



Cite this article: Succi S, Coveney PV. 2019

Big data: the end of the scientific method?

Phil. Trans. R. Soc. A **377**: 20180145.

<http://dx.doi.org/10.1098/rsta.2018.0145>

Accepted: 25 July 2018

One contribution of 11 to a theme issue
'Multiscale modelling, simulation and
computing: from the desktop to the exascale'.

Subject Areas:

computer modelling and simulation,
computational physics, artificial intelligence

Keywords:

Big data, multiscale modelling, simulation,
artificial intelligence

Author for correspondence:

Peter V. Coveney

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

Big data: the end of the scientific method?

Sauro Succi^{1,2} and Peter V. Coveney^{3,4}

¹Center for Life Nano Sciences at La Sapienza, Istituto Italiano di Tecnologia, viale R. Margherita, 265, 00161, Roma, Italy

²Institute for Applied Computational Science, J. Paulson School of Engineering and Applied Sciences, Harvard University, 29 Oxford Street, Cambridge, USA

³Centre for Computational Science, Department of Chemistry, University College London, London, UK

⁴Yale University, New Haven, USA

 PVC, 0000-0002-8787-7256

For it is not the abundance of knowledge, but the interior feeling and taste of things, which is accustomed to satisfy the desire of the soul.
(Saint Ignatius of Loyola).

We argue that the boldest claims of big data (BD) are in need of revision and toning-down, in view of a few basic lessons learned from the science of complex systems. We point out that, once the most extravagant claims of BD are properly discarded, a synergistic merging of BD with big theory offers considerable potential to spawn a new scientific paradigm capable of overcoming some of the major barriers confronted by the modern scientific method originating with Galileo. These obstacles are due to the presence of nonlinearity, non-locality and hyperdimensions which one encounters frequently in multi-scale modelling of complex systems.

This article is part of the theme issue 'Multiscale modelling, simulation and computing: from the desktop to the exascale'.

1. Introduction

Our current society is characterized by an unprecedented ability to produce and store breathtaking amounts of data and, much more importantly, by the ability to navigate across them in such a way as to distil from them useful information, hence knowledge. This has

now reached the point of spawning a separate discipline, so-called big data (BD), which has taken the scientific and business domains by storm. Like all technological revolutions, the import of BD goes far beyond the scientific realm, reaching down into deep philosophical and epistemological questions, not to mention societal ones. One of the most relevant is: Are we facing a new epoch in which the power of data renders obsolete the use of the scientific method as we have known it since Galileo? That is, insight gained through a self-reinforcing loop between experimental data and theoretical analysis, based on the use of mathematics and modelling?

For if, as sometimes appears true today, anything can be inferred by detecting patterns within huge databases, what is the point of modelling anymore? This extreme stance is summarized in Anderson's provocative statement: 'With enough data, the numbers speak for themselves, correlation replaces causation, and science can advance even without coherent models or unified theories'. In a nutshell, it is a data-driven version of Archimedes' fulcrum: give me enough data and I shall move the world. As radical as this new empiricism is, it brings up an intriguing point: is understanding overrated? Could it be that smart algorithmic searching through oceans of data can spare us the labour (and the joys) of learning how the world works [1]?

'Why learn, if you can look it up?' is another way of articulating the same idea. At least among intellectuals, the retort is that of C.S. Lewis, 'Once you have surrendered your brain, you've surrendered your life' (paraphrased) [2]. In the sequel, we shall offer rational arguments in support of this instinctive reaction whilst recognizing the perspectives opened up by BD approaches.

2. Why is big data so sexy?

BD flourishes upon four main observations, namely

- (i) The explosive growth of data production/acquisition/navigation capabilities.
- (ii) Reading off patterns from complex datasets through smart search algorithms may be faster and more revealing than modelling the underlying behaviour, i.e. using theory.
- (iii) It applies to *any* discipline, including those traditionally not deemed suitable for mathematical treatment, including life sciences (another way of putting this is to suggest that these domains are too complex to be modelled).
- (iv) Its involvement in immediate application to business and politics, 'opinion dynamics', 'sentiment analysis' and so on, furnishes another set of domains which raise many ethical questions.

While the four points above hold disruptive potential for science and society, in the following we shall illustrate how and why, based on basic findings within the modern science of complexity, all of them may lead to false expectations and, at their nadir, even to dangerous social, economical and political manipulation.

The four points we shall make in response are the following:

- (i) Complex systems are strongly correlated, hence they do not (generally) obey Gaussian statistics.
- (ii) No data are big enough for systems with strong sensitivity to data inaccuracies.
- (iii) Correlation does not imply causation, the link between the two becoming exponentially fainter at increasing data size.
- (iv) In a finite-capacity world, too much data is just as bad as no data.

Far from being exceptional, our four assertions apply to most complex systems of relevance to modern science and society, such as far from equilibrium nonlinear physics, finance, wealth distribution and many social phenomena as well. So, there can be no excuse for ignoring them.

3. Complex systems do not (generally) obey Gaussian statistics

BD radicalism draws heavily upon a fairly general fact of life: the Law of Large Numbers, the main content of which is that, with enough samples, call it N , errors (uncertainty) are bound to surrender to certainty. The most famous aspect of this fact is the square-root law of the noise/signal ratio:

$$\frac{\sigma}{m} \propto \frac{1}{\sqrt{N}}, \quad (3.1)$$

where m is the mean value and σ its root-mean-square associated with a given stochastic process X . In other words, let $m_N = (x_1 + \dots + x_N)/N$ be the mean value of a given quantity X as obtained from N measurements, i.e. the data. It is well known that m_N approaches the correct mean, m , in the limit of $N \rightarrow \infty$. Even better, one can estimate how fast such convergence is attained by inspecting the mean square departure from the mean, also known as the variance, namely

$$\sigma_N^2 = \frac{1}{N} \sum_{i=1}^N (x_i - m_N)^2. \quad (3.2)$$

Under fairly general assumptions, it can be shown that the root-mean-square (rms) departure from the mean decays like $1/\sqrt{N}$. With enough measurements, uncertainty surrenders: this is the triumph of BD [3].

