Computational modelling and sustainable development

Jonathan Dawes
Department of Mathematical Sciences, University of Bath, Bath BA2 7AY, UK

January 2023

Abstract: This article discusses themes of quantitative modelling relevant to 'world modelling' in the sense of The Limits to Growth and its recent successors (in particular the Earth4All model) and attempts to place them within a more general context of mathematical and computational modelling. Since the 1960s, computational modelling has been employed widely across science and engineering and increasingly in social and public policy. System Dynamics represents a specific kind of computational modelling and has its particular limitations. Directions for possible future development of the Earth4All model are set out.

Keywords: System Dynamics; dynamical systems; complex systems; machine learning; artificial intelligence

1. Introduction

The pioneering work in the early 1970s by the authors of The Limits to Growth, sponsored by the Club of Rome, was novel in several ways, notably (i) its attempts to generate a coherent and overarching description of the interactions between population, pollution, and resource use, and (ii) its use of what was at the time substantial computational resources to formulate and run the model for the six scenarios that form the basis for The Limits to Growth report and its predictions.

In the decades that followed not only were the results, and modelling assumptions (both explicit and implicit) discussed widely, but computational modelling itself has become a widely used tool employed to help clarify, understand and predict complex phenomena across a wide range of fields, not just in scientific and industrial research but also in social science and public policy defined broadly. This breadth of application has evolved alongside huge advances in computational power, data processing and storage, and algorithm design.

In the light of both of these issues - the increased breadth of modelling approaches and areas of application and the technological advances - it is of interest to reflect on The Limits to Growth and its contemporary successor, the Earth for All model, to review how these models are situated in the wider computational modelling landscape and to speculate as to what future modelling directions, or research questions, might look like.

The commonly-heard wisdom, often attributed to the statistician George Box, that "All models are wrong, but some are useful" is a convenient starting point for any discussion of mathematical or statistical modelling, and is indeed close to being a tautology. For, in order to be a model of a process, at least one (and usually many) aspects of that process must be discarded, or approximated, or assumed to take a particular form. A complete and exact copy of a process is not a model of a process: it is the same process itself. The point of a model is to obtain a simplified description that enables key concepts and their interaction to be understood. The level of simplification may vary, but it is always driven by the need to convey truth to others: for example specialists in other domains, decision makers, or the general public. The focus on interactions between elements of the model puts the activity of modelling immediately very close to the language and philosophy of 'complex systems': how do complicated phenomena arise through the interaction of a number of simpler parts? And what outcomes could not be envisaged even through very careful consideration of the parts of the system in isolation?

1.1 Computational modelling

The report 'Computational Modelling: Technological Futures' (abbreviated here to CMTF for convenience) published in 2018 by the UK Government Office for Science and commissioned by the Prime Minister's Council for Science and Technology attempts to survey the contemporary modelling landscape in non-technical language. For example, it reminds readers that models are tools to help understand a part of reality and, by definition, omit a certain level of detail. Moreover, models are designed to address specific questions and may not provide helpful (or even valid) answers when they are re-purposed in order to address new questions. This sense of the bounded scope of any particular model is essential. One could indeed go further and seek always to establish the conditions under which a model fails to give answers that agree with reality, thus establishing clearly a sense of the boundaries, or constraints, under which the model is operating. This is often in practice an extremely useful exercise, although it is rarely reported on and rarely discussed. This failure to acknowledge the operational boundaries of a model is perhaps one of our most significant cultural failings around modelling as a discipline. We do not like to write, or talk, about failure but it is scientifically essential that we understand when and how models fail to capture the underlying reality.

In terms of the purpose of modelling, the CMTF report sets out five purposes that might lie behind the construction of a model. These are summarised briefly in Table 1 below.

Modelling purpose	Advantages	Disadvantages
Analogy	Provides new insights into complex systems.	Does not capture full range of influences or outcomes
Illustration or visualisation	Assists in the communication of new ideas. Usually rapid to run the model	Restricts worldview to only the model ingredients
Understanding theory	Enables testing of ideas and hypotheses.	Theoretical assumptions may be validated in the model but not be (sufficiently) true in the real world
Explanation or exploration of future scenarios	Can be validated quantitatively. Able to demonstrate the range of likely outcomes, even allowing probabilities to be assigned across these	Increased model complexity obscures conditions relating to different outcomes. Higher computational complexity required to explore the full range of scenarios
Prediction or forecasting	Can predict outcomes unknown to the model builder	Model complexity inhibits transparency and explanatory power

Table 1: Five purposes for computational modelling, in ascending order of modelling detail, together with brief notes on their relative advantages and disadvantages. Author's adaption from *Computational Modelling: Technological Futures* (UK Government Office for Science, 2018).

The level of modelling detail required increases as one moves down the list set out in Table 1. Indeed, to move between different purposes might require a fundamental change in the modelling approach. For example, a model built to visualise a particular process, or to validate a theoretical point of view should probably not be used without further development as a prediction tool since the model will, probably by construction, omit one or more factors that would allow the investigation of the level of uncertainty in the model outputs, and hence the level of uncertainty present in any predictions or forecasts. As the authors of CMTF remark "*It is dangerous to interpret an exploration of theory as a conclusion about how the real world works.*" (page 20). To an extent, of course, this division into a classification of five purposes and their advantages and disadvantages is rather artificial; different modelling approaches will be constrained in different ways, and some advantages and disadvantages are common to more than one row in Table 1, although not indicated there explicitly.

