



ELSEVIER

Physica A 302 (2001) 89–99

PHYSICA A

www.elsevier.com/locate/physa

Complexity through nonextensivity

William Bialek^{a*,1}, Ilya Nemenman^{a,2}, Naftali Tishby^{a,b}

^a*NEC Research Institute, 4 Independence Way, Princeton, NJ 08540, USA*

^b*School of Computer Science and Engineering, and Center for Neural Computation, Hebrew University, Jerusalem 91904, Israel*

Abstract

The problem of defining and studying complexity of a time series has interested people for years. In the context of dynamical systems, Grassberger has suggested that a slow approach of the entropy to its extensive asymptotic limit is a sign of complexity. We investigate this idea further by information theoretic and statistical mechanics techniques and show that these arguments can be made precise, and that they generalize many previous approaches to complexity, in particular, unifying ideas from the physics literature with ideas from learning and coding theory; there are even connections of this statistical approach to algorithmic or Kolmogorov complexity. Moreover, a set of simple axioms similar to those used by Shannon in his development of information theory allows us to prove that the divergent part of the subextensive component of the entropy is a unique complexity measure. We classify time series by their complexities and demonstrate that beyond the “logarithmic” complexity classes widely anticipated in the literature there are qualitatively more complex, “power-law” classes which deserve more attention. © 2001 Published by Elsevier Science B.V.

The problem of quantifying complexity is very old. Interest in the field has been fueled by three sorts of questions. First, one would like to make precise the impression that some systems, such as life on earth or a turbulent fluid flow, evolve toward a state of higher complexity, and one would like to classify these states; this is the realm of dynamical systems theory. Second, in choosing among different models that describe an experiment, one wants to quantify a preference for simpler explanations or, equivalently, provide a penalty for complex models that can be weighed against the more conventional goodness-of-fit criteria; this type of question usually is investigated in statistics. Finally, there are questions about how hard it is to compute or to describe

* Corresponding author.

¹ Present address: Department of Physics, Princeton University, Princeton, NJ 08544, USA.

² Present address: Institute for Theoretical Physics, University of California at Santa Barbara, Santa Barbara, CA 93105, USA.

the state of a complex system; this is the area of formal mathematics and computer science.

Research in each of these three directions has given birth to numerous definitions of complexity. The usual objective is to make these definitions focused enough to be operational in particular contexts but general enough to connect with our intuitive notions. For many years, the dominant candidate for a universal measure has been the mathematically rigorous notion of *Kolmogorov* or *algorithmic complexity* that measures (roughly) the minimum length of a computer program that can recreate the observed time series [1]. Unfortunately, there is no algorithm that can calculate the Kolmogorov complexity of all data sets. Therefore, for applications to statistics, Rissanen [2,3] and others have developed a new concept: *stochastic complexity* of the data with respect to a particular class of models, which measures the shortest total description of the data and the model within the class, but cannot rule out the possibility that a different model class could generate a shorter code.

The main difficulty of all these approaches is that the Kolmogorov complexity is closely related to the Shannon entropy, which means that it measures something closer to our intuitive concept of randomness than to the intuitive concept of complexity [4]. A true random string cannot be compressed and hence requires a long description, yet the physical process that generates this string may be very simple. As physicists, our notions of complexity correspond to statements about the underlying process, and not directly to the description length or Kolmogorov complexity: a dynamics with a predictable constant output (small algorithmic complexity) is as trivial as one for which the output is completely unpredictable and random (large algorithmic complexity), while really complex processes lie somewhere in between.

The two extreme cases, however, have one feature in common: the entropy of the output strings (or, equivalently, the Kolmogorov complexity of a typical one) either is a fixed constant or grows exactly linearly with the length of the strings. In both cases, corrections to the asymptotic behavior *do not grow with the size of the data set*. This allowed Grassberger [5] to identify the slow approach of the entropy to its extensive limit as a sign of complexity. He has proposed several functions to analyze this slow approach and studied systems that exhibited a broad range of complexity properties.

