The Relationship between Numerical Precision and Forecast Lead Time in the Lorenz'95 System

FENWICK C. COOPER

University of Oxford, Oxford, United Kingdom

PETER D. DÜBEN

University of Oxford, Oxford, and European Centre for Medium-Range Weather Forecasts, Reading, United Kingdom

CHRISTOPHE DENIS

ENS Paris-Saclay, Center for Mathematical Studies and Their Applications, Cachan, France

ANDREW DAWSON

European Centre for Medium-Range Weather Forecasts, Reading, United Kingdom

PETER ASHWIN

University of Exeter, Exeter, United Kingdom

(Manuscript received 6 June 2018, in final form 29 May 2019)

ABSTRACT

We test the impact of changing numerical precision upon forecasts using the chaotic Lorenz'95 system. We find that in comparison with discretization and numerical rounding errors, the dominant source of errors are the initial condition errors. These initial condition errors introduced into the Lorenz'95 system grow exponentially at a rate according to the leading Lyapunov exponent. Given this information we show that the number of bits necessary to represent the system state can be reduced linearly in time without significantly affecting forecast skill. This is in addition to any initial reduction in precision to that of the initial conditions and also implies the potential to reduce some storage costs. An approach to vary precision locally within simulations, guided by the direction of eigenvectors of the growth and decay of forecast error (the “singular vectors”), did not show a satisfying impact upon forecast skill in relation to cost savings that could be achieved with a uniform reduction of precision. The error in a selection of ECMWF forecasts as a function of the number of bits used to store them indicates that precision might also be reduced in operational systems.

1. Introduction

Several studies have shown that numerical precision can be reduced significantly in atmospheric modeling and that such a reduction in precision promises large savings in computational costs (e.g., Düben et al. 2014; Düben and Palmer 2014; Jeffress et al. 2017). For example, reducing numerical precision reduces the amount of data that must be exchanged between a computer's CPU and main memory, between nodes of a cluster and, depending upon the hardware, can increase the number of

calculations per CPU clock cycle. All of which are potential bottlenecks in a model's performance. If computational cost is reduced, savings can be reinvested to achieve model simulations at higher resolution or model complexity, or more ensemble members in ensemble forecasts to improve future predictions (e.g., Palmer 2012). The idea is to reduce numerical precision to the minimal value that can be justified by the level of model uncertainty and information content within model simulations to enable models to run as efficiently as possible.

Numerical precision can be reduced more strongly at later forecast lead times compared to a reduction in

Corresponding author: Fenwick C. Cooper, fenwick@littlestick.com

DOI: 10.1175/MWR-D-18-0200.1

© 2020 American Meteorological Society. For information regarding reuse of this content and general copyright information, consult the [AMS Copyright Policy](https://www.ametsoc.org/PUBSReuseLicenses) (www.ametsoc.org/PUBSReuseLicenses).

numerical precision at the beginning of a forecast (Düben et al. 2015). This is consistent with the idea that one can reduce precision to a level that can be justified by model uncertainty since model error will grow with forecast lead time. For simulations of a nonlinear and chaotic system that start from imperfect initial conditions, the expansion or contraction of errors is governed over short time scales by the so-called singular values and vectors (Molteni and Palmer 1993), and for typical perturbations, at longer time scales the mean error will grow exponentially, with the rate of error growth called the leading Lyapunov exponent of the system (e.g., Wolf et al. 1985; Vannitsem 2017). Wolf et al. 1985 define the leading Lyapunov exponent λ by

$$\lambda = \lim_{t \rightarrow \infty} \frac{1}{t} \log \frac{p(t)}{p(0)},$$

where $p(t)$ is the separation of trajectories in phase space as a function of time and $p(0)$ is sufficiently small. In practice the limit as $t \rightarrow \infty$ is never reached and $p(0)$ is never sufficiently small, so an approximation, perhaps estimated from an ensemble average, is used. Note that the limit as $t \rightarrow \infty$ indicates that the Lyapunov exponent is not necessarily well approximated by the initial growth of perturbations. After some time, the average distance between the truth and a model trajectory approaches the limits of the system and the mean error approaches a constant. This is often loosely referred to as the forecast reaching “climatology”. In real weather forecasts, the growth of mean forecast error is assumed to be a superposition of an exponential growth rate due to the growth of imperfect initial conditions and a linear growth due to the use of an imperfect model with limited accuracy, for example due to limited resolution and errors in model formulation (e.g., Magnusson and Källén 2013).

