Climate Modeling in Low Precision: Effects of Both Deterministic and Stochastic Rounding

E. ADAM PAXTON,^a MATTHEW CHANTRY,^a MILAN KLÖWER,^a LEO SAFFIN,^b AND TIM PALMER^a

^a *University of Oxford, Oxford, United Kingdom*

^b *University of Leeds, Leeds, United Kingdom*

(Manuscript received 4 May 2021, in final form 17 October 2021)

ABSTRACT: Motivated by recent advances in operational weather forecasting, we study the efficacy of low-precision arithmetic for climate simulations. We develop a framework to measure rounding error in a climate model, which provides a stress test for a low-precision version of the model, and we apply our method to a variety of models including the Lorenz system, a shallow water approximation for flow over a ridge, and a coarse-resolution spectral global atmospheric model with simplified parameterizations (SPEEDY). Although double precision [52 significant bits (sbits)] is standard across operational climate models, in our experiments we find that single precision (23 sbits) is more than enough and that as low as half precision (10 sbits) is often sufficient. For example, SPEEDY can be run with 12 sbits across the code with negligible rounding error, and with 10 sbits if minor errors are accepted, amounting to less than $0.1 \text{ mm (6 h)}^{-1}$ for average gridpoint precipitation, for example. Our test is based on the Wasserstein metric and this provides stringent nonparametric bounds on rounding error accounting for annual means as well as extreme weather events. In addition, by testing models using both round-to-nearest (RN) and stochastic rounding (SR) we find that SR can mitigate rounding error across a range of applications, and thus our results also provide some evidence that SR could be relevant to next-generation climate models. Further research is needed to test if our results can be generalized to higher resolutions and alternative numerical schemes. However, the results open a promising avenue toward the use of low-precision hardware for improved climate modeling.

SIGNIFICANCE STATEMENT: Weather and climate models provide vital information for decision-making, and will become ever more important in the future with a changed climate and more extreme weather. A central limitation to improved models are computational resources, which is why some weather forecasters have recently shifted from conventional 64-bit to more efficient 32-bit computations, which can provide equally accurate forecasts. Climate models, however, still compute in 64 bits, and adapting to lower precision requires a detailed analysis of rounding errors. We develop methods to quantify rounding error in a climate model, and find similar precision acceptable across weather and climate models, with even 16 bits often sufficient for an accurate climate. This opens a promising avenue for computational efficiency gains in climate modeling.

KEYWORDS: Error analysis; Numerical analysis/modeling; Climate models

1. Introduction

Modern numerical Earth system models require enormous amounts of computational resources and place significant demand on the world's most powerful supercomputers. As such, operational forecasting centers are stretched to make best use of resources and seek ways of reducing unnecessary computation wherever possible.

One idea to improve computational efficiency that has gained attention in recent years is to utilize low-precision arithmetic—in place of conventional 64-bit arithmetic—for computationally intensive parts of the code. This has been accompanied by parallel trends in deep learning where low precision is deployed routinely and for which novel hardware is now emerging (Gupta et al. 2015; Micikevicius et al. 2018).

Whether such hardware can be exploited for weather and climate, however, ultimately depends on the cumulative effect of rounding error. In fact, a number of studies have shown that much numerical weather prediction, at least on the short time scales relevant for forecasts, can be optimized for low precision (Chantry et al. 2019; Hatfield et al. 2019; Jeffress et al. 2017; Klöwer et al. 2020; Saffin et al. 2020; Prims et al. 2019) and forecasting centers are already exploiting this in operations. Indeed, the European Centre for Medium-Range Weather Forecasts recently ported the atmospheric component of its flagship Integrated Forecast System to single precision (Maass 2021; Vána et al. 2017), while MeteoSwiss and the U.K. Met Office have previously implemented single- and mixed-precision codes respectively (Rüdisühli et al. 2014; Gilham 2018).

As operational weather forecasters experiment with more efficient low-precision hardware, it is natural to ask whether low precision is suitable for climate modeling (i.e., long time scales) and this is the question addressed by the current paper. Compared to weather forecasting, where research in low precision has focused to date, climate modeling presents a different problem requiring some new techniques. While an

Denotes content that is immediately available upon publication as open access.

Corresponding author: E. Adam Paxton, edmund.paxton@physics.ox.ac.uk

FIG. 1. An example of the Wasserstein distance (WD). The WD from g_1 to f is 1, where an optimal transport strategy is to shift all of the probability mass up by one unit. By contrast, it is easily seen that $\text{WD}(f, g_2) > 1$, reflecting the intuition that g_1 is a better approximation of f than is g_2 . The figure is motivated by [Robin et al. \(2017\)](#).

ensemble weather forecast seeks a relatively localized probability distribution over the possible states of the atmosphere at a given time, the exact state is understood to be totally unpredictable on long time scales due to chaos, and a climate model seeks instead to approximate the statistics of states over a long time period (in the language of ergodic theory, the climatological object of interest is the invariant probability distribution). Thus the test for a low-precision climate model should be whether it has the same statistics (invariant distribution) as its high-precision counterpart.

We develop such a test based on the Wasserstein distance (WD) from optimal transport theory, which provides a natural notion of closeness between probability distributions. The WD is defined as the cost of an optimal strategy for transporting probability mass between distributions with respect to a cost $c(x, y)$ of transporting unit mass from $x \in \mathbb{R}^d$ to $y \in \mathbb{R}^d$, where throughout this paper we take $c(x, y) = |x - y|$ so that cost has the same units as the underlying field.

The WD is an appropriate metric for this study because 1) it is nonparametric, 2) it has favorable geometrical properties (see [Fig. 1](#)), 3) it is interpretable in appropriate physical units, and 4) it bounds a range of expected values covering both the mean response and extreme weather events. The metric is popular in machine learning ([Arjovsky et al. 2017](#)) and has recently been suggested as an appropriate measure of skill in climate modeling ([Robin et al. 2017](#); [Vissio et al. 2020](#); [Vissio and Lucarini 2018](#)); however, since it is not so well known within the community, a survey of the WD, including a rigorous definition and discussion of computational techniques, is given in [appendix A](#). In particular, see [sections a and e of appendix A](#) for further discussion of points 1–4.