Now let us ask ourselves: what are the ‘general assumptions’ we alluded to above? The answer is that the variables x_i must: (i) be uncorrelated, i.e. each outcome x_i is independent of the previous one and does not affect the next one either, (ii) exhibit a finite variance. As we shall see, neither of the two should be taken for granted.

With these two premises, the central limit theorem pertaining to the Law of Large Numbers shows that the sum X_N obeys Gaussian statistics, i.e. a bell-shaped curve

$$p_G(y) = \frac{1}{2\pi\sigma} e^{-y^2/2}, \quad (3.3)$$

where $y = (x - m)/\sigma$ is the normalized (de-trended) version of x (here x stands for any generic stochastic variable).

The Gaussian distribution exhibits many important properties, but here we shall focus on the following one: *Outliers stand very poor chances of manifesting themselves and, precisely because they are the carriers of uncertainty, uncertainty is heavily suppressed.* Then the numbers indeed speak for themselves: the probability of finding an event one-sigma away from the mean is about 30%, a number which goes down to just 4.5% at two-sigma. The demise of uncertainty is dramatic: at five-sigma we find just about one in half a million and at six-sigma less than two in a billion! This adumbrates a very comfortable world, where uncertainty has no chances because outliers are heavily suppressed; fluctuations recede and are absorbed within the mean, an overly powerful attractor. A comfortable, if somewhat grey, world of stable and reassuring conformity.

The Gaussian distribution plays an undeniable role across all walks of science and society, to the point of still being regarded by many as a universal descriptor of uncertainty. The truth is that, for all its monumental importance, the Gaussian distribution is far from being universal. In fact, it fails to describe most phenomena where complexity holds sway.

Why? Because complex systems, almost by definition, are correlated! When a turbulent whiff is ejected from the wind-shield of our car, it affects the surrounding air flow, so that the next whiff will meet with an environment which is not the same it would have met in the absence of the previous whiff. The system affects its environment, the two are correlated, the statistics of whiffs (turbulence) is not Gaussian. A similar story goes for most complex states of flowing matter [4] and complex systems in general.

This is a far cry from the ‘fair coin’, in which head or tail now has no effect on head or tail at the next toss. In complex systems, the coin is hardly fair. As a result, the statistics of correlated events is much more tolerant towards outliers, with the consequence of a much higher, sometimes even

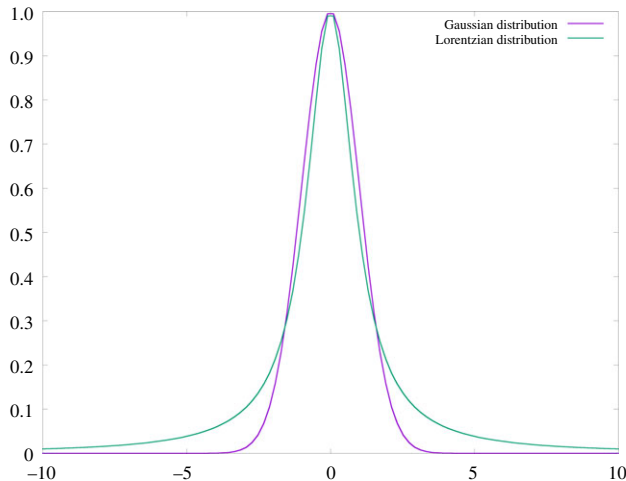


Figure 1. A comparison of Gaussian and Lorentzian distributions. Note the persistence in the latter distribution of far larger events from the mean. To emphasize the different decay away from the peak (most probable) value, such peak is taken to be the same for both distributions, which means that they are not normalized to the same value. (Online version in colour.)

unbounded, variance. In a nutshell, it is a world much more full of (good and bad) surprises, just as is real life!

The prototypical example is the Lorentz distribution:

$$p_L(x) = \frac{a/\pi}{a^2 + (x - m)^2}. \quad (3.4)$$

For inliers, $|x - m| \ll a$, this is virtually indistinguishable from a Gaussian, and all is fine and well. But for events far in excess of a , outliers or *rare events* in the following, the difference is dramatic: the Lorentz distribution decays much more slowly than the Gaussian, which is the reason why its variance is formally infinite (figure 1). As an example, if human height measurements were distributed according to Lorentz, with an average $m = 1.75$ (metres) and $a = 10$ cm, the probability of finding a human 2.75 m high, i.e. at 10σ , would be of the order of three per cent. The decay is so slow as to sound ridiculous: going to 20σ just halves the number, which means that walking along the street of any city in the world, humans 3.75 m high would be commonplace!¹ Ridiculous as it seems for human heights (which are indeed Gaussian-distributed), such slow decay of outliers is a regular occurrence in complex systems, be they natural, financial or social. The dire consequences of treating the financial world as a Gaussian-regulated one are most compellingly (and often hilariously) discussed in Taleb's books *Fooled by randomness* [5] and *The black swan* [6].

Infinite variance is a bit far-fetched, since in the real-world signals and measurements are necessarily finite, but the message comes across loud and clear: the mean and the variance are no longer sufficient to capture the statistical nature of the phenomenon. For this purpose, higher-order *moments* must be inspected.

By moments, we mean sums (integrals in continuum space, with due care) of the form

$$M_q = \sum_{i=1}^N x_i^q, \quad (3.5)$$

where q is usually (but not necessarily) a positive integer. By normalizing with $M_0 = N$, it is clear that M_1 is the mean and $M_2 - M_1^2$ is the variance. For a Gaussian distribution, this is all

¹For the detail-thirsty, this probability is computed by means of the cumulative distribution of the Lorentz function, namely $P_L(x) = 1/2 + 1/\pi \arctan(x - m/a)$.

we need to know, because all higher-order moments with $q > 2$ follow directly from these two. But for a generic distribution, this is no longer the case and more moments need to be specified in particular, from the very definition, it is readily appreciated that large q 's give increasing weight to large x 's, i.e. the aforementioned rare events play an increasing role. Thus, inspection of M_q with $q > 2$ is paramount to the understanding of complex processes, an utterly non-Gaussian world trailblazed by the turbulence community and now widespread in most walks of the science of complexity. The bottomline is that, in the presence of correlations, Gaussian statistics no longer hold, and uncertainty does not give in so easily under data pressure, in that the convergence to zero uncertainty is much slower than the inverse square-root law. For a moment of order q , it is likely to be $N^{-1/q}$, which is a nearly flat in practice for $q \gg 2$. For instance, with $q = 8$, cutting down uncertainty by a factor of 2 takes $2^8 = 256$ times more data.