To deal with the same problem, Rissanen has emphasized strongly that fitting a model to data represents an encoding of those data, or predicting future data. Shorter encodings generally mean better prediction or generalization. However, much of the code usually describes the meaningless, nongeneralizable “noise”—statistical fluctuations within the model. Only model description is relevant to prediction, and this part of the code has been termed the *model complexity* [2,3]. While systems with model complexity of very different types are known, the two extreme examples above are similar; it only takes a fixed number of bits to code either a call to a random number generator or to a constant function.

The present work may be viewed as expanding on the notions of subextensivity and effective prediction. We construct a coherent theory that brings these ideas together in an intuitive way, but nonetheless is sufficiently general to be applied in many different

contexts. We will show that with only a little bit of work, Grassberger’s definitions may be made as mathematically precise as they are aesthetically pleasing. Finally, we will argue that the definitions are unique if one accepts a set of simple axioms in the spirit of Shannon’s original work, and that these definitions relate to the usual Kolmogorov complexity in a straightforward way. Much of this paper follows closely a more detailed analysis in Ref. [6], to which we refer for calculation details and a thorough discussion of the (large) relevant literature.

Our path to connecting the various complexity measures begins by noticing that the subextensive components of entropy identified by Grassberger in fact determine the information available for making predictions. This also suggests a connection to the importance or value of information, especially in a biological or economic context: information is valuable if it can be used to guide our actions, but actions take time and hence observed data can be useful only to the extent that those data inform us about the state of the world at later times. It would be attractive if what we identify as “complex” in a time series were also the “useful” or “meaningful” components.

While prediction may come in various forms, depending on context, information theory allows us to treat all of them on the same footing. For this we only need to recognize that all predictions are probabilistic, and that, even before we look at the data, we know that certain futures are more likely than others. This knowledge can be summarized by a prior probability distribution for the futures. Our observations on the past lead us to a new, more tightly concentrated distribution, the distribution of futures conditional on the past data. Different kinds of predictions are different slices through or averages over this conditional distribution, but information theory quantifies the “concentration” of the distribution without making any commitment as to which averages will be most interesting.

Imagine that we observe a stream of data $x(t)$ over a time interval $-T < t < 0$; let all of these past data be denoted by the shorthand x_{past} . We are interested in saying something about the future, so we want to know about the data $x(t)$ that will be observed in the time interval $0 < t < T'$; let these future data be called x_{future} . In the absence of any other knowledge, futures are drawn from the probability distribution $P(x_{\text{future}})$, while observations of particular past data x_{past} tell us that futures will be drawn from the conditional distribution $P(x_{\text{future}}|x_{\text{past}})$. The greater concentration of the conditional distribution can be quantified by the fact that it has smaller entropy than the prior distribution, and this reduction in entropy is Shannon’s definition of the information that the past provides about the future. We can write the average of this *predictive information* as

$$\mathcal{I}_{\text{pred}}(T, T') = \left\langle \log_2 \left[\frac{P(x_{\text{future}}|x_{\text{past}})}{P(x_{\text{future}})} \right] \right\rangle \quad (1)$$

$$= - \langle \log_2 P(x_{\text{future}}) \rangle - \langle \log_2 P(x_{\text{past}}) \rangle - [- \langle \log_2 P(x_{\text{future}}, x_{\text{past}}) \rangle], \quad (2)$$

where $\langle \dots \rangle$ denotes an average over the joint distribution of the past and the future, $P(x_{\text{future}}, x_{\text{past}})$.

Each of the terms in Eq. (2) is an entropy. Since we are interested in predictability or generalization, which are associated with some features of the signal persisting forever, we may assume stationarity or invariance under time translations. Then the entropy of the past data depends only on the duration of our observations, so we can write $-\langle \log_2 P(x_{\text{past}}) \rangle = S(T)$, and by the same argument $-\langle \log_2 P(x_{\text{future}}) \rangle = S(T')$. Finally, the entropy of the past and the future taken together is the entropy of observations on a window of duration $T + T'$, so that $-\langle \log_2 P(x_{\text{future}}, x_{\text{past}}) \rangle = S(T + T')$. Putting these equations together, we obtain

$$\mathcal{I}_{\text{pred}}(T, T') = S(T) + S(T') - S(T + T'). \quad (3)$$

In the same way that the entropy of a gas at fixed density is proportional to the volume, the entropy of a time series (asymptotically) is proportional to its duration, so that $\lim_{T \rightarrow \infty} S(T)/T = \mathcal{S}_0$; entropy is an extensive quantity. But from Eq. (3) any extensive component of the entropy cancels in the computation of the predictive information: *predictability is a deviation from extensivity*. If we write

$$S(T) = \mathcal{S}_0 T + S_1(T), \quad (4)$$

then Eq. (3) tells us that the predictive information is related *only* to the nonextensive term $S_1(T)$.