Any test can only bound the effects of rounding error at low precision relative to the variability of probability distributions generated by a corresponding high-precision experiment. Experiments must thus be carefully designed to minimize such variability in order to isolate the effects of rounding error. For example, by taking an ensemble of sufficiently long integrations one can reduce initial condition variability, and by keeping external factors such as greenhouse

gas emissions annually periodic one can reduce variability due to nonstationarity. A choice of metric with strong properties (e.g., points 1–4 listed above) is then crucial for interpretation of the resulting bounds on rounding error in order to have confidence in the reliability of a low-precision model. Although we developed our methods to measure rounding error, we hope they might also be of interest to the broader climate modeling community.

a. Structure of the document

[Section 2](#) covers the Lorenz system, intended to illustrate our methodology using a low-dimensional and well-known chaotic dynamical system. [Section 3](#) then extends the analysis to a high-dimensional dynamical system via a finite-difference scheme for a shallow-water model. [Section 4](#) presents a short interlude into a simple heat diffusion model, the main aim of which is to illustrate clearly the effect of stochastic rounding in preventing stagnation in time-stepping schemes. [Section 5](#) analyses the SPEEDY model in low precision, both through an idealized El Niño experiment and through the detailed climatological methods developed in [sections 2 and 3](#). [Section 6](#) provides conclusions and poses areas for further research. Acknowledgments are then given along with links to our source code. [Appendix A](#) provides relevant background on the Wasserstein distance, which is central to our work, while [appendix B](#) documents the different arithmetic formats and rounding modes that are referenced in the paper.

b. Notation

In this paper we write Float64, Float32, and Float16 to denote the IEEE-754 standard formats for double, single, and half precision respectively, and BFloat16 for Google's Brain floating point format (see [Fig. B1 in appendix B](#)), with round-to-nearest used implicitly as the rounding scheme. We write Float32sr, Float16sr, and BFloat16sr whenever stochastic rounding is used instead (see [appendix B](#)). All low-precision formats (i.e., everything except Float64) have been implemented in software, rather than in hardware; see the data availability statement (following the acknowledgments) for details.

FIG. 2. The Lorenz system integrated and data-binned at different precisions to approximate the invariant probability distribution. Warmer color corresponds to more probability mass.

2. The Lorenz system

To illustrate our methods, which will later be applied to more complex climate models, we first consider the Lorenz system:

$$\begin{aligned} \frac{dx}{dt} &= 10(y - x), \\ \frac{dy}{dt} &= \left(\frac{8}{3} - z\right)x - y, \\ \frac{dz}{dt} &= xy - 28z. \end{aligned} \quad (1)$$

Derived first in [Lorenz \(1963\)](#) in a study of convection, (1) exhibits features of nonlinear dynamics representative of the real atmosphere ([Molteni and Kucharski 2018](#)). The system has a fractal attractor $\mathcal{A} \subseteq \mathbb{R}^3$ and the dynamics on $\mathcal{A}$ is chaotic, rendering the precise numerical approximation of any specific orbit futile. On the other hand, a study of asymptotics on $\mathcal{A}$ reveals statistics common across orbits, which one may hope to approximate. Indeed, there exists a unique invariant probability distribution ν supported on $\mathcal{A}$ such that for almost any orbit $\mathbf{x}(t) = [x(t), y(t), z(t)]$ initiated in the basin of attraction of $\mathcal{A}$ and any bounded continuous function ϕ

$$\lim_{T \rightarrow \infty} \left\{ \frac{1}{T} \int_0^T \phi[\mathbf{x}(t)] dt \right\} = \int \int \int \phi(\mathbf{x}) d\nu(\mathbf{x}) \quad (2)$$

(cf. [Tucker 1999](#)). Thus ν encodes the long-time statistics of the system. For example, taking $\phi(\mathbf{x}) = 1$ for $\mathbf{x} \in B$ and $\phi(\mathbf{x}) = 0$ outside of a neighborhood of B , from (1) we see the average time an orbit spends in a region B is the

probability mass $\nu(B)$. In the context of climate modeling, the test for a low-precision integration of (1) is whether it produces approximately the same ν as its high-precision counterpart.

We first sampled 10 initial conditions $i_0, \dots, i_9 \in \mathbb{R}^3$ from a normal distribution with unit variance centered at the center of mass of the attractor $\mathcal{A}$. We then integrated (1) at high precision (Float64) initialized at initial conditions $i_0, \dots, i_4$ for 220 000 model time units (mtu) each and discarded the first 20 000 mtu as spinup to allow for any orbits initially perturbed off $\mathcal{A}$ to return to $\mathcal{A}$, and we labeled these five runs as the control ensemble e_i^{control} . Next, for each comparison arithmetic (including Float64 as a control) we integrated (1) initialized at initial conditions $i_5, \dots, i_9$ for 220 000 mtu, discarded the first 20 000 mtu, and labeled these as the competitor e_j . Each integration used the Runge–Kutta fourth-order scheme with a time step of $dt = 0.002$. For background on the different arithmetic formats and stochastic rounding, see [appendix B](#).

In general, results will be sensitive to both the choice of numerical scheme and the time step; however, we will not dwell on such issues since our aim is to develop a method to measure rounding error in a climatological context, rather than to obtain an optimal integration.

Integrations are binned and plotted in [Fig. 2](#). While the Float32, Float32sr, and Float16sr integrations appear to approximate the high-precision attractor well, Float16 suffers from the small number of available states forcing the evolution into an early periodic orbit, while we found that BFloat16 collapsed on a point attractor with this fine time step, although both of these integrations were notably improved by stochastic rounding. For each arithmetic we computed the ensemble mean Wasserstein distance (WD; see [appendix A](#))

FIG. 3. Measuring rounding error in the climate of the Lorenz system. Plot shows Wasserstein distances between the probability distributions generated by low-precision and high-precision ensembles (bin width 6.0) where black is high precision vs high precision for reference.

between the five probability distributions generated by e_i and the five distributions generated by e_j^{control} , the evolution of this quantity with time is plotted in Fig. 3 with a log-log scale. Note that the Float64 competitor (black in Fig. 3) shows the mean WD between a pair of high-precision integrations initiated at different initial conditions, and thus gives a measure of the variability of the experiment at high precision, which is important to be able to draw conclusions. Figure 3 confirms quantitatively what is suggested by Fig. 2 in that the Float32, Float32sr, and Float16sr curves closely follow Float64 in the approach to statistical equilibrium, showing that rounding error is small relative to high-precision variability, while for BFloat16sr rounding error has notably perturbed the dynamics. The size of the high-precision variability after 100 000 mtu (Fig. 3) is less than 0.1 model space units (msu), which is small in the context of distributions supported on the Lorenz attractor, which has characteristic length scale of approximately 50 msu (Fig. 2).

To compute WDs we approximated the probability distributions by data binning with cubed bins and a bin width of 6 msu. This is a coarse estimation but we found results were not sensitive to decreasing bin width [in agreement with Vissio and Lucarini (2018)] and we also performed the same computation approximating by the empirical distributions generated by 2500 samples as well as approximating by Sinkhorn divergences, and by marginalizing onto one-dimensional distributions (see section c in appendix A), all of which gave analogous results. The methods based upon sampling and marginalization provide alternatives to data-binning in higher dimensional settings, as we will come to in the next section.