In this paper, we investigate the extent to which one can take advantage of error growth by making a steady reduction in numerical precision with forecast lead time. We consider the impact of reduced precision upon both the dynamics and forecast skill of the chaotic Lorenz’95 system (Lorenz 1995). Although the Lorenz’95 “atmosphere” is unrealistic, it was created to share some aspects of the predictability of the real atmosphere, as explained in detail in Lorenz (1995). It is a chaotic system with solutions that decorrelate in time, which is the property necessary to link forecast error to numerical precision, and it is a high-dimensional system with propagating solutions that decay in space, which is the property that relates it to a more realistic atmospheric model. Integrating the Lorenz’95 system to equilibrium requires a tiny fraction of the computational effort required to do the same with a numerical weather model.

This enables us to complete the large number of ensemble integrations required to distinguish between small model differences.

Today, it is common to store model output in weather forecasts at the same level of numerical precision over the entire length of the forecast. It is straightforward to use a reduction in precision to reduce the storage costs of weather model output. Savings can be used to store more model fields or to simply increase the rate of model outputs. A similar reduction in precision with forecast lead time could also be performed for floating-point arithmetic within model simulations when switching from double (64 bits) to single (32 bits) or other precisions, for example in the network data communicated between nodes, or when using hardware that allows the use of flexible floating-point precision such as Field Programmable Gate Arrays (FPGAs), see Jeffress et al. (2017).

2. Simulations with the Lorenz model

We study the single layer Lorenz’95 system, (sometimes referred to as Lorenz’96), that is defined by

$$\frac{\partial x_j}{\partial t} = x_{j-1}(x_{j+1} - x_{j-2}) - x_j + f, \quad (1)$$

where f is constant in time, $j = 1, 2, 3, \dots, K$ and the system is periodic $x_j = x_{j+K} = x_{j-K}$ (Lorenz 1995). In a very loose sense, the K variables are said to represent an atmospheric quantity at equally spaced sectors around the equator with a time unit of around 5 days. The quadratic terms represent advection, the decay term represents some kind of dissipation and f an external forcing. In this paper, we choose $K = 40$ and $f = 20$ making the system chaotic. Equation (1) is integrated using a fourth-order Runge–Kutta method (e.g., Press et al. 2007) with a time step of $\Delta t = 0.01$.

We investigate a reduction in precision using two different methods:

1) Precision analysis using a software emulator

The 64 bits in an IEEE standard double precision floating-point number are made up of 52 significant bits, 11 exponent bits and a sign bit (Zuras and Cowlshaw 2008). We perform integrations that use different numbers of bits for the significant of floating-point numbers. To control the precision of calculations, we use the reduced-precision-emulator (*rpe*, see Dawson and Düben 2017). This library enables us to simulate the code as if it was run on reduced precision hardware.

2) Precision analysis using Verificarlo

We perform a precision analysis using the Verificarlo compiler (Denis et al. 2016). Verificarlo is a compiler

FIG. 1. Logarithm of the mean-squared error (MSE) of reduced precision integrations with respect to a double precision integration as a function of time. The lowest line in dark green represents the MSE when using a 51 bit significand, the red line above it corresponds to a 50 bit significand, the cyan line above that 49 bits, and so on. The top line corresponds to a 3 bit significand. The different starting points in the plot are caused by the reduction of precision in the initial conditions.

that automatically applies Monte Carlo arithmetic. By adding random noise to input variables, Monte Carlo arithmetic tracks rounding and cancellation errors throughout model integrations and enables identification of their effects.