An increase in the level of modelling detail is almost always accompanied by a decrease in the transparency of the modelling approach. This is made even worse by some (but not all) data-driven approaches within what is collectively referred to as 'machine learning', the algorithmic engine behind much of what is colloquially referred to as Artificial Intelligence. We will return to this issue in section 3.

To conclude this introduction we remark that the CMTF report includes in fact a concise discussion of The Limits to Growth (CMTF, pages 20-22), as part of its discussion of the second level in the hierarchy: 'illustration or visualisation'. The CMTF report attempts rather explicitly to justify its inclusion at that level in the hierarchy by including graphs that summarise the general shape of the overshoot and collapse scenarios that The Limits to Growth warned of. But clearly the way in which The Limits to Growth was able to communicate new ideas was a key part of its success in moving the debate around economic growth and environmental damage into new ground in the 1970s and 1980s. It could certainly also be argued that the formulation of The Limits to Growth embedded various problematic features into the debate that continue to resonate. For example the role of social and cultural norms, the role of the inherent tensions between commercial law and national legal systems, and the need to identify which parts of civil society, business and industry, or national governments are responsible for actions to address the challenges described by The Limits to Growth. Models that had included these features more prominently might have directed the discussion differently in the following decades.

The broad outline of this article is as follows; in general terms the article starts from more theoretical perspectives and moves towards specific applications in later sections. In section 1 we discuss System Dynamics as the paradigm within which both the Limits to Growth and the Earth for All models are framed and compare it with the sub-field of mathematics known as dynamical systems theory. Section 2 presents a broader review of the landscape of modern computational modelling and examples of where specific modelling frameworks have recently been used (including other examples of system dynamics models). Section 3 contains a discussion and conclusions focussing specifically on the current Earth for All programme and Earth4All model.

2. System Dynamics and dynamical systems

In this section we give brief overviews of two quantitative modelling areas that seems to be less closely related in practice, and driven by rather separate communities, than their names might suggest: System Dynamics and the mathematical field known as dynamical systems. The similarity in the names provides an easy provocation and motivation for such a discussion. After brief, and incomplete, overviews of each area, a discussion of the similarities and differences in the aims and approaches then follows.

2.1 An overview of System Dynamics

Although many substantial overviews of System Dynamics are available, not least the book *Thinking in Systems* by Donella Meadows (Meadows, 2008), it is useful here to summarise at least a few key elements. Meadows begins with a definition of a system as "an interconnected set of elements that is coherently organised in a way that achieves something." From this she abstracts three kinds of entity: elements, interconnections and a (sense of an overall) purpose. Implicit in this definition is a sense of a system as evolving through time (hence 'Dynamics') rather than being a fixed framework. The state of a system is given in terms of the levels of various elements, referred to as *stocks.* Levels of stocks change in time due to flows. The flow rates themselves in general depend on many things: other flow rates, the levels of other stocks, fixed influences beyond the boundaries of the system, and system parameters. System Dynamics models can then be set out in two ways: causal loop diagrams (CLDs) and stock-and-flow diagrams (or simply 'flow diagrams').

A very simple example would be a system to control the temperature in a room. The warmth of the room is the stock, with its temperature being the level of the stock. The inflow of heat into the room and the outflow into the open air through the windows are the flow rates. The flow rates themselves are controlled by system parameters such as the desired temperature of the room, the outside air temperature, and the efficiency of the heating system.

As even this example demonstrates, there may be many more rates, and many more influences on those rates, than stocks. It is also clear that System Dynamics is inherently concerned with measurable quantities that can be reasonably be described as continuously varying in time, for example levels of food supply, or demographics. At a conceptual level it is less easy to capture more abstract notions such as the level of social unrest, or the degree to which governance is democratic, or more discrete ones (for example whether a country has, or does not have, nuclear weapons).

Complexity in System Dynamics models arises out of feedbacks: situations in which rates eventually influence themselves as part of closed loops containing other rates (or stocks). Feedback loops can be either stabilising or destabilising. In the simplest case, as observed by Malthus and Ricardo, population sizes are subject to destabilising feedback since a larger population results in more births, and hence a further increase in population unless other effects (such as food scarcity or disease) are present. In System Dynamics terms, this exponential growth in population is due to a *reinforcing loop* that describes the destabilisation. Stabilising influences are due to negative feedbacks, termed *balancing loops*.