3. A shallow water model

We next consider the shallow water model from Klöwer et al. (2020), which describes turbulent flow in a rectangular ocean basin, driven by a steady zonally symmetric wind forcing over a meridionally symmetric ridge. The equations are

$$\begin{aligned} \frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} + f \hat{\mathbf{z}} \times \mathbf{u} &= -g \nabla \eta + \mathbf{D} + \mathbf{F}, \\ \frac{\partial \eta}{\partial t} + \nabla \cdot (\mathbf{u} h) &= 0, \end{aligned}$$

where $\mathbf{u} = (u, v)$ is velocity, η is surface elevation, $\mathbf{D} = -(c_D/h)|\mathbf{u}| - \nu \nabla^4 \mathbf{u}$ is a nonlinear diffusive term with coefficients c_D and ν (bottom friction and biharmonic viscosity coefficient), $\mathbf{F}$ is wind forcing, f is the Coriolis parameter, $h = H + \eta$ is layer thickness, and H is the time-independent depth of the water at rest describing the ridge at the fluid base. The ocean basin dimensions were taken as 2000 km $\times$ 1000 km with average depth of 500 m. We integrated the equations using the scheme from Klöwer et al. (2020), which uses finite differences on an Arakawa C-grid and fourth-order Runge–Kutta in time combined with a semi-implicit scheme for the dissipative terms, with a time step of 6 h, and refer to Klöwer et al. (2020) for more details on the numerics.

Following the methodology developed in section 2 we integrated for 20 years discarding the first year of each run as spinup, taking a five-member ensemble for each arithmetic. Snapshots of the evolution in each case are plotted in Fig. 4a. We computed ensemble mean pairwise WDs between the distributions generated by high- and low-precision ensembles and plot this quantity evolving with time in Fig. 4b. We found

that for Float16 and BFloat16sr rounding error is significant while for Float32 and Float16sr rounding error is small relative to high-precision variability. In particular, our results show that rounding errors at half-precision are successfully mitigated by stochastic rounding in this climate experiment. Again, we refer to [appendix B](#) for background on these different number formats.

The main difference between the methods of this section and [section 2](#) is that we considered here a dynamical system that is high dimensional, so that approximating probability distributions is nontrivial. For [Fig. 4b](#) we approximated the invariant distributions by taking 2500 uniformly distributed random samples in time and computed WDs between the corresponding empirical distributions (see [section c](#) of [appendix A](#)). This approximation method does not give readily interpretable results due to a curse of dimensionality; however, we obtained analogous results by marginalizing onto one-dimensional subspaces. We save discussion of such marginalized results for [section 5](#) in the context of a global atmospheric model.

4. Interlude: Heat diffusion in a soil column

In this section we briefly consider a very simplified land surface component of a global climate model. This is a trivial case of climatology since all solutions converge upon a constant equilibrium temperature and so there is no need to use the WD in this setting. We include this simple example, however, because it clearly illustrates a major advantage of stochastic rounding (SR) over round-to-nearest (RN) in preventing stagnation.

This section was partially motivated by [Harvey and Verseghy \(2016\)](#) and [Dawson et al. \(2018\)](#). In [Harvey and Verseghy \(2016\)](#) the authors observed that the Canadian Land Surface Scheme (CLASS) could not be run effectively at single precision in large part because of an issue of stagnation. They argued that single-precision arithmetic was not appropriate for climate modeling with the scheme, which relies crucially on accurate representations of slowly varying processes such as permafrost thawing, and that double precision or even quadruple precision should be adopted instead. The setup considered here was introduced in [Dawson et al. \(2018\)](#) as a toy model that retained some features of CLASS, most crucially the stagnation at single precision with RN. The authors of [Dawson et al. \(2018\)](#) proposed mixed precision to avoid stagnation, while the results of this section indicate that SR provides an alternative approach.

Following [Dawson et al. \(2018\)](#), we consider an idealized soil column that is heated from the top and thermally insulated from the bottom:

$$\frac{\partial T}{\partial t} = D \frac{\partial^2 T}{\partial z^2},$$

with $T(t, 0) = 280$, $(\partial T / \partial z)(t, H) = 0$, and $T(0, z) = 273$, where $T(t, z)$ is temperature in kelvin, $H = 60$ m is soil depth, and $D = 7 \times 10^{-7} \text{ m}^2 \text{ s}^{-1}$ is the coefficient of diffusivity, and discretize as

$$T_j^{n+1} = T_j^n + D\Delta t \left[\frac{T_{j+1}^n - 2T_j^n + T_{j-1}^n}{(\Delta z)^2} \right], \quad (3)$$

with $\Delta z = 1$ m and $\Delta t = 1800$ s.

We integrated for 100 years, and the results are plotted in [Fig. 5](#). Stagnation is apparent for Float32 and Float16 where the small tendency term in (3) is repeatedly rounded down to zero, so that heat does not diffuse effectively through the soil column. This is mitigated by SR, however, which assigns a nonzero probability of rounding up after the addition in (3) (see [section b](#) of [appendix B](#)). Rounding error is negligible with Float32sr; while visible as noise in Float16sr, the solution shares the large-scale pattern of Float64.

To be clear, this section is not intended to imply that SR is necessary in the low-precision integration of the heat equation. There are other ways to avoid stagnation, such as increasing the time step, which is extremely small in this example and well below what is necessary for stability, or implementing a compensated summation for the time-stepping. Rather, this section aims to illustrate an interesting advantage of SR in mitigating stagnation by means of a clear and visual example. For more analysis of SR in the numerical solution of the heat equation see [Croci and Giles \(2020\)](#).

5. A global atmospheric model

Finally, we proceed to a global atmospheric circulation model: the Simplified Parameterizations Primitive Equation Dynamics version 41 (SPEEDY). SPEEDY is a coarse-resolution model employing a T30 spectral truncation with a 48×96 latitude–longitude grid, eight vertical levels, and a 40-min time step, and is forced by annually periodic fields obtained from ERA reanalysis together with a prescribed sea surface temperature anomaly [Kucharski et al. \(2013\)](#). For this section, in order to isolate the effects of numerical precision we truncated only the significant bits, so that when we speak of half precision, for example, we refer to 10 significant bits (sbits) and 11 exponent bits rather than the IEEE-754 standard 5 exponent bits (see [section a](#) of [appendix B](#)).

As a first test, we constructed a constant-in-time SST anomaly field to crudely simulate an El Niño event and ran SPEEDY both with and without to investigate the mean field response. This anomaly field was constructed [partially following [Dogar et al. \(2017\)](#)] by taking the Pearson correlation coefficients between an ERA reanalysis time series at each grid point and the Niño-3.4 index over 1979–2019 and multiplying by a factor of 4 in an attempt to produce temperature anomalies in Kelvin roughly of the magnitude of the 2015 El Niño. [Figure 6](#) shows the El Niño response for precipitation and geopotential height at 500 hPa (Z500) for both double and half precision and it is seen that the latter certainly simulates a similar response to the El Niño. The area-weighted Pearson correlations between the double

FIG. 4. The shallow water model integrated at different precision levels. (a) A snapshot of the flow speed (m s^{-1}), initiated from the same initial condition, after 50 days. (b) The effect of rounding error on climatology.

precision and half precision mean El Niño responses for the northern extratropics, southern extratropics, and tropics were calculated as (0.99, 0.99, 0.99) for precipitation and (0.98, 0.99, 0.93) for Z500.