This explains why the BD trumpets should be toned down: when rare events are not so rare, convergence rates can be frustratingly slow even in the face of petabytes of data.

(a) To gauss or not to gauss: nonlinear correlations

It is natural to ask if there is a qualitative criterion to predict whether a given system would or would not obey Gaussian statistics. While the present authors are not aware of any rigorous 'proof' in this direction, robust heuristics are certainly available. We have mentioned before that the Law of Large Numbers rests on the assumption that the sequence of stochastic events be uncorrelated, that is, the occurrence of a given realization does not depend on the previous occurrences and does not affect the subsequent ones. This is obviously an idealization, but one which works eminently well, as long as the system in point can be treated as isolated from its environment and not subject to any form of nonlinearity. Yet, by definition, most complex systems do interact with their environment and they affect it in various ways. Since the environment couples back to the system, it is clear that self-reinforcing or self-destroying loops get set up in the process. Self-reinforcing loops imply that a given occurrence affects the environment in such a way as to make such an occurrence more likely to happen again in the future. This is the basic mechanism giving rise to persistent correlations, the unfair coin we alluded to earlier on in this paper. And persistent correlations are a commonplace in most complex systems, be they natural, financial, political, psychological or social.

4. Sensitivity to data inaccuracies

The main goal of BD is to extract patterns from data, i.e. to unveil correlations between apparently disconnected phenomena. Given two processes, say $X = \{x_i\}$ and $Y = \{y_i\}$, $i = 1, \dots, N$, the standard measure of their correlation, is the *covariance*, defined as

$$C(X, Y) = \frac{\sum_{i=1}^N x_i y_i}{\sigma_x \sigma_y}, \quad (4.1)$$

where the sequences x_i and y_i are assumed to be de-trended, i.e. of zero-mean. The most perfect correlation is $Y = X$, which delivers $C = 1$, its opposite being perfect anti-correlation, $Y = -X$, yielding $C = -1$. Also interesting is the case of plain indifference, zero-correlation $C = 0$, which means that sets of positively and negatively correlated events are in perfect balance.

In geometrical language, $C = 0$ implies that the two N -dimensional vectors X and Y are orthogonal. Adhering to this language, the correlation C can be thought of as the cosine of the angle between the vectors X and Y , in which case we write $C = (X, Y) / (\sigma_x \sigma_y)$, where (\cdot) denotes the scalar product in Euclidean space. Needless to say, under Gaussian statistics, correlation coefficients converge to the 'exact' values in the limit of infinite N .

But again, this is not necessarily the case if X and Y stem from complex processes. Moreover, quadratic correlators, such as the covariance in equation (4.1), which originate directly from the notion of Euclidean distance $d(X, Y) = \sum_i (y_i - x_i)^2$, may not be adequate to capture the complex nature of the phenomena, just as mean and variance are by no means the full story in the presence

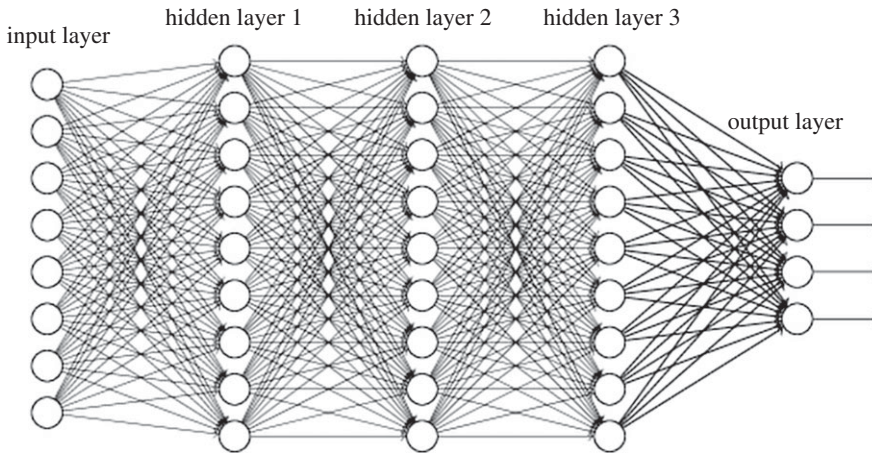


Figure 2. Sketch of a deep neural network. Note the presence of multiple hidden layers within which specific processing activity takes place, between the input and output layers of processing elements, called neurons.

of rare events! In particular, higher order ‘distances’, possibly not even Euclidean ones, should be inspected, their resilience to data pressure being inevitably much higher than the one offered by the variance. Once again, error convergence might be a very slow function of data size.

A similar argument goes for more sophisticated forms of learning, such as the currently all-popular *machine learning*. Here, a neural net is trained to recognize patterns within a given set of data, by adjusting the weights of the connections in such a way as to minimize a given error functional (the cost function in machine-learning jargon). Given a set of input data x_i , the neural net produces a corresponding output y_i of the form

$$y_i = f \left(\sum_j W_{ij} x_j \right),$$

where W_{ij} is the connecting weight between nodes i and j belonging to two subsequent layers of the network and $f()$ a suitable transfer function, typically a sigmoid or variations thereof (figure 2).

The output signal is then compared to the target data Y_i to form the loss (error) function

$$E\{W\} = \sum_{i=1}^N d(y_i, Y_i),$$

where $d(x, y)$ is some metric distance in data space, usually, but not necessarily, the standard Euclidean one.

