

Opinion piece   Check for updates

Cite this article: Coveney PV, Highfield RR. 2021 When we can trust computers (and when we can't). *Phil. Trans. R. Soc. A* **379**: 20200067. <https://doi.org/10.1098/rsta.2020.0067>

Accepted: 26 June 2020

One contribution of 15 to a theme issue 'Reliability and reproducibility in computational science implementing verification, validation and uncertainty quantification *in silico*'.

Subject Areas:

computational chemistry, medicinal chemistry, artificial intelligence, computer modelling and simulation, medical computing

Keywords:

validation, verification, uncertainty quantification, big data, machine learning, artificial intelligence

Author for correspondence:

Peter V. Coveney
e-mail: p.v.coveney@ucl.ac.uk

When we can trust computers (and when we can't)

Peter V. Coveney^{1,2} and Roger R. Highfield³

¹Centre for Computational Science, University College London, Gordon Street, London WC1H 0AJ, UK

²Institute for Informatics, Science Park 904, University of Amsterdam, 1098 XH Amsterdam, The Netherlands

³Science Museum, Exhibition Road, London SW7 2DD, UK

 PVC, 0000-0002-8787-7256

With the relentless rise of computer power, there is a widespread expectation that computers can solve the most pressing problems of science, and even more besides. We explore the limits of computational modelling and conclude that, in the domains of science and engineering which are relatively simple and firmly grounded in theory, these methods are indeed powerful. Even so, the availability of code, data and documentation, along with a range of techniques for validation, verification and uncertainty quantification, are essential for building trust in computer-generated findings. When it comes to complex systems in domains of science that are less firmly grounded in theory, notably biology and medicine, to say nothing of the social sciences and humanities, computers can create the illusion of objectivity, not least because the rise of big data and machine-learning pose new challenges to reproducibility, while lacking true explanatory power. We also discuss important aspects of the natural world which cannot be solved by digital means. In the long term, renewed emphasis on analogue methods will be necessary to temper the excessive faith currently placed in digital computation.

This article is part of the theme issue 'Reliability and reproducibility in computational science: implementing verification, validation and uncertainty quantification *in silico*'.

1. Introduction

Scientists and engineers take reproducibility seriously [1] for reasons that, though obvious, are well worth restating. Research relies on a never-ending dialogue between hypothesis and experiment, a conversation that advances more quickly towards understanding phenomena when not distracted by false leads.

Reproducibility and quantifying uncertainty are vital for the design of buildings, bridges, cars, spacecraft and aircraft, along with weather forecasting and a host of other applications where lives depend on performing calculations correctly. Reproducibility is also a welcome corollary of the open science movement that seeks to make all aspects of research transparent and accessible—whether publications, data, physical samples, algorithms, software or documentation. Many have referred to reproducibility as being a ‘cornerstone of science’ [2–4].

However, problems of reproducibility arise in all fields of science, from biomedicine, migration and climate to advanced materials, fusion derived energy and high energy physics. There are many reasons why ‘one-off’ observations cannot be reproduced, even with the same methods, owing to the aleatoric, or random, nature of many phenomena. In such instances, reliable statistical measures are essential to convert measurements into robust findings. False correlations are usually washed away when, over time, they are scrutinized by more systematic, bigger and better-designed studies.

The extent to which reproducibility is an issue for computer modelling is more profound and convoluted, however, depending on the domain of interest, the complexity of the system, the power of available theory, the customs and practices of different scientific communities, together with practical concerns, such as when commercial considerations are challenged by scientific findings [5].

For research on microscopic and relatively simple systems, such as those found in physics and chemistry, for example, theory—both classical and quantum mechanical—offers a powerful way to curate the design of experiments and weigh up the validity of results. In these and other domains of science that are grounded firmly on theory, computational methods more easily help to confer apparent objectivity, with the obvious exceptions of pathological science [6] and fraud [7]. For the very reason that the underlying theory is established and trusted in these fields, there is perhaps less emphasis than there should be on verification and validation (‘solving the equations right’ and ‘solving the right equations’, respectively [8]) along with uncertainty quantification—collectively known by the acronym VVUQ. By comparison, in macroscopic systems of interest to engineers, applied mathematicians, computational scientists and technologists and others who have to design devices and systems that actually work, and which must not put people’s lives in jeopardy, VVUQ, is a way of life—in every sense—to ensure that simulations are credible.

This VVUQ philosophy underpins advances in computer hardware and algorithms that improve our ability to model complex processes using techniques such as finite-element analysis, and computational fluid dynamics for end-to-end simulations in virtual prototyping and to create digital twins [9]. There is a virtuous circle in VVUQ, where experimental data hone simulations, while simulations hone experiments and data interpretation. In this way, the ability to simulate an experiment influences validation by experiment.

In other domains, however, notably biology and biomedical sciences, theories have rarely attained the power and generality of physics. The state space of biological systems tends to be so vast that detailed predictions are often elusive, and VVUQ is less well established, though that is now changing rapidly as, for example, models and simulations begin to find clinical use [10].

Despite the often stated importance of reproducibility, researchers still find various ways to unwittingly fool themselves and their peers [11]. Data dredging—also known as blind big data, data fishing, data snooping and p-hacking—seeks results that can be presented as statistically significant, without any knowledge of the structural characteristics of the problem, let alone first devising a hypothesis about the underlying mechanistic relationships. While corroboration or indirect supporting evidence may be reassuring, when taken too far it can lead to the

interpretation of random patterns as evidence of correlations, and to the conflation of these correlations with causative effects.

Spurred on by the current reward and recognition systems of academia, it is easier and very tempting to quickly publish one-off findings which appear transformative, rather than invest additional money, energy and time to ensure that these one-off findings are reproducible. As a consequence, a significant number of ‘discoveries’ turn out to be unreliable because they are more likely to depend on small populations, weak statistics and flawed analysis [12–16]. There is also a temptation to carry out *post hoc* rationalization or HARKing, ‘Hypothesizing After the Results are Known’ and to invest more effort into explaining away unexpected findings than validating expected results. Most contemporary research depends heavily on computers which generate numbers with great facility. Ultimately, though, computers are themselves tools that are designed and used by people. Because human beings have a capacity for self-deception [17], the datasets and algorithms that they create can be subject to unconscious biases of various kinds, for example in the way data are collected and curated in data dredging activities, a lack of standardized data analysis workflows [18], or the selection of tools that generate promising results, even if their use is not appropriate in the circumstances.