To explore the full climatology, we next followed the WD calculations of sections 2 and 3. We generated initial conditions $i_0, \dots, i_9$ by integrating from rest for 11 years at 51 sbits of precision (effectively double precision plus a tiny rounding error) before discarding the first year as spinup and taking the initial conditions from the starts of each of the 10 subsequent years. This method was intended to emulate sampling from the high-precision invariant distribution while avoiding overlap in the high-precision ensemble. We then constructed our control ensemble e_i^{control} and competitor ensembles e_j by integrating for 10 years from the initial conditions $i_0, \dots, i_4$ and $i_5, \dots, i_9$ respectively. The SST

anomaly was turned off so that boundary conditions were annually periodic.

To circumvent issues of dimensionality (see section c of appendix A) we first marginalized onto the distributions spanned by individual spatial grid points and measure error by WDs between these 1D distributions. We call these grid-point Wasserstein distances (GPWDs) and note that this is the approach adopted in Vissio et al. (2020). To address correlations between grid points, we then checked our GPWD results against approximate WDs between the full distributions, which were obtained via a Monte Carlo sampling approach as was done in section 3. While such results are harder to interpret quantitatively (see appendix A) we found that they were analogous to the GPWD results. In particular, no errors were detected by this method which were not present in the GPWDs.

FIG. 5. Heat diffusion in a soil column with different number formats and rounding modes.

The gridpoint mean and 95th percentile GPWD results for total precipitation, Z500, and horizontal wind speed at 500 hPa are plotted in Fig. 7 as they evolve with time. To give insight into the spatial distribution of rounding errors, Fig. 8 shows maps of both the absolute error

$$\text{mean}\left[\text{WD}\left(e_i^{\text{low-precision}}, e_j^{\text{control}}\right) - \text{WD}\left(e_i^{\text{high-precision}}, e_j^{\text{control}}\right)\right], \quad (4)$$

and the log relative error

$$\log_{10}\left\{\frac{\text{mean}\left[\text{WD}\left(e_i^{\text{low-precision}}, e_j^{\text{control}}\right)\right]}{\text{mean}\left[\text{WD}\left(e_i^{\text{high-precision}}, e_j^{\text{control}}\right)\right]}\right\}, \quad (5)$$

for precipitation (convective and large-scale combined) across grid points after the 10-yr integrations have completed.

For both geopotential height and horizontal wind speed we found that rounding error was negligible relative to high-precision for 12 sbits and above, while a small rounding error emerged at 10 sbits. For precipitation the picture was similar, except with a very small rounding error emerging at 12 sbits. Figure 8 reveals those grid points at which rounding error becomes significant relative to high-precision variability for precipitation at 10 sbits. Rounding error is negligible relative to high-precision variability across all grid points for 14 sbits and above, and the mean high-precision variability for precipitation after 10 years is around 0.04 mm

$(6 \text{ h})^{-1}$, which is very small (see section e of appendix A for interpretation of WDs in terms of expected values). Moreover, the rounding error at 10 sbits is small, with gridpoint mean values of $0.07 \text{ mm } (6 \text{ h})^{-1}$, 5 m, and 0.3 m s^{-1} for precipitation, geopotential height, and horizontal wind speed respectively, and with the worst affected grid points seeing errors on the order of $1 \text{ mm } (6 \text{ h})^{-1}$, 25 m, and 1 m s^{-1} respectively (recall that these values provide bounds on annual means as well as extreme weather events; see section e of appendix A). To give more intuition behind the size of rounding error at 10 sbits, the probability distributions for precipitation at some of the worst affected tropical grid points (coastal Suriname at 5.56°N , 56.25°W and western Nigeria at 9.27°N , 3.75°E) after 10 years are plotted in Fig. 9. It is clearly seen that the difference between double and half precision, even at these worst affected grid points, is slight. It may also be noted from Figs. 7 and 8 that stochastic rounding partially mitigates rounding error at half-precision.

We also computed differences in annual mean precipitation and found that, by and large, these were of the same order as the GPWDs, indicating that precipitation error was largely accounted for by differences in the means. In general, however, WD bounds are much stronger than mean bounds (see section e of appendix A) so we can be confident that our estimates give stringent bounds on rounding error.

To summarize this section, with external forcings held annually periodic we found that a well-defined large-to-medium scale structure of the invariant probability distribution of SPEEDY

FIG. 6. Annual mean response to El Niño: double vs half precision.

emerged after 10 years. This statement is quantified by the high-precision variability as measured by the gridpoint mean GPWD (black curve in Fig. 7), which after 10 years was about $0.05 \text{ mm (6 h)}^{-1}$ for precipitation, for example. The finer-scale structures, which account for less than $0.05 \text{ mm (6 h)}^{-1}$ in gridpoint mean GPWD, remain ill defined, so we cannot conclude that there is no rounding error, but only that any potential rounding error must be smaller than $0.05 \text{ mm (6 h)}^{-1}$. If we were to increase the integration time, we would expect that the high-precision variability would decrease as finer-scale structures in the invariant distribution emerge, which would give sharper bounds on rounding error. In fact, our empirical results indicate an approximate power law in the rate of decrease of the high-precision variability, as seen in the linear structure of Fig. 7, which gives some indication of how long a modeler might have to integrate for to obtain a desired high-

precision variability. It is up to the climate model developer to determine what is an acceptable bound on rounding error. For the case of SPEEDY, we felt that the measured high-precision variabilities after 10 years were small relative to existing model biases, and thus provided an appropriate bound. Moreover, Fig. 9 shows that even at the worst affected grid points, the effects of rounding error at half-precision are slight.

6. Conclusions

While there is now convincing evidence that low-precision arithmetic can be suitable for accurate numerical weather prediction, before this work there had not been a detailed study of the effects of rounding error on climate simulations, and we have set out to address this imbalance.

FIG. 7. Measuring rounding error in SPEEDY climatology with grid point Wasserstein distances. Double precision variability in black provides a reference.

We have argued that an appropriate metric to measure rounding error in the context of a chaotic climate model is the Wasserstein distance (WD), an intuitive and nonparametric metric that provides bounds on a range of expected values including those relevant for extreme weather events. By constructing experiments minimizing the variability between probability distributions at high precision and comparing WDs against low precision we have obtained stringent bounds on rounding error, and we have found that error is typically insignificant until truncating as low as half precision in our climate experiments.

