



Real patterns and indispensability

Abel Suñé¹ · Manolo Martínez¹

Received: 1 March 2019 / Accepted: 24 July 2019 / Published online: 26 October 2019
© Springer Nature B.V. 2019

Abstract

While scientific inquiry crucially relies on the extraction of patterns from data, we still have a far from perfect understanding of the metaphysics of patterns—and, in particular, of what makes a pattern real. In this paper we derive a criterion of real-patternhood from the notion of conditional Kolmogorov complexity. The resulting account belongs to the philosophical tradition, initiated by Dennett (*J Philos* 88(1):27–51, 1991), that links real-patternhood to data compressibility, but is simpler and formally more perspicuous than other proposals previously defended in the literature. It also successfully enforces a non-redundancy principle, suggested by Ladyman and Ross (*Every thing must go: metaphysics naturalized*, Oxford University Press, Oxford, 2007), that aims to exclude from real-patternhood those patterns that can be ignored without loss of information about the target dataset, and which their own account fails to enforce.

Keywords Kolmogorov complexity · Real patterns · Structure functions · Algorithmic information theory · Metaphysics of science

1 Introduction

Scientific inquiry often depends on the extraction of patterns from data. The huge datasets and corpora typical of many contemporary scientific projects have only made this dependence more obvious and central. Genomics (Barrett et al. 2013), connectomics (Sporns et al. 2005) and astronomy (Feigelson and Babu 2012) are obvious examples, but the trend is quite general. Philosophers of science have been paying increasing attention to patterns, so as to keep up with this trend, but also in the hope

Abel Suñé and Manolo Martínez have contributed equally to this work.

✉ Manolo Martínez
manolomartinez@ub.edu
Abel Suñé
martinezsune@ub.edu

¹ LOGOS Research Group in Analytic Philosophy, Barcelona Institute of Analytic Philosophy, Department of Philosophy, Universitat de Barcelona, Montalegre, 6, 08001 Barcelona, Catalonia

that these entities will fruitfully supplement other entities, such as natural kinds, in the metaphysics of science (Andersen 2017; Ladyman and Ross 2007; McAllister 2003a, b, 2011, 2015, 2017; Petersen 2013).

On the other hand, as is often the case when a new theoretical tool starts to gain prominence, we still have a far from perfect understanding of the notion of a pattern itself. In this paper we propose a way to spell this out that builds upon Ladyman and Ross's theory of *real patterns* (RP henceforth), the most sophisticated account currently on offer. RP substantially extends and refines the idea (first proposed by Dennett (1991) and prefigured, more or less explicitly, by Bogen and Woodward (1988), and Rissanen (1998), among many others) that there are patterns in a dataset D insofar as one can describe a computer program that outputs D while being shorter than D—a program, that is, that *compresses* D. The underlying insight is that patterns correspond to redundancies in the dataset, and it is these redundancies that are exploited by the algorithm implemented by the program.

A further question is what, precisely, is a pattern. The “underlying insight” just mentioned does not take a stand on this: we are invited to conclude that there are patterns present in D if D is compressible, but we have been given no guidance, for any entity P, as to whether it is warranted to claim that P is a pattern in D. This is the question that Don Ross and James Ladyman take up in a series of papers culminating with their 2007 book, and which results in their RP account of real patterns. While RP represents substantial progress towards the development of a metaphysics of patterns, somewhat surprisingly its concrete formulation has not been closely scrutinized in the literature—perhaps, we speculate, because it is built with a computer-theoretic toolbox that is comparatively alien to many metaphysicians of science.

Close scrutiny shows that RP is not without problems: one of the goals of this theory is to provide a criterion of *non-redundancy* (also, interchangeably, *indispensability*) for patterns—the idea, roughly, being that one should count as real all and only those patterns such that ignoring them results in an incomplete description of the world. We will show that RP does not succeed in providing such a criterion. We will also offer an alternative definition of a real pattern, simpler and more formally perspicuous than RP, that provides a workable criterion of indispensability.

Before getting on to that, though, it is important to clarify the scope of our discussion in this paper. In the literature we are engaging with, “real” is taken with a grain of salt: e.g. Dennett refuses to answer the question of whether his account of real patterns is “a sort of instrumentalism or a sort of realism” (1991, p. 51), and Ladyman and Ross (also L&R henceforth) explicitly endorse “a version of instrumentalism about all propositions referring to self-subsistent individual objects, chairs and electrons alike” but also “realism about the domain of scientific description” (2007, p. 198). For our current purposes, this is the sort of realism that matters. We are currently interested in developing a coherent notion of a real pattern, in the comparatively soft understanding of realism typical of the literature on real patterns. That is, we are after a principled way to ground the claim that some patterns are, but some other patterns are not, present in a certain dataset. Whether the criterion for real-patternhood to be developed here is also able to underpin stronger versions of scientific realism (having to do, for example, with the existence of certain objects or kinds of objects) is an extremely interesting topic that will have to be taken up on another occasion.

A point of terminology: L&R use “pattern” to refer to arbitrary strings of symbols