The weights are then updated according to some dynamic minimization schedule, so as to achieve the minimum error. It is then clear that if the functional $E\{W\}$ is smooth, the search is easy and robust against data inaccuracies. If, on the other hand, the error landscape is corrugated, the expected case for complex systems in which higher order moments carry most of the relevant information, even small inaccuracies can result in the wrong set of weights, where wrong means that such weights are likely to fail when applied to new data beyond those they had been trained for (figure 3).

Of course, failsafe scenarios also exist, whereby different sets of weights work even though they differ considerably from each other, because all local minima are basically equivalent quasi-solutions of the optimization problem. We suspect, without proof, that such a form of benevolence lies at the basis of the most remarkable successes of machine-learning and modern artificial intelligence applications.

But this cannot be assumed to be the universal rule: a big red light is there in general, the name of the game being ‘overfitting’, that is a stiff solution exists which reproduces very well a given set

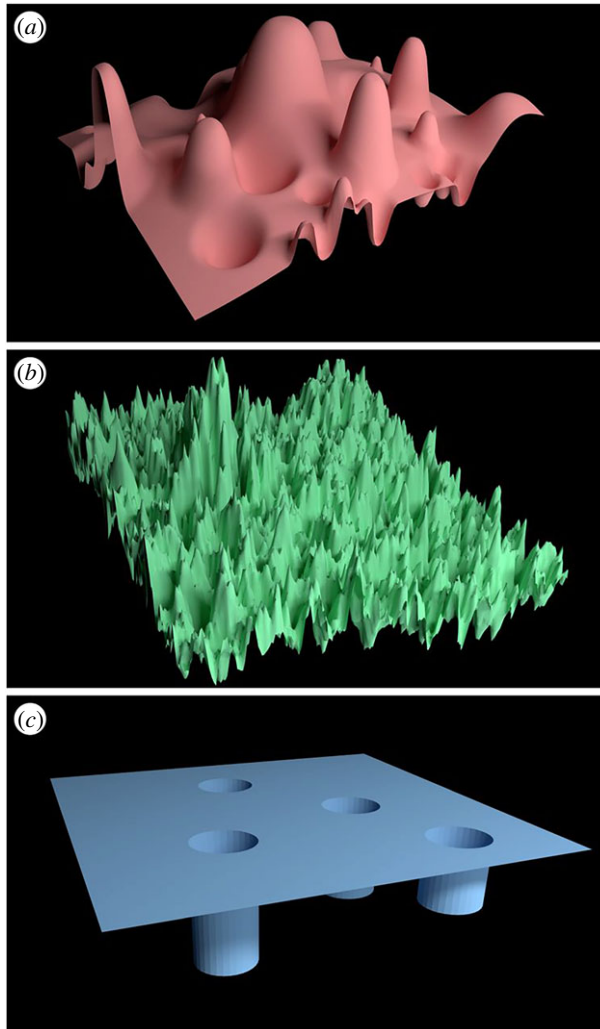


Figure 3. Examples of three classes of landscapes that may be encountered in connecting input to output variables, shown in three dimensions but typically, in fact, arising in much higher dimensions which cannot be drawn. (a) A relatively smooth (i.e. continuous) landscape over which machine-learning algorithms might be expected to perform well; (b) a fractal landscape which is not differentiable and contains structure on all length scales; (c) another discontinuous landscape with no gradients. Neither (b,c) would be expected to perform reliably in the context of machine-learning algorithms. Reproduced from Coveney *et al.* [7].

of data, but fails grossly as soon as the dataset is enlarged, if only slightly. Although well known to the machine-learning community, these problems are typically swept under the carpet by the most ardent BD aficionados.

5. The two distant sisters: correlation and causation

The fact that correlation does not imply causation is such a well-known topic that we only mention it for completeness.

It is indeed well recognized that even if two signals manage to register a very high correlation coefficient C (close to 1), this does not necessarily imply that they are mechanistically related. They may be false correlations (FC), as opposed to true correlations (TC), the latter signalling a

true causal connection. The matter lends itself to hilarious observations: the rate of drowning by falling in a pool appears tightly correlated with Nicolas Cage's movies: unless one assumes that Cage's movies are so badly received as to induce some to drown themselves, there is little question that this is an FC. This case is trivial, but the general problem is not: distinguishing between TC's and FC's is an art, as the problem is both hard and important.

The embarrassing fact is that FCs grow much more rapidly with size of dataset under investigation than the true ones (the nuggets). As recently proven by Calude & Longo [8] the TC/FC ratio is a very steeply decreasing function of data size. Meng, on the other hand, has shown that to be able to make statistically reliable inferences one needs to have access to a very substantial (i.e. greater than 50%) fraction of the data on which to perform one's machine learning [9].

Once again, how big is big enough to make reliable machine learning prediction, remains a very open question. To be sure, we are a very far cry from the comfortable inverse square root law of Gaussian statistics. What is clearly required in the field of BD and machine learning is many more theorems that reliably specify the domain of validity of the methods and the amounts of data to produce statistically reliable conclusions. One recent paper that sets out the way forward is by Karbalayghareh *et al.* [10]

6. Life in a finite world: too much data is like no data

Finally, there is an argument that hinges directly on epistemology and society, since it has to do with a most prized human attribute: wisdom, the ability to take the right decision. Wisdom is often represented as the top level of a pyramid of four, the DIKW (data-information-knowledge-wisdom) chain, the one enabling us to take well-informed decisions. From data we extract information, from information we extract knowledge and finally from knowledge we distil the ultimate goal: Wisdom, the ability to do the 'right thing' (figure 4).

BD-driven decision theory is obviously of paramount importance to science, business and society, as it is to each of our private lives. But, as a matter of fact, the 'constitutive relation' between Data and Information, Information versus Knowledge and Knowledge versus Wisdom is not well known, to put it mildly. In the following, we shall argue that the pyramid representation is deceptive, for it conveys the idea that the layers stand in a simple linear relationship to one another, which is by no means the case. More importantly, it suggests that by expanding the basis (data) all upper-lying layers will expand accordingly, whence the mantra: more data, more wisdom.