We cannot conclude from our results that a state-of-the-art Earth system model can be run with equally low precision, since such codes are hugely complex, which can make low-precision issues difficult to overcome. However, given that the unit round-off error scales exponentially with the number of sbits, it would appear that the current industry standard of double precision across all model components is likely overkill. In terms of acceptable precision, our results for SPEEDY are similar to those found in an analysis of the initial-value problem, suggesting that a level of precision suitable on weather time scales might also be suitable for climate for this model—something that is not obvious a priori. In light of recent operational successes with single-precision weather forecasting, this is a promising result in the direction of potential single-precision climate modeling; however, further research will be required to assess the generalizability of our results.

Regarding stochastic rounding (SR), although not currently in hardware, interest from machine learning together with a number of recently released patents suggest that it might become available soon (Crocì and Giles 2020). Rounding error is present in all numerical models, but with deterministic rounding schemes it can be hard to identify and may contribute to systemic biases. With SR, however, potential rounding error is appropriately treated as another source of model uncertainty, which is then sampled by an ensemble of models runs and reflected in probabilistic predictions. In addition, our experiments have shown that SR can make models more resilient to rounding error, especially at low precision. While some of the advantages of SR

are well understood, such as in the context of solving linear diffusive equations (Crocì and Giles 2020), in other settings its benefits are more obscure. Further research is called for to shed more light on the contexts in which SR and other low-precision formats can benefit weather and climate models. In addition, the community is encouraged to engage now with chip-makers in order to influence hardware development for next-generation models.

Acknowledgments. The first author would like to thank Lorenzo Pacchiardi, Stephen Jeffress, Sam Hatfield, Peter Düben, Peter Weston, and Dimitar Pashov for interesting discussions, as well as Lucy Harris for a careful reading of a first draft of the manuscript. We thank the three anonymous referees for their careful reading and helpful comments. E. A. Paxton, M. Klöwer, and T. Palmer were supported by the European Research Council Grant 741112, M. Klöwer was supported by the Natural Environmental Research Council Grant NE/L002612/1, M. Chantry and T. Palmer were supported by a grant from the Office of Naval Research Global, and T. Palmer holds a Royal Society Research Professorship.

Data availability statement. Throughout our work low precision was emulated in software. For this we found the type system of the Julia language to be well suited, and we made use of the github.com/milankl/StochasticRounding.jl and github.com/JuliaMath/BFloat16s.jl packages for stochastic rounding and BFloat16 formats. For Fortran code we used the reduced precision emulator of Dawson and Düben (2017), for which we developed a custom branch with stochastic rounding. For the Lorenz system integration we made use of github.com/milankl/Lorenz63.jl; for the shallow water model github.com/milankl/ShallowWaters.jl version 0.4; and for SPEEDY we used a branch primarily developed by Saffin for which some changes were made to optimize for low precision, which may be found at github.com/eapax/speedy. To compute optimal transport metrics in one dimension we used the `scipy.stats` package for Python while for higher-dimensional computations including Monte Carlo methods we built a custom solver at github.com/eapax/EarthMover.jl.

FIG. 8. Spatial distribution of precipitation error after 10 years [cf. (4) and (5)].

FIG. 9. Precipitation climatology at two grid points that are representative of the largest rounding errors at half precision.

APPENDIX A

The Wasserstein Distance

The Wasserstein distance (WD) defines distance between probability distributions μ and ν as the lowest cost at which one can transport all probability mass from μ to ν with respect to a cost function $c(x, y) = |x - y|^p$ that sets the cost to transport unit mass from position x to position y , where in our work we have taken $p = 1$ so that cost has the units of the underlying field. In this appendix we will first motivate the WD as a tool for the analysis of climate data by listing some of its favorable properties, before giving the formal definition of the WD, discussing methods for its computation, comparing it with other common metrics, and highlighting its interpretability through a useful dual formulation.

a. Properties of the WD

As stated earlier, the WD is defined (at least informally) as the smallest cost required to transport one probability distribution into another. Before giving the formal details of this definition, let us first motivate the WD by listing some of its favorable properties.

First, the WD is nonparametric and versatile. It does not require any specific structure of the distributions such as Gaussianity and it can be used to compare both singular and continuous distributions. This is an important point for climate modeling which presents a wide range of probability distributions. For example, the climatological distributions corresponding to South Asian rainfall or the subtropical jet stream latitude are multimodal, while for the Lorenz

system the object of interest is a singular probability distribution supported on a fractal attractor. The ability to consider singular distributions is also useful since it accommodates working directly with the empirical distributions corresponding to a sample of data, rather than first binning the data into a histogram, for example.

Second, the WD is intuitive. It may be interpreted as the minimum amount of work required to transport one distribution into the other, an idea which is readily conceptualized, and it takes the units of the underlying field. For example, for distributions of rainfall measured in millimeters per day (mm day^{-1}), a WD of 1 can be thought of as a difference of 1 mm day^{-1} . Moreover, this figure provides bounds on differences in mean rainfall as well as differences in extreme rainfall, for example (see section e of appendix A).

Third, the WD takes into account geometry. To illustrate this, consider three simple distributions with densities:

$$f(x) = \begin{cases} 1 & x \in [0, 1] \\ 0 & \text{otherwise} \end{cases},$$

$$g_1(x) = \begin{cases} 1 & x \in [1, 2] \\ 0 & \text{otherwise} \end{cases},$$

$$g_2(x) = \begin{cases} 1 & x \in [9, 10] \\ 0 & \text{otherwise} \end{cases}.$$

Now if f is taken as the true distribution and g_1 and g_2 as approximations of f , then clearly g_1 gives the better approximation, due to its proximity to f . This is reflected with $\text{WD}(f, g_1) < \text{WD}(f, g_2)$. By contrast, considering instead

the L^p distances between densities, for example, would give $\|f - g_1\|_{L^p} = \|f - g_2\|_{L^p} = 2$ for all $p \geq 1$, which does not reflect the geometry, and this is only intensified in higher dimensions as is illustrated in Fig. 1. This is not just a shortcoming of the L^p metric but is shared by measures such as the Kolmogorov–Smirnov test or the Kullback–Liebler divergence.^{A1}

More generally, the WD metrizes the space of probability distributions with respect to weak convergence, which means that closeness in the sense of the WD corresponds to closeness with respect to a natural topology (Villani 2003, theorem 7.12).

b. Formal definition of the Wasserstein distance

There are two alternative formulations of optimal transport due to Monge and Kantorovich and it is helpful to consider both when computing WDs.