No field of science is immune to these issues, but they are particularly challenging in domains where systems are complex and many dimensional, weakly underpinned by theoretical understanding, and exhibit nonlinearity, chaos and long-range correlations. With the rise of digital computing power, approaches predicated on big data, machine learning (ML) and artificial intelligence (AI) are frequently deemed to be indispensable. ML and AI are increasingly used to sift experimental and simulation data for otherwise hidden patterns that such methods may suggest are significant. Reproducibility is particularly important here because these forms of data analysis play a disproportionate role in producing results and supporting conclusions.

Some even maintain that big data analyses can do away with the scientific method [19]. However, as datasets increase in size, the ratio of false to true correlations increases very rapidly, so one must be able to reliably distinguish false from true if one is able to find robust correlations. That is difficult to do without a reliable theory underpinning the data being analysed. We, like others [20], argue that the faith placed in big data analyses is profoundly misguided: to be successful, big data methods must be more firmly grounded on the scientific method [21]. Far from being a threat to the scientific method, the weaknesses of blind big data methods serve as a timely reminder that the scientific method remains the most powerful means we have to understand our world.

2. Reproducibility

In science, unlike politics, it does not matter how many people say or agree about something: if science is to be objective, it has to be reproducible (within the error bars). Observations and ‘scientific facts and results’ cannot depend on who is reporting them but must be universal. The consensus is the business of politics and the scientific equivalent only comes after the slow accumulation of unambiguous pieces of empirical evidence (albeit most research and programmes are still funded on the basis of what the majority of people on a review panel thinks is right, so that scientists who have previously been successful are more likely to be awarded grants [22,23]).

There is some debate about the definition of reproducibility [24]. Some argue that replicability is more important than reproducibility. Others maintain that the gold standard of research should be ‘re-testability’, where the result is replicated rather than the experiment itself, though the degree to which the ‘same result’ can emerge from different setups, software and implementations is open to question [25].

By reproducibility we mean the repetition of the findings of an experiment or calculation, generally by others, providing independent confirmation and confidence that we understand what was done and how, thus ensuring that reliable ideas are able to propagate through the scientific community and become widely adopted. When it comes to computer modelling,

reproducibility means that the original data and code can be analysed by any independent, sceptical investigator to reach the same conclusions. The status of all investigators is supposedly equal and the same results should be obtained regardless of who is performing the study, within well-defined error bars—that is, reproducibility must be framed as a statistically robust criterion because so many factors can change between one set of observations and another, no matter who performs the experiment.

The uncertainties come in two forms: (i) ‘epistemic’, or systematic errors, which might be due to differences in measuring apparatus; and (ii) ‘aleatoric’, caused by random effects. The latter typically arise in chaotic dynamical systems which manifest extreme sensitivity to initial conditions, and/or because of variations in conditions outside of the control of an experimentalist.

By seeking to control uncertainty in terms of a margin of error, reproducibility means that an experiment or observation is robust enough to survive all manner of scientific analysis. Note, of course, that reproducibility is a necessary but not a sufficient condition for an observation to be deemed scientific. In the scientific enterprise, a single result or measurement can never provide a definitive resolution for or against a theory. Unlike mathematics, which advances when a proof is published, it takes much more than a single finding to establish a novel scientific insight or idea. Indeed, in the Popperian view of science, there can be no final vindication of the validity of a scientific theory: they are all provisional, and may eventually be falsified. The extreme form of the modern machine-learners’ pre-Baconian view stands in stark opposition to this: there is no theory at all, only data, and success is measured by how well one’s learning algorithm performs at discerning correlations within these data, even though many of these correlations will turn out to be false, random or meaningless.

Moreover, in recent years, the integrity of the scientific endeavour has been open to question because of issues around reproducibility, notably in the biological sciences. Confidence in the reliability of clinical research has, for example, been under increasing scrutiny [5]. In 2005, Ioannidis wrote an influential article about biomedical research, entitled ‘Why Most Published Research Findings are False’, in which he assessed the positive predictive value of the truth of a research finding from values such as threshold of significance and power of the statistical test applied [26]. He found that the more teams were involved in studying a given topic, the less likely the research findings from individual studies turn out to be true. This seemingly paradoxical corollary follows because of the scramble to replicate the most impressive ‘positive’ results and the attraction of refuting claims made in a prestigious journal, so that early replications tend to be biased against the initial findings. This ‘Proteus phenomenon’ has been observed as an early sequence of extreme, opposite results in retrospective hypothesis-generating molecular genetic research [27], although there is often a fine line to be drawn between contrarianism, wilful misrepresentation and the scepticism (*nullius in verba*) that is the hallmark of good science [28].

Such lack of reproducibility can be troubling. An investigation of 49 medical studies undertaken between 1990 and 2003—with more than 1000 citations in total—found that 16% were contradicted by subsequent studies, 16% found stronger effects than subsequent studies, 44% were replicated and 24% remained largely unchallenged [29]. In psychological science, a large portion of independent experimental replications did not reproduce evidence supporting the original results despite using high-powered designs and original materials when available [30]. Even worse performance is found in cognitive neuroscience [13].

Scientists more widely are routinely confronted with issues of reproducibility: a May 2016 survey in *Nature* of more than 1576 scientists reported that more than 70% had tried and failed to reproduce another scientist’s experiments, and more than half had failed to reproduce their own experiments [31]. This lack of reproducibility can be devastating for the credibility of a field.

3. Modelling and simulation

Computers are critical in all fields of data analysis and computer simulations need to be reliable—validated, verified and their uncertainty quantified—so that they can feed into real-world applications and decisions be they governmental policies dealing with pandemics, for

the global climate emergency, the provision of food and shelter for refugee populations fleeing conflicts, creation of new materials, the design of the first commercial fusion reactor or to assist doctors to test medication on a virtual patient before a real one.

