

Low Complexity, Low Probability Patterns and Consequences for Algorithmic Probability Applications

Mohamed Alaskandarani¹ and Kamaludin Dingle^{1,2,3,*}

¹Centre for Applied Mathematics and Bioinformatics (CAMB),

Department of Mathematics and Natural Sciences, Gulf University for Science and Technology, Kuwait

²Department of Chemical Engineering and Biotechnology, Cambridge University, UK

³Department of Computing and Mathematical Sciences, Caltech, USA

(Dated: July 20, 2022)

Developing new ways to estimate probabilities can be valuable for science, statistics, and engineering. By considering the information content of different output patterns, recent work invoking algorithmic information theory has shown that a priori probability predictions based on pattern complexities can be made in a broad class of input-output maps. These algorithmic probability predictions do not depend on a detailed knowledge of how output patterns were produced, or historical statistical data. Although quantitatively fairly accurate, a main weakness of these predictions is that they are given as an upper bound on the probability of a pattern, but many low complexity, low probability patterns occur, for which the upper bound has little predictive value. Here we study this low complexity, low probability phenomenon by looking at example maps, namely a finite state transducer, natural time series data, RNA molecule structures, and polynomial curves. Some mechanisms causing low complexity, low probability behaviour are identified, and we argue this behaviour should be assumed as a default in the real world algorithmic probability studies. Additionally, we examine some applications of algorithmic probability and discuss some implications of low complexity, low probability patterns for several research areas including simplicity in physics and biology, a priori probability predictions, Solomonoff induction and Occam's razor, machine learning, and password guessing.

Keywords: Algorithmic probability; Kolmogorov complexity; prediction; induction

I. INTRODUCTION

Many problems in science, statistics, and engineering revolve around estimating probabilities of events. This is especially true in the current climate, where machine learning and data science are enjoying broad applications. Hence developing new methods for calculating, bounding, or predicting probabilities is valuable. One direction for such predictions is in theoretical computer science where *algorithmic information theory* [1–4] (AIT) provides a theoretical framework for studying randomness, probability, and complexity. The central quantity of AIT is *Kolmogorov complexity*, $K(x)$, which measures the complexity of an individual object or pattern x via the amount of information required to describe or generate x . The object could be a binary string, integer, graph, or indeed anything that can be represented as a binary string. The AIT coding theorem [5] establishes a fundamental connection between complexity and probability predictions in very general settings. It states that the chances that some output x is produced via some generic computation mechanism is directly related to the complexity of x . More formally, the coding theorem states that $P(x) \sim 2^{-K(x)}$ where $P(x)$ is the probability that an output x is generated by a (prefix optimal) universal Turing machine fed with a random binary program input. $P(x)$ is known as the *algorithmic probability* of x , and the associated probability distribution is known as the *universal distribution* [6].

Algorithmic probability and AIT results are typically difficult to apply in real-world settings due to the fact that $K(x)$ is uncomputable, the theorems assume the presence of uni-

versal Turing machines (UTMs), and the results are asymptotic and stated with accuracy to within an unknown constant. Despite these theoretical difficulties, in practice many successful applications of AIT have been made based, for example in bioinformatics [10, 11], physics [12], signal denoising [13], among many other applications [4]. Mostly these applications use standard compression algorithms to approximate $K(x)$, sometimes combined with various forms of theorem approximation. Algorithmic probability estimates have been made numerically by random sampling [14] and enumeration [15–17] of computer programs.

From a very different perspective, algorithmic probability estimates have also been made via deriving a weaker form of the coding theorem, applicable in real-world contexts [9], taking the form of an upper bound. This weaker form upper bound was applied in a range of input-output maps to make a priori predictions of the probability of different shapes and patterns, such as the probability of different RNA shapes appearing on a random choice of genetic sequence, or the probability of differential equation solution profile shapes, on random choice of input parameters, and several other examples [7, 9]. Surprisingly, it was found that probability estimates could be made directly from the complexities of the shapes themselves, without recourse to the details of the map or reference to how the shapes were generated. The authors of [9] termed this phenomenon of an inverse relation between complexity and probability *simplicity bias* (SB). One important drawback of this work is that only an upper bound prediction was made, rather than a direct probability value prediction. In contrast to original coding theorem, in practice it has been observed that many simple patterns have low probability [9]. Such low Kolmogorov complexity, low probability (LKLP) outputs present a weakness in the predictive ability of the upper bound, because for these outputs their com-

* Correspondence: dinglek@caltech.edu

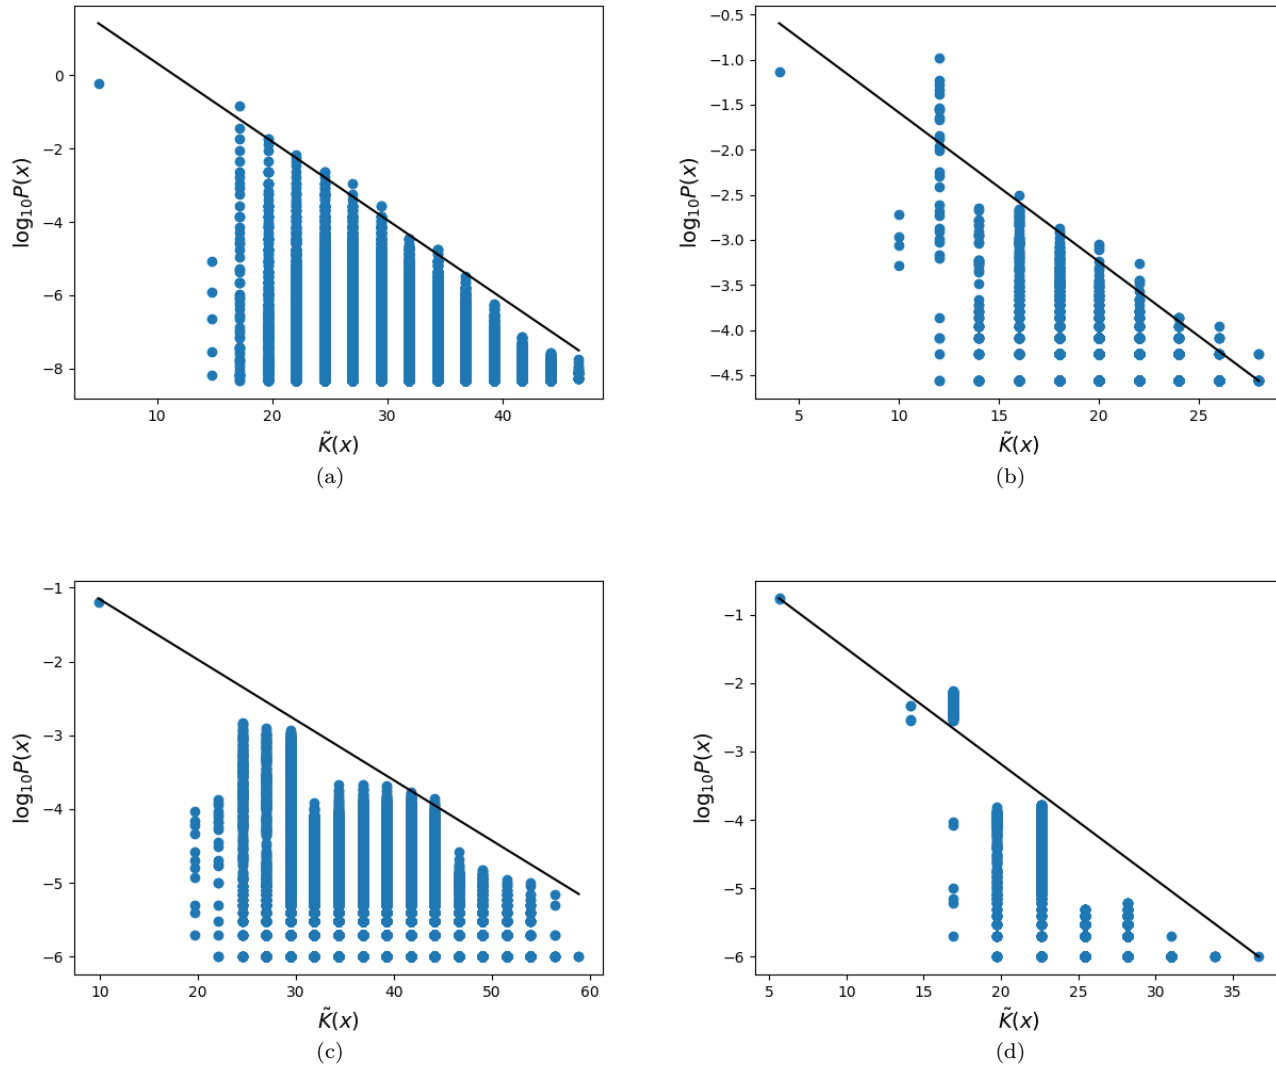


FIG. 1. **Simplicity bias plots with low complexity, low probability outputs.** Blue dots represent different binary string output patterns. The black line is the upper bound of Eq. (3). Probabilities $P(x)$ are the fraction of random inputs that map to output x . $\tilde{K}(x)$ are estimated complexity values for each output x . There are many data points far below the upper bound; these points are low Kolmogorov complexity, low probability (LKLP) output patterns. The maps studied are: (a) A finite state transducer (FST) with data taken from ref. [7]; (b) time series patterns from the World Bank Open Data, with data taken from [8]; (c) computationally predicted RNA secondary structures from 1 million randomly sampled input sequences; and (d) discretised polynomials curves generated with $x \in (0,1)$ and random coefficients, data taken from [9].

plexity and probability values are largely disconnected and hence predicting one from the other is more challenging. Understanding the causes and properties of LKLP outputs may help to improve the accuracy of applications of algorithmic probability, such as better a priori probability predictions and induction. In this work we investigate LKLP behaviour, and its implications for applied algorithmic probability via some examples from previously published works.