Results

To evaluate model quality at different levels of numerical precision relative to a double precision control, we average the mean-squared error (MSE),

$$\text{MSE}(t) = \sum_{j=1}^K [x_j(t) - \hat{x}_j(t)]^2,$$

in the reduced precision x_j variables with respect to the double precision control $\hat{x}_j$ variables, of forecasts made using 400 independent integrations between $t = 0$, after spinup, and $t = 7$, see Fig. 1. Prior to each forecast we performed a “spinup” integration from $t = -20$ to $t = 0$ from random initial conditions, so that our forecast initial conditions start from on, or near to, the attractor. The kink in the curves in Fig. 1 around $t = 1$ is reproduced in an independent set of integrations and is also visible in the autocorrelation of x_j . To check that our conclusions are not simply an artifact of the Runge–Kutta scheme, we verified the results using a second-order backward differentiation formula with an adaptive time step using the *sundials ida* package (Hindmarsh et al. 2005).

FIG. 2. The gradient of each line in the MSE plot in Fig. 1. The average gradient is estimated between $t = 0.1$ and a $\log(\text{MSE})$ of 0 with the circle indicating the average over all 400 forecasts and the dashed lines indicating the uncertainty in this average (the standard deviation multiplied by $2/\sqrt{400}$). Note that the y axis covers a change in the gradient of about 10% of the overall value.

In Fig. 1 the mean-squared error increases at an approximately exponential rate toward a climatological mean value, which is somewhat independent of precision. This rate approximates the leading Lyapunov exponent of the system. In Fig. 1 it appears that the exponent for each line is comparable, while the lines are separated by approximately equal values of $\log(\text{MSE})$. This indicates that MSE is dominated by errors in the initial condition rather than any other errors due to reduced precision, such as inaccurate additions or multiplications. For example, consider a 50 bit and a 30 bit significand. The initial condition error in the 30 bit system is much larger than in the 50 bit system, but these errors increase at approximately the same rate in time in both systems.

The gradient of each of the $\log(\text{MSE})$ lines in Fig. 1 is estimated and plotted in Fig. 2. The gradient (or leading Lyapunov exponent) hardly changes between integrations with a 30 to 51 bit significand. The gradient does change as the number of bits is reduced below 30, indicating the impact of reduced precision on the chaotic dynamics of the system; however, a reduction to a 13 bit significand alters the leading Lyapunov exponent by less than 10%. The final climatological error seems to be very consistent over a large range of precision. Note that $\log(\text{MSE})$ in the Runge–Kutta time step, estimated by a fifth-order method (Cash and Karp 1990, as applied in Press et al. 2007), is approximately -29 corresponding to a significand of about 28 bits.

The limited impact of reducing precision upon the Lyapunov exponent suggests that both data storage and

FIG. 3. (a) $\log(\text{MSE})$ of a 51 bit model (solid black line) and a 23 bit model (dashed black line) as a function of time. The curves are identical to their equivalents in Fig. 1. At $t = 3.5$ (vertical gray line) the MSE of the 51 bit model is about the same value as the initial condition error of the 23 bit model. The 23 bit model is initialized with the 51 bit model state at this point (solid red line). (b) Logarithm of the ensemble variance for a model simulation integrated with Verifcarlo. The solid black curve represents the ensemble variance when the least significant bit is randomly flipped. The dashed black curve represents the ensemble variance when the 29 least significant bits are randomly flipped, leaving $52 - 29 = 23$ bits of precision remaining. The solid red curve is identical to the dashed black curve, except it has been moved to start at $t = 3.5$.

precision in the model integration can be reduced as error grows. To test this, the 51 bit model is integrated until its MSE, averaged over all 400 forecasts, is slightly larger than the initial condition MSE of the 23 bit model, (which corresponds to IEEE single precision floating-point format). At this point ($t = 3.5$) the precision is reduced to 23 bits and the forecast continues. When compared with simply continuing the 51 bit integration, the 23 bit forecast does have a slightly larger average MSE after $t = 3.5$, but the difference is small (see Fig. 3a). To check that the MSE is in fact caused by finite precision rounding error, these integrations were repeated using Verifcarlo with 50 ensemble members, Fig. 3b. These results with Verifcarlo confirm those found using *rpe*, that the logarithm of the variance due to initial rounding errors, increases linearly in time.