Reproducibility in computer simulations would seem trivial to the uninitiated: enter the same data into the same program on the same architecture and you should get the same results. In practice, however, there are many barriers to overcome to ensure the fidelity of a model in a computational environment [32]. Overall, it can be challenging if not impossible to test the claims and arguments made by authors in published work without access to the original code and data, and even in some instances the machines the software ran on. One study of what the authors dubbed ‘weak repeatability’ examined 402 papers with results backed by code and found that, for one-third, they were able to obtain the code and build it within half an hour, while for just under half they succeeded with significant extra effort. For the remainder, it was not possible to verify the published findings. The authors reported that some researchers are reluctant to share their source code, for instance for commercial and licensing reasons, or because of dependencies on other software, whether due to external libraries or compilers, or because the version they used in their paper had been superseded, or had been lost due to lack of backup. Many detailed choices in the design and implementation of a simulation never make it into published papers. Frequently, the principal code developer has moved on, the code turns out to depend on exotic hardware, there is inadequate documentation, and/or the code developers say that they are too busy to help [33]. There are some high-profile examples of these issues, from the disclosure of climate codes and data,¹ to delays in sharing codes for COVID-19 pandemic modelling.² If the public are to have confidence in computing models that could directly affect them, transparency, openness and the timely release of code and data are critical.

In response to this challenge, there have been various proposals to allow scientists to openly share code and data that underlie their research publications: RunMyCode [runmycode.org] and, perhaps better known, GitHub [github.com]; SHARE, a web portal to create, share and access remote virtual machines that can be cited from research papers to make an article fully reproducible and interactive [34]; PaperMâché, another means to view and interact with a paper using virtual machines [35]; various means to create ‘executable papers’ [36,37]; and a verifiable result identifier (VRI), which consists of trusted and automatically generated strings that point to publicly available results originally created by the computational process [38].

In addition to external verification, there are many initiatives to incorporate verification and validation into computer model development, along with uncertainty quantification techniques to verify and validate the models [39]. In the United States, for example, the American Society of Mechanical Engineers has a Standards Committee for the development of verification and validation (V&V) procedures for computational solid mechanics models, guidelines and recommended practices have been developed by the National Aeronautics and Space Administration (NASA)³; the US Defense Nuclear Facilities Safety Board backs model V&V for all safety-related nuclear facility design, analyses and operations, while various groups within the DOE laboratories (including Sandia, Los Alamos and Lawrence Livermore) are conducting research in this area [40]. In Europe, the VECMA (Verified Exascale Computing for Multiscale Applications) project⁴ is developing software tools that can be applied to many research domains, from the laptop to the emerging generation of exascale supercomputers in order to validate, verify and quantify the uncertainty within highly diverse applications.

¹The disclosure of climate data from the Climatic Research Unit at the University of East Anglia—Science and Technology Committee: <https://publications.parliament.uk/pa/cm200910/cmselect/cmsctech/387/38703.htm> (accessed 29 December 2020).

²See <https://www.imperial.ac.uk/mrc-global-infectiousdisease-analysis/covid-19/report-9-impact-of-npis-on-covid-19/> and <https://github.com/ImperialCollegeLondon/covid19model> (accessed 29 December 2020).

³The NASA Langley UQ Challenge on Optimization Under Uncertainty: <https://uqtools.larc.nasa.gov/nasa-uq-challenge-problem-2020/> (accessed 29 December 2020).

⁴VECMA Verified Exascale Computing for Multiscale Applications: <https://www.vecma.eu/> (accessed 29 December 2020).

The major challenge faced by the state of the art is that many scientific models are multiphysics in nature, combining two or more kinds of physics, for instance to simulate the behaviour of plasmas in tokamak nuclear fusion reactors [41], electromechanical systems [42] or in food processing [43]. Even more common, and more challenging, many models are also multiscale, which require the successful convergence of various theories that operate at different temporal and/or spatial scales. They are widespread at the interface between various fields, notably physics, chemistry and biology. The ability to integrate macroscopic universality and molecular individualism is perhaps the greatest challenge of multiscale modelling [44]. As one example, we certainly need multiscale models if we are to predict the biology and medicine that underpin the behaviour of an individual person. Digital medicine is increasingly important and, as a corollary of this, there have been calls for steps to avoid a reproducibility ‘crisis’ of the kind that has engulfed other areas of biomedicine [45].

Although there are many kinds of multiscale modelling, there now exist protocols to enable the verification, validation, and uncertainty quantification of multiscale models [46]. The VECMA toolkit,⁵ which is not only open source but whose development is also performed openly, has many components: FabSim3, to organize and perform complex remote tasks; EasyVVUQ, a Python library designed to facilitate verification, validation and uncertainty quantification for a variety of simulations [47,48]; QCG Pilot Job, to provide the efficient and reliable execution of large number of computational jobs; QCG-Now, to prepare and run computational jobs on high-performance computing machines; QCG-Client, to provide support for a variety of computing jobs, from simple ones to complex distributed workflows; EasyVVUQ-QCGPilotJob, for efficient, parallel execution of demanding EasyVVUQ scenarios on high-performance machines; and MUSCLE 3, to make creating coupled multiscale simulations easier, and to then enable efficient uncertainty quantification of such models.

The VECMA toolkit is already being applied in several circumstances: climate modelling, where multiscale simulations of the atmosphere and oceans are required; forecasting refugee movements away from conflicts, or as a result of climate change, to help prioritize resources and investigate the effects of border closures and other policy decisions [49]; for exploring the mechanical properties of a simulated material at several length and time scales with verified multiscale simulations; and multiscale simulations to understand the mechanisms of heat and particle transport in fusion devices, which is important because the transport plays a key role in determining the size, shape and more detailed design and operating conditions of a future fusion power reactor, and hence the possibility of extracting almost limitless energy; and verified simulations to aid in the decision-making of drug prescriptions, simulating how drugs interact with a virtual version of a patient’s proteins, [50] or how stents will behave when placed in virtual versions of arteries [51]. The toolkit has also been used to demonstrate the very considerable uncertainty in the predictions arising from the CovidSim code used to make predictions of death rates caused by the COVID-19 pandemic [52,53].

4. Big data, machine learning and reproducibility

Recent years have seen an explosive growth in digital data accompanied by the rising public awareness that their lives depend on ‘algorithms’, though it is plain to all that any computer code is based on an algorithm, without which it will not run. Under the banner of AI and ML, many of these algorithms seek patterns in those data. Some—emphatically not the authors of this paper—even claim that this approach will be faster and more revealing than modelling the underlying behaviour by the use of conventional theory, modelling and simulation [54]. This approach is particularly attractive in disciplines traditionally not deemed suitable for mathematical treatment because they are so complex, notably life and social sciences, along with the humanities.