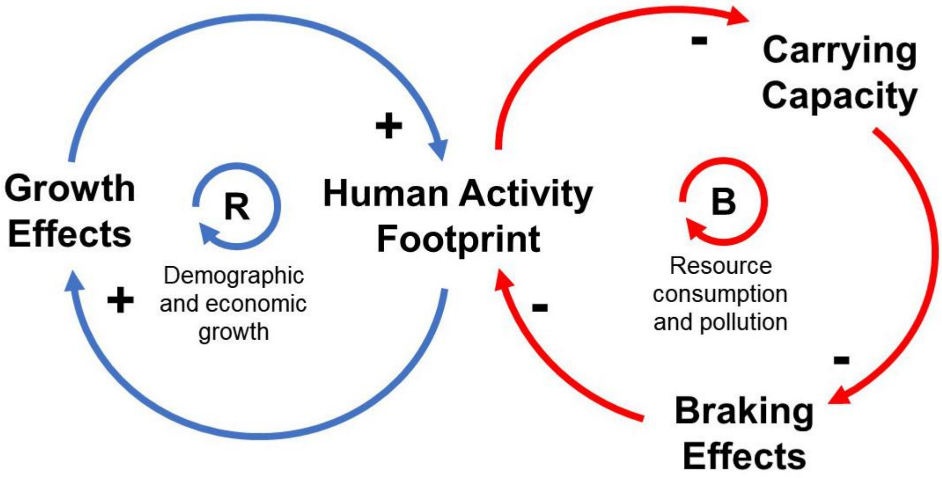

Figure 1: A highly simplified version of the central argument of The Limits to Growth in which the self-reinforcing nature of economic and population growth (the left-hand side loop) drives an ever-increasing footprint of human activity, balanced by the consequential reduction in resources which acts to suppress human activity, shown by the balancing loop on the right hand side. A positive sign at the head of an arrow indicates that a change in the quantity at the tail of the arrow produces a change in the same direction in the quantity at the arrowhead. A negative sign indicates that changes in the head and tail quantities will be in opposite directions. Adapted from CTMGF (2018), page 22.

Causal loop diagrams provide a useful summary of these feedbacks within a System Dynamics model, as illustrated in Figure 1 which shows a CLD that abstracts the main argument of The Limits to Growth. Figure 1 contains one reinforcing loop, denoted ***R***, on the left hand side, in which economic and population growth serve to expand the global footprint of human activity. The size of the global footprint is also influenced by the Earth's carrying capacity. Indeed, an increase in the Earth's carrying capacity would allow an increase in the human activity footprint (since the two negative signs on the red arrows in the lower half of the right hand side would combine to produce a positive influence overall). But because an increase in the human activity footprint leads directly (through the generation of pollution and the depletion of natural resources) to a reduction in the Earth's carrying capacity, the loop on the right hand side, denoted ***B***, is overall a balancing loop (three negative influences combining to generate an overall negative influence). The larger number of arrows in the balancing loop also hints at the fact that the balancing effect takes longer to occur, although this is not explicit. This longer response time allows the reinforcing loop to dominate earlier and to drive the system into overshoot, before the balancing loop is able to bring about a correction, or in extreme cases a collapse, mode of response.

Stock-and-flow diagrams for the World 2 and World 3 models are in themselves not particularly large, and are reproduced over two pages each, for example, in Forrester (1973) on pages 20-21, and in Cole et al (1973) on pages 16-17 and 20-21. The complexity in these models arises from the number of system parameters, for example multipliers that combine to produce birth and death rates, the addition of time delays, and the need to estimate many functional relationships (described as 'look-up tables') that describe the nonlinear relations between rates and levels.

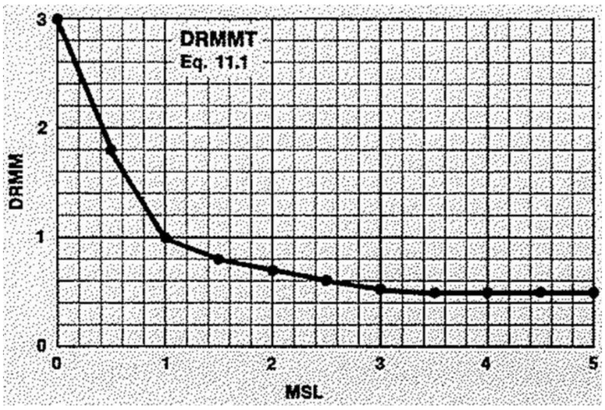


```
DRMM.K=TABHL(DRMMT,MSL.K,0,5,.5)                         11, A
DRMMT=3/1.8/1/.8/.7/.6/.53/.5/.5/.5/.5                   11.1, T
   DRMM     - DEATH-RATE-FROM-MATERIAL MULTIPLIER
                (DIMENSIONLESS)
   TABHL    - LOGICAL FUNCTION, TABLE LOOK UP AND
                INTERPOLATION
   DRMMT    - DEATH-RATE-FROM-MATERIAL-MULTIPLIER TABLE
   MSL      - MATERIAL STANDARD OF LIVING (DIMENSIONLESS)
```

Figure 2: The Death-Rate-from-Material-Multiplier that enhances or reduces the basic death rate if the Material Standard of Living (MSL) falls or rises above a value of one, respectively. Left: a plot of the graph of DRMM (vertical axis) as a function of MSL (horizontal axis). Right: the DYNAMO program code that implements the piecewise-linear function shown in the graph on the left. The look-up table comprises the dots shown in the graph which are defined by the numerical values in the first two lines of code (MSL values in the first line, DRMM values in the second line). Reproduced from Forrester (1973) with permission.