We know two general facts about the behavior of $S_1(T)$. First, the corrections to extensive behavior are positive, $S_1(T) \geq 0$. Second, the statement that entropy is extensive is the statement that the limit $\lim_{T \rightarrow \infty} S(T)/T = \mathcal{S}_0$ exists, and for this to be true we must also have $\lim_{T \rightarrow \infty} S_1(T)/T = 0$. Thus, the nonextensive terms in the entropy must be *subextensive*, that is they must grow with T less rapidly than a linear function. Taken together, these facts guarantee that the predictive information is positive and subextensive. Further, if we let the future extend forward for a very long time, $T' \rightarrow \infty$, then we can measure the information that our sample provides about the entire future,

$$I_{\text{pred}}(T) = \lim_{T' \rightarrow \infty} \mathcal{I}_{\text{pred}}(T, T') = S_1(T), \quad (5)$$

and this is precisely equal to the subextensive entropy.

If we have been observing a time series for a (long) time T , then the total amount of data we have collected in is measured by the entropy $S(T)$, and at large T this is given approximately by $\mathcal{S}_0 T$. But the predictive information that we have gathered cannot grow linearly with time, even if we are making predictions about a future which stretches out to infinity. As a result, of the total information we have taken in by observing x_{past} , only a vanishing fraction is of relevance to the prediction:

$$\lim_{T \rightarrow \infty} \frac{\text{Predictive information}}{\text{Total information}} = \frac{I_{\text{pred}}(T)}{S(T)} \rightarrow 0. \quad (6)$$

In this precise sense, most of what we observe is irrelevant to the problem of predicting the future. Since the average Kolmogorov complexity of a time series is related to its (total) Shannon entropy, this result means also that most of the algorithm that is required to encode the data encodes aspects of the data that are useless for prediction or for guiding our actions based on the data. This is a strong indication that the

usual notions of Kolmogorov complexity in fact do not capture anything at all like the (intuitive) utility of the data stream.

Consider the case where time is measured in discrete steps, so that we have seen N time points x_1, x_2, \dots, x_N . How much is there to learn about the underlying pattern in these data? In the limit of large number of observations, $N \rightarrow \infty$ or $T \rightarrow \infty$, the answer to this question is surprisingly universal: predictive information may either stay finite, or grow to infinity together with T ; in the latter case the rate of growth may be slow (logarithmic) or fast (sublinear power).

The first possibility, $\lim_{T \rightarrow \infty} I_{\text{pred}}(T) = \text{constant}$, means that no matter how long we observe we gain only a finite amount of information about the future. This situation prevails in both extreme cases mentioned above. For example, when the dynamics are too regular, as for a purely periodic system, complete prediction is possible once we know the phase, and if we sample the data at discrete times this is a finite amount of information; longer period orbits intuitively are more complex and also have larger I_{pred} , but this does not change the limiting behavior $\lim_{T \rightarrow \infty} I_{\text{pred}}(T) = \text{constant}$.

Similarly, the predictive information can be small when the dynamics are irregular but the best predictions are controlled only by the immediate past, so that the correlation times of the observable data are finite. This happens, for example, in many physical systems far away from phase transitions. Imagine, for example, that we observe $x(t)$ at a series of discrete times $\{t_n\}$, and that at each time point we find the value x_n . Then we always can write the joint distribution of the N data points as a product,

$$P(x_1, x_2, \dots, x_N) = P(x_1)P(x_2|x_1)P(x_3|x_2, x_1) \cdots . \tag{7}$$

For Markov processes, what we observe at t_n depends only on events at the previous time step t_{n-1} , so that

$$P(x_n | \{x_1 \leq i \leq n-1\}) = P(x_n | x_{n-1}) \tag{8}$$

and hence the predictive information reduces to

$$I_{\text{pred}} = \left\langle \log_2 \left[\frac{P(x_n | x_{n-1})}{P(x_n)} \right] \right\rangle . \tag{9}$$

The maximum possible predictive information in this case is the entropy of the distribution of states at one time step, which in turn is bounded by the logarithm of the number of accessible states. To approach this bound, the system must maintain memory for a long time, since the predictive information is reduced by the entropy of the transition probabilities. Thus, systems with more states and longer memories have larger values of I_{pred} .