This flies in the face of a very general fact of life: sooner or later, all finite systems hit their ceiling, the technical name of the game being *nonlinear saturation*, another well-known concept in the science of complex systems. This is the very general competition-driven phenomenon by which increasing data supply leads to saturation and sometimes even loss of information; adding further data actually destroys information.

But let us discuss saturation first. A well-known example of nonlinear saturation is logistic growth in population dynamics. Let x be the number of individuals of a given species which reproduce at a rate $a > 0$, say a births per year per individual. In differential terms $dx/dt = ax$, leading to untamed exponential growth. But obviously in a finite environment, with finite space and a finite amount of food, such untamed growth cannot last forever for the environmental finiteness will necessarily generate competition, hence a depletion term. Assuming competition is only between two individuals at a time, this results in the famous logistic equation

$$\frac{dx}{dt} = ax - bx^2, \quad (6.1)$$

where b measures the strength of competition. The right-hand side is the epitome of what we mean, the rate of change growing linearly with x but beyond a certain threshold, $x^* = a/2b$, decreasing until it comes to a halt at $x = a/b$. By then the population stops growing, and the number of individuals left at that point is $x = a/b$, also known as the *capacity* of the system. It is

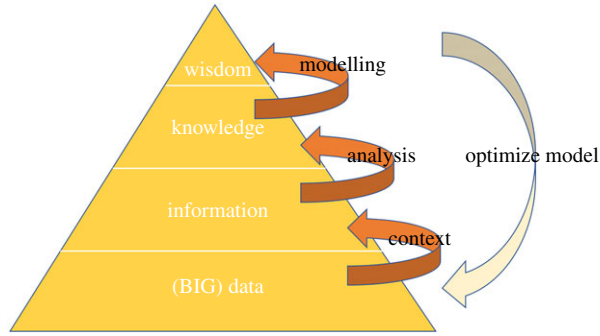


Figure 4. A depiction of the DIKW pyramid to show the cooperation between big data and modelling. It displays how ‘data’, when put in context, leads to ‘information’; analysing the ‘information’ yields ‘knowledge’; ‘knowledge’ gained can be deeply understood by hypothesizing a model for its underlying cause leading to ‘wisdom’ which can be used to optimize the model by repeating the process.

readily seen that the capacity goes inversely with the competition rate: the fiercer the competitors, the higher their needs and the smaller their number. As expected, big consumers present a threat, as is well known to those not driving four-by-four vehicles in the crowded streets of Rome and London.

The right-hand side of (6.1) can be cast in a more informative form as follows:

$$R(x) = bx(c - x), \quad (6.2)$$

where R indicates the effective rate and $c = a/b$ is the capacity.

This shows a nice symmetry (duality) between the population x and the co-population $\bar{x} = c - x$, namely the gap between the actual number of individuals and the system’s capacity. Such symmetry is further exposed by writing the equations in dual form

$$\frac{dx}{dt} = bx\bar{x}, \quad (6.3)$$

$$\frac{d\bar{x}}{dt} = -bx\bar{x}. \quad (6.4)$$

Writing the equation in this manner highlights the dual process of generating population (matter) and annihilating co-population (co-matter). In passing, we note that the above system is invariant under the exchange $x \leftrightarrow \bar{x}$, in combination with time inversion $t \rightarrow -t$, which means that the backward-time evolution of the co-population is the same as the forward-time evolution of the population. Such types of dual relations are typical of finite-size systems hosting nonlinear cooperative/competitive interactions, in the generalized form

$$R(x) = bx^\alpha (c - x)^\beta,$$

where the exponents α and β , as well as the coefficients, may change depending on the specific phenomenon at hand. But the dual structure remains because it reflects the existence of a finite capacity. The above example refers to population growth in time, which is not necessarily related to the information growth with data. Nevertheless, it is our everyday experience that, beyond a certain threshold, further data does not add any information, simply because additional data contain less and less new information, and ultimately no new information at all.

This is quite common in complex systems: for instance, the number of degrees of freedom of a turbulent flow (Information) grows like $R^{9/4}$, where R is the Reynolds number, a dimensionless group measuring the strength of nonlinearity of fluid equations, whereas the volume of space (data) hosting the turbulent flow grows like R^3 (because R scales like the linear size of the volume). Hence the *information density*, i.e. the physical information per unit volume, scales like $I/V = R^{9/4-3} = R^{-3/4}$ a very steep decay at increasing Reynolds number. Given that R of the

order of a million or more is a commonplace in real-life, it is clear that adding volume provides increasingly less return on investment in terms of gain of physical information. We speculate, without proof, that this a general rule in the natural world.

Let us now come to the worst-case scenario: data which destroy information. Eventually, additional data may even *contradict* previous data, perhaps because of inaccuracy but more devious scenarios are not hard to imagine, thereby *destroying* information, because the new and the old data annihilate each other. In the latter scenario, information gain turns into information loss: seeing too much starts to be like not seeing enough, to borrow from C.S. Lewis again.

We argue that such a dual trend applies to the DIKW chain as well, at least to the two lower layers: *too much data is just like no data at all*. In fact, this is possibly still more general: in a finite world, close to capacity, competitive interactions arise which either annihilate the return on investment (information per data unit) or even make it negative, thereby destroying information and productivity, over-communication being a well-known example in point.²

Of course one can argue that, in actual practice, this depends on where the capacity is, so that BD can move it upwards and shift the problem away. We maintain that such threshold shifting without due insight is purely chimeric. Unless we can apprehend the logical structures underlying any given phenomenon, we may just keep generating data conflicts that data accumulation alone will not be able to resolve: in fact, quite the opposite.

In science, we strive to go from data-starved to data-rich, yet a blind data-driven procedure, as often advocated by the most enthusiastic BD neophytes, may well take us from data-rich to data-buried science, unless a just dose of theoretical reasoning is used as an antidote [11].