Figure 2 shows, as an example, the Death-Rate-from-Material Multiplier (DRMM) from the World 2 model which is proposed to be a sharply decreasing function of the material standard of living (MSL), normalised using the value in 1970, i.e. if the MSL increases above the 1970 reference point then the DRMM falls below unity, to perhaps a value as low as one half. However if the MSL decreases below the 1970 value then the DRMM rises rapidly to values of two or even three, causing the death rate to rise rapidly also. Three other multipliers affect the death rate computed in World 3; these describe increases (or reductions) in the normal death rate due to levels of pollution, food availability and crowding. The effect of each of these is applied independently, which Forrester acknowledged immediately was potentially problematic, but hoped that the model did not enter states which the death rate had varied so substantially that the overall effect became unrealistic.

Nevertheless, the level of complexity and 'internal parametrisation' illustrated by the above discussion of the death rate fuelled one substantial avenue of criticism of The Limits to Growth, namely that with quite so many parameters, and look-up tables of this kind, for which the underlying data could not be reconstructed empirically, the use of the model for any kind of quantitative analysis was extremely speculative. The uncertainty in these proposed relationships, and the sensitivity of the model's results to variations in these details, let alone changes in the structural assumptions in the model (for example that changes in the population death rate could be described by multiplying together factors such as DRMM relating to changes in overcrowding, pollution, food supply and material standard of living, all treated independently of each other) were seen as evidence of the fragility of the foundations of The Limits to Growth approach. It is interesting to see also how clearly Cole et al (1973) expressed their concern at what they saw as a replacement of mental models by computer models, and "computer fetishism" (their term) which they define as "[endowing] the computer model with a validity and an independent power which altogether transcends the mental models which are its essential basis." (Cole et al 1973, page 8). From this distance it feels as though their concerns about the use of computational models have been both clearly acknowledged and largely overcome by the huge growth in computational modelling that has influenced every scientific field.

To summarise, the aim of System Dynamics remains to capture the structural interactions between system parts and to express these in mathematical form. As exemplified by the World 2 and World 3 models, parameter values are selected carefully in order to calibrate the model outputs so that they always start in a state as close as possible to the current state of the world; there is a single initial condition that is relevant to the model, and the parameters are fitted so that past trends (for example deducing underlying values to take for birth and death rates by fitting the model to reproduce correctly the growth in world population from 1900 to 1970) are reproduced as well as possible (Forrester, 1973). The outputs of the model are framed as a number of scenarios over a pre-defined future time period (1970 to 2100), in order to generate a range of possible futures which can be discussed in qualitative terms rather than being seen as specific predictions of what will occur.

2.2 An overview of dynamical systems theory

In mathematical terms, a dynamical system is loosely defined as a rule that is applied repeatedly to evolve a system through a set of different states, starting from a known initial state (Strogatz, 2014). The initial state and the evolution rule define the dynamics. Given this generality, the idea of a dynamical system certainly includes every deterministic model described using differential equations (i.e. equations that specify the rate of change in time of the system in terms of the current state of the system). With relatively small adjustments to the definition of dynamical system we can include more complicated elements such stochastic effects, time delays, dependencies on the whole history of the system. The evolution of the system, through the repeated application of an evolution rule, naturally makes one think in terms of the kind of time-discrete update rule as applied in The Limits to Growth, and effectively there are only two choices for how the rule is applied: either one considers time to be discrete, or to be continuous. For computational reasons a continuous update rule must be discretised into a set of small time increments, which is then precisely how The Limits to Growth operates, computationally.

At this point it might seem as though dynamical systems theory is general enough to cover much of applied mathematics. What sets dynamical systems theory apart from other sub-fields is its philosophical focus on questions of long-term trends and behaviour. Dynamical systems as a sub-field asks two broad questions: (i) what are the possible kinds of behaviour that a system settles down to at long times? and (ii) how might those kinds of behaviour change qualitatively as the system parameters are varied. In contrast, initial transient parts of the system behaviour are ignored since the precise form taken by the initial behaviour would depend significantly on the choice of initial state for the system. Moreover, the initial state is often not known precisely, and so the features of the long-term behaviour that are most valuable, and certainly the features that will be most robust, are probably those that are to a large extent independent of the choice of initial state. Hence the initial transient, being closely determined by the initial state, is also ignored.

To respond to the first question: long-time behaviours are often very simple: a relaxation to equilibrium, or to a periodic oscillation in time. More complicated examples include quasiperiodic oscillations in time (where the system returns arbitrarily often to a state that is arbitrarily close to a chosen previous state, but without being exactly periodic), or deterministically chaotic, in which small differences between two initial states grow until the future behaviours are uncorrelated, but both remain bounded and exploring the same collection of future states, in a statistically equivalent fashion (Stewart, 2002). These four kinds of behaviour are qualitatively distinct, in that each has its own distinctive features that enable one to distinguish between them through system measurements that do not depend on the choice of coordinates used to specify the evolution rule or the state of the system.