For the Monge formulation, suppose we have discrete probability distributions

$$\mu = \frac{1}{N} \sum_{i=1}^N \delta_{x_i}, \quad \eta = \frac{1}{N} \sum_{i=1}^N \delta_{y_i}, \quad (\text{A1})$$

where $x_i, y_i \in \mathbb{R}^d$ and we visualize each as a distribution of equal masses on $\mathbb{R}^d$. The masses might be books on a bookshelf $\mathbb{R}$ or shipping crates on a dockside $\mathbb{R}^2$. We are tasked with transporting μ to η . The masses cannot be split so a transport strategy is a permutation σ of N objects. Introducing a cost function $c(x, y)$ defining the cost to move unit mass from position x to position y , the cost of σ is $(1/N) \sum_{i=1}^N c(x_i, y_{\sigma(i)})$ and the optimal cost is the cost of an optimal strategy $\min_{\sigma \in S_N} \left\{ (1/N) \sum_{i=1}^N c(x_i, y_{\sigma(i)}) \right\}$ where S_N is the set of permutations. The special case $c(x, y) = |x - y|$ defines (Monge's version of) the WD:

$$W_1(\mu, \eta) = \min_{\sigma \in S_N} \frac{1}{N} \sum_{i=1}^N |x_i - y_{\sigma(i)}|. \quad (\text{A2})$$

Kantorovich's formulation is a relaxation of Monge's in that masses are viewed as continuous rather than discrete (think piles of sand rather than shipping crates) so mass can be subdivided in infinitely many ways. Suppose now a pair of distributions

$$\mu = \sum_{i=1}^{M_1} P_i \delta_{x_i}, \quad \eta = \sum_{j=1}^{M_2} Q_j \delta_{y_j}, \quad (\text{A3})$$

where $\sum_i P_i = \sum_j Q_j = 1$ and $P_i, Q_j \geq 0$ (i.e., the $\mathbf{P}$ and $\mathbf{Q}$ terms are probability vectors). In applications μ and ν may represent discrete probability histograms, where the points x_i and y_j are the midpoints of bins and P_i and Q_j are weights.^{A2}

^{A1} The KL divergence is particularly ill suited to comparing probability distributions on $\mathbb{R}^d$ and this example gives $\text{KL}(f, g_1) = \text{KL}(f, g_2) = \infty$.

^{A2} Kantorovich's formulation extends to more general distributions, but we consider (A3) to simplify things.

A transport strategy is now defined as a nonnegative valued matrix $\pi \in \mathbb{R}_{\geq 0}^{M_1 \times M_2}$ where π_{ij} denotes the amount of mass transported from x_i to y_j and for conservation of mass we impose $\sum_j \pi_{ij} = P_i$, $\sum_i \pi_{ij} = Q_j$. Write c_{ij} for the cost to move unit mass from x_i to y_j so the cost of a strategy is $\sum_{i,j} c_{ij} \pi_{ij}$ and the optimal cost is $\min_{\pi \in \Pi(\mathbf{P}, \mathbf{Q})} \sum_{i,j} c_{ij} \pi_{ij}$, where

$$\Pi(\mathbf{P}, \mathbf{Q}) = \left\{ \pi \in \mathbb{R}_{\geq 0}^{M_1 \times M_2} : \sum_j \pi_{ij} = P_i, \sum_i \pi_{ij} = Q_j \right\} \quad (\text{A4})$$

is the set of possible transport strategies. Note that the space $\Pi(\mathbf{P}, \mathbf{Q})$ is also the space of joint distributions with marginals $\mathbf{P}$ and $\mathbf{Q}$, and $\Pi(\mathbf{P}, \mathbf{Q})$ is nonempty as can be seen by considering the independence distribution $\pi_{ij} = P_i Q_j$. The special case $c_{ij} = |x_i - y_j|$ defines (Kantorovich's version of) the WD:

$$W_1(\mu, \nu) = \min_{\pi \in \Pi(\mathbf{P}, \mathbf{Q})} \sum_{i,j} |x_i - y_j| \pi_{ij}. \quad (\text{A5})$$

Note that (A5) gives a linear optimization problem while (A2) has no obvious linear structure. If $M_1 = M_2 = N$ and $P_i = Q_j = 1/N$ it may be shown there is a minimizing π for (A5), which is an optimal strategy in the sense of Monge so (A2) and (A5) are consistent (Villani 2003, pp. 5–6).

c. Computing the Wasserstein distance

Suppose one has samples $\{x_i\}_{i=1}^N$ and $\{y_i\}_{i=1}^N$ drawn from a pair of distributions μ and ν on $\mathbb{R}^d$. Then one can either compute the Monge WD [(A2)] between the empirical distributions $\mu_n = (1/N) \sum_i \delta_{x_i}$ and $\eta_n = (1/N) \sum_i \delta_{y_i}$ directly or one can first perform a data-binning of the data into M bins and compute the Kantorovich WD [(A5)] between the resulting histograms. The complexity of the former scales with the sample size N while the latter scales with the number of bins M .

The computation of (A2) is a special case of the assignment problem from economics, which can be solved in $O(N^3)$ by the well-known Hungarian algorithm. On the other hand, the Kantorovich formulation (A5) is an example of a problem in linear programming. The set $\Pi(\mathbf{P}, \mathbf{Q})$ is a convex polytope and as the cost function is linear it follows that the minimum must be attained on a vertex of this polytope. A minimizing vertex can be found, for example, via the simplex algorithm.

When $d = 1$ the WD can be computed easily as there is an explicit formula. Indeed, for two 1D distributions with cumulative distribution functions (CDFs) F and G , respectively, the 1-WD is (Villani 2003, p. 75)

$$W_1(F, G) = \int_{-\infty}^{\infty} |F(x) - G(x)| dx. \quad (\text{A6})$$

When d is large data-binning is infeasible and it is more natural to work with the empirical distributions μ_n, ν_n directly; however, there is a curse of dimensionality in this context. Indeed, one has

$$\mathbb{E}\left[\left|W_1(\mu_n, \nu_n) - W_1(\mu, \nu)\right|\right] = O(n^{-1/d})$$

and this bound is sharp in general, which gives very slow convergence in high dimension in some cases (Dudley 1969). Construction of a metric to rival the WD that does not suffer a curse of dimensionality is an open problem and focus of active research.

In our work we have found that, despite the curse of dimensionality, computing the WD between empirical distributions with a modest sample size is computationally feasible and provides a useful checksum, usually in agreement with results obtained for example by marginalizing on one-dimensional subspaces. We also note that interesting recent work has shown a regularized form of the WD called the Sinkhorn divergence (SD) (Cuturi 2013) has improved sample complexity with a dimension agnostic convergence rate of $O(n^{-1/2})$ for appropriate regularizing parameters (Genevay et al. 2019) and in our work we found that many WD computations could be corroborated by SDs provided a suitable choice of regularizing parameter was chosen to ensure convergence of the Sinkhorn algorithm.

d. Comparing other metrics

It is interesting to note the similarity between (A6) and the continuous rank probability score often used for weather forecast skill

$$\text{CRPS}(F, G) = \int_{-\infty}^{\infty} |F(x) - G(x)|^2 dx, \quad (\text{A7})$$

and with the Kolmogorov–Smirnov test

$$\text{KS}(F, G) = \sup_{x \in (-\infty, \infty)} |F(x) - G(x)|, \quad (\text{A8})$$

for CDFs F and G . For the WD with $p = 2$ cost there is the explicit formula in 1D

$$W_2(F, G) = \int_{-\infty}^{\infty} |F^{-1}(x) - G^{-1}(x)|^2 dx \quad (\text{A9})$$

where F^{-1} and G^{-1} are generalized inverses (Villani 2003, p. 75). Note that (A6), (A7), and (A9) take account of the geometry of $\mathbb{R}$ (cf. section a of appendix A) while (A8) does not.

e. Duality

Suppose we have a pair of distributions μ and ν representing, for example, rainfall at a fixed location in mm $(6 \text{ h})^{-1}$. How can we interpret a WD of, say, 1 between μ and ν ?

