM Turing tests. In each of (a)–(c) 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, $1Sv = 10^6 m^3 s^{-1}$, negative means equatorwards). Adapted from Haine (2010). (b) Salinity (colors) on a section north of Denmark Strait (annual average; the heavy contour is the 27.80 σ_0 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 1000 m) ocean temperature mea-625 surements. The colored dots show different instruments and observing platforms. (b) History of sea level 626 measurements from satellite altimetry expressed by the cumulative number of days of measurement. The satel-627 lite missions and their durations are indicated with the colored bars. (c) History of cutting-edge global OGCM 628 and IPCC ocean model resolution expressed by the lengthscale of the horizontal grid and the number of model 629 grid points. Each dot represents one ocean model and the OGCMs are from Bryan and Lewis (1979); Semtner 630 and Chervin (1992); Maltrud et al. (1998); Maltrud and McClean (2005), and Rocha et al. (2016). (d) History 631 of top supercomputers using Rmax speed (FLOPS = floating point operations per second) for fastest machines 632 (open circles) and ECMWF and NCAR machines (closed circles). The lines show best fit exponential growth in 633 each panel ($\tau_{2\times}$ is the doubling time). 634



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.