II. BACKGROUND THEORY

Some brief accessible background theory is given here; more formal and detailed presentations are available in, eg refs. [4, 18, 19]. A universal Turing machine (UTM) [20]

is an abstract general computing device which can simulate any other Turing machine. A UTM has the highest computational capacity, and can implement any conceivable algorithm which could in principle be run on a computer. A programming language is called *Turing complete* (or computationally universal) if the language is sufficiently expressive to be able to simulate a UTM, and therefore implement any algorithm. Common languages like Python, C, FORTRAN, etc are Turing complete. If a function can be computed by a finite mechanical procedure, then it is a computable function. For computable functions, all inputs or programs eventually halt, and cannot (for example) run on forever in an infinite loop. Common functions like polynomials, exponentials, trigonometric functions, etc are computable.

The Kolmogorov complexity $K_U(x)$ of a string x with re-

spect to U , is defined [1–3] as

$$K_U(x) = \min_p \{|p| : U(p) = x\} \quad (1)$$

where p is a binary program for a prefix optimal UTM U , and $|p|$ indicates the length of the binary program p in bits. Due to the invariance theorem [4] for any two optimal UTMs U and V , $K_U(x) = K_V(x) + O(1)$ so that the complexity of x is independent of the machine, up to additive constants. Hence we conventionally drop the subscript U in $K_U(x)$, and speak of ‘the’ Kolmogorov complexity $K(x)$. Informally, $K(x)$ can be defined as the length of a shortest program that produces x , or simply as the size in bits of the compressed version of x . If x contains repeating patterns like $x = 10101010101010$ then it is easy to compress, and hence $K(x)$ will be small. On the other hand, a randomly generated bit string of length n is highly unlikely to contain any significant patterns, and hence can only be described via specifying each bit separately without any compression, so that $K(x) \approx n$ bits. Other names for $K(x)$ are *descriptive complexity*, *algorithmic complexity*, and *program-size complexity*. Fundamentally, $K(x)$ measures the amount of information to describe or generate x precisely and unambiguously. Note that Shannon information and Kolmogorov complexity are closely related to each other [21], but also differ fundamentally because Shannon information quantifies the information content of a random source, while Kolmogorov complexity quantifies the information of individual sequences or objects.

Solomonoff invented algorithmic probability [22], but it was later formalised and extended by Levin [5], who proved the AIT *coding theorem* in 1974 which states that

$$P(x) = 2^{-K(x)+O(1)} \quad (2)$$

where $P(x)$ is the probability that UTM U generates output string x on being fed random bits as a program (again we have dropped the subscript U). Thus, high complexity outputs have exponentially low probability, and simple outputs must have high probability.

Coding theorem-like behaviour in real-world input-output maps was studied, leading to the observation of a phenomenon called *simplicity bias* (SB) [9]. SB is captured mathematically as

$$P(x) \leq 2^{-a\tilde{K}(x)-b} \quad (3)$$

where $P(x)$ is the (computable) probability of observing output x on random choice of inputs, and $\tilde{K}(x)$ is the estimated Kolmogorov complexity of the output x : complex outputs from input-output maps have lower probabilities, and high probability outputs are simpler. The constants $a > 0$ can often be estimated without recourse to sampling, but just by knowing or estimating the total number of different possible outputs [9]. Using $b = 0$ is the default guess for this constant, but it can also be fit to the data via partial sampling.

A complete understanding of exactly which maps will show SB has not been developed, but SB is expected to appear in many input-output maps, under fairly general conditions. Importantly, the map should be ‘simple’ (technically of $O(1)$ complexity) to prevent the map itself from dominating over inputs in defining output patterns. If an arbitrarily complex map was permitted, outputs could have arbitrary complexities and probabilities, and thereby remove

any connection between probability and complexity. Strong bias is a prerequisite for SB, typically the relative value of the largest and smallest probabilities should be more than the value of the number of outputs, ie for N_O outputs, $\max(P(x))/\min(P(x)) > N_O$.

III. RESULTS

A. Low complexity, low probability behaviour is very common

Based on numerical experiments using the upper bound of Eq. (3), Dingle et al [9] reported that many outputs x have probability values far below their predicted upper bounds, ie $P(x) \ll 2^{-a\tilde{K}(x)-b}$ for many x . This observation was not due to the bound being trivially loose, because it is a tight bound for many outputs. Significantly, the small fraction of outputs which were close to the upper bound absorbed most of the probability mass, or in other words most of the inputs map to outputs for which the bound is tight. As a consequence it was shown [7, 9] analytically and numerically that for an output x generated by a random input, with high probability $P(x) \approx 2^{-a\tilde{K}(x)-b}$.

To illustrate the LKLP phenomenon, in Figure 1 we show four probability-complexity plots. The probabilities $P(x)$ are calculated as the fraction of random inputs which produce output x . The complexity $\tilde{K}(x)$ denotes an estimate of the Kolmogorov complexity of each output, using a slightly adapted [9] version of the famous Lempel-Ziv [23] 1976 complexity measure. The black lines are fitted upper-bounds, depicting the SB upper bound of Eq. (3). All four maps have been studied in earlier works, and are: (a) a finite state transducer (FST), which is a simple generic model of computation, but unlike a UTM has a very limited computational capacity (being the lowest on the Chomsky hierarchy). Here the outputs are length 30 bits, and were obtained from thorough sampling of binary string inputs, and the data was taken from ref. [7]. (b) Time series data taken from the World Bank Open Data project (<https://data.worldbank.org>), which have been discretised to binary strings (of length 16 bits) while studying SB [8]. (c) Computationally predicted [24] RNA secondary structures obtained from randomly sampling 1 million sequences of length $L = 35$ nucleotides. Following the protocol of ref. [9] in which SB in computationally generated RNA structures was first reported, predicted dot-bracket structures were converted to binary strings and complexity values were thereby estimated. (d) Polynomial curves $y = \sum_{i=1}^{14} \alpha_i x^i$ with Gaussian random coefficients α_i and $x \in (0, 1)$, with data from [9]. The curves were discretised to binary strings by the up/down method [25, 26]. More details of these maps and relevant analyses can be found in the original cited papers.

It is apparent that in each panel of Figure 1 the data points show a similar ‘triangle’ shape, with some points closely following the upper bound (black line), but at the same time for each complexity value a large variation in probability values is observed, inferred from the many points far below the bound. Occupying the bottom left corner of the ‘triangles’ are the outputs which exhibit the strongest LKLP behaviour, because they have very low probability, but at the same time

have relatively low complexity values. This LKLP ‘triangle’ depicted in all maps of Figure 1 appears in essentially all examples of SB studied so far, including in biology [27], machine learning [28], and other contexts [9] (Supp. Info.).

In Figure 1 and in most other SB examples cited, the previously described complexity estimator $\hat{K}(x)$ based on the Lempel-Ziv complexity measure was used. However, LKLP outputs have also been observed when using other complexity measures, such as those used in protein quaternary structures and polyominoes [27], so LKLP outputs are not merely an artefact of the Lempel-Ziv complexity measure. From another angle, nearly all the cited complexity-probability plots showing SB were generated via random sampling of inputs, and so it might be suggested that with more sampling, or full enumeration of inputs, LKLP outputs would disappear or at least become much less pronounced. However, this suggestion is countered by the observation that LKLP behaviour is also observed in maps for which probabilities were directly calculated by enumerating every possible input, for example L-systems [9] and RNA length $L = 15$ nucleotides [7]. Perhaps the only clear example of a map which does not show strong LKLP behaviour is in a small polyomino system using path complexity (see Fig 4.2(a) in Chapter 4 of ref. [29]).

To avoid confusion, we stress that LKLP behaviour is not due to a failing of complexity measures to detect patterns. If, say, the Lempel-Ziv measure failed to detect an important pattern and thereby assigned a high complexity to an output which is actually simple, then the opposite behaviour would be observed in the complexity-probability plots, where some outputs would be far *above* the upper bound, not far below the bound.

A well known limitation of the original coding theorem is choice of UTM, the problem being that algorithmic probability estimates from different machines will differ [30], and there is no known way to remove this dependence [31]. However, due to the invariance theorem, this machine dependence is only a problem for (typically small) outputs with low complexities, for which $O(1)$ terms due to translating between UTMs may dominate, but this is not a problem for large complexity strings. This machine dependence limitation may not be of great concern perhaps, because many real-world phenomenon are highly complex. In contrast, LKLP outputs do not disappear for large complexity outputs, and so LKLP can be seen as a more serious problem for applied algorithmic probability than machine dependence.

Finally, in many maps which exhibit SB, some patterns are actually impossible to make, such that $P(x) = 0$, so this could be viewed as an extreme form of LKLP behaviour.