However, to build a machine-learning system, you have to decide what data you are going to choose to populate it. That choice is frequently made without any attempt to first try to

⁵VECMA Toolkit: <https://www.vecma-toolkit.eu/> (accessed 29 December 2020).

understand the structural characteristics that underlie the system of interest, with the result that the 'AI system' produced strongly reflects the limitations or biases (be they implicit or explicit) of its creators.

Moreover, there are four fundamental issues with big data that are frequently not recognized by practitioners [54]: complex systems are strongly correlated, so they do not generally obey Gaussian statistics; no datasets are large enough for systems with strong sensitivity to rounding or inaccuracies; correlation does not imply causality; and too much data can be as bad as no data: although computers can be trained on larger datasets than the human brain can absorb, there are fundamental limitations to the power of such datasets (as one very real example, mapping genotype to phenotype is far from straightforward), not least due to their digital character.

All machine-learning algorithms are initialized using (pseudo) random number generators and have to be run vast numbers of times to ensure that their statistical predictions are robust. However, they typically make plenty of other assumptions, such as smoothness (i.e. continuity) between data points. The problem is that nonlinear systems are often anything but smooth, and there can be jumps, discontinuities and singularities.

Not only the smoothness of behaviour but also the forms of distribution of data regularly assumed by machine learners are frequently unknown or untrue in complex systems. Indeed, many such approaches are distribution free, in the sense that there is no knowledge provided about the way the data being used is distributed in a statistical sense [54]. Often, a Gaussian (normal) distribution is assumed by default; while this distribution plays an undeniable role across all walks of science it is far from universal. Indeed, it fails to describe most phenomena where complexity holds sway because, rather than resting on randomness, these typically have feedback loops, interactions and correlations.

ML is often used to seek correlations in data. But in a real-world system, for instance, in a living cell that is a cauldron of activity of 42 million protein molecules [55], can we be confident that we have captured the right data? Random data dredging for complex problems is doomed to fail where one has no idea which variables are important. In these cases, data dredging will always be defeated by the curse of dimensionality—there will simply be far too much data needed to fill in the hyperdimensional space for blind ML to produce correlations to any degree of confidence. On top of that, as mentioned earlier, the ratio of false to true correlations soars with the size of the dataset, so that too much data can be worse than no data at all [55].

There are practical considerations too. Machine-learning systems can never be better than the data they are trained on, which can contain biases 'whether morally neutral as toward insects or flowers, problematic as toward race or gender, or even simply veridical, reflecting the *status quo* distribution of gender with respect to careers or first names' [56]. In healthcare systems, for example, where commercial prediction algorithms are used to identify and help patients with complex health needs, significant racial bias has been found [57]. Cathy O'Neil's book, *Weapons of Math Destruction* is replete with examples of this kind, covering virtually all walks of life, and their harmful impact on modern societies [58].

Machine-learning systems are black boxes, even to the researchers that build them, making it hard for their creators, let alone others, to assess the results produced by these glorified curve-fitting systems. Precise replication would be nearly impossible given the natural randomness in neural networks and variations in hardware and code. That is one reason why blind ML is unlikely to ever be accepted by regulatory authorities in medical practice as a basis for offering drugs to patients. To comply with the regulatory authorities such as the US Food and Drug Administration and the European Medicines Agency, the predictions of a ML algorithm are not enough and it is essential that an underlying mechanistic explanation is also provided, one which can explain not only when a drug works but also when it fails, and/or produces side effects.

There are even deeper problems of principle in seeking to produce reliable predictions about the behaviour of complex systems of the sort one encounters frequently in the most pressing problems of twenty-first-century science. We are thinking particularly, in life sciences, medicine, healthcare and environmental sciences, where systems typically involve large numbers of variables and many parameters. The question is how to select these variables and parameters

to best fit the data. Despite the constant refrain that we live in the age of 'Big Data', the data we have available is never enough to model problems of this degree of complexity. Unlike more traditional reductionist models, where one may reasonably assume one has sufficient data to estimate a small number of parameters, such as a drug interacting with a nerve cell receptor, this ceases to be the case in complex and emergent systems, such as modelling a nerve cell itself. The favourite approach of the moment is of course to select ML, which involves adjustments of large numbers of parameters inside the neural network 'models' used; these can be tuned to fit the data available but have little to no predictability beyond the range of the data used because they do not take into account the structural characteristics of the phenomenon under study. This is a form of overfitting [59]. As a result of the uncertainty in all these parameters, the model itself becomes uncertain as testing it involves an assessment of probability distributions over the parameters and, with nowhere near adequate data available, it is not clear if it can be validated in a meaningful manner [60]. For some related issues of a more speculative and philosophical nature in the study of complexity, see Succi [61]. The recently trumpeted announcement that DeepMind's AlphaFold algorithm solved the 'grand challenge' of predicting three-dimensional protein x-ray crystal structures from their one dimensional amino acid sequences, according to the Critical Assessment of Protein Structure Prediction [62], would certainly amount to a major development if and when it is fully confirmed scientifically; but it is notable that it produces poor matches to similar protein structures determined by another experimental method, nuclear magnetic resonance [63,64].

Compounding all this, there is a fundamental problem that undermines our faith in simulations which arises from the digital nature of modern computers, whether classical or quantum. Digital computers make use of four billion rational numbers that range from plus to minus infinity, the so-called 'single-precision IEEE floating-point numbers', which refers to a technical standard for floating-point arithmetic established by the Institute of Electrical and Electronics Engineers in the 1950s; they also frequently use double-precision floating-point numbers, while half-precision has become commonplace of late in the running of machine-learning algorithms.

However, digital computers only use a very small subset of the rational numbers—so-called dyadic numbers, whose denominators are powers of 2 because of the binary system underlying all digital computers—and the way these numbers are distributed is highly nonuniform. Moreover, there are infinitely more irrational than rational numbers, which are ignored by all digital computers because to store any one of them, typically, one would require an infinite memory. Manifestly, the IEEE floating-point numbers are a poor representation even of the rational numbers. Recent work by one of us (PVC), in collaboration with Bruce Boghosian and Hongyan Wan at Tufts University, demonstrates that there are major errors in the computer-based prediction of the behaviour of arguably the simplest of chaotic dynamical systems, the generalised Bernoulli map, for single precision floating point numbers. For a subset of values of the model's solitary parameter, very large errors accrue that cannot be mitigated by any increase in the precision of the numerical representation. For other parameter values, double precision reduces the sizeable errors substantially (Milan Kloewer, private communication with PVC). However, this leaves open the question as to whether double precision floating point numbers are themselves sufficient to handle the far more exquisite complexity of real world molecular dynamics and fluid turbulence, which originate in dynamical systems that are many orders of magnitude more complicated. The spectrum of the unstable periodic orbits of the map is badly damaged regardless of the precision of the floating point numbers [65].