Since cost is defined as $c(x, y) = |x - y|$ a nice property of the WD is that it inherits the units of rainfall so that we can interpret the difference in mm $(6 \text{ h})^{-1}$. Heuristically, this difference tells us that a cost of at least 1 mm $(6 \text{ h})^{-1}$ must be spent to transport μ to ν , and this takes into account both mean and extreme rainfall.

Moreover, this distance gives bounds on a range of expected values by the Kantorovich–Rubinstein duality (Villani 2003, theorem 1.14), which states

$$W_1(\mu, \nu) = \sup_{f \in \text{Lip}_1} \left| \mathbb{E}[f(X_\mu)] - \mathbb{E}[f(Y_\nu)] \right|,$$

where X_μ and Y_ν are random variables with laws μ and ν , and Lip_1 is the space of functions satisfying $|f(x) - f(y)| \leq |x - y|$ (the 1-Lipschitz functions). Taking $f(x) = x$ and duality gives

$$\left| \mathbb{E}[X_\mu] - \mathbb{E}[Y_\nu] \right| \leq W_1(\mu, \nu),$$

which shows that a WD of 1 mm $(6 \text{ h})^{-1}$ implies a difference in expected rainfall of less than 1 mm $(6 \text{ h})^{-1}$ (note that this bound is sharp when μ and ν are Dirac masses). But duality also gives bounds on expected extreme rainfall. To see this, suppose extreme rainfall is defined as any rainfall that falls in excess of r_c mm $(6 \text{ h})^{-1}$ where r_c is some critical value. Then taking $f(x) = 0$ for $x < r_c$ and $f(x) = x - r_c$ for $x \geq r_c$ gives

$$\left| \mathbb{E}[f(X_\mu)] - \mathbb{E}[f(Y_\nu)] \right| \leq W_1(\mu, \nu),$$

which shows a difference in expected extreme rainfall of $\leq 1 \text{ mm } (6 \text{ h})^{-1}$.

Understanding such heuristics is helpful in interpreting the bounds on rounding error derived in this paper.

APPENDIX B

Number Formats

a. Floating-point arithmetic

The standard arithmetic format for scientific computing is the floating-point number (float). The bits in a float are divided into three groups: a sign bit, the exponent bits, and the significant bits. A nonzero exponent specifies an interval $I = [2^{e-\text{bias}}, 2^{e+1-\text{bias}})$ with a bias to allow for negative exponents while the exponent bits are interpreted as an unsigned integer e . By convention the bias is taken as $2^{k-1} - 1$ where k is the number of exponent bits. For $e = 0$, floats are defined on an interval $I = (0, \pm 2^{1-\text{bias}})$ called the subnormal range. The significant bits specify a point on I from an evenly spaced partition of I . Thus, the bias together with the number of exponent bits determines the range of representable normal numbers, while the subnormal range, and therefore the smallest representable number, is determined by the bias together with the number of significant bits. Some different float formats available in hardware are shown in Fig. B1.

The IEEE-754 Float64 format is called double precision, and Float32 and Float16 are called single and half precision, respectively.

b. Rounding

The default rounding mode for floats is round-to-nearest tie-to-even (RN), which rounds an exact result x to the

(a) A visualisation of floats.

(b) The distribution of bits in (a).

	Float64	Float32	Float16	BFloat16	TensorFloat32
sign	1	1	1	1	1
exponent	11	8	5	8	8
significant	52	23	10	7	10

FIG. B1. Floating-point number formats currently available in hardware. Float64, Float32, and Float16 are the IEEE-754 standards for double, single, and half precision, while BFloat16 and TensorFloat32 are custom formats developed by Google Brain and Nvidia respectively. (a) A visualization of floats. (b) The distribution of bits in (a).

nearest representable number x_i . In case x is halfway between two representable numbers, the result will be tied to the even float, whose significand ends in a zero bit. These special cases are therefore alternately round up or down, which removes a bias that would otherwise persist.

For stochastic rounding (SR) rounding of x down to a representable number x_1 or up to x_2 occurs at probabilities that are proportional to the respective distances. Specifically, if u is the distance between x_1 , x_2 , then x will be rounded to x_1 with probability $1 - u^{-1}(x - x_1)$ and to x_2 with probability $u^{-1}(x - x_1)$.

The introduced absolute rounding error for SR is always at least as big as for RN and when low-probability round away from nearest occurs, it can be up to $\pm u$, twice as large as for round-to-nearest. However, by construction, SR is exact in expectation and thus in particular by the law of large numbers one has

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \text{stochastic round}(x) = x,$$

with the limit obtained in the strong sense. Moreover, by sometimes rounding small remainders up, rather than always rounding them down as in RN, systemic errors can sometimes be avoided with SR, such as in stagnation (see section 4 for an example).

While the law of large numbers may plausibly be invoked in simple additive numerical schemes as in section 4, in other numerical schemes such as for nonlinear evolution equations its applicability is less clear.

It is worth noting that SR at low precision requires computation at a higher precision in order to generate the probabilities for rounding; however, all numbers are written, read, and communicated at low precision. It is also interesting to note that SR can easily be implemented with a random number sampled from the uniform distribution. This

means that random samples can be computed in advance of or in parallel to the arithmetic.