B. Estimating which of two outputs is more likely

Another way to examine and quantify LKLP behaviour is in terms of predicting which of two outputs has higher probability. In the absence of LKLP outputs, if $K(x) < K(y)$ for two strings x and y then it follows from Eq. (2) that $P(x) > P(y)$, ignoring $O(1)$ terms. With LKLP behaviour, the direct connection between probability and complexity is disrupted, and $K(x) < K(y)$ does not imply that $P(x) > P(y)$ necessarily. Note that predicting which of two strings has higher probability does not depend on estimating or fit-

ting the parameters a and b in Eq. (3); only the complexity values are required. This kind of calculation was done earlier by [8, 9].

The manner in which the strings x and y are chosen also affects the strength of the connection between complexity and probability. If x and y result from randomly chosen inputs, then they will be sampled with weights according to their probabilities, and in this case $P(x) \approx 2^{-a\hat{K}(x)-b}$ and $P(y) \approx 2^{-a\hat{K}(y)-b}$ [7]. In this probability-weighted sampling scenario, for randomly generated outputs, the probability-complexity connection is quite strong. Hence we should expect to be able to predict whether $P(x) > P(y)$ or $P(x) < P(y)$ just from the complexity values, with quite high accuracy. If instead x and y are chosen uniformly randomly from the full set of possible outputs, then we can expect the probability-complexity connection to be less strong, and predicting whether $P(x) > P(y)$ or $P(x) < P(y)$ just from the complexity values will be less accurate.

Here we computationally study the question of which of two strings has higher probability, both using the randomly generated pairs x, y and also uniformly sampled x, y . The protocol we employ is: predict $P(x) > P(y)$ if $K(x) < K(y)$, $P(x) < P(y)$ if $K(x) > K(y)$, and we randomly guess which has higher probability if $K(x) = K(y)$. Note that there will be relatively few unique complexity values (eg only $\sim n$ values for strings of length n), and hence it happens quite commonly that two random strings have the same complexity, while it is relatively rare for two strings to have exactly the same probability value.

For both probability-weighted sampling and uniform sampling, the protocol just described has a null prediction accuracy level of 50% (assuming a null hypothesis that output complexity has no predictive value). The calculated accuracy values (after 10,000 sampled pairs) are: FST probability-weighted sampling achieved an accuracy of 79% and uniform sampling an accuracy of 63%; time series 81% and 80%; RNA 78% and 71%; polynomial 82% and 73%. It is quite striking that high levels of accuracy can be achieved just by complexity estimates, which does not even require estimating the values of a or b . As expected, the accuracy values for probability-weighted sampling are higher, while substantially lower for uniform sampling, but still substantially above the 50% null model mark. Overall, the effect of LKLP outputs is clear: while some % accuracy values are (surprisingly) high, they are still substantially below 100% which is what we would expect in the absence of LKLP outputs.

Note that true uniform sampling strictly requires knowing the full set of outputs, which might entail fully enumerating all inputs, or very thorough input sampling. In the four maps studied here, only the FST map was thoroughly sampled whereas the other maps were sampled partially (eg for RNA, only 10^6 sampled inputs out of a possible $4^{30} \sim 10^{18}$ RNA sequence inputs were made). Hence our stated accuracy values for ‘uniform sampling’ will be somewhat overestimates for RNA, time series, and polynomials, but for the FST map the accuracy value of 63% is likely to be close to the true value.

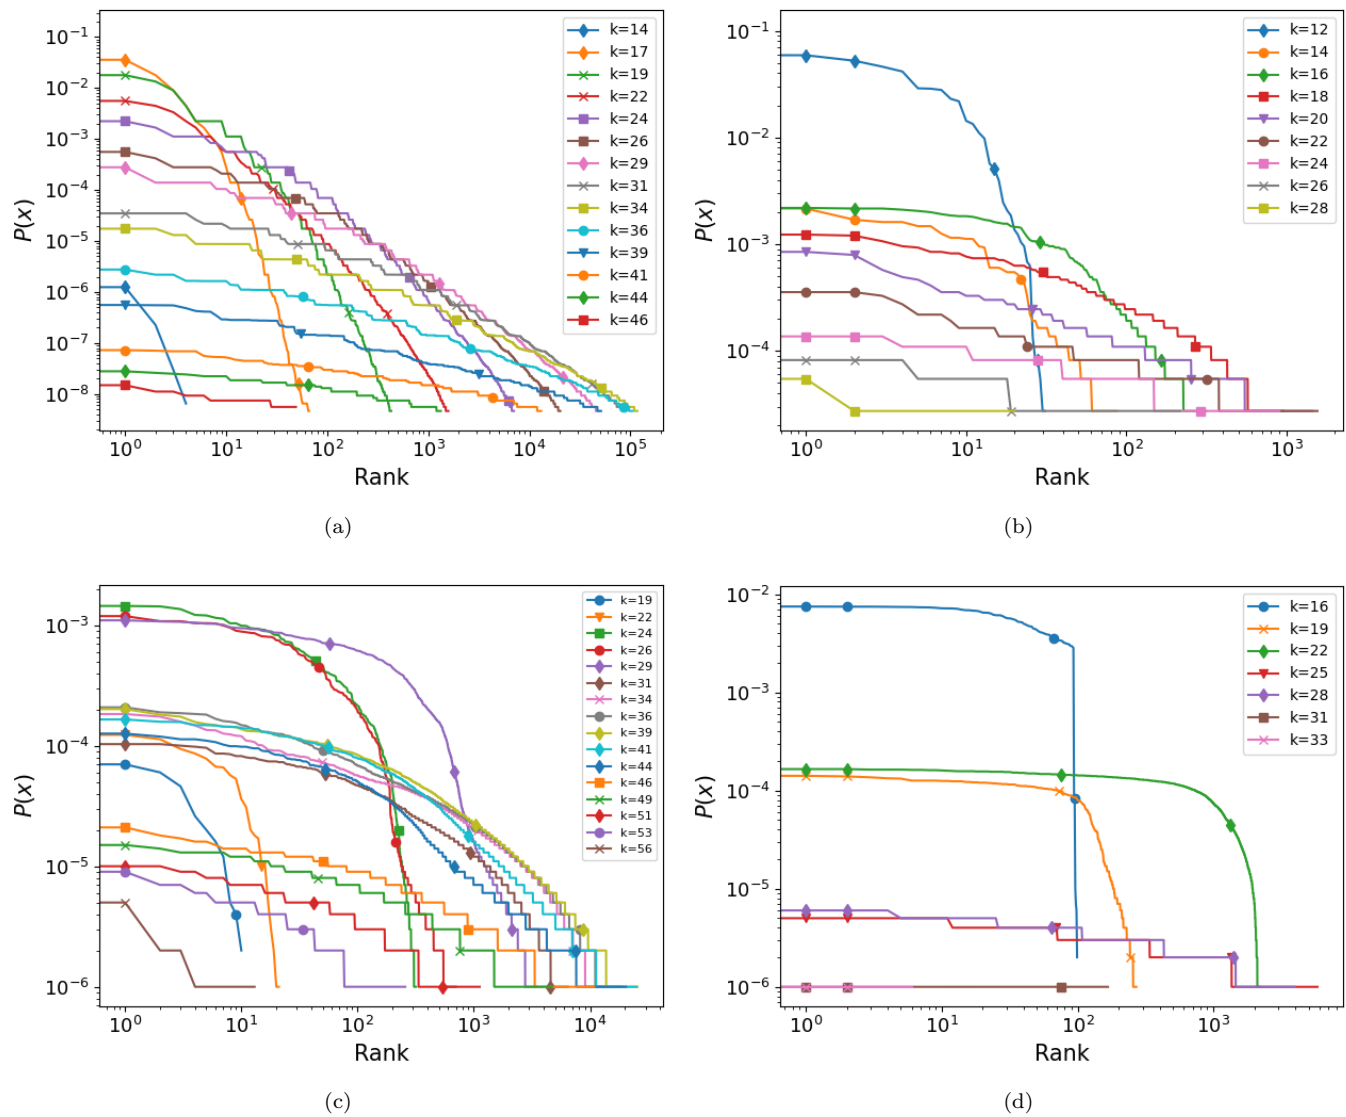


FIG. 2. **Rank plots of probabilities, separated by complexity values.** Each coloured line depicts the decay in probability of outputs with the same complexity value k . The complexity values k are shown in the figure legends. The plots present the same data as in Figure 1, but now plotted as rank plots separated by complexity value. (a) FST; (b) time series; (c) RNA structures; and (d) random polynomials. In each case, a similar pattern of decay is observed, and almost all rank plots decay to the lowest probability for that map.

C. Why do low complexity, low probability outputs exist?

We have seen LKLP outputs in essentially all the maps we have studied. Why are they so common? In the original AIT coding theorem given in Eq. (2), there are no LKLP output patterns, except in the rather modest sense that the $O(1)$ term in the exponent allows for relatively small variation of probabilities within a constant, mathematically $|K(x) - \log_2(1/P(x))| < c$ for some constant c . The reason simple outputs cannot have very low probability in the original AIT coding theorem is that the theorem is based on random prefix codes run on a UTM. In a UTM an output x always has a program p of length $K(x)$ by definition of Kolmogorov complexity. The probability that the specific program p appears as the first p bits of the random binary program is $2^{-K(x)}$ and therefore $P(x) \geq 2^{-K(x)}$ implying

that all simple strings must have high probability. Hence in this UTM setting, it is not possible to have LKLP outputs with $P(x) \ll 2^{-K(x)+O(1)}$.