Given the approximations involved in digital simulations of chaotic systems, found in models used to predict the weather, climate, molecular dynamics, chemical reactions, fusion energy and much more, not to speak of the various sources of measurement errors, it is never possible to obtain exact agreement with experimental results. In short, the use of floating-point numbers instead of real numbers can contribute additional systematic errors in numerical schemes that have not so far been fully assessed [61].

5. The solution

For modelling, we need to tackle both epistemic and aleatoric sources of error. To deal with these challenges, a number of countermeasures have been put forward: documenting detailed methodological and statistical plans of an experiment ahead of data collection (preregistration); demanding that studies are thoroughly replicated before they are published; [66,67] insisting on collaborations to double-check findings; [68] explicit consideration of alternative hypotheses, even processing all reasonable scenarios [69]; the sharing of methods, data, computer code and results in central repositories, such as the Open Science Framework,⁶ a free, open platform to support research, enable collaboration and ‘team science’ [70]; and blind data analysis, where data are shifted by an amount known only to the computer, leaving researchers with no idea what their findings implied until everyone agrees on the analyses and the blindfold is lifted. The role of universities, as in the Brazilian Reproducibility Initiative, [71] is important, along with conferences, such as the World Conferences on Research Integrity [<https://wcrif.org/>], and the actions of funding agencies, such as the US National Institutes of Health, [72] the UK research councils and the Wellcome Trust, [73] along with the French National Center for Scientific Research (CNRS), which has launched CASCAD, Certification Agency for Scientific Code and Data [www.cascad.tech], the first public laboratory specialized in the certification of the reproducibility of scientific research.

ML requires special consideration. A survey of 400 algorithms presented in papers at two top AI conferences (the 2013 and 2016 International Joint Conferences on Artificial Intelligence, IJCAI, and the 2014 and 2016 Association for the Advance of Artificial Intelligence, AAAI, conferences [74]) found that only 6% of the presenters shared the algorithm’s code [75]. The most commonly used machine-learning platforms provided by big tech companies have poor support for reproducibility [76]. Studies have shown that even if the results of a deep learning model could be reproduced, a slightly different experiment would not support the findings—yet another example of overfitting—which is common in machine-learning research. In other words, unreproducible findings can be built upon supposedly reproducible methods [77].

Rather than continuing to simply fund, pursue and promote ‘blind’ big data projects, more resources should be allocated to the elucidation of the multiphysics, multiscale and stochastic processes controlling the behaviour of complex systems, such as those in biology, medicine, healthcare and environmental science [21]. Finding robust predictive mechanistic models that provide explanatory insights will be of particular value for ML when dealing with sparse and incomplete sets of data, ill-posed problems, exploring vast design spaces to seek correlations and then, most importantly, for identifying correlations. Where ML provides a correlation, multiscale modelling can test if this correlation is causal.

There are also demands in some fields for a reproducibility checklist, [78] to make AI reproducibility more practical, reliable and effective. Another suggestion is the use of so-called ‘Model Cards’ – documentation that accompanies trained machine-learning models which outline the application domains, the context in which they are being used and their carefully benchmarked evaluation in a variety of conditions, such as across different cultural, demographic and phenotypic groups; [79] and proposals for best practice in reporting experimental results which permit for robust comparison [80].

Despite the caveat that computers are made and used by people, there is also considerable interest in their use to design and run experiments, for instance using Bayesian optimization methods, such as in the field of cognitive neuroscience [81] and to model infectious diseases and immunology quantitatively [82].

When it comes to the limitations of digital computing, research is still required to determine how reliably we can compute the properties of such chaotic dynamical systems on digital computers. Among possible solutions, one that seems guaranteed to succeed is analogue

⁶Open Science Framework (OSF). <https://osf.io/>.

computing, an older idea, able to handle the numerical continuum of reality in a way that digital computers can only approximate [83].

6. Synthesis and conclusion

In the short term, notably in the biosciences, better data collection, curation, validation, verification and uncertainty quantification procedures of the kind described here, will make computer simulations more reproducible, while ML will benefit from a more rigorous and transparent approach. The field of big data and ML has become extremely influential but without big theory, it remains dogged by a lack of firm theoretical underpinning ensuring its results are reliable [21]. Indeed, we have argued that in the modern era in which we aspire to describe really complex systems, involving many variables and vast numbers of parameters, there is not sufficient data to apply these methods reliably. Our models are likely to remain uncertain in many respects, as it is so difficult to validate them.

In the medium term, AI methods may, if carefully produced, improve the design, objectivity and analysis of experiments. However, this will always require the participation of people to devise the underlying hypotheses and, as a result, it is important to ensure that they fully grasp the assumptions on which these algorithms are based and are also open about these assumptions.

It is already becoming increasingly clear that ‘artificial intelligence’ is a digital approximation to reality. Moreover, in the long term, when we are firmly in the era of routine exascale and perhaps eventually also quantum computation, we will have to grapple with a more fundamental issue. Even though there are those who believe the complexity of the universe can be understood in terms of simple programs rather than by means of concise mathematical equations [84,85] digital computers are limited in the extent to which they can capture the richness of the real world [83,86]. Freeman Dyson, for example, speculated that for this reason the downloading of a human consciousness into a digital computer would involve ‘a certain loss of our finer feelings and qualities’ [87]. In the quantum and exascale computing eras, we will need renewed emphasis on the analogue world and analogue computational methods if we are to trust our computers [83]. A photon-based computer, Jiuzhang, recently demonstrated a very substantial quantum computational advantage over classical computers, providing a glimpse of the potential of analogue computing, being both extremely fast and having very low power consumption [88].

Data accessibility. This article has no additional data.

Authors’ contributions. All authors contributed to the concept and writing of the article.