More importantly, we all live in a finite world, and sooner or later its finite capacity is going to be noticed. Even though such a basic reality is a taboo for many commercially inspired promoters of BD, we had better prepare for that. Aggressive BD distracts attention from our limits, which is pretty dangerous for society, long before it is for science. Indeed, whatever philosophical stance one may adopt, it is clear that treating our resources as if they were unlimited is a sure way to disasters of various sorts, be they environmental, financial or social.

7. Knowledge for business: big data and big lies

Good business is beneficial, and so is any healthy driver of economic wealth. Hailing the pursuit of knowledge and innovation as a Trojan horse for business, is not.

BD has a potentially enormous bearing on society, as is reflected by Pentland's recent book *Social physics* [12], the discipline which endeavours to explain, predict and *influence* human behaviour (for good) based on physical-mathematical principles, treating individuals as 'thinking molecules'. The concept is technically appealing, although one clearly walking on a very thin tightrope between good-willed science and social manipulation. In principle, one can think of society as the 'material made up of thinking molecules' and ask how to best design such 'material', so as to optimize moral values while being bereft of cheating, with flourishing economies, social equity and so on along a rosy-carpeted avenue. In fact, such models provide scientific underpinning to Bauman's illuminating metaphor of a liquid society as one where the rules change faster than most individuals can adjust to, leaving the majority behind [13]. Influencing human behaviour through 'healthy' social pressure and flow of ideas is a noble goal, albeit one which walks on a high wire, the name of the (bad) game being plain manipulation for profit. More precisely, what is hailed as pursuit of knowledge and innovation is, in fact, a very different goal, named *zero sales resistance*.

Zero sales resistance is an obvious goal for the most rapacious forms of capitalism, spinning around the 'money for money's sake' paradigm, instead of more enlightened forms of capitalism, in which money is the just and natural follow-on from healthy innovation, filling genuine societal gaps.

²Incidentally, the reader should appreciate that away from capacity, the logistic equation (6.2) equation is linear because when $x \ll c$, then $c - x \sim c$.

The damage done by BD brainware at the service of rapacious capitalism is all too evident: acquisition of private data in return for the dream of ‘celebrity for everyone’, is a super-clever strategy, as it hits at the very roots of human weakness.

Pentland, being well aware of the danger, invokes a new deal on data. But best intentions are easily fooled; so, here again C.S. Lewis appears apposite: ‘When man proclaims conquest of power of nature, what it really means is conquest of power of *some* men over other men’. When social media hit at human weaknesses, such as the desperate need for fame through a growing list of ‘followers’, collecting money for the disaster brought about by a tsunami, good as it is, does not change the final balance sheet: mankind loses anyway.

A similar story applies to the big claims that cross the border into big lies, such as the promises of the so-called *Master Algorithm*, allegedly capable of extracting *all* the information from the data, doing *everything*, just everything we want, even before we ask for it [14].³

In this essay, we hope we have made it clear why some of the boldest claims of BD are in fact little or nothing short of big lies, that is:

- (i) Complex systems support uncertainty to a much stronger degree than the Law of Large Numbers (Gaussian statistics) would have us believe. The implication is that error decay with data volume is considerably slower, to the point of becoming impractically slow even in the face of zettabytes;
- (ii) No system is infinite, but when operating to their maximal extent, complex systems support the onset of competitive interactions, in turn leading to data conflicts, which may either saturate the return on investment (in terms of the information gained per unit of data) or even make it negative by supplying more data than a finite-capacity system can process. Those BD aficionados who promise us ‘all we want and more’ simply choose to ignore this, and it is not hard to see why.
- (iii) In the end, most of BD comes down to more or less sophisticated forms of curve fitting based on error minimization. Such minimization procedures fare well if the error landscape is smooth, but they exhibit fragility towards corrugated ones in other situations, which are the rule in complex systems (figure 3).

Given these properties of nonlinear systems, the idea of replacing understanding with glorified curve fitting, no matter how ‘clever’, appears a pretty questionable bargain, to put it mildly. In this context, it is amusing to recall a conversation between Enrico Fermi and Freeman Dyson in 1952 [15]:

Fermi: How many parameters did you use in your calculations?

Dyson: Four.

Fermi: My friend John von Neumann used to say with four parameters, I can fit an elephant and with five, I can make him wiggle his trunk.

And, with that, the conversation was over.

8. What can be done?

Once radical empiricism, hype-blinded high-tech optimism and the most rapacious forms of business motivation are filtered out, what remains of BD are nonetheless a serious and promising scientific methodology. In the end, however, it is nothing other than an elaborate form of curve fitting, but this is not intended as a dismissive statement: sophisticated forms of inference, search and optimization are involved in this activity, which deserve credit and respect, although certainly not awe.

There is no doubt that the ‘big data/machine learning/artificial intelligence’ (BD/ML/AI) approach has plenty of scope to play a constructive and important role in addressing major scientific problems. Among the applications, pattern recognition is particularly powerful in

³This sentence appears in the marketing introduction (Italian version) of the book, not in the book itself.

detecting patterns which might otherwise remain hidden indefinitely (modulo the problem of false positives mentioned earlier). Possibly the most important role is likely to be in establishing patterns which then demand further explanation, where scientific theories are required to make sense of what is discovered. We have written elsewhere [7] of the fact that rapid ‘successes’ of BD approaches take far longer to turn into sources of scientific insight.

In passing, however, we cannot refrain from commenting on the resurgence of use of the term ‘artificial intelligence’ in this context, more than forty years after Marvin Minsky’s unfortunate claim that computers were just a few years away from emulating human intelligence. That wild claim led to a decades-long ‘AI winter’ from which one observes not only a thaw but the extravagant hype accompanying any claimed successes of the BD approach. We do not propose to digress into a discussion of AI, other than to point out that the concept has been subjected to penetrating analysis among others by Roger Penrose [16], who argues cogently that no digital computer will ever be capable of matching the human brain in terms of its ability to resolve problems such as those that reside in the class of the Gödelian undecidable. It matters not one iota that a so-called ‘AI machine’ has the capability of assimilating the contents of staggeringly vast numbers texts (Tegmark [17]).

In the final part of this essay, we focus rather on the positive aspects, namely how BD might assist with the struggle of the human mind to overcome three notorious barriers: nonlinearity, non-locality and hyperdimensional spaces.