Computable maps lack some computational power as compared to UTMs, which means that there are some algorithms which they cannot implement (for example, ones that don't halt). As suggested earlier [7] the weaker computational ability of these computable maps suggests that they may not be able to compute some output patterns which are nonetheless simple, or at least cannot compute them with efficient short programs. Hence we can expect some type of LKLP in computable maps, and the lower bound $P(x) \geq 2^{-K(x)}$ does not necessarily hold.

Looking more closely at the lower bound vs. upper bound contrast, it appears that the upper bound on $P(x)$ is more fundamental. The justification is that if the upper bound was violated, and some complex outputs had high probabil-

ues will be combined together into one grouping, but their true complexities will actually be varied, and hence we can expect their probabilities to vary also, apparently yielding LKLP outputs.

The fact that the LKLP ‘triangle’ is common to essentially all maps studied so far in the literature suggests there may be a general explanation. Burkat & Dingle [33] studied SB in the Bernoulli process with parameter $\neq 0.5$, and reported that both the single highest and single lowest probability outputs are also the simplest, and pointed out the cause of LKLP outputs in the Bernoulli map is easy to understand. As an extension, we propose here that if an output is (even roughly) made up of statistically independent parts, then the output may be approximated as a (biased) Bernoulli process, for which we expect LKLP behaviour [33]. As an example, for a time series of some fixed length n , the non-adjacent values are typically correlated only weakly, hence roughly independent, especially if separated by longer intervals. Even for RNA secondary structure for which there are definite correlations between distant parts of the sequence due to bonding interactions, but nonetheless much of the structure is roughly independent of other structure provided the sequence is not very short. This line of argument might help to explain some instances of LKLP outputs.

D. Rank plots of probability values, separated by complexity

We now study the distribution of LKLP probabilities for the four maps shown in Figure 1 above. Within Figure 1, LKLP outputs appear as a ‘column’ of overlapping blue dots stretching down from the black line upper bound, but the distribution of those LKLP probability values is not easy to discern. Instead in Figure 2 for each of those maps, a loglog plot is given showing the rank and probability of each output, coloured separately by complexity value k , so that the distribution is more easily visualised.

Another observation is that the probability decay follows a similar form, with the lowest complexity curves starting at a high probability and quickly decaying, and the higher complexity curves starting lower probability and slowly decaying. This observation can be rationalised by noting that for low complexities, there are exponentially few possible outputs (ie few simple patterns), yet some of these have exponentially high probability as expected from SB. As the complexity increases, it is well known from AIT that there are exponentially more possible patterns for a given complexity value, yet the maximum probability decays exponentially (consistent with the upper bound of Eq. 3), which explains why the rank plots become wider and less tall for higher complexities.

Interestingly, in the FST map, which has the cleanest data, the profile of the different curves appear to depict a straight line decay on the loglog plot, suggesting some kind of Zipf's power law upper bound on the rank decay profiles. Although

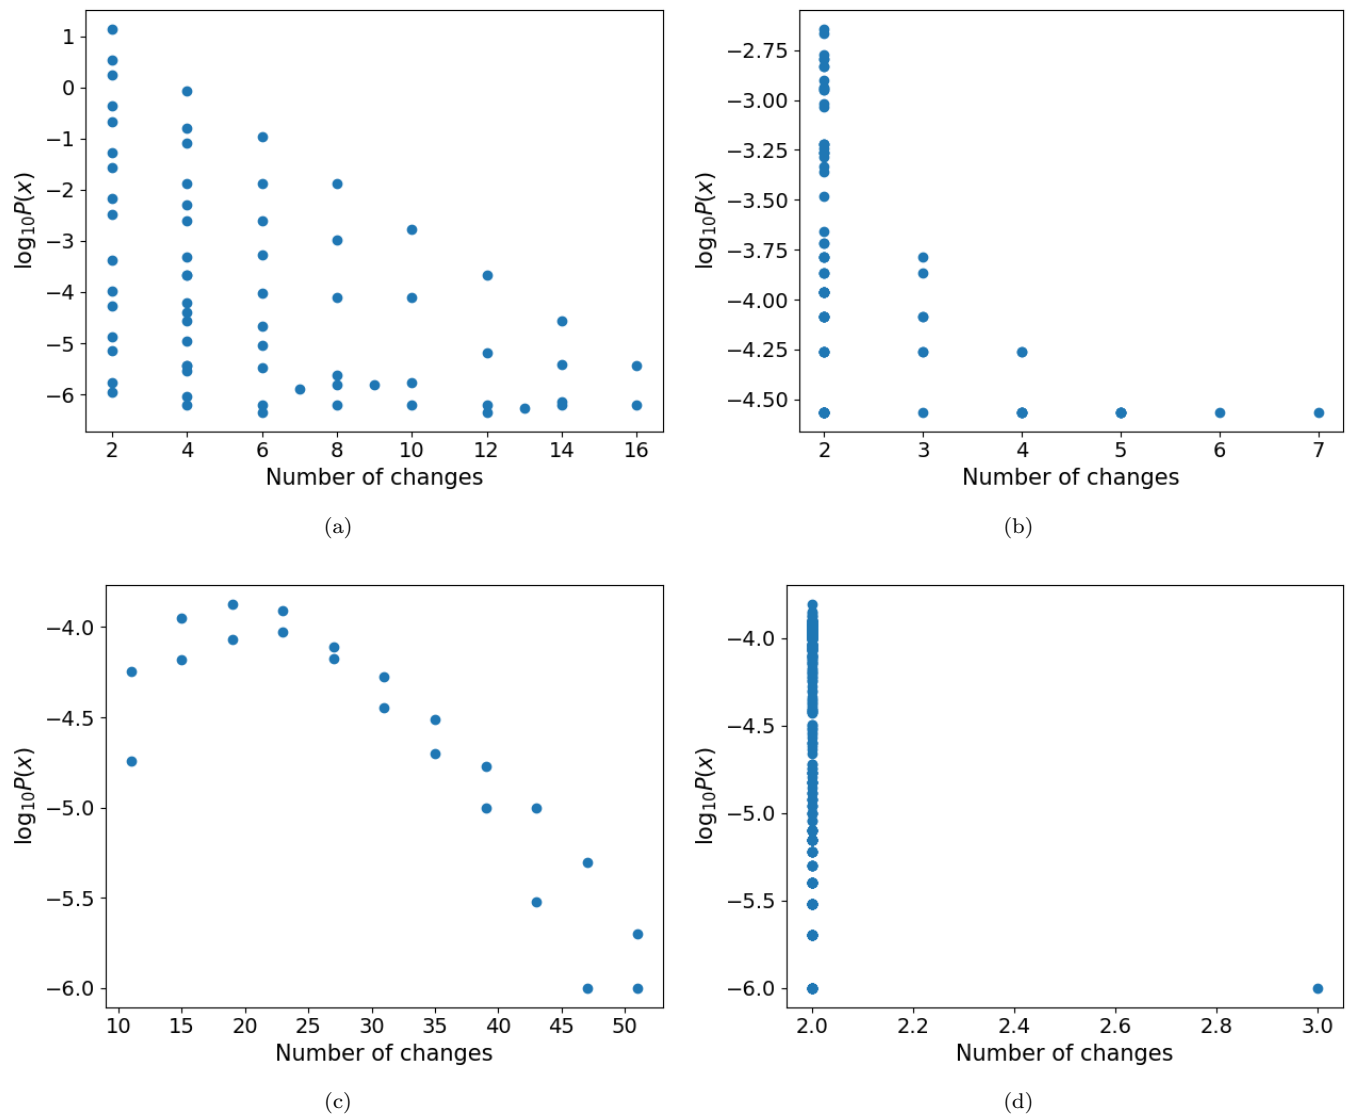


FIG. 3. **Number of 0/1 of changes in an output string vs the output probability, for just one complexity value k .** (a) For the FST map, there is a clear decay in probability with increasing numbers of 0/1 changes, using complexity $k = 17$; (b) similarly for the time series data, there is a clear decay, using $k = 14$; (c) For the RNA there is a strong decay after an initial increase, using $k = 22$; (d) no trend for the polynomials data, using $k = 19$.

the other maps do not show this power-law behaviour clearly, they are also more noisy datasets. We conjecture that some kind of Zipf's law may be a more general feature of these kinds of plots.

It was shown [7] that the sum of the probability mass of outputs with \log_2 probability that is at least Δ bits below the upper bound is no more than $2^{-\Delta+1+O(1)}$. This implies that if there is decay in probabilities, it should be at least exponential. However this does not explain why there is, for all maps, a roughly smooth exponential decay to the same lowest probability. This remains an open question for future work.