We may compute an error covariance matrix for the short time evolution of perturbations to the linearized governing equations (Molteni and Palmer 1993). If an ensemble of integrations is started, each member with a small normally distributed perturbation to its initial condition, then the error covariance matrix describes the initial expansion and contraction in time of the distribution of ensemble member states. We might suppose that precision may be reduced in the direction of the contracting eigenvectors of this matrix (sometimes loosely referred to as the “singular vectors”), because any errors will get smaller. However, we found that in the Lorenz’95 system, the sum of the first three eigenvectors, where the error covariance is contracting, looks rather like the sum of the final three eigenvectors, where the error covariance

is expanding, see Fig. 4a. The overall growth or contraction of error at a grid point depends upon all of the singular vectors and may be small relative to the impact of only the leading singular vectors.

To test if a reduction in precision in contracting error directions is beneficial, we project the contribution of all of the singular vectors onto the x_j variables. At a particular time step, this contribution at a particular x_j might be above one, indicating expanding error, or below one indicating contracting error. For this time step we then choose the significand (separately for each x_j) to be either 18 bits for contracting or 19 bits for expanding errors. The MSE of this integration lies in between the 18 bit and 19 bit curves in Fig. 1, see Fig. 4b. Repeating the experiment with 17 and 20 bits for the contracting and expanding directions, yields a MSE curve between the 17 and 18 bit curves in Fig. 1, also in Fig. 4b. The reduction in forecast error due to an increase in significand size from 17 bits to 18 bits is larger than the reduction in forecast error due to increasing the significand size from 17 to 20 bits restricted to when the error is increasing. These results summarize our conclusions arrived at by comparing several other choices of precision, (not shown here).

3. Evaluation of stored forecast data

We use data output from the high-resolution, operational weather forecasts of the European Centre for Medium-Range Weather Forecasts (ECMWF). Figure 5a shows the MSE between 20° and 90°N for geopotential

FIG. 4. (a) For a particular point in time of a single integration, the sum of the first three (expanding error) singular vectors (black) and the last three (contracting error) singular vectors (red). The similarity of the black and red lines is typical in the Lorenz'95 system. (b) Difference between the 19 bit $\log(\text{MSE})$ and the 18 bit (solid black line) and 17 bit (dashed black line) curves in Fig. 1. The red curve illustrates the case when contracting errors are represented by 18 bits and expanding by 19 bits. The blue curve illustrates the case when contracting errors are represented by 17 bits and expanding by 20 bits. A $\log(\text{MSE})$ difference of zero indicates the same error as the 19 bit integration.

height at 500 hPa (Z500). The errors are averaged for 12 forecasts that were initialized on the first day of each month in 2016. The MSE is calculated in gridpoint space for data stored at varying levels of precision on a reduced Gaussian grid with 640 grid points between pole and equator. At ECMWF, model data are stored in the form of GRIB files (Dey et al. 2003) in a fixed point format, with an equal spacing of the interval between the minimal and maximal field value within a vertical model level. Precision of output files can be adjusted by the user who can adjust the number of bits that are used per variable. In Fig. 5a, it is visible that the impact of a precision reduction

is reducing with forecast lead time, albeit over a much smaller range of precision and much larger initial condition uncertainty than that considered here. The analysis of the Lorenz'95 system here provides a very basic potential explanation of why this is the case. As explained in the next section, in particular Eq. (4), it is a property of all chaotic systems. Being a much more complex model, it is not certain, however, that the time-dependent impact of stored precision upon the MSE here is due to the same chaotic processes as those observed with the Lorenz'95 system. Or even that other nonchaotic processes dominate the error growth. Compare with Fig. 5b, which shows