Competing interests. The authors have no competing interests.

Funding. P.V.C. is grateful for funding from the UK EPSRC for the UKCOMES UK High-End Computing Consortium (grant no. EP/R029598/1), from MRC for a Medical Bioinformatics (grant no. MR/L016311/1), the European Commission for the CompBioMed, CompBioMed2 and VECMA (grant nos. 675451, 823712 and 800925, respectively) and special funding from the UCL Provost.

Acknowledgements. The authors are grateful for many stimulating conversations with Bruce Boghosian, Daan Crommelin, Ed Dougherty, Derek Groen, Alfons Hoekstra, Robin Richardson and David Wright.

References

1. Whitaker K. 2020 ‘The Turing Way’ - A handbook for reproducible data science. Accessed June 24, 2020. <https://www.turing.ac.uk/research/research-projects/turing-way-handbook-reproducible-data-science>.
2. Sonnenburg S *et al.* 2007 The need for open source software in machine learning. *J. Mach. Learn. Res.* **8**, 2443–2466.
3. Wren K. 2014 As concerns about non-reproducible data mount, some solutions take shape. Washington, DC: American Association for the Advancement of Science. Accessed 30 December, 2020. <https://www.aaas.org/news/concerns-about-non-reproducible-data-mount-some-solutions-take-shape>
4. Crocker J, Cooper ML. 2011 Addressing scientific fraud. *Science* **334**, 1182. (doi:10.1126/science.1216775)

5. Goldacre B. 2012 *Bad pharma: how drug companies mislead doctors and harm patients*. London, UK: Fourth Estate.
6. Langmuir I. 1989 Pathological science. *Res. Manag.* **32**, 11–17. (doi:10.1080/08956308.1989.11670607)
7. Broad W, Wade N. 1983 *Betrayers of the truth: fraud and deceit in the halls of science*. New York, NY: Century Publishing.
8. Roache PJ. 1997 Quantification of uncertainty in computational fluid dynamics. *Annu. Rev. Fluid. Mech.* **29**, 123–160. (doi:10.1146/annurev.fluid.29.1.123)
9. Assessing the Reliability of Complex Models. 2012 *Mathematical and statistical foundations of verification, validation, and uncertainty quantification*. Washington, DC: National Academies Press.
10. Steinman DA, Migliavacca F. 2018 Editorial: special issue on verification, validation, and uncertainty quantification of cardiovascular models: towards effective VVUQ for translating cardiovascular modelling to clinical utility. *Cardiovasc. Eng. Technol.* **9**, 539–543. (doi:10.1007/s13239-018-00393-z)
11. Nuzzo R. 2015 How scientists fool themselves – and how they can stop. *Nature* **526**, 182–185. (doi:10.1038/526182a)
12. Dumas-Mallet E, Button KS, Boraud T, Gonon F, Munafò MR. 2017 Low statistical power in biomedical science: a review of three human research domains. *R. Soc. Open Sci.* **4**, 160254. (doi:10.1098/rsos.160254)
13. Szucs D, Ioannidis JPA. 2017 Empirical assessment of published effect sizes and power in the recent cognitive neuroscience and psychology literature. *PLoS Biol.* **15**, e2000797. (doi:10.1371/journal.pbio.2000797)
14. Jennions MD, Moller AP. 2003 A survey of the statistical power of research in behavioral ecology and animal behavior. *Behav. Ecol.* **14**, 438–445. (doi:10.1093/beheco/14.3.438)
15. Fraley RC, Vazire S. 2014 The N-pact factor: evaluating the quality of empirical journals with respect to sample size and statistical power. *PLoS ONE* **9**, e109019. (doi:10.1371/journal.pone.0109019)
16. Button KS, Ioannidis JPA, Mokrysz C, Nosek BA, Flint J, Robinson ES, Munafò MR. 2013 Power failure: why small sample size undermines the reliability of neuroscience. *Nat. Rev. Neurosci.* **14**, 365–376. (doi:10.1038/nrn3475)
17. von Hippel W, Trivers R. 2011 The evolution and psychology of self-deception. *Behav. Brain Sci.* **34**, 1–16. (doi:10.1017/S0140525X10001354)
18. Botvinik-Nezer R *et al.* 2020 Variability in the analysis of a single neuroimaging dataset by many teams. *Nature* **582**, 84–88. (doi:10.1038/s41586-020-2314-9)
19. Anderson C. 2008 The end of theory: the data deluge makes the scientific method obsolete. *Wired Magazine* **16**, 16-07.
20. Forde JZ, Paganini M. 2019 The scientific method in the science of machine learning. Published online April 2019. (arXiv:1904.10922)
21. Coveney PV, Dougherty ER, Highfield RR. 2016 Big data need big theory too. *Phil. Trans. R. Soc. A* **374**, 20160153. (doi:10.1098/rsta.2016.0153)
22. Merton RK. 1968 The Matthew effect in science: the reward and communication systems of science are considered. *Science* **159**, 56–63. (doi:10.1126/science.159.3810.56)
23. Bol T, de Vaan M, van de Rijt A. 2018 The Matthew effect in science funding. *Proc. Natl Acad. Sci. USA* **115**, 4887–4890. (doi:10.1073/pnas.1719557115)
24. Baker M. 2016 Muddled meanings hamper efforts to fix reproducibility crisis. *Nature* Published online June 2016. (doi:10.1038/nature.2016.20076)
25. Drummond C. 2018 Reproducible research: a minority opinion. *J. Exp. Theor. Artif. Intell.* **30**, 1–11. (doi:10.1080/0952813X.2017.1413140)
26. Ioannidis JPA. 2005 Why most published research findings are false. *PLoS Med.* **2**, e124. (doi:10.1371/journal.pmed.0020124)
27. Ioannidis JPA, Trikalinos TA. 2005 Early extreme contradictory estimates may appear in published research: the proteus phenomenon in molecular genetics research and randomized trials. *J. Clin. Epidemiol.* **58**, 543–549. (doi:10.1016/j.jclinepi.2004.10.019)
28. de Winter J, Happee R. 2013 Why selective publication of statistically significant results can be effective. *PLoS ONE* **8**, e66463. (doi:10.1371/journal.pone.0066463)
29. Ioannidis JPA. 2005 Contradicted and initially stronger effects in highly cited clinical research. *J. Am. Med. Assoc.* **294**, 218–228. (doi:10.1001/jama.294.2.218)