E. Is there a bias against changes in natural data?

We saw that map-specific idiosyncrasies can bias output probabilities via patterns which are hard to make for the map but are not in themselves very complex. But are there any general trends which are likely to be common across different real world maps? The AIT coding theorem in Eq. (2) and the related upper bound in Eq. (3) are based on fundamental information content of patterns assumed to be generated by UTMs, but in the natural world many patterns may not be physically easy to make. We suspect there is a disconnect between simple patterns with low information content, and those which are easy to make by real world systems. As an example, the patterns $x = 00000001111111$ and $y = 01010101010101$ have similar complexity from an information perspective, eg they have the same entropy and same size. But it may be that oscillations of y are 'harder' to

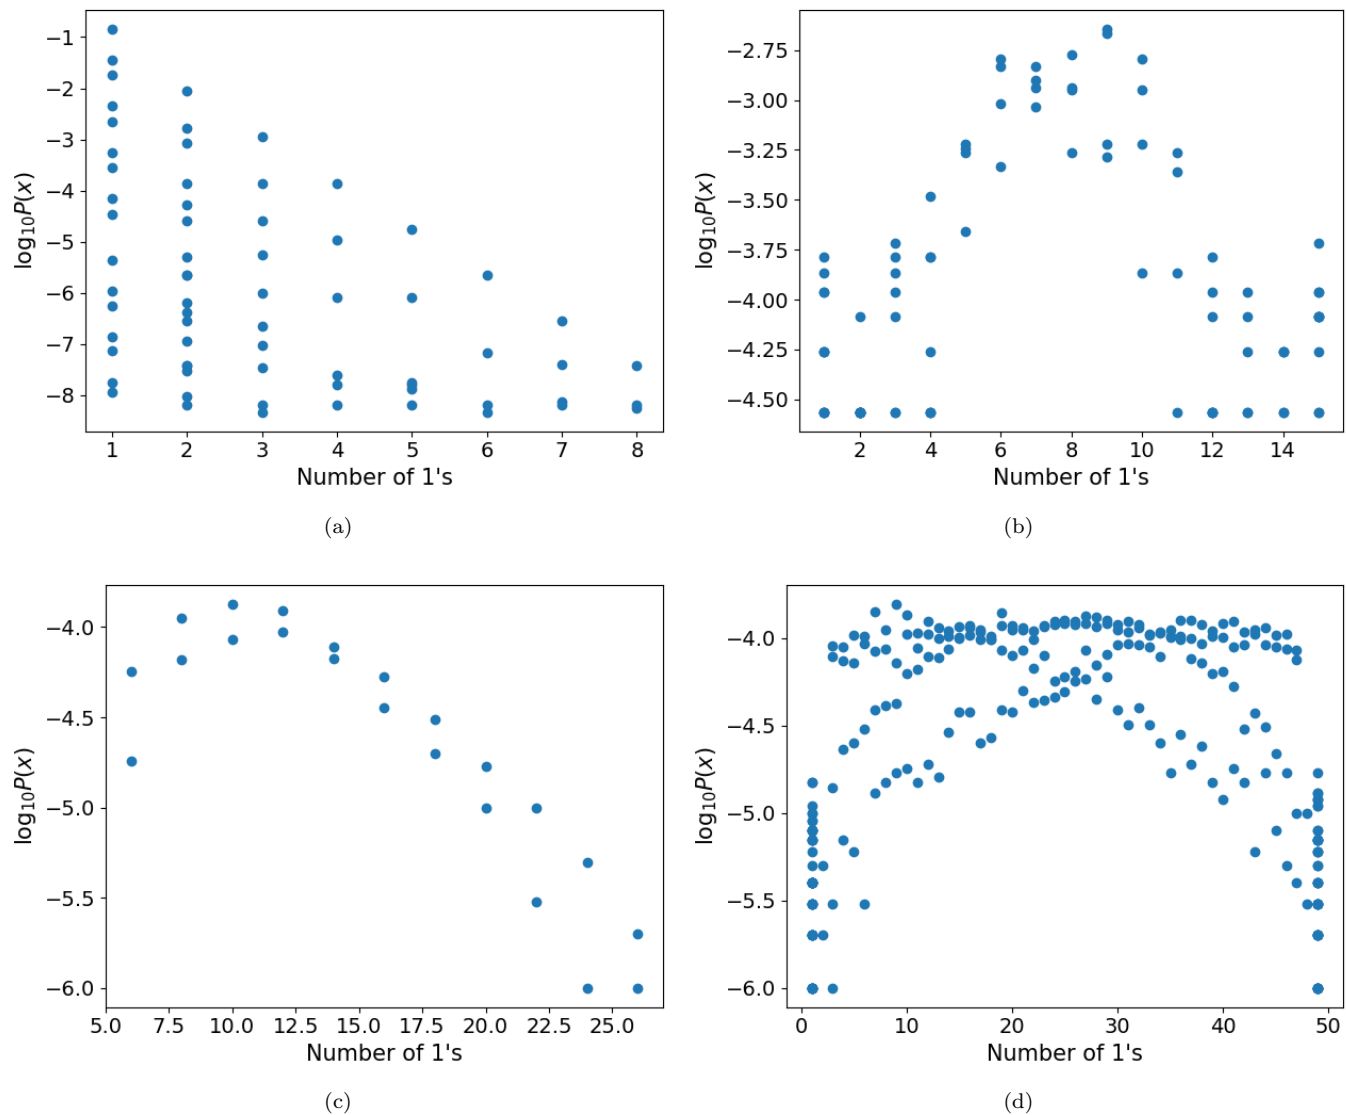


FIG. 4. **The number of 1's in an output string vs the string's probability.** For each example map a plot is given for just one complexity value k , the second simplest, chosen to aid illustration. (a) The FST example shows a linear trend in log probability with increasing number of 1's, using complexity $k=17$; (b) Large and small numbers of 1's are not favoured in the time series, using $k=14$; (c) RNA shows an overall negative relation, using $k=22$; (d) polynomials, similar to time series, large and small numbers of 1's are not favoured, using $k=19$.

maintain in the physical world than the one simple change in x , because change is often expensive in terms of energy or mechanics (with notable exceptions like pendulums, but even they can be easily upset). A more extreme example is the Champernowne's number 0.123456789101112.... which again is a very simple in terms of information content, but it is hard to see how any natural system in the physical world could make such a pattern. Potentially these kinds of patterns may be LKLP in a range of different computable, physically relevant maps.

As a first brief investigation into whether change is 'hard' to generate for natural systems and hence associated to LKLP outputs, we study just the frequency of 0/1 changes in outputs. By "changes" we mean the number of times a 1 is followed by a 0, or a 0 followed by a 1, eg the string 00111010 has four changes, and the string 0000 has no changes. We ask: for two strings of the same complexity value k , does the

probability of an output decrease with increasing numbers of 0/1 changes?

Figure 3 shows plots for the number of changes vs probability for each map. To aid illustration in the figure, for each map we chose just a single complexity value k , which was the second lowest of the output complexity value. The reason for this choice is that LKLP behaviour is most pronounced for very low complexities. For the FST, time series, and RNA there is a clear trend showing an exponential decay in probability with increasing number of 0/1 changes. The polynomial case does not show a trend, but on the other hand there are only two value (2 and 3 changes) which makes observing a trend difficult.

The linear correlation and p -values for the trends are given here in the format of complexity value, correlation coefficient, and p -value. For brevity, values are only reported if either the p -value was statistically significant (taken

as < 0.05) or the correlation coefficient r was non-trivial, taken here as $r < -0.4$. FST: $k=14$, $r=-0.95$, $p\text{-value}=0.01$; $k=17$, $r=-0.49$, $p\text{-value}=3e-05$. Time series: $k=12$, $r=-0.80$, $p\text{-value}=3e-08$; $k=14$, $r=-0.44$, $p\text{-value}=1e-05$; $k=16$, $r=-0.42$, $p\text{-value}=1e-14$; $k=18$, $r=-0.46$, $p\text{-value}=1e-50$. RNA: $k=19$, $r=-0.82$, $p\text{-value}=0.002$; $k=22$, $r=-0.84$, $p\text{-value}=1e-06$; $k=24$, $r=-0.46$, $p\text{-value}=3e-18$; $k=26$, $r=-0.46$, $p\text{-value}=6e-38$; $k=29$, $r=-0.45$, $p\text{-value}=2e-168$. Polynomial: $k=16$, $r=-0.95$, $p\text{-value}=6e-51$; $k=22$, $r=-0.70$, $p\text{-value}=6e-321$.

We conclude that a bias for or against changes is a simple mechanism that can quickly lead to exponentially large variation in probabilities, for outputs with the same informational complexity. Further, we tentatively conjecture that in the natural data the bias is more likely to be against changes.

F. Bias for 1s or 0s

Another simple mechanism which might conceivably strongly affect probabilities is a bias for or against 1's or 0's. In terms of information theory, 1's and 0's are given equal weighting, and all common complexity measures would assign equal complexity to a string and to the same string but with the 0's and 1's flipped, eg 010110 flipped to become 101001. Despite this, it is easy to imagine that a natural system may have a bias for 1's or 0's, and hence the symmetry in information between a string and its flipped counterpart would not extend to a similarity in the probabilities of the two strings (Cf. [33]). Further, if a discretisation such as converting a real-valued time series into a binary string was performed, then how exactly this was done could bias for 0's or 1's.

We investigate this potential bias in the four maps described above. Figure 4 shows the scatter plots for each map, with the number of 1's in a string vs the probability of the string. Again here just one complexity value k for each map is used; if there was no LKLP behaviour, then all strings would have the same probability. The FST example in Figure 4(a) has clear bias against 1's, with exponentially decaying probability with more and more 1's. The time series and polynomial data do not display a linear trend, but rather large and small fractions of 1's are associated with lower probabilities, which for the time series can be expected from the way the discretisation was done. The reason is less clear for the polynomial example. The RNA example has an overall bias against 1's, but it is not quite a linear trend. Based on these few examples, there is no obvious common relation between the number of 1's in a string and its probability.