The way in which the long-time system behaviour changes when system parameters are changed, raised in the second question above, is often known as bifurcation theory. It hinges on being able to tell when two dynamical systems are qualitatively equivalent to each other, and when they are not. Bifurcation theory enables us to build a complete catalogue that describes the typical, or generic, kinds of transition between, for example, the four long-time behaviours listed above: equilibrium, periodic oscillation, quasiperiodic oscillation, and chaos. Whether these transitions are desirable or not depends significantly on the modelling context. For example, in structural engineering the transition between an equilibrium and a periodic oscillation is usually bad: an 'instability' in which a structure begins to sway or buckle. However, in biological systems that rely on periodic oscillation to perform their function, such as electrical activity in the heart, circadian rhythms governing daily sleeping and wakefulness, or even the cell cycle, the maintenance of that oscillation is essential to the proper function of the system, and to life itself.

2.3 Contrasts and comparisons

Both System Dynamics and dynamical systems theory aim to understand the development from initial states of a complex system, described essentially as the solution to a set of nonlinear differential equations, i.e. equations that relate the rates of change of a set of stocks, whose levels define the state of the system, to the inflows and outflows that influence them, and which are in turn modified by the levels themselves. There is a common underlying language and set of concepts.

The central difference is one of purpose. Dynamical systems theory, as a mathematical enterprise, lies firmly within the scope of purpose 3, as described in Table 1, i.e. 'understanding theory', since by taking the view that transients are less important than ultimate states, dynamical systems theory essentially ignores quantitative changes in individual solution trajectories unless they are accompanied by a qualitative change. For example, a simple pendulum comprising a mass at the end of a rigid rod that is fixed at its other end and allowed to swing freely around the pivot, when released from a starting position that is a small deflection to the vertical, oscillates periodically in time around the vertical position. If the angle of release is increased, the oscillation is qualitatively the same; only the amplitude of the motion has increased. This remains true until the starting position of the mass is set to be vertically above the pivot. The vertical position is another equilibrium of the pendulum and so its motion is now qualitatively different (and in this idealised case there is no motion at all!). There is also a subtle distinction between this initial vertical state and an initial state in which the pendulum is started close to vertical but with a small impulse that is just enough for the pendulum to perform one revolution and arrive, in the limit of a very long time, precisely at the vertical position. Dynamical systems theory gives us a precise vocabulary to distinguish between these kinds of behaviour, based in large part on the ultimate fate of system trajectories.

In contrast, System Dynamics is concerned only with a finite time horizon: past behaviour is used for calibration and to optimise appropriate system parameters (and indeed structural modelling assumptions) while future behaviour then provides the sense of "explanation or exploration of future scenarios" that falls into purpose 4 in the CMTF classification shown in Table 1. Indeed, it is entirely possible that, since a System Dynamics model is constructed with a particular time horizon in mind, effects omitted from the model mean that outside this time window it becomes manifestly unreliable (for example quantities that must be positive become unphysically negative). Indeed Forrester noted this in relation to the death rate discussion referred to in Section 2.1 earlier; see Forrester (1973), pages 43-44. This would be unsatisfactory from a dynamical systems viewpoint – the behaviour at long times should be a robust feature of the system - but allowable in System Dynamics since such long time horizons fall outside the scope of the modelling exercise.

An interesting final observation concerns model robustness and genericity, which dynamical systems aims to characterise. Even when the focus of modelling is over a finite time horizon, it is possible that understanding of the ultimate behaviour of solutions is of use since it may point to the kinds of additional effect that might be seen within the finite time window of interest, as parameters are varied. If the dynamical systems viewpoint demonstrates, for example, that oscillations were not possible at any point on future trajectories, then this information would be reassuring that even on finite time windows such behaviour should not emerge even if parameter values are changed significantly.

3. Trends in computational modelling

3.1 Digital Twins

A digital twin refers to a computation model, and usually a detailed one, that is continually updated in the light of new external data arriving from sensors, or other information sources. In this sense, a digital twin is perhaps more concerned with 'nowcasting' rather than 'forecasting': it aims to give a sense of what the current state of the system is rather than any future prediction. The term 'digital twin' implies a high level of detail so such modelling work usually falls into one of the later categories described in Table 1. An obvious physical example is the sophisticated telemetry systems used in Formula 1 racing, in which real-time data from the car is not only recorded for later analysis but used to inform decision making during the race itself.

In the context of earth systems and global modelling such digital twins have progressed furthest in the context of fixed assets and urban infrastructure. As the CMTF report notes, the elaboration of precise information about the construction of complex structures such as oil platforms has paved the way to digital twins that provide detailed representations of the built environment and national (critical) infrastructure systems. Building Information Management (BIM) systems should become increasingly routine ways to provide the backbone of the data on the complete construction of a building and the utilities and services that it uses. Probably the most advanced project of this kind is the *Virtual Singapore* project directed by Singapore's National Research Foundation. To date this project has focussed on compiling 3D mappings of Singapore's infrastructure. This allows new, but static (i.e. not in real time) analyses, such as the modelling of the optimal locations for new solar panels based on building heights and orientations. However it also opens up new dynamical modelling possibilities such as real-time control of vehicular and pedestrian traffic flows around the city. Singapore's centralised governance allows such systems to overcome a number of privacy concerns that might prohibit such plans elsewhere, but the challenges of security, data reliability and standardisation are also considerable.