30. Nosek BA, Aarts AA, Anderson CJ, Joanna E, Kappes H. 2015 Estimating the reproducibility of psychological science. *Science* **349**, aac4716. (doi:10.1126/science.aac4716)
31. Baker M. 2016 1,500 scientists lift the lid on reproducibility. *Nature* **533**, 452–454. (doi:10.1038/533452a)
32. Ivie P, Thain D. 2018 Reproducibility in scientific computing. *ACM Comput. Surv.* **51**, 1–36. (doi:10.1145/3186266)
33. Collberg C, Proebsting T, Moraila G, Shankaran A, Shi Z, Warren AM. 2014 Measuring reproducibility in computer systems research. *Dep. Comput. Sci. Univ. Arizona, Tech. Rep.* **51**, 1–37. (doi:10.1145/3186266)
34. Van Gorp P, Mazanek S. 2011 SHARE: a web portal for creating and sharing executable research papers. *Procedia Comput. Sci.* **4**, 589–597. (doi:10.1016/j.procs.2011.04.062)
35. Brammer GR, Crosby RW, Matthews SJ, Williams TL. 2011 Paper Mâché: creating dynamic reproducible science. *Procedia Comput. Sci.* **4**, 658–667. (doi:10.1016/j.procs.2011.04.069)
36. Koop D *et al.* 2011 A provenance-based infrastructure to support the life cycle of executable papers. *Procedia Comput. Sci.* **4**, 648–657. (doi:10.1016/j.procs.2011.04.068)
37. Müller W, Rojas I, Eberhart A, Haase P, Schmidt M. 2011 A-R-E: the author-review-execute environment. *Procedia Comput. Sci.* **4**, 627–636. (doi:10.1016/j.procs.2011.04.066)
38. Gavish M, Donoho D. 2011 A universal identifier for computational results. *Procedia Comput. Sci.* **4**, 637–647. (doi:10.1016/j.procs.2011.04.067)
39. Roy CJ, Oberkampf WL. 2011 A comprehensive framework for verification, validation, and uncertainty quantification in scientific computing. *Comput. Methods Appl. Mech. Eng.* **200**, 2131–2144. (doi:10.1016/j.cma.2011.03.016)
40. Thacker BH, Doebling SW, Hemez FM, Anderson MC, Pepin JE, Rodriguez EA. 2004 Concepts of model verification and validation. (doi:10.2172/835920)
41. Williamson RL, Hales JD, Novascone SR, Tonks MR, Gaston DR, Permann CJ, Andrs D, Martineau RC. 2012 Multidimensional multiphysics simulation of nuclear fuel behavior. *J. Nucl. Mater.* **423**, 149–163. (doi:10.1016/j.jnucmat.2012.01.012)
42. Dede EM, Lee J, Nomura T. 2014 *Multiphysics simulation*. London, UK: Springer.
43. Knoerzer K, Juliano P, Versteeg C (eds) 2011 *Innovative food processing technologies: advances in multiphysics simulation*. Hoboken, NJ: Wiley-Blackwell.
44. Coveney PV, Boon JP, Succi S. 2016 Bridging the gaps at the physics–chemistry–biology interface. *Phil. Trans. R. Soc. A* **374**, 20160335. (doi:10.1098/rsta.2016.0335)
45. Stupple A, Singerman D, Celi LA. 2019 The reproducibility crisis in the age of digital medicine. *npj Digit. Med.* **2**, 2. (doi:10.1038/s41746-019-0079-z)
46. Karabasov S, Nerukh D, Hoekstra A, Chopard B, Coveney PV. 2014 Multiscale modelling: approaches and challenges. *Phil. Trans. R. Soc. A* **372**, 20130390. (doi:10.1098/rsta.2013.0390)
47. Wright DW *et al.* 2020 Building confidence in simulation: applications of EasyVVUQ. *Adv. Theory Simul.* **3**, 1900246.
48. Richardson RA, Wright DW, Edeling W, Jancauskas V, Lakhili J, Coveney PV. 2020 EasyVVUQ: a library for verification, validation and uncertainty quantification in high performance computing. *J. Open Res. Softw.* **8**, 11. (doi:10.5334/jors.303)
49. Suleimenova D, Bell D, Groen D. 2017 A generalized simulation development approach for predicting refugee destinations. *Sci. Rep.* **7**, 13377. (doi:10.1038/s41598-017-13828-9)
50. Potterton A, Hussein FS, Southey MWY, Bodkin MJ, Heifetz A, Coveney PV, Townsend-Nicholson A. 2019 Ensemble-based steered molecular dynamics predicts relative residence time of a 2A receptor binders. *J. Chem. Theory Comput.* **15**, 3316–3330. (doi:10.1021/acs.jctc.8b01270)
51. Nikishova A, Veen L, Zun P, Hoekstra AG. 2019 Semi-intrusive multiscale metamodelling uncertainty quantification with application to a model of in-stent restenosis. *Phil. Trans. R. Soc. A* **377**, 20180154. (doi:10.1098/rsta.2018.0154)
52. Edeling W *et al.* 2020 Model uncertainty and decision making: predicting the impact of COVID-19 using the CovidSim epidemiological code. See <https://www.researchsquare.com/article/rs-82122/v3> (accessed 29 December 2020)
53. Adam D 2020 Simulating the pandemic: what COVID forecasters can learn from climate models. See <https://www.nature.com/articles/d41586-020-03208-1> (accessed 30 December 2020)