More interesting are those cases in which $I_{\text{pred}}(T)$ diverges at large T . In physical systems we know that there are critical points where correlation times become infinite, so that optimal predictions will be influenced by events in the arbitrarily distant past. Under these conditions, the predictive information can grow without bound as T becomes large; for many systems the divergence is logarithmic, $I_{\text{pred}}(T \rightarrow \infty) \propto \log T$.

Long-range correlation also are important in a time series where we can learn some underlying rules. Suppose a series of random vector variables $\{\vec{x}_i\}$ are drawn

independently from the same probability distribution $Q(\vec{x}|\alpha)$, and this distribution depends on a (potentially infinite dimensional) vector of parameters α . The parameters are unknown, and before the series starts they are chosen randomly from a distribution $\mathcal{P}(\alpha)$. In this setting, at least implicitly, our observations of $\{\vec{x}_i\}$ provide data from which we can learn the parameters α . Here we put aside (for the moment) the usual problem of learning—which might involve constructing some estimation or regression scheme that determines a “best fit” α from the data $\{\vec{x}_i\}$ —and treat the ensemble of data streams $P[\{\vec{x}_i\}]$ as we would any other set of configurations in statistical mechanics or dynamical systems theory. In particular, we can compute the entropy of the distribution $P[\{\vec{x}_i\}]$ even if we cannot provide explicit algorithms for solving the learning problem.

As is shown in Ref. [6], the crucial quantity in such analysis is the density of models in the vicinity of the target $\bar{\alpha}$ —the parameters that actually generated the sequence. For two distributions, a natural distance measure is the Kullback–Leibler divergence $D(\bar{\alpha}|\alpha) = \int d\vec{x} Q(\vec{x}|\bar{\alpha}) \log[Q(\vec{x}|\bar{\alpha})/Q(\vec{x}|\alpha)]$, and the density is

$$\rho(D; \bar{\alpha}) = \int d^K \alpha \mathcal{P}(\alpha) \delta[D - D_{\text{KL}}(\bar{\alpha}|\alpha)]. \quad (10)$$

If ρ is large as $D \rightarrow 0$, then there are many models almost identical to the target one. Therefore, specifying the model is not very informative and cannot be done to a high precision. The predictive information and the complexity of the model specification both are small in this case. On the other hand, if ρ decreases very rapidly near $D=0$, the situation is reversed. It is then clear that the density, but not the number of parameters or any other simplistic measure, characterizes predictability and the complexity of prediction. If, as often is the case for $\dim \alpha < \infty$, the density behaves in the way common to finite-dimensional systems of the usual statistical mechanics,

$$\rho(D \rightarrow 0, \bar{\alpha}) \approx AD^{(K-2)/2}, \quad (11)$$

then the predictive information to the leading order is [6]

$$I_{\text{pred}}(N) \approx K/2 \log N. \quad (12)$$

The modern theory of learning is concerned in large part with quantifying the complexity of a model class, and in particular with replacing a simple count of parameters with a more rigorous notion of dimensionality for the space of models; for a general review of these ideas see Ref. [7], and for discussion close in spirit to ours see Ref. [8]. The important point here is that the dimensionality of the model class, and hence the complexity of the class in the sense of learning theory, emerges as the coefficient of the logarithmic divergence in I_{pred} . Thus, a measure of complexity in learning problems can be derived from a more general dynamical systems or statistical mechanics point of view, treating the data in the learning problem as a time series or one-dimensional lattice. The logarithmic complexity class that we identify as being associated with finite-dimensional models also arises, for example, at the Feigenbaum accumulation point in the period-doubling route to chaos [5].