The linear correlation and p -value are given now, but only for correlations that were both non-trivial and statistically significant. FST: $k=14$, $r=-0.95$, $p\text{-value}=0.01$; $k=17$, $r=-0.50$, $p\text{-value}=2e-05$; $k=22$, $r=-0.41$, $p\text{-value}=4e-65$. Time series: no linear relations. RNA: $k=19$, $r=-0.82$, $p\text{-value}=0.002$; $k=22$, $r=-0.84$, $p\text{-value}=1e-06$; $k=24$, $r=-0.48$, $p\text{-value}=6e-20$; $k=26$, $r=-0.50$, $p\text{-value}=6e-45$; $k=29$, $r=-0.47$, $p\text{-value}=2e-189$. Polynomial: No linear correlations.

It seems that exponentially varying probability, unrelated to complexity, can easily arise in natural systems from biases for 1's or 0's and/or changes and oscillations. This helps to understand LKLP behaviour, but a full understanding of

the ubiquitous 'triangle' scatter plots we see arising from SB remains to be determined.

IV. APPLYING ALGORITHMIC PROBABILITY

A. Systems effectively equivalent to a UTM

Algorithmic probability has been applied or invoked in many different areas of science and mathematics, where implicitly or explicitly the original coding theorem involving a UTM is assumed. No physical computational system can actualise a UTM in the strictest sense, because a UTM requires infinite computational resources, including infinite memory and computation time, while this is not possible in practice. Nonetheless, if a system is at least Turing complete, and if a large memory resource is available which can be expanded as needed, and 'long' computations are permitted, then the system can be considered effectively equivalent to a UTM.

We have seen that for computable maps, low Kolmogorov complexity, low probability (LKLP) outputs are ubiquitous and their existence can be understood from a number of angles. While the original coding theorem in Eq. (2) states a direct connection between Kolmogorov complexity and probability, in applications of algorithmic probability to real world scenarios, if it is known that the system under consideration is not effectively equivalent to a UTM, then LKLP outputs should be expected.

Additionally, even if a system has this effective equivalence property and hence has the ability to implement arbitrary algorithms, another issue is whether or not it is reasonable to assume that the system is processing purely random input programs (hence any conceivable algorithm might be implemented), or instead is merely implementing some random inputs to some computable functions for which LKLP outputs might be expected. (Indeed, whether the input programs being processed are reasonably modelled as random at all is another question.) For example, a desktop computer with a Turing complete language compiler like Python is effectively equivalent to a UTM, but if the computer is used to run RNA structure prediction with random RNA sequences, then we can still expect that the outputs we are observing will show LKLP behaviour like in Figure 1(c). In this scenario, the UTM-equivalent computational capacity of the underlying system (ie the desktop computer) does not preclude LKLP outputs, and is in a sense irrelevant. Even though the computer is effectively a UTM, we do not expect to see patterns like the algorithmically simple digits of π (or perhaps even some non-pseudo random patterns) appearing with high probability, because the programs sampled are not relevant to that type of pattern.

B. Two questions for physical systems

The preceding discussion suggests two questions in reference to when we expect LKLP outputs: Is the physical system of interest likely to be effectively equivalent to a UTM? If effectively UTM equivalent, is it reasonable to assume that the physical system of interest is running random programs for which all manner of algorithms might be computed, or

instead computing only a small fraction of computable algorithms?

Regarding the first question, the physical universe does support universal computation, because it can (and is) performed within the universe. Moreover, Wolfram (Chapter 12 of [34]) has proposed the Principle of Computational Equivalence which states that systems in nature which are not obviously simple, eg the weather, have maximally possible computational power, implying that many or even most natural systems operate effectively equivalent to a UTM. Other work on undecidability in physics [35–38] and the computational capacity of physical world [39–41] may tend to support the possibility of high-level computation in the natural world. Despite these points, the Principle has not been proven to hold and it is not clear that it does actually hold very commonly in nature.

Regarding the second question, there are many examples in nature of maps which are clearly not performing arbitrary algorithms, but instead only performing a narrow set of computable functions. In biology, the mapping from DNA sequence programs to biological forms can often be modelled with computable mathematical functions, and hence we can expect LKLP behaviour. Another example: the time series pattern describing the daily temperature in London, UK, over the year has been ‘computed’ from a variety of input factors such as global air currents and cloud movements. But this computation is described by computable mathematical functions (and some uncomputable random noise), and hence we can expect LKLP behaviour if there is any SB in this system. It may be that the weather system could in principle perform arbitrary computations, but it is far from clear that it is actually doing this, and hence LKLP behaviour might be expected.

Because computable maps abound in nature, and UTM-equivalent systems are not obviously very common, and even if they are, do not commonly implement the full range of possible algorithms, we suggest that LKLP behaviour should be the default assumption in applications of algorithmic probability in real-world settings. We now highlight a selection of applications of algorithmic probability and see how or if LKLP behaviour is relevant.

C. Explaining symmetry and simplicity in biology

It has recently been argued [27] that a significant factor causing the symmetry and simplicity in many biological forms (eg in large symmetric biomolecules, petal arrangements in flowers) is due to simplicity bias described by the upper bound in Eq. (3). This application of algorithmic probability is unproblematic, because even in the presence of LKLP outputs, there remains a bias towards simpler outputs.

D. Time series prediction

The authors of ref. [8] invoked algorithmic probability and the upper bound in Eq. (3) to make predictions about natural time series data taken from the World Bank Open Data project (the example time series data set described above

was taken from their study). They proposed to make a kind of ‘forecasting without historical data’. In this prediction context, LKLP outputs do cause a problem for prediction, because although the predicted upper bound closely matches the upper bound of the data, the presence of LKLP outputs imply that many series extrapolations will be predicted to have high probability (on account of being simple) but in fact have low probability. This amounts to a significant weakness in the prediction due to LKLP outputs. A similar challenge to prediction accuracy may also affect the prediction-by-compression of time series studied by Ryabko et al [42].

E. What will I see next?

Müller [43] has recently proposed a novel application of algorithmic probability as a framework for addressing fundamental problems in theoretical physics, such as why the universe obeys simple laws. He points out that many problems in science and philosophy reduce to the question “*What will I see next?*”, and suggests that this single unified approach to framing questions about the world also has a single unified approach to answering them in all contexts uniformly, namely (conditional) algorithmic probability. To see how, we briefly recap some relevant theory [44]: the continuous universal a priori probability $M(x)$ is defined as

$$M(x) = \sum_{p: V(p)=x*} 2^{-|p|} \quad (4)$$

where p is a program that runs on some UTM V and produces an output string which starts with x and then continues (perhaps never halting), denoted by $x*$. So $M(x)$ is the probability that an infinite string begins with x . Now, $M(y|x) := M(xy)/M(x)$ provides a general method for predicting the extrapolation y given a history x . Further, Solomonoff showed that even if μ is some computable (eg real world) probability distribution over binary strings, then the conditional probability prediction problem $\mu(y|x)$ can be estimated by $M(y|x)$, with μ -probability 1 if $|x| \rightarrow \infty$ and μ is low complexity computable function [45]. We point out that predictions based on $M(y|x)$ may be very inaccurate, ie $M(y|x) \gg \mu(y|x)$, if y may be a LKLP pattern, however.

Returning to the unified approach to framing questions about the world, Müller takes as a starting point the assumption that there is only the state of the observer (which is not fundamentally embedded into anything), and then postulates that what happens next to that observer is dictated by algorithmic probability. He goes on to show that this looks to the observer as if he/she was embedded in some computable probabilistic world. Finally, assuming there is an external world with computable laws within which the observer is embedded, then it is proposed that algorithmic probability gives approximately correct predictions for what the observer sees next, as indicated by the close quantitative relation between $\mu(y|x)$ and $M(y|x)$.

Do LKLP outputs impact the stated goal to address a unified question in a unified manner? Given that algorithmic probability was directly postulated [43] to be the best way to approach the question of “*What will I see next?*”, it is perhaps inappropriate to consider whether it is reasonable

to assume the presence of a UTM. Nonetheless it is worth noting that LKLP outputs may be relevant if the manner in which the state of the observer is updated is not in fact reasonably modelled as the result of (random) programs fed into a system which is effectively equivalent to a UTM, but instead as outputs from computable maps.

Müller also suggested that algorithmic probability helps to explain why we see a “simple” world and physics laws. Even with LKLP behaviour, this argument would still hold, because the simplicity bias upper bound still favours simple (compressible) outcomes and extrapolations of historical patterns.

Somewhat related to the preceding, Lloyd [46] has argued that quantum fluctuations act a random programs which are computed by the universe, to produce complexity in the universe. Lloyd’s argument (which invokes AIT) explicitly proposes that the universe acts as a computer; it would be interesting to consider the implications (if any) for LKLP outputs this perspective on the physical universe.