54. Succi S, Coveney PV. 2019 Big data: the end of the scientific method? *Phil. Trans. R. Soc. A* **377**, 20180145. (doi:10.1098/rsta.2018.0145)
55. Ho B, Baryshnikova A, Brown GW. 2018 Unification of protein abundance datasets yields a quantitative *Saccharomyces cerevisiae* proteome. *Cell Syst.* **6**, 192–205.e3. (doi:10.1016/j.cels.2017.12.004)
56. Caliskan A, Bryson JJ, Narayanan A. 2017 Semantics derived automatically from language corpora contain human-like biases. *Science* **356**, 183–186. (doi:10.1126/science.aal4230)
57. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. 2019 Dissecting racial bias in an algorithm used to manage the health of populations. *Science* **366**, 447–453. (doi:10.1126/science.aax2342)
58. O’Neil C. 2016 *Weapons of math destruction*. New York, NY: Crown Random House.
59. Hawkins DM. 2004 The problem of overfitting. *J. Chem. Inf. Comput. Sci.* **44**, 1–12. (doi:10.1021/ci0342472)
60. Dougherty ER. 2016 *The evolution of scientific knowledge: from certainty to uncertainty*. Bellingham, WA: SPIE Press.
61. Succi S. 2019 Of naturalness and complexity. *Eur. Phys. J. Plus* **134**, 97. (doi:10.1140/epjp/i2019-12576-3)
62. Alphafold. 2020 AlphaFold: a solution to a 50-year-old grand challenge in biology. See <https://deepmind.com/blog/article/alphafold-a-solution-to-a-50-year-old-grand-challenge-in-biology> (accessed 30 December 2020)
63. Kinch L, Pei J, Schaeffer D, Grishin N. 2020 CASP14 tertiary structure prediction assessment topology (FM) category. See https://predictioncenter.org/casp14/doc/presentations/2020_11_30_Topology_assessment1_Kinch_Updated.pdf.
64. Pei J. 2020 Examples of FM/TBM target analysis. See https://predictioncenter.org/casp14/doc/presentations/2020_11_30_Topology_assessment2_JiminPei.pdf.
65. Boghosian BM, Coveney PV, Wang H. 2019 A new pathology in the simulation of chaotic dynamical systems on digital computers. *Adv. Theory Simul.* **2**, 1900125. (doi:10.1002/adts.201900125)
66. Schooler JW. 2014 Metascience could rescue the ‘replication crisis’. *Nature* **515**, 9. (doi:10.1038/515009a)
67. Wilson EO. 1998 Consilience: the unity of knowledge. In: *Consilience: the unity of knowledge*. (ed. Alfred A. Knopf), p. 332. <http://www.wtf.tw/ref/wilson.pdf>.
68. Yong E. 2013 Psychologists strike a blow for reproducibility. *Nature*. Published online November 2013. (doi:10.1038/nature.2013.14232)
69. Steegen S, Tuerlinckx F, Gelman A, Vanpaemel W. 2016 Increasing transparency through a multiverse analysis. *Perspect. Psychol. Sci.* **11**, 702–712. (doi:10.1177/1745691616658637)
70. Stokols D, Hall KL, Taylor BK, Moser RP. 2008 The science of team science. *Am. J. Prev. Med.* **35**, S77–S89. (doi:10.1016/j.amepre.2008.05.002)
71. de Oliveira Andrade R. 2019 Brazilian biomedical science faces reproducibility test. *Nature* **569**, 318–319. (doi:10.1038/d41586-019-01485-z)
72. Collins FS, Tabak LA. 2014 Policy: NIH plans to enhance reproducibility. *Nature* **505**, 612–613. (doi:10.1038/505612a)
73. The Academy of Medical Sciences. 2015 Reproducibility and reliability of biomedical research: improving research practice. In Symp. Rep. 2015 October, pp. 1–77.
74. Gundersen OE, Kjensmo S. 2018 State of the art: Reproducibility in artificial intelligence. In *32nd AAAI Conf. Artif Intell AAAI 2018*, pp. 1644–1651. Published online 2018.
75. Hutson M. 2018 Missing data hinder replication of artificial intelligence studies. *Science*. Published online February 2018. (doi:10.1126/science.aat3298)
76. Isdahl R, Gundersen OE. 2019 Out-of-the-Box Reproducibility: A Survey of Machine Learning Platforms. In *2019 15th Int. Conf. on EScience (EScience)*, pp. 86–95. IEEE. (doi:10.1109/eScience.2019.00017)
77. Bouthillier X, Laurent C, Vincent P. 2019 Unreproducible research is reproducible. In *36th Int. Conf. Mach Learn ICML 2019*. 2019-June, pp. 1150–1159.
78. Pineau J. 2018 The machine learning reproducibility checklist (Version 1.0). Published online 2018. (<https://www.cs.mcgill.ca/~jpineau/ReproducibilityChecklist.pdf>)
79. Mitchell M, Wu S, Zaldívar A, Barnes P, Vasserman L, Hutchinson B, Spitzer E, Raji ID, Gebru T. 2019 Model cards for model reporting. In *Proceedings of the Conference on Fairness*,

- Accountability, and Transparency, Atlanta, GA, 29–31 January 2019.* New York, NY: Association for Computing Machinery. (doi:10.1145/3287560.3287596)
80. Dodge J, Gururangan S, Card D, Schwartz R, Smith NA. 2019 Show your work: improved reporting of experimental results. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), Hong Kong, 3–7 November 2019*, pp. 2185–2194. Stroudsburg, PA: Association for Computational Linguistics. (doi:10.18653/v1/D19-1224)
 81. Lorenz R, Hampshire A, Leech R. 2017 Neuroadaptive Bayesian optimization and hypothesis testing. *Trends Cogn. Sci.* **21**, 155–167. (doi:10.1016/j.tics.2017.01.006)
 82. Eccleston RC, Coveney PV, Dalchau N. 2017 Host genotype and time dependent antigen presentation of viral peptides: predictions from theory. *Sci. Rep.* **7**, 14367. (doi:10.1038/s41598-017-14415-8)
 83. Coveney PV, Highfield RR. 2020 From digital hype to analogue reality: Universal simulation beyond the quantum and exascale eras. *J. Comput. Sci.* **46**, 101093. (doi:10.1016/j.jocs.2020.101093)
 84. Wolfram S. The Wolfram Physics Project: a project to find the fundamental theory of physics. See <https://www.wolframphysics.org/>.
 85. Wolfram S. 2002 A New Kind of Science. Wolfram Media. See <https://www.wolframscience.com/nks/>.
 86. Wolfram S. 2020 A class of models with the potential to represent fundamental physics. *Complex Systems* **29**, 107–536. (doi:10.25088/ComplexSystems.29.2.107)
 87. Dyson F. 2020 Is life analog or digital? Published 2001. Accessed June 24, 2020. See https://www.edge.org/conversation/freeman_dyson-is-life-analog-or-digital.
 88. Zhong H-S *et al.* 2020 Quantum computational advantage using photons. *Science* **370**, 1460–1463. (doi:10.1126/science.abe8770)