As noted by Grassberger in his original discussion, there are time series for which the divergence of I_{pred} is stronger than a logarithm. We can construct an example by looking at the density function ρ in our learning problem above: finite-dimensional models are associated with algebraic decay of the density as $D \rightarrow 0$, and we can imagine that there are model classes in which this decay is more rapid, for example

$$\rho(D \rightarrow 0) \approx A \exp[-B/D^\mu], \quad \mu > 0. \quad (13)$$

In this case it can be shown that the predictive information diverges very rapidly, as a sublinear power-law, [6]

$$I_{\text{pred}}(N) \sim N^{\mu/(\mu+1)}. \quad (14)$$

One way that this scenario can arise is if the distribution $Q(\vec{x})$ that we are trying to learn does not belong to any finite parameter family, but is itself drawn from a distribution that enforces a degree of smoothness [9]. Understandably, stronger smoothness constraints have smaller powers μ (less to predict) than the weaker ones (more to predict). For example, a rather simple case of predicting a one-dimensional variable that comes from a continuous distribution produces $I_{\text{pred}}(N) \sim \sqrt{N}$.

As with the logarithmic class, we expect that power-law divergences in I_{pred} are not restricted to the learning problems that we have studied in detail. The general point is that such behavior will be seen in problems where predictability over long scales, rather than being controlled by a fixed set of ever more precisely known parameters, is governed by a progressively more detailed description—effectively increasing the number of parameters—as we collect more data. This seems a plausible description of what happens in language, where rules of spelling allow us to predict forthcoming letters of long words, grammar binds the words together, and compositional unity of the entire text allows us to make predictions about the subject of the last page of the book after reading only the first few. Indeed, Shannon’s classic experiment on the predictability of English text (by human readers!) shows this behavior [10,11], and more recently several groups have extracted power-law subextensive components from the numerical analysis of large corpora of text (see, for example, [12,13]).

Interestingly, even without an explicit example, a simple argument ensures existence of exponential densities and, therefore, power-law predictive information models. If the number of parameters in a learning problem is not finite then in principle it is impossible to predict anything unless there is some appropriate regularization. If we let the number of parameters stay finite but become large, then there is *more* to be learned and correspondingly the predictive information grows in proportion to this number. On the other hand, if the number of parameters becomes infinite without regularization, then the predictive information should go to zero since nothing can be learned. We should be able to see this happen in a regularized problem as the regularization weakens: eventually the regularization would be insufficient and the predictive information would vanish. The only way this can happen is if the predictive information grows more and more rapidly with N as we weaken the regularization, until finally it becomes extensive (equivalently, drops to zero) at the point where prediction becomes impossible. To

realize this scenario we have to go beyond $I_{\text{pred}} \propto \log T$ with $I_{\text{pred}} \propto N^{\mu/(\mu+1)}$; the transition from increasing predictive information to zero occurs as $\mu \rightarrow \infty$.

This discussion makes it clear that the predictive information (the subextensive entropy) distinguishes between problems of intuitively different complexity and thus, in accord to Grassberger's definitions [5], is probably a good choice for a universal complexity measure. Can this be made more precise?

First, we need to decide whether we want to attach measures of complexity to a particular signal $x(t)$ or whether we are interested in measures that are defined by an average over the ensemble $P[x(t)]$. One problem in assigning complexity to single realizations is that there can be atypical data streams. Grassberger [5] in particular has argued that our visual intuition about the complexity of spatial patterns is an ensemble concept, even if the ensemble is only implicit. The fact that we admit probabilistic models is crucial: even at a colloquial level, if we allow for probabilistic models then there is a simple description for a sequence of truly random bits, but if we insist on a deterministic model then it may be very complicated to generate precisely the observed string of bits. Furthermore, in the context of probabilistic models it hardly makes sense to ask for a dynamics that generates a particular data stream; we must ask for dynamics that generate the data with reasonable probability, which is more or less equivalent to asking that the given string be a typical member of the ensemble generated by the model. All of these paths lead us to thinking not about single strings but about ensembles in the tradition of statistical mechanics, and so we shall search for measures of complexity that are averages over the distribution $P[x(t)]$.