F. Universal gambling

A universal gambling scheme which considers specific individual outcomes (and explicitly was based on Kolmogorov complexity) was introduced by Cover [47] in 1974, in which it was suggested that an investment portfolio should be constructed while respecting probability predictions essentially similar to algorithmic probability. In more detail, an investor having observed a binary string financial time series $x_1x_2\dots x_i$ for some index i , in predicting whether the next bit will be a 0, the formula Cover suggested was $p(0) = 2^{l(x_1x_2\dots x_i0)} / 2^{l(x_1x_2\dots x_i)}$, and $p(1)$ is similarly defined. In this formula $l(z)$ measures the minimum codelength for some string z , which is fundamentally the same notion as the Kolmogorov complexity of z , which appears in the coding theorem.

It seems reasonable to assume that the natural context for applying this universal gambling strategy, that is time series in financial markets, are not the result of UTMs but rather some computable processes. Hence it seems likely that LKLP behaviour would be a challenge for the strategy, due to the possibility that for example $x_1x_2\dots x_i0$ might be simpler than $x_1x_2\dots x_i1$ but due to LKLP behaviour, the latter is more likely. We are not saying that the gambling scheme is invalid, just that its efficacy would be reduced somewhat in the presence of LKLP outputs.

G. Solomonoff induction and Occam’s razor

Occam’s razor is a fundamental principle of scientific reasoning, philosophy, and model selection [25] stating that simpler explanations or models should take preference over more complex ones [48]. Despite its wide application and common sense appeal, a formal grounding for this principle has been a challenge for philosophy. Solomonoff introduced the idea of algorithmic probability in the 1960’s [4, 22] as part of an investigation into induction, and it has been argued that his formal method of induction has solved the long-standing philosophical problem of why simpler explanations

should be preferred [49–51]. Solomonoff’s basic argument is that because simple hypotheses/explanations are a priori more likely to appear from a random program running on a UTM, then given some observed data the simplest hypothesis/explanation should have the highest probability and hence should be preferred (assuming it explains some observed data as well as another competing hypothesis).

Does LKLP behaviour affect the applicability of Solomonoff induction to the real world? If we assume that the contexts within which induction is to be made — presumably the physical world — result from random programs fed into a UTM, then LKLP do not even exist. However, as discussed above it seems reasonable to assume that the physical world is the result of computable physical laws and therefore we can expect (many) LKLP outputs. Solomonoff induction is premised on the fact that a given simple pattern (hypothesis) is a priori more likely than some other given more complex one, and interestingly this property still holds albeit in a weaker sense, even with LKLP outputs.

Recall that we saw above (Section IIIB) when predicting whether x or y is more likely, that for both probability-weighted and uniform sampling simpler strings tend to have higher probability. However, because uniform sampling the prediction accuracy is low, eg 63% for uniform sampling with the FST, this suggests that while Solomonoff induction is still valid in LKLP settings, it is less likely to lead to correct induction as compared to in the UTM setting for which complexity and probability is much more closely connected. Therefore a challenge to applying Solomonoff induction in the real world is faced because when weighing up two hypotheses H_1 and H_2 that both explain the observed data, H_1 may be simpler than H_2 , but perhaps H_1 is a LKLP output and hence much less likely than H_2 . By extension, there is a challenge to justifying Occam’s razor in the real physical world — it is still valid even with LKLP, but the argument for preferring simpler hypotheses is somewhat weakened.

Relatedly, Hutter has formed a universal theory of artificial intelligence based on algorithmic probability and Solomonoff induction [52] (see also Legg and Veness [14] for a numerical implementation of the theory). The implications of LKLP outputs for this research project will be similar to those for Occam’s razor, namely if the intelligent agent is making observations resulting from an environment that is known to be generated by a computable map, then it is still true that often a given simpler hypothesis is a priori more likely than some given complex hypothesis, but not that rarely the reverse is true. Hence in the computable setting it may be that this form of induction is less likely to lead to correct predictions as compared to the UTM setting.

H. DNN generalisation

Another recent and important application of simplicity bias is in machine learning, where it has been argued [28] that the surprising generalisation ability of deep neural networks (DNN) is due in part to the fact that DNN are biased towards simple functions (by invoking Eq. (3)), and natural functions are also biased towards simple functions, so the problem of learning functions is significantly easier than it would be if there was no simplicity bias in functions. Be-

cause LKLP behaviour is also found in DNN [28], and LKLP has been observed in so many other natural settings, this raises a question: does LKLP behaviour create a challenge for this invocation of algorithmic probability? One way to look at this is to ask, do the simple functions in nature coincide with the simple functions towards which the DNN are biased? Without LKLP outputs in either DNN or nature, they would automatically coincide, but this is no longer automatic given that LKLP outputs exist. If the natural functions are highly probable simple functions but different to the highly probable functions that DNN produce (or at least not a subset of these functions) then the argument for why DNN generalise is weaker. On the other hand, if highly probable natural functions do coincide (or are at least a subset) of the highly probable functions generated by DNN then it would be interesting to consider why this is. This question relates to our earlier question regarding whether there might be common patterns to which types of outputs are LKLP: if different maps or systems have completely unrelated LKLP patterns, then an overlap between different systems is less likely as compared to if there are more general typical patterns across different systems.

I. Password guessing

Although not based on Kolmogorov complexity and Levin's Eq. (2), an essentially very similar probability prediction method has been derived by workers in information theory. Merhav and Feder [53] point out in an influential review of *universal prediction* that $2^{-LZ(x)}$ is a universal probability assignment for prediction, citing the work of ref. [54] and others. In this context, $LZ(x)$ is the Lempel-Ziv compression complexity measure, which is very similar to the measure used in most studies of SB. Merhav and Cohen [55] have recently used this $2^{-LZ(x)}$ universal probability predictor in a cryptography setting, where they suggested that it forms an optimal method for guessing passwords. It would be interesting to investigate any examples of LKLP passwords. If they did occur, then they might represent a source of wasted guesses of the guessing strategy, because the strategy would assign a high probability of $\sim 2^{-LZ(x)}$ to some password x , which is in fact very rarely used as a password by people.

V. DISCUSSION

We have investigated the occurrence of low Kolmogorov complexity, low probability (LKLP) outputs in computable functions with randomly sampled inputs. The central messages are that (a) LKLP outputs have been observed in essentially all maps for which simplicity bias (SB) has been studied; (b) LKLP outputs are expected in computable maps for theoretical reasons; (c) there appears to be some common statistical patterns to the distributions of the LKLP outputs which form a kind of 'triangle' shape in probability-complexity plots; and (d) when applying algorithmic probability in real world applications, LKLP outputs should be the default expectation unless there is good reason to expect that the outputs are indeed generated by purely ran-

dom programs fed into a universal Turing machine (UTM), which we suggest is probably an uncommon scenario in science, engineering, finance, etc. Further, we briefly surveyed some works in which algorithmic probability in some form has been invoked, and discussed some possible implications of LKLP outputs, including studies of a priori predictions, Solomonoff induction and Occam's razor. The main LKLP implication is that the connection between complexity and probability is considerably less strong, as compared to in the original algorithmic information theory (AIT) coding theorem.

A main motivation for this study is developing theory for improved a priori probability predictions. The AIT coding theorem states that output probabilities can be directly found via the Kolmogorov complexity of the output, rather than, say, using historical frequency statistics to estimate probabilities. In earlier work [9] a practical weaker version of this AIT coding theorem was presented, in the form of an upper bound on probabilities, rather than a direct estimate of probabilities. Understanding the causes and nature of LKLP outputs which fall far below the upper bound may help to improve predictions of the probabilities of those outputs, and hence a stronger theory of a priori probability predictions may be laid out.

Although algorithmic probability was originally formulated in the context of random algorithms/programs that generate outputs via a computer, we stress that algorithmic probability estimates — especially the upper bound of Eq. (3) — are not limited to what are usually understood as algorithms per se. Instead these estimates can be applied to a wide range of problems for which output patterns result from mathematical functions with some form of input parameters. For example, the parameters of a large ordinary differential equation system are not usually understood as an 'algorithm' for the solution profile, nonetheless the upper bound has been shown to predict the probability of outputs in such systems [9, 27]. Even more distant from a computer running a program, the upper bound has been shown to work well in predicting natural time series patterns [8] for which both the notion of program and computer is much less clearly defined.

There has been a lot of discussion in the statistics and philosophy communities regarding how to choose a Bayesian prior, and AIT promises to be one way to address this [56] (but see also ref. [57] for a critique of this approach). Our work here is directly relevant to making practical implementations of this AIT answer to the Bayesian prior question.

Several open questions remain. In general, how can we better understand the causes and nature of LKLP outputs? Given a small sample of outputs, can these be used to predict which outputs are likely to be LKLP, and which are likely to be close to the bound? Why is it that the data points form 'triangles' in Figure 1 and other simplicity bias studies? Are there common patterns across different systems dictating which outputs will be LKLP? How can we best incorporate simplicity bias probability predictions into other probability estimation approaches, such as machine learning?

Acknowledgements: KD acknowledges financial support from the Gulf University for Science and Technology Seed Grant (grant number 234271). We thank Paris Flood, Iain Johnston, Ard Louis, Nora Martin, Christopher Mingard, and Markus Müller for valuable discussions and suggestions related to this work.

Data availability: The data sets generated during and analysed during the current study are available from the corresponding authors on request.

Author contributions: MA performed the numerical calculations. KD conceived the study and wrote the paper.