(a) Nonlinearity

Nonlinearity is a notoriously tough cookie for theoretical modelling, for various reasons, primarily because nonlinear systems do not respond in proportion to the extent to which they are prompted. The most spectacular and popular metaphor of nonlinearity is the well known ‘butterfly’ effect, namely the little butterfly beating her wings in Cuba and triggering a hurricane in Miami in the process. This is the ominous side of nonlinearity, the one that hits straight at our ability to predict the future, the harbinger of uncertainty.

Less widely known perhaps is the sunny side of nonlinearity, its *constructive* power, which is most apparent in biology where it underlies spatio-temporal organization. We shall not delve any further into this vital edifice of modern science [18]. Up to half a century ago, nonlinearity was hidden under the carpet of science for two good reasons. First, many systems under comparatively small loads, do indeed respond linearly (consider, for example, the logistic equation away from capacity). Second, linear systems are incomparably easier to deal with on mathematical grounds. It was only in the 1960s, with the birth of chaos theory, that nonlinearity started to be fully embraced by the scientific method, and it has continued to advance across all fields of science ever since.

While BD can certainly be of assistance in tackling some of the vagaries of nonlinear systems, the fractal nature of many nonlinear dynamical systems utterly defies any notion of the smooth mappings upon which essentially all machine-learning algorithms are based, rendering them nugatory from the outset in such contexts. Indeed, the discontinuous nature of many nonlinear systems is simply not amenable to approaches based on machine learning’s common assumptions that relationships are smooth and differentiable.

(b) Non-locality

A further source of difficulty for scientific investigation is non-locality understood as meaning the presence of long-range correlations, by which we mean that interactions between entities (such as particles or fluid domains) decay very slowly with increasing distance. The classical example is the gravitational many-body problem, in which the force decays with the square of the inverse distance between two bodies. Non-locality is a problem because it generates an all-to-all interaction scenario in which the computational complexity grows quadratically with the number of interacting units. The problem is far more acute in the quantum context, where non-locality

takes the form of the once-dreaded ‘action at a distance’, or more precisely to entanglement, meaning that different parts of a system remain causally connected even when they are arbitrarily far apart. This challenges our basic intuition that things interact most when they are in close proximity.

Leaving aside abundant metaphysical and science fiction based ramifications, entanglement stands as a highly counterintuitive and difficult phenomenon to deal with by the current methods of theoretical science. Addressing quantum correlations with machine learning is plainly a major challenge [19].

(c) Hyper-dimensions

We are used to living in three spatial dimensions, plus time, and we are clearly often in difficulty in going further. In fact, visualizing objects, not to mention dynamic phenomena, in just three dimensions seems to be complicated enough, as everyone dealing with visualization software knows all too well. Yet most problems inhabit a much larger domain, known as *phase-space*.

Macroscopic systems consist of a huge number of individual components, typically in the order of the Avogadro number $A_v \sim 6 \times 10^{23}$ for the case of standard quantities of matter that we encounter. If each component is endowed with just six degrees of freedom, say its position in space and its velocity, this makes six times Avogadro’s number of variables, namely a mathematical problem in six times Avogadro’s number of dimensions. Such is the monster-dimensionality that matters for many modelling purposes.

Thanks to a gracious gift ‘we neither understand, nor deserve’ in Eugene Wigner’s words [20], much can be learned about these systems by solving problems in a much lower number of dimensions using the methods of statistical mechanics [21]. Even so, the task of modelling complex systems, say in weather forecasting and in protein folding, to name but two outstanding problems in modern science, remains very hard [7].

Calculating the electronic structure of molecules is firmly in the class of computationally intractable problems. Accurate calculations scale factorially in the size of the basis sets used and render the highest levels of theory/accuracy essentially unattainable for anything other than the smallest of molecular systems. Here, considerable hype if not expectation has been focused on the construction of working quantum computers which would exhibit a special form of ‘quantum parallelism’ that allows even such kinds of classically intractable problems to be solved on feasible time scales. There is absolutely no chance of BD/ML/AI being applicable here: each problem is in a class of its own, and there are not going to be sufficient examples of solved problems available any time soon on which inference-based approaches could even begin to be contemplated.

What is curious about the current fad for quantum computing is that, as with BD etc., the hype is at its peak in the big corporations, such as Microsoft, Google, IBM, *et al.*, who make claims we will have a working quantum computer in five years (we are excluding the D-Wave adiabatic quantum variant). These are the very same corporations which inundate us with reminders of the power of BD/ML/AI. And remarkably, one of the applications they say would be a ‘killer app’ is in quantum chemistry, for the discovery of future drugs, at the same time as they promote BD methods to do the same thing.

9. A new scientific deal

It would be highly desirable if BD and particularly machine-learning techniques could help surmount the three basic barriers to our understanding described above.

For now, however, in hard-core physical science at least, there is little evidence of any major BD-driven breakthroughs, at least not in fields where insight and understanding rather than zero sales resistance is the prime target: physics and chemistry do not succumb readily to the seduction of BD/ML/AI. It is extremely rare for specialists in these domains to simply go out and collect vast quantities of data, bereft of any guiding theory as to why it should be done. There are some exceptions, perhaps the most intriguing of which is astronomy, where sky scanning telescopes

scrape up vast quantities of data for which machine learning has proved to be a powerful way of both processing it and suggesting interpretations of recorded measurements. In subjects where the level of theoretical understanding is deep, it is deemed aberrant to ignore it all and resort to collecting data in a blind manner. Yet, this is precisely what is advocated in the less theoretically grounded disciplines of biology and medicine, let alone social sciences and economics. The oft-repeated mantra of the life sciences, as the pursuit of ‘hypothesis driven research’, has been cast aside in favour of large data collection activities [7].