FIG. 5. (a) MSE of 500 hPa geopotential height (Z500) data, between 20° and 90°N, of ECMWF's operational deterministic forecast (a single integration at 9 km resolution for each forecast) with respect to the operational reanalysis (an estimate of the true state using measurements). The errors are averaged for 12 forecasts that were initialized on the first day of each month in 2016. The MSE is calculated in gridpoint space for data stored at varying levels of precision (indicated) on a reduced Gaussian grid with 640 grid points between pole and equator. The 10 bit and 16 bit lines are approximately coincident. (b) MSE in the 51 bit Lorenz'95 forecast when the 51 bit output is stored at a range of precisions. The inset plot is a zoom of the full plot, illustrating a plausible range of behavior that the ECMWF forecast is displaying.

the equivalent plot for the Lorenz'95 system where double precision model output has been truncated to a range of precisions for storage.

4. Discussion

In our Lorenz'95 experiments, forecast error due to a reduction in precision is dominated by the reduction in accuracy of the initial condition. The growth rate of errors, that is described by the Lyapunov exponent, is very similar for integrations utilizing a range of precisions. Since rounding errors grow exponentially, the number of bits of the systems state that are not perturbed by rounding errors reduces linearly with forecast lead time if only unperturbed bits contain relevant information regarding the model state.

To make the final point above explicit we write the forecast error

$$E \approx E_0 \exp(\lambda t), \quad (2)$$

where E_0 is the initial condition error and λ is the leading Lyapunov exponent. If data are stored in an integer format so that the space of possible values is evenly divided (as in GRIB), then the smallest distance D that can be represented is given by

$$D = D_0 2^{-n}, \quad (3)$$

where D_0 is the maximum range for which a physical quantity is meaningful (e.g., for temperature in kelvin $230 < T < 320 \text{ K} \rightarrow D_0 = 90$) and n is the number of bits used. If we want the numerical precision to match the MSE, we can set D and E to be equal and obtain an equation for the number of bits that should be used as a function of forecast lead time:

$$n \approx \frac{\log(D_0) - \log(E_0)}{\log(2)} - \frac{\lambda}{\log(2)} t. \quad (4)$$

The constants E_0 and λ might be estimated by curve fitting, for example that performed for each point in Fig. 2. Other, more sophisticated methods are also available (e.g., Wolf et al. 1985). A similar relationship to Eq. (4) has been found by Wang and Li (2014) in a different context. The equation suggests that the number of bits required is approximately linearly related to the forecast lead time, with the gradient determined by the leading Lyapunov exponent, which indicates that storage requirements can be halved (see the dark gray region in Fig. 6). Note that this reduction is in addition to the initial reduction in precision to the initial value uncertainty (see the white region in Fig. 6). In the future, customized computing hardware that works

FIG. 6. An illustration of approximate storage (and potentially computational) cost savings possible, as suggested by the Lorenz'95 experiments described in this paper. The x axis represents forecast time, the y axis represents the number of bits used in a computation and the integral of lines on this plot represents cost. Integrating at standard precision for an extended time is the most expensive approach. By reducing precision to the initial condition error ($D = E_0$) throughout the forecast, the white area may be eliminated from the cost. Stopping the forecast when climatology has been reached removes the light gray area. Reducing the precision optimally as a function of time ($D = E$) can potentially yield a further halving of storage cost (elimination of the dark gray area). The area between the dashed horizontal line and the origin indicates the additional cost of continuing the forecast at low precision, perhaps for the collection of climatological or seasonal statistics.

at a reduced precision may be designed. At this hardware level, both the area of silicon and electrical power required for a calculation (and hence monetary cost) are approximately proportional to the number of bits of precision (e.g., Jeffress et al. 2017). It therefore seems likely that these costs can ultimately also be approximately halved. However, we have not considered the exponent.