Once we focus on average quantities, we can provide an axiomatic proof (much in the spirit of Shannon's [14] arguments establishing entropy as a unique information measure) that links I_{pred} to complexity. We can start by adopting Shannon's postulates as constraints on a measure of complexity: if there are N equally likely signals, then the measure should be monotonic in N ; if the signal is decomposable into statistically independent parts, then the measure should be additive with respect to this decomposition; and if the signal can be described as a leaf on a tree of statistically independent decisions, then the measure should be a weighted sum of the measures at each branching point. We believe that these constraints are as plausible for complexity measures as for information measures, and it is well known from Shannon's original work that this set of constraints leaves the entropy as the only possibility. Since we are discussing a time dependent signal, this entropy depends on the duration of our sample, $S(T)$. We know of course that this cannot be the end of the discussion, because we need to distinguish between randomness (entropy) and complexity. The path to this distinction is to introduce other constraints on our measure.

First, we notice that if the signal x is continuous, then the entropy is not invariant under transformations of x that do not mix point at different times (reparameterizations). It seems reasonable to ask that complexity be a function of the process we are observing and not of the coordinate system in which we choose to record our observations. However, that it is not the whole function $S(T)$ which depends on the coordinate system for x ; it is only the extensive component of the entropy that has this

noninvariance. This can be seen more generally by noting that subextensive terms in the entropy contribute to the mutual information among different segments of the data stream (including the predictive information defined here), while the extensive entropy cannot; mutual information is coordinate invariant, so all of the noninvariance must reside in the extensive term. Thus, any measure complexity that is coordinate invariant must discard the extensive component of the entropy.

If we continue along these lines, we can think about the asymptotic expansion of the entropy at large T . The extensive term is the first term in this series, and we have seen that it must be discarded. What about the other terms? In the context of predicting in a parameterized model, most of the terms in this series depend in detail on our prior distribution in parameter space, which might seem odd for a measure of complexity. More generally, if we consider transformations of the data stream $x(t)$ that mix points within a temporal window of size τ , then for $T \gg \tau$, the entropy $S(T)$ may have subextensive terms which are constant, and these are not invariant under this class of transformations. On the other hand, if there are divergent subextensive terms, these *are* invariant under such temporally local transformations.³ So if we insist that measures of complexity be invariant not only under instantaneous coordinate transformations, but also under temporally local transformations, then we can discard both the extensive and the finite subextensive terms in the entropy, leaving only the divergent subextensive terms as a possible measure of complexity.

To illustrate the purpose of these two extra conditions, we may think of the following example: measuring velocity of a turbulent fluid flow at a given point. The condition of invariance under reparameterizations means that the complexity is independent of the scale used by the speedometer. On the other hand, the second condition ensures that the temporal mixing due to the finiteness of the inertia of the speedometer's needle does not change the estimated complexity of the flow.

In our view, these arguments (or their slight variation also presented in Ref. [6]) settle the question of the unique definition of complexity. Not only is the divergent subextensive component of the entropy the unique complexity measure, but it is also a universal one since it is connected in a straightforward way to many other measures that have arisen in statistics and in dynamical systems theory. A bit less straightforward is the connection to the Kolmogorov's definition that started the whole discussion, but even this can also be made.

To make this connection, we follow the suggestion of Standish [15] that one should focus not on the complexity of particular strings but of equivalence classes. In the present case it is natural to define an equivalence class of data $x(-T < t \leq 0)$ as those data that generate indistinguishable conditional probability distributions for the future, $P[x(t > 0) | x(-T < t \leq 0)]$. If this conditional distribution has sufficient statistics, then there exists a compression of the past data $x(-T < t \leq 0)$ into exactly $I_{\text{pred}}(T)$ bits while preserving all of the mutual information with the future. But this means that

³ Throughout this discussion we assume that the signal x at one point in time is finite dimensional. There are subtleties if we allow x to represent the configuration of a spatially infinite system.

the ensemble of data in an equivalence class can be described, on average, using exactly this many bits. Thus, for dynamics such that the prediction problem has sufficient statistics, the average Kolmogorov complexity of equivalence classes defined by the indistinguishability of predictions is equal to the predictive information. By the arguments above, prediction is *the* useful thing which we can do with a data stream, and so in this case it makes sense to say that the Kolmogorov complexity of representing the useful bits of data is equal to the predictive information. Note also that Kolmogorov complexity is defined only up to a constant depending on the computer used [1]. A computer independent definition requires ignoring constant terms and focusing only on asymptotic behavior. This agrees very well with our arguments above that identified only the divergent part of the predictive information with the complexity of a data stream.