REFERENCES

- Arjovsky, M., S. Chintala, and L. Bottou, 2017: Wasserstein generative adversarial networks. *Proc. 34th Int. Conf. on Machine Learning*, Vol. 70, 214–223, <http://proceedings.mlr.press/v70/arjovsky17a.html>.
- Chantry, M., T. Thornes, T. Palmer, and P. Düben, 2019: Scale-selective precision for weather and climate forecasting. *Mon. Wea. Rev.*, **147**, 645–655, <https://doi.org/10.1175/MWR-D-18-0308.1>.
- Croci, M., and M. B. Giles, 2020: Effects of round-to-nearest and stochastic rounding in the numerical solution of the heat equation in low precision. ArXiv, <https://arxiv.org/abs/2010.16225>.
- Cuturi, M., 2013: Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in Neural Information Processing Systems*, Vol. 26, 2292–2300, <http://papers.nips.cc/paper/4927-sinkhorn-distances-lightspeed-computation-of-optimal-transport.pdf>.
- Dawson, A., and P. Düben, 2017: rpe v5: An emulator for reduced floating-point precision in large numerical simulations. *Geosci. Model Dev.*, **10**, 2221–2230, <https://doi.org/10.5194/gmd-10-2221-2017>.
- , —, D. A. MacLeod, and T. N. Palmer, 2018: Reliable low precision simulations in land surface models. *Climate Dyn.*, **51**, 2657–2666, <https://doi.org/10.1007/s00382-017-4034-x>.
- Dogar, M. M., F. Kucharski, and S. Azharuddin, 2017: Study of the global and regional climatic impacts of ENSO magnitude using SPEEDY AGCM. *J. Earth Syst. Sci.*, **126**, 30, <https://doi.org/10.1007/s12040-017-0804-4>.
- Dudley, R. M., 1969: The speed of mean Glivenko-Cantelli convergence. *Ann. Math. Stat.*, **40**, 40–50, <https://doi.org/10.1214/aoms/1177697802>.
- Genevay, A., L. Chizat, F. Bach, M. Cuturi, and G. Peyré, 2019: Sample complexity of Sinkhorn divergences. *PMLR*, **89**, 1574–1583, <https://arxiv.org/abs/1810.02733>.

- Gilham, R., 2018: 32-bit physics in the Unified Model. Met Office Tech. Rep. 626, 16 pp., https://digital.nmla.metoffice.gov.uk/IO_951e52e5-6698-485e-ad33-54d0a2b0ce99/.
- Gupta, S., A. Agrawal, K. Gopalakrishnan, and P. Narayanan, 2015: Deep learning with limited numerical precision. *PMLR*, **37**, 1737–1746, <https://proceedings.mlr.press/v37/gupta15.html>.
- Harvey, R., and D. L. Verseghy, 2016: The reliability of single precision computations in the simulation of deep soil heat diffusion in a land surface model. *Climate Dyn.*, **46**, 3865–3882, <https://doi.org/10.1007/s00382-015-2809-5>.
- Hatfield, S., M. Chantry, P. Düben, and T. Palmer, 2019: Accelerating high-resolution weather models with deep-learning hardware. *Proc. Platform for Advanced Scientific Computing Conference*, ACM, Zurich, Switzerland, <https://doi.org/10.1145/3324989.3325711>.
- Jeffress, S., P. Düben, and T. Palmer, 2017: Bitwise efficiency in chaotic models. *Proc. Roy. Soc.*, **A473**, 20170144, <https://doi.org/10.1098/rspa.2017.0144>.
- Klöwer, M., P. D. Düben, and T. N. Palmer, 2020: Number formats, error mitigation, and scope for 16-bit arithmetics in weather and climate modeling analyzed with a shallow water model. *J. Adv. Model. Earth Syst.*, **12**, e2020MS002246, <https://doi.org/10.1029/2020MS002246>.
- Kucharski, F., F. Molteni, M. P. King, R. Farneti, I.-S. Kang, and L. Feudale, 2013: On the need of intermediate complexity general circulation models: A “SPEEDY” example. *Bull. Amer. Meteor. Soc.*, **94**, 25–30, <https://doi.org/10.1175/BAMS-D-11-00238.1>.
- Lorenz, E. N., 1963: Deterministic nonperiodic flow. *J. Atmos. Sci.*, **20**, 130–141, [https://doi.org/10.1175/1520-0469\(1963\)020<0130:DNF>2.0.CO;2](https://doi.org/10.1175/1520-0469(1963)020<0130:DNF>2.0.CO;2).
- Maass, C., 2021: ECMWF implementation of IFS cycle 47r2. ECMWF, <https://confluence.ecmwf.int/display/FCST/Implementation+of+IFS+Cycle+47r2>.
- Miciekevicius, P., and Coauthors, 2018: Mixed precision training. Poster, *Int. Conf. on Learning Representations*, Vancouver, BC, Canada, ICLR, <https://openreview.net/forum?id=r1gs9JgRZ>.
- Molteni, F., and F. Kucharski, 2018: A heuristic dynamical model of the North Atlantic Oscillation with a Lorenz-type chaotic attractor. *Climate Dyn.*, **52**, 6173–6193, <https://doi.org/10.1007/s00382-018-4509-4>.
- Prims, O. T., M. C. Acosta, A. M. Moore, M. Castrillo, K. Serradell, A. Cortés, and F. J. Doblas-Reyes, 2019: How to use mixed precision in ocean models: Exploring a potential reduction of numerical precision in NEMO 4.0 and ROMS 3.6. *Geosci. Model Dev.*, **12**, 3135–3148, <https://doi.org/10.5194/gmd-12-3135-2019>.
- Robin, Y., P. Yiou, and P. Naveau, 2017: Detecting changes in forced climate attractors with Wasserstein distance. *Nonlinear Processes Geophys.*, **24**, 393–405, <https://doi.org/10.5194/npg-24-393-2017>.
- Rüdisühli, S., A. Walser, and O. Fuhrer, 2014: COSMO in single precision. *COSMO Newsletter*, No. 14, Consortium for Small-Scale Modeling, Offenbach, Germany, 70–87, http://www.cosmo-model.org/content/model/documentation/newsLetters/newsLetter14/cnl14_09.pdf.
- Saffin, L., S. Hatfield, P. Düben, and T. Palmer, 2020: Reduced-precision parametrization: Lessons from an intermediate-complexity atmospheric model. *Quart. J. Roy. Meteor. Soc.*, **146**, 1590–1607, <https://doi.org/10.1002/qj.3754>.
- Tucker, W., 1999: The Lorenz attractor exists. *C. R. Acad. Sci.*, **328**, 1197–1202, [https://doi.org/10.1016/S0764-4442\(99\)80439-X](https://doi.org/10.1016/S0764-4442(99)80439-X).
- Váňa, F., P. Düben, S. Lang, T. Palmer, M. Leutbecher, D. Salmond, and G. Carver, 2017: Single precision in weather forecasting models: An evaluation with the IFS. *Mon. Wea. Rev.*, **145**, 495–502, <https://doi.org/10.1175/MWR-D-16-0228.1>.
- Villani, C., 2003: *Topics in Optimal Transportation*. American Mathematical Society, 370 pp., <https://books.google.co.uk/books?id=GqRXYFxe0l0C>.
- Vissio, G., and V. Lucarini, 2018: Evaluating a stochastic parametrization for a fast–slow system using the Wasserstein distance. *Nonlinear Processes Geophys.*, **25**, 413–427, <https://doi.org/10.5194/npg-25-413-2018>.
- , V. Lembo, V. Lucarini, and M. Ghil, 2020: Evaluating the performance of climate models based on Wasserstein distance. *Geophys. Res. Lett.*, **47**, e2020GL089385, <https://doi.org/10.1029/2020GL089385>.