In addition to the exponential error growth due to the chaotic dynamics, the missing accuracy of a weather model itself is a large and linearly growing source of error (Magnusson and Källén 2013). Predictability is also state dependent (e.g., Dawson and Palmer 2015) and a reduction of precision, solely justified by mean errors over all weather regimes, is likely to reduce skill for the most predictable regimes. On the other hand, the reduction in forecast error obtainable by increasing the number of significant bits by one everywhere in the Lorenz'95 system was found to be much larger than the reduction due to changing precision according to directions of increasing and decreasing error, according to the “singular values.” Given the dominance of the initial condition error, and the difficulty of making use of the contracting singular values, it seems that it would be difficult to obtain further reductions in precision of the Lorenz'95 system through considering the contracting local Lyapunov covariant vectors (see, e.g., Ginelli et al. 2013).

Seasonal forecasts have been shown to hold useful information for several months (e.g., Weisheimer and Palmer 2014). At seasonal time scales, the forecast error will be close to climatology and Eq. (4) needs to be modified to retain this information. Reduction in precision should be related to the meteorological information we are interested in and the growth behavior of forecast mean error will depend on the error metric and the diagnostic that is considered (Orrell 2002). For example, slightly displacing a front within forecasts will have a strong negative impact on the MSE but not necessarily destroy all predictive information for weather forecasters.

The points above can be addressed by using a slower reduction in precision than that suggested by the leading Lyapunov exponent and by limiting the minimal level of precision toward the end of a forecast. We have shown evidence that it might be possible for operational forecast model fields to be stored at a low precision (~ 10 bits) and still allow for their appropriate representation of arbitrary derived quantities, given the huge uncertainties. Of course, if only a small quantity of output data are required, for example in order to plot 500 hPa synoptic-scale weather charts of a few model variables, then there is no advantage to storage at low precision.

Acknowledgments. This research was funded by EPSRC project ReCoVER EP/M008495/1 via a pilot study “Optimal use of reduced precision hardware in weather and climate models.” (RFFLP 019) awarded to Tim N. Palmer, Peter Ashwin, Fenwick C. Cooper, and Peter D. Düben. Peter D. Düben, Andrew Dawson, and Tim Palmer received funding from an ERC grant (Towards the Prototype Probabilistic Earth-System Model for Climate Prediction, project reference 291406). Peter Düben was also supported by the ESIWACE project funded from the European Union’s Horizon 2020 programme under Grant 675191. The research materials supporting this publication can be accessed by contacting Fenwick Cooper (e-mail: Fenwick@LittleStick.com).