And, if the best minds are employed in large corporations to work out how to persuade people to click on online advertisements instead of cracking hard-core science problems, not much can be expected to change in the years to come. An even more delicate story goes for social sciences and certainly for business, where the burgeoning growth of BD, more often than not fuelled by bombastic claims, is a compelling fact, with job offers towering over the job market to an astonishing extent. But, as we hope we have made clear in this essay, BD is by no means the panacea its extreme aficionados want to portray to us and, most importantly, to funding agencies. It is neither Archimedes’ fulcrum, nor the end of insight.

Therefore, instead of rendering theory, modelling and simulation obsolete, BD should and will ultimately be used to complement and enhance it. Examples are flourishing in the current literature, with machine-learning techniques being embedded to assist large-scale simulations of complex systems in materials science, chaotic systems, turbulence [22–25] and also to provide major strides towards personalized medicine [11], a prototypical problem for which statistical knowledge will never be a replacement for patient-specific modelling [7]. It is not hard to predict that major progress may result from an inventive blend of the two, perhaps emerging as a new scientific methodology.

Acknowledgements. This essay grew out of the *Lectio Magistralis* ‘Big Data Science: the End of the Scientific Method as We Know It?’ given by S.S. at the University of Bologna and various talks by P.V.C. on the need of big theory for big data. S.S. appreciates enlightening discussions with S. Strogatz and G. Parisi. P.V.C. thanks E. Dougherty, F. Alexander and R. Highfield for valuable discussions, along with J. Dagley for stylistic advice. Both authors thank A. Hoekstra and P. Sloom for many illuminating discussions over the years.

Data accessibility. This article has no additional data.

Competing interests. We declare we have no competing interests.

Funding. S.S. wishes to acknowledge financial support from the European Research Council under the European Union’s Horizon 2020 Framework Programme (no. FP/2014–2020)/ERC Grant Agreement no. 739964 (COPMAT). P.V.C. is grateful for funding from the MRC Medical Bioinformatics project (MR/L016311/1), EU H2020 CompBioMed and VECMA (Grant nos. 675451 and 800925) and from the UCL Provost.

References

1. Strogatz S. 2006 *The end of insight, John Brockman series ‘What is your dangerous idea?’*, Ch 26. New York, NY: Edge org Brockman inc.
2. Lewis CS. 2009 *The abolition of man*. New York, NY: Harper Collins.
3. Silver N. 2012 *The signal and the noise: why so many predictions fail, but some don’t*. The Penguin Press: New York.
4. Succi S. 2018 *The lattice Boltzmann equation (for complex states of flowing matter)*. Oxford, UK: Oxford University Press.
5. Taleb N. 2005 *Fooled by randomness: the hidden role of chance in life and in the markets*. Random House: New York.
6. Taleb N. 2007 *The black swan: the impact of the highly improbable*. Random House: New York.
7. Coveney PV, Dougherty ER, Highfield R. 2016 Big data need big theory too. *Phil. Trans. R. Soc. A* 374, 20160153. (doi:10.1098/rsta.2016.0153)
8. Calude CS, Longo G. 2017 The deluge of spurious correlations in big data. *Found. Sci.* 22, 595. (doi:10.1007/s10699-016-9489-4)
9. Meng XL. 2014 A trio of inference problems that could win you a nobel prize in statistics (if you help fund it). In *Past, Present, and Future of Statistical Science* (eds X Lin, C Genest, DL Banks, G Molenberghs, DW Scott, J-L Wang), pp. 537–562. Boca Raton, FL: CRC Press.

10. Karbalayghareh A, Qian X, Dougherty ER. 2018 Optimal Bayesian transfer learning. *IEEE Trans. Signal Process* **66**, 3724–3739. (doi:10.1109/TSP.2018.2839583)
11. Eccleston RC, Wan S, Dalchau N, Coveney PV. 2017 The role of multiscale protein dynamics in antigen presentation and T lymphocyte recognition. *Front. Immunol.* **8**, 797. (doi:10.3389/fimmu.2017.00797)
12. Pentland A. 2015 *Social physics: how good ideas spread—the lessons of a new science*. London, UK: Penguin Random House.
13. Bauman Z. 2000 *Liquid modernity*. Polity Press: Cambridge.
14. Domingos P. 2015 *The Master Algorithm*. Philadelphia, PA: Basic Books.
15. Dyson F. 2004 A meeting with Enrico Fermi, how one intuitive physicist rescued a team from fruitless research. *Nature* **427**, 297. (doi:10.1038/427297a)
16. Penrose R. 1989 *The emperor's new mind*. Oxford, UK: Oxford University Press.
17. Tegmark M. 2017 *Life 3.0: being human in the age of artificial intelligence*. Allen Lane: London.
18. Coveney PV, Highfield R. 1995 *Frontiers of complexity: the search for order in a chaotic world*. Faber & Faber: London.
19. Biamonte J, Wittek P, Pancotti N, Rebentrost N, Wiebe N, Lloyd S. 2017 Quantum machine learning. *Nature* **454**, 195–202. (doi:10.1038/nature23474)
20. Wigner EP. 1990 The unreasonable effectiveness of mathematics in the natural sciences. *Math. Sci.* 291–306. (doi:10.1142/1104)
21. Coveney PV, Highfield R. 1990 *The arrow of time*. W. H. Allen: London.
22. Schoenholz SS, Cubuk ED, Sussman DM, Kaxiras E, Liu AJ. 2016 A structural approach to relaxation in glassy liquids. *Nat. Phys.* **12**, 469–471. (doi:10.1038/nphys3644)
23. Ling J, Kurzawski A, Templeton J. 2016 Reynolds averaged turbulence modelling using deep neural networks with embedded invariance. *J. Fluid Mech.* **807**, 155–166. (doi:10.1017/jfm.2016.615)
24. Pathak J, Hunt B, Girvan M, Lu Z, Ott E. 2018 Model-free prediction of large spatiotemporally chaotic systems from data: a reservoir computing approach. *Phys. Rev. Lett.* **120**, 024102. (doi:10.1103/PhysRevLett.120.024102)
25. Fang R, Sondak D, Protopapas P, Succi S. 2018 Using deep learning for physical insights into turbulence modelling. In *Proc. 31th Conf. on Discrete Simulation of Fluid Dynamics*.