A natural extension of this kind of project for the future is to use satellite Earth Observation data to extend the coverage of such data gathering, at the expense probably of spatial resolution. That would allow digital twin modelling for Earth, in the sense of real-time data collection and visualisation at a global scale, which in turn might allow the construction of computational models that were continually improved and corrected by the real-time data.

3.2 Artificial Intelligence and Machine Learning

The modern history of Artificial Intelligence is short (70 years) but illuminating. After at least two unsuccessful waves of interest and funding (both ending in disappointment and unfulfilled ambition), it appears as though there is currently a much better alignment of the three fundamental ingredients that computational modelling in the AI sense requires: raw computational power, data availability, and algorithm design. Previous waves of interest have lacked at least one of these. Common to all previous waves and the present time is, in the view of this author, a keen sense of self-belief in many commercial and research organisations and marketing activity that continually sets out ambitions for the field that remain out of reach, and indeed may in some sense be getting further away as we

Machine Learning refers to the algorithmic engines that lie beneath the current wave of AI success. Although many mathematical techniques have been collected together and are routinely referred to together as 'machine learning', one focal point of machine learning lies in the application of neural networks to learn patterns in data and then to make predictions based on that past experience. Immediately from that definition we can see the inevitability, without conscious human intervention, of the maintenance and propagation into the future of existing and historic biases. In its most common form, an artificial neural network (ANN) is optimised to reproduce the results contained in a training dataset, and then tested on another, separate, dataset (the 'train-and-test' paradigm). By designing the neural network as a collection of initially identical nonlinear units one can usually take advantage of various universal approximation theorems in the mathematical literature that guarantee that an arbitrarily large network exposed to arbitrarily large amounts of training data will be able, for example, to classify images by their content (one long-running example being to distinguish images of chihuahuas from images of blueberry muffins, as in the 2017 article by Mariya Yao).

The rapid recent increase in the accuracy of such systems is well known but it is worth quantitatively illustrating in order to reinforce the point that this is such a recent phenomenon. Figures 3 and 4 illustrate the historical improvement in machine learning classification algorithms taking as examples firstly general image classification, and secondly a more specific task concerning the standardised handwritten samples of the digits 0–9 known as the Modified National Institute of Standards and Technology (MNIST) database. As with images, and also with Natural Language Processing and OpenAI's ChatGPT system, the current state of the art is close to, or better than, the error that human beings would typically have with a (narrowly-defined) task.

Figure 3: Machine learning algorithms now regularly exceed human performance (red horizontal dashed line) on the specific task of image classification. Reproduced from the Electronic Frontier Foundation website at https://www.eff.org/ai/metrics#Vision under the Creative Commons Attribution Licence CC-BY-3.0.

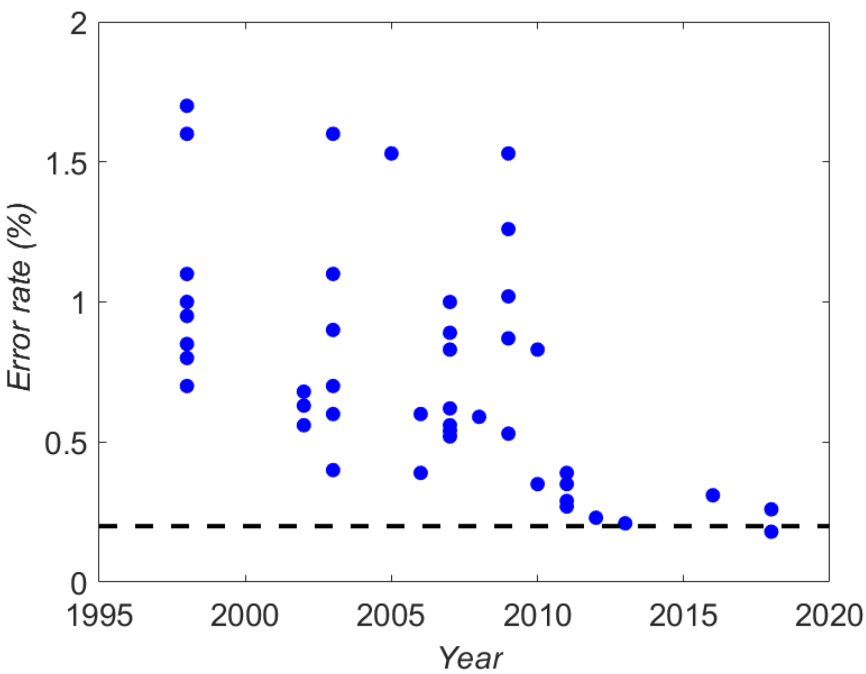

Figure 4: Machine learning algorithms also now regularly exceed human performance on the specific task of recognising digits from the MNIST dataset. The horizontal dashed line indicates typical human-level performance on this task. Plot produced by the author from data available at https://en.wikipedia.org/wiki/MNIST_database .