REFERENCES

- Cash, J. R., and A. H. Karp, 1990: A variable order Runge–Kutta method for initial value problems with rapidly varying right-hand sides. *ACM Trans. Math. Software*, **16**, 201–222, <https://doi.org/10.1145/79505.79507>.
- Dawson, A., and T. N. Palmer, 2015: Simulating weather regimes: Impact of model resolution and stochastic parameterization. *Climate Dyn.*, **44**, 2177–2193, <https://doi.org/10.1007/s00382-014-2238-x>.
- , and P. D. 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>.
- Denis, C., P. de Oliveira Castro, and E. Petit, 2016: Verificarlo: Checking floating point accuracy through Monte Carlo arithmetic. *2016 IEEE 23rd Symp. on Computer Arithmetic (ARITH)*, Santa Clara, CA, IEEE, <https://doi.org/10.1109/ARITH.2016.31>.
- Dey, C. H., C. Sanders, J. Clochard, J. Hennessy, and S. Elliott, 2003: Guide to the WMO table driven code form used for the representation and exchange of regularly spaced data in binary form: FM 92 GRIB Edition 2. WMO, Geneva, Switzerland, 103 pp., http://www.wmo.int/pages/prog/www/WMOCodes/Guides/GRIB/GRIB2_062006.pdf.
- Düben, P. D., and T. N. Palmer, 2014: Benchmark tests for numerical weather forecasts on inexact hardware. *Mon. Wea. Rev.*, **142**, 3809–3829, <https://doi.org/10.1175/MWR-D-14-00110.1>.
- , H. McNamara, and T. N. Palmer, 2014: The use of imprecise processing to improve accuracy in weather and climate prediction. *J. Comput. Phys.*, **271**, 2–18, <https://doi.org/10.1016/j.jcp.2013.10.042>.
- , F. P. Russell, X. Niu, W. Luk, and T. N. Palmer, 2015: On the use of programmable hardware and reduced numerical precision in earth-system modeling. *J. Adv. Model. Earth Syst.*, **7**, 1393–1408, <https://doi.org/10.1002/2015MS000494>.
- Ginelli, F., H. Chate, R. Livi, and A. Politi, 2013: Covariant Lyapunov vectors. *J. Phys. A: Math. Theor.*, **46**, 254005, <https://doi.org/10.1088/1751-8113/46/25/254005>.
- Hindmarsh, A. C., P. N. Brown, K. E. Grant, S. L. Lee, R. Serban, D. E. Shumaker, and C. S. Woodward, 2005: SUNDIALS: Suite of nonlinear and differential/algebraic equation solvers. *ACM Trans. Math. Software*, **31**, 363–396, <https://doi.org/10.1145/1089014.1089020>.
- Jeffress, S. A., P. D. Düben, and T. N. Palmer, 2017: Bitwise efficiency in chaotic models. *Proc. Roy. Soc. London*, **473A**, <https://doi.org/10.1098/rspa.2017.0144>.
- Lorenz, E. N., 1995: Predictability: A problem partly solved. *Proc. ECMWF Seminar on Predictability*, Reading, United Kingdom, ECMWF, 1–18.
- Magnusson, L., and E. Källén, 2013: Factors influencing skill improvements in the ECMWF forecasting system. *Mon. Wea. Rev.*, **141**, 3142–3153, <https://doi.org/10.1175/MWR-D-12-00318.1>.
- Molteni, F., and T. N. Palmer, 1993: Predictability and finite-time instability of the northern winter circulation. *Quart. J. Roy. Meteor. Soc.*, **119**, 269–298, <https://doi.org/10.1002/qj.49711951004>.
- Orrell, D., 2002: Role of the metric in forecast error growth: How chaotic is the weather? *Tellus*, **54A**, 350–362, <https://doi.org/10.3402/tellusa.v54i4.12159>.
- Palmer, T. N., 2012: Towards the probabilistic Earth-system simulator: A vision for the future of climate and weather prediction. *Quart. J. Roy. Meteor. Soc.*, **138**, 841–861, <https://doi.org/10.1002/qj.1923>.
- Press, W. H., S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, 2007: *Numerical Recipes: The Art of Scientific Computing*. 3d ed. Cambridge University Press, 1256 pp.
- Vannitsem, S., 2017: Predictability of large-scale atmospheric motions: Lyapunov exponents and error dynamics. *Chaos*, **27**, 032101, <https://doi.org/10.1063/1.4979042>.
- Wang, P., and J. Li, 2014: On the relation between reliable computation time, float-point precision and the Lyapunov exponent in chaotic systems. ArXiv:1410.4919.
- Weisheimer, A., and T. N. Palmer, 2014: On the reliability of seasonal climate forecasts. *J. Roy. Soc. Interface*, **11**, <https://doi.org/10.1098/rsif.2013.1162>.
- Wolf, A., J. B. Swift, A. H. Swinney, and J. A. Vastano, 1985: Determining Lyapunov exponents from a time series. *Physica D*, **16**, 285–317, [https://doi.org/10.1016/0167-2789\(85\)90011-9](https://doi.org/10.1016/0167-2789(85)90011-9).
- Zuras, D., and M. Cowlshaw, 2008: IEEE standard for floating-point arithmetic. IEEE Computer Society, <https://doi.org/10.1109/IEEESTD.2008.4610935>.