In the terminology suggested by Grassberger, the statement that the prediction problem has sufficient statistics means that the *True Measure Complexity* is equal to the *Effective Measure Complexity* [5]; similarly, the *statistical complexity* defined by Crutchfield and coworkers [16] then also is equal to predictive information defined here. These are strong statements, and it is likely that they are not true precisely for most natural data streams. More generally one can ask for compressions that preserve the maximum fraction of the relevant (in this case, predictive) information, and our intuitive notion of data being “understandable” or “summarizable” is that these selective compressions can be very efficient [17]—here efficiency means that we can compress the past into a description with length not much larger than $I_{\text{pred}}(T)$ while preserving a finite fraction of the (diverging) information about the future; an example is when we summarize data by the parameters of the model that describes the underlying stochastic process. The opposite situation is illustrated by certain cryptographic codes, where the relevant information is accessible (at best) only from the entire data set. Thus, we can classify data streams by their predictive information, but additionally by whether this predictive information can be represented efficiently. For those data where efficient representation is possible, the predictive information and the mean Kolmogorov complexity of future-equivalent classes will be similar; with more care we can guarantee that these quantities are proportional as $T \rightarrow \infty$. Perhaps Wigner’s famous remarks about the unreasonable effectiveness of mathematics in the natural sciences could be rephrased as the conjecture that the data streams occurring in nature—although often complex as measured by their predictive information—nonetheless belong to this efficiently representable class.

References

- [1] M. Li, P. Vitányi, *An Introduction to Kolmogorov Complexity and its Applications*, Springer, New York, 1993.
- [2] J. Rissanen, *Stochastic Complexity and Statistical Inquiry*, World Scientific, Singapore, 1989.
- [3] J. Rissanen, *IEEE Trans. Inf. Theory* 42 (1996) 40–47.
- [4] C. Bennett, in: W.H. Zurek (Ed.), *Complexity, Entropy and the Physics of Information*, Addison-Wesley, Redwood City, 1990, pp. 137–148.

- [5] P. Grassberger, *Int. J. Theor. Phys.* 25 (1986) 907–938.
- [6] W. Bialek, I. Nemenman, N. Tishby, *Neural Comput.* 13 (2001) 2409–2463. E-print: [arXiv:physics/0007070](#).
- [7] V. Vapnik, *Statistical Learning Theory*, Wiley, New York, 1998.
- [8] V. Balasubramanian, *Neural Comput.* 9 (1997) 349–368.
- [9] W. Bialek, C. Callan, S. Strong, *Phys. Rev. Lett.* 77 (1996) 4693–4697.
- [10] C.E. Shannon, *Bell Sys. Tech. J.* 30 (1951) 50–64.
- [11] W. Hilberg, *Frequenz* 44 (1990) 243–248.
- [12] W. Ebeling, T. Pöschel, *Europhys. Lett.* 26 (1994) 241–246.
- [13] T. Schurmann, P. Grassberger, *Chaos* 6 (1996) 414–427.
- [14] C.E. Shannon, *Bell Sys. Tech. J.* 27 (1948) 379–423, 623–656.
- [15] R.K. Standish, *Complexity Int.*, submitted for publication. E-print: [arXiv:nlin.AO/0101006](#).
- [16] C.R. Shalizi, J.P. Crutchfield, *J. Stat. Phys.* (2001) in press. E-print: [arXiv:cond-mat/9907176](#).
- [17] N. Tishby, F. Pereira, W. Bialek, in: B. Hajek, R.S. Sreenivas (Eds.), *Proceedings of the 37th Annual Allerton Conference on Communication, Control and Computing*, University of Illinois, 1999, pp. 368–377. E-print: [arXiv:physics/0004057](#).