It is therefore clear that machine learning, as a combination of algorithm development combined with data availability and raw computational power, is finally becoming the technology that futurists were hoping for a generation ago. This sense of optimism must of course be tempered by the challenges that remain related to transparency, regulation, ownership, privacy, bias, and discrimination. The use of AI as 'black boxes' will increasingly be challenged both legally and by public opinion.

Rolnick et al (2022) present a review of uses of AI in generating solutions to climate change challenges across thirteen domains, including data-rich and engineering-intensive areas such as energy systems and transport, but also looking at how AI could contribute to challenges in the social impacts of climate adaption, in promoting individual behaviour change, and in collective decision making. This provides a response to calls, as in the Villani Report (2018) for "innovations in AI [to] be used to optimize energy consumption ... and achieve a better understanding of the effects of human activity on the environment." (Villani, 2018, page 6). Section 4 of the Villani Report also highlights the complex relationship between the current growth in AI, and its expanding footprint in terms of energy consumption and natural resources (particularly rare earth elements used in hardware manufacturing), and environmental and biodiversity crises. AI could provide many solutions to these climate-related problems but the resulting expansion in the global technological footprint could also end up making the underlying problems worse at the same time.

What perhaps is even more uncomfortable is the knowledge that at least some part of the expansion of the technology footprint (data centres, small satellites, growth of internet and mobile phone infrastructure, energy supply and distribution) would need to occur *before* the collection of data that would allow AI-based solutions to assist with averting environmental problems. This perhaps points to the usefulness, where possible, of less resource-intensive modelling based on a priori physical and biochemical principles in order to identify how to bring data collection and data-intensive AI methods to bear most effectively. That might point to synergies between physically-based and data-driven approaches in computational modelling.

4. Discussion and conclusions

This short essay has sought to provide a brief summary of modern themes in computational modelling in its widest sense. To conclude we discuss six possible ways in which these various themes might contribute to contemporary successors to The Limits to Growth, in particular the Earth for All project and the Earth4All model inspired by, and the modern successor to, the World 3 model (Dixson-Declève et al, 2022). It is not claimed that incorporating any of these potential contributions into Earth4All is either easy, or indeed actually fundamental to the continued success of the Earth4All programme. They are proposed as contributions to the future discussion and development of 'world modelling' in the spirit started by The Limits to Growth, as challenges for the future development of these and similar models.

1. Participatory modelling.

One enormous advantage of the school of thought represented by practitioners of System Dynamics in general is a participatory approach to modelling. That is, to involve from the very beginning a group of people who work within the system under study and who can contribute directly their experiences into the modelling process. This has clearly been beneficial particularly for Systems Dynamics approaches to complex issues where many different kinds of factor are at play. An example of this is the study of the causes of obesity in the UK (UK Government Office for Science, 2007) which commissioned an evidence base from experts across disciplines that included food production and consumption, social and individual psychology, physical activity and the environment, and physiology and then tested and compared these initial responses in a series of workshops.

A similar theme of participation from stakeholders is clear in the increasing use of Citizens' Assemblies to discuss and build confidence in national proposals to tackle climate change (UK House of Commons, 2020).

A further sense in which modelling is a participatory process is the sense of it being a scientific method open to critique from the community. The current drive towards greater transparency in science, and open-access publication, links to a need in System Dynamics as in all other computational modelling to be precise about the details of the model and its implementation.

2. Regional disaggregation and international trade.

A regionalised version of the Earth4All model divided the world into ten regions each driven by the same system dynamics structure but allowing for variations in parametrisations (for example the look-up tables and multipliers discussed in section 2.1 above). This has led to a greater understanding of the similarities in development across these different regions, as presented by Collste et al (2020). What is less clear is the role that international trade plays in such a regionalised model. Substitution effects between labour, or goods produced in different regions may allow inequalities between regions to grow even if overall global trends are positive. Adding a model for international trade does of course substantially increase complexity in the modelling but it might enable the further testing of the model assumptions and again generate additional future scenarios for discussion.

3. The role of transnational corporations and the strength of international law.
As Picciotto (1997) notes " *The main problem for international law is the extremely weak nature of the international political system*". It follows that international law is limited by the ability of nation states to agree on common objectives and approaches, especially to questions concerning the global commons which are central to this kind of world modelling. It therefore feels essential in models such as Earth4All to attempt to take account of the degree of international consensus and the strength of international frameworks, as compared to the private interests of transnational corporations and the extent to which global resources may become reserved to private interests.

Capturing these concepts in a quantitative model is of course much more demanding than physical elements of modelling (e.g. population sizes, energy usage) but together with point 2 above it may help to indicate other possible future scenarios based on higher or lower future degrees of international agreement, and therefore bring new dimensions to the modelling challenge.

4. A digital twin approach.
The 'digital twin' idea, that is, to make an annual comparison of global trends (as reported for example by UN or World Bank statistical agencies) with model predictions and subsequently to adjust the model, is natural and appealing. The disadvantage perhaps is that each adjustment to the model effectively invalidates previous predictions. However, by tracking the parameter changes that are made on each occasion one might learn valuable information about which structural features of the model were most robust, and which most in need of future clarification. The model will always 'be wrong' but identifying where more attention should be focussed would again be valuable in the development of further modelling approaches in the future.