- [1] R. J. Solomonoff. A preliminary report on a general theory of inductive inference (revision of report v-131). *Contract AF*, 49(639):376, 1960.
- [2] A.N. Kolmogorov. Three approaches to the quantitative definition of information. *Problems of information transmission*, 1(1):1–7, 1965.
- [3] Gregory J Chaitin. A theory of program size formally identical to information theory. *Journal of the ACM (JACM)*, 22(3):329–340, 1975.
- [4] M. Li and P.M.B. Vitányi. *An introduction to Kolmogorov complexity and its applications*. Springer-Verlag New York Inc, 2008.
- [5] L.A. Levin. Laws of information conservation (nongrowth) and aspects of the foundation of probability theory. *Problemy Peredachi Informatsii*, 10(3):30–35, 1974.
- [6] Ray J Solomonoff. The kolmogorov lecture the universal distribution and machine learning. *The Computer Journal*, 46(6):598–601, 2003.
- [7] Kamaludin Dingle, Guillermo Valle Pérez, and Ard A Louis. Generic predictions of output probability based on complexities of inputs and outputs. *Scientific reports*, 10(1):1–9, 2020.
- [8] Kamaludin Dingle, Rafiq Kamal, and Boumediene Hamzi. A note on a priori forecasting and simplicity bias in time series. *arXiv preprint arXiv:2203.05391*, 2022.
- [9] Kamaludin Dingle, Chico Q Camargo, and Ard A Louis. Input-output maps are strongly biased towards simple outputs. *Nature communications*, 9(1):761, 2018.
- [10] R. Cilibrasi and P.M.B. Vitányi. Clustering by compression. *Information Theory, IEEE Transactions on*, 51(4):1523–1545, 2005.
- [11] P. Ferragina, R. Giancarlo, V. Greco, G. Manzini, and G. Valiente. Compression-based classification of biological sequences and structures via the universal similarity metric: experimental assessment. *BMC bioinformatics*, 8(1):252, 2007.
- [12] Ram Avinery, Micha Kornreich, and Roy Beck. Universal and accessible entropy estimation using a compression algorithm. *Physical review letters*, 123(17):178102, 2019.
- [13] Paul MB Vitányi. Similarity and denoising. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 371(1984), 2013.
- [14] Shane Legg and Joel Veness. An approximation of the universal intelligence measure. *arXiv preprint arXiv:1109.5951*, 2011.
- [15] J.P. Delahaye and H. Zenil. Numerical evaluation of algorithmic complexity for short strings: A glance into the innermost structure of algorithmic randomness. *Appl. Math. Comput.*, 219:63–77, 2012.
- [16] Fernando Soler-Toscano, Hector Zenil, Jean-Paul Delahaye, and Nicolas Gauvrit. Calculating Kolmogorov complexity from the output frequency distributions of small Turing machines. *PloS one*, 9(5):e96223, 2014.
- [17] Hector Zenil, Liliana Badillo, Santiago Hernández-Orozco, and Francisco Hernández-Quiroz. Coding-theorem like behaviour and emergence of the universal distribution from resource-bounded algorithmic probability. *International Journal of Parallel, Emergent and Distributed Systems*, 34(2):161–180, 2019.
- [18] C.S. Calude. *Information and randomness: An algorithmic perspective*. Springer, 2002.
- [19] P. Gács. *Lecture notes on descriptive complexity and randomness*. Boston University, Graduate School of Arts and Sciences, Computer Science Department, 1988.
- [20] Alan Mathison Turing. On computable numbers, with an application to the entscheidungsproblem. *J. of Math*, 58(345-363):5, 1936.
- [21] Peter Grunwald and Paul Vitányi. Shannon information and Kolmogorov complexity. *arXiv preprint cs/0410002*, 2004.
- [22] Ray J Solomonoff. Algorithmic probability: Theory and applications. In *Information theory and statistical learning*, pages 1–23. Springer, 2009.
- [23] A. Lempel and J. Ziv. On the complexity of finite sequences. *Information Theory, IEEE Transactions on*, 22(1):75–81, 1976.
- [24] Ronny Lorenz, Stephan H Bernhart, Christian Höner Zu Siederdisen, Hakim Tafer, Christoph Flamm, Peter F Stadler, and Ivo L Hofacker. Viennarna package 2.0. *Algorithms for molecular biology*, 6(1):26, 2011.
- [25] Mark H Hansen and Bin Yu. Model selection and the principle of minimum description length. *Journal of the American Statistical Association*, 96(454):746–774, 2001.
- [26] K. Willbrand, F. Radvanyi, J.P. Nadal, J.P. Thiery, and T.M.A. Fink. Identifying genes from up-down properties of microarray expression series. *Bioinformatics*, 21(20):3859–3864, 2005.
- [27] Iain G Johnston, Kamaludin Dingle, Sam F Greenbury, Chico Q Camargo, Jonathan PK Doye, Sebastian E Ahnert, and Ard A Louis. Symmetry and simplicity spontaneously emerge from the algorithmic nature of evolution. *Proceedings of the National Academy of Sciences*, 119(11):e2113883119, 2022.
- [28] Guillermo Valle-Perez, Chico Q Camargo, and Ard A Louis. Deep learning generalizes because the parameter-function map is biased towards simple functions. *arXiv preprint arXiv:1805.08522*, 2018.
- [29] Kamaludin Dingle. *Probabilistic bias in genotype-phenotype maps*. PhD thesis, University of Oxford, 2014.
- [30] Jan Leike and Marcus Hutter. Bad universal priors and notions of optimality. In *Conference on Learning Theory*, pages 1244–1259. PMLR, 2015.
- [31] Markus Müller. Stationary algorithmic probability. *Theoretical Computer Science*, 411(1):113–130, 2010.
- [32] SF Greenbury and Sebastian Edmund Ahnert. The organization of biological sequences into constrained and unconstrained parts determines fundamental properties of genotype-phenotype maps. *Journal of The Royal Society Interface*, 12(113):20150724, 2015.
- [33] Jędrzej Burkat and Kamaludin Dingle. A note on simplicity bias in the bernoulli process. *In preparation*, 2022.
- [34] S. Wolfram. *A New Kind of Science*. Wolfram Media, 2002.
- [35] Christopher Moore. Unpredictability and undecidability in dynamical systems. *Physical Review Letters*, 64(20):2354,

- 1990.
- [36] Karl Svozil. *Randomness & undecidability in physics*. World Scientific, 1993.
 - [37] Toby S Cubitt, David Perez-Garcia, and Michael M Wolf. Undecidability of the spectral gap. *Nature*, 528(7581):207–211, 2015.
 - [38] Anthony Aguirre, Zeeya Merali, and David Sloan. Undecidability, uncomputability, and unpredictability, 2021.
 - [39] Scott Aaronson. Guest column: Np-complete problems and physical reality. *ACM Sigact News*, 36(1):30–52, 2005.
 - [40] Hector Zenil. *A computable universe: understanding and exploring nature as computation*. World Scientific, 2013.
 - [41] Seth Lloyd. *Programming the universe: a quantum computer scientist takes on the cosmos*. Vintage, 2006.
 - [42] Boris Ryabko, Jaakko Astola, and Mikhail Malyutov. *Compression-based methods of statistical analysis and prediction of time series*. Springer, 2016.
 - [43] Markus P Müller. Law without law: from observer states to physics via algorithmic information theory. *Quantum*, 4:301, 2020.
 - [44] Marcus Hutter, Shane Legg, and Paul MB Vitanyi. Algorithmic probability. *Scholarpedia*, 2(8):2572, 2007.
 - [45] Ray Solomonoff. Complexity-based induction systems: comparisons and convergence theorems. *IEEE transactions on Information Theory*, 24(4):422–432, 1978.
 - [46] Seth Lloyd. The computational universe. *Information and the nature of reality: From physics to metaphysics*, pages 92–103, 2010.
 - [47] TM Cover. Universal gambling schemes and the complexity measures of kolmogorov and chaitin. rep. no. 12, statistics dep, 1974.
 - [48] Alan Baker. Simplicity. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Winter 2016 edition, 2016.
 - [49] Aron Vallinder. Solomonoff induction: A solution to the problem of the priors? 2012.
 - [50] WW Kirchherr, M Li, PMB Vitanyi, et al. The miraculous universal distribution. *Mathematical Intelligencer*, 1997.
 - [51] Samuel Rathmanner and Marcus Hutter. A philosophical treatise of universal induction. *Entropy*, 13(6):1076–1136, 2011.
 - [52] Marcus Hutter. *Universal artificial intelligence: Sequential decisions based on algorithmic probability*. Springer Science & Business Media, 2004.
 - [53] Neri Merhav and Meir Feder. Universal prediction. *IEEE Transactions on Information Theory*, 44(6):2124–2147, 1998.
 - [54] Eli Plotnik, Marcelo J Weinberger, and Jacob Ziv. Upper bounds on the probability of sequences emitted by finite-state sources and on the redundancy of the lempel-ziv algorithm. *IEEE transactions on information theory*, 38(1):66–72, 1992.
 - [55] Neri Merhav and Asaf Cohen. Universal randomized guessing with application to asynchronous decentralized brute-force attacks. *IEEE Transactions on Information Theory*, 66(1):114–129, 2019.
 - [56] Marcus Hutter. On universal prediction and bayesian confirmation. *Theoretical Computer Science*, 384(1):33–48, 2007.
 - [57] Sven Neth. A dilemma for solomonoff prediction. *Philosophy of Science*, pages 1–25, 2022.