5. Technological tipping points.
One exogeneous input into the Earth4All model is a general level of technological advance of around 1% per year. As well as the sensitivity of the model to this rate of improvement, one should also explore the sensitivity to more abrupt increases in technological advance. Figures 3 and 4 illustrate that, since there is a threshold that machine performance needs to meet in order to be able to replace human performance, there is an abrupt change in the sense that there is a specific time at which machines overtake humans at the task. The same is true in other narrow AI tasks such as playing chess or Go. So also with combinations of algorithms and hardware, as in automation in manufacturing or warehouse operations. Hence, at heart, technological advances are often disruptive, although their diffusion into markets does provide a mechanism for smoothing their influences, as captured by

Everett Rogers' "Diffusion of innovation" theory in 1962 (Rogers, 2003) and its mathematization by Bass (1969). But it would be useful to explore, as future scenarios, the effects of much more rapid changes, perhaps to one sector alone, of much more substantial technological advances. This would also provide incentives perhaps for directing further publically-funded R&D into different sectors, for example food supply.

6. Uncertainty quantification.
Linked to the previous point but of a more general nature is the general sense of the level of uncertainty in the model. Being able to talk about ranges of outcomes may be more useful in some policy formulation situations than single point estimates for the future. Although Earth4All is of course founded on the need to formulate contrasting scenarios of the future, some sense of the general level of uncertainty contained within the 'Too Little Too Late' and 'Giant Leap' scenarios would be valuable, both in terms of validating the robustness of the modelling and in policy debates. It aligns with the sense elaborated on earlier that understanding the limitations of a model is a necessary part of understanding where it succeeds.

To conclude, the six directions above, and perhaps the overarching theme of this article, is advocacy for the wider discussion and analysis of world models, and for the participation of different parts of the academic community in their construction and analysis. This use of a number of separate research groups and then a collation of their results, lies behind some of the most successful global modelling projects of recent years such as the IPCC Assessment Reports. Here a set of General Circulation Models that clearly differ in their underlying parametrisations and approximations are combined in order to test our understanding of complex climate phenomena. The unanimity of their results provides compelling conclusions that are needed, as we have seen in successive climate COPs, to drive the necessary political agreement and progress towards a better world for all.

Bibliography

BASS, F. (1969). «A new product growth for model consumer durables ». *Management Science* **15**(5), 215–227.

COLE H.S.D., FREEMAN C., JAHODA M. and PAVITT K.LR. (1973) (Eds), *Thinking About the Future: A Critique of The Limits to Growth.* Sussex University Press.

COLLSTE D, CORNELL S.E., RANDERS J, ROCKSTRŐM J, STOKNES P.E. (2021) « Human well-being in the Anthropocene: limits to growth ». *Global Sustainability* **4**, e30, 1–9.

DIXSON-DECLÈVE et al (2022), *Earth For All: A Survival Guide for Humanity*. New Society Publishers.

FORRESTER J.W. (1973), *World Dynamics*. Second Edition. Wright-Allen Press, Cambridge, MA.

MEADOWS D.H. (2008), *Thinking in Systems.* Chelsea Green.

PICCIOTTO S. (1997), « International law: the legitimation of power in world affairs », In P. Ireland and P. Laleng (Eds.), *The Critical Lawyers' Handbook 2*, pp 13-29. Pluto Press.

ROGERS, E. (2003) *Diffusion of Innovations*. Fifth Edition. Simon and Schuster.

ROLNICK D. et al (2022), « Tackling climate change with machine learning », *ACM Comput. Surv.* **55**(2) Article 42.

SINGAPORE NATIONAL RESEARCH FOUNDATION, *Virtual Singapore* project. https://www.nrf.gov.sg/programmes/virtual-singapore. Accessed on 24 January 2023.

STEWART I. (2002), *Does God Play Dice?: The New Mathematics of Chaos*. Second Edition. Wiley-Blackwell.

STROGATZ S.H. (2014), *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering*. Second Edition. Westview Press.

UK GOVERNMENT OFFICE FOR SCIENCE (2007). *Tackling Obesities – Project Report.* (Second Edition). UK Government Department of Innovation, Universities, and Skills. Reference: DIUS/PUB 8654/2K/12/07/AR.

UK GOVERNMENT OFFICE FOR SCIENCE (2018). Report: *Computational Modelling: Technological Futures.* UK Government.

UK HOUSE OF COMMONS (2020). *The path to net zero: Climate Assembly UK report*. House of Commons. https://www.climateassembly.uk/. Accessed on 28 January 2023.

VILLANI C. (2018). « *Donner un Sens à L'Intelligence Articificielle: Pour une Stratégie Nationale et Européene » (For a Meaningful Artificial Intelligence: Towards a French and European Strategy).* Conseil national du numérique (French Digital Council). https://www.aiforhumanity.fr/pdfs/MissionVillani_Report_ENG-VF.pdf

YAO M. (2017), ««*Chihuahua or Muffin? Searching For The Best Computer Vision API.* » https://www.topbots.com/chihuahua-muffin-searching-best-computer-vision-api/. Accessed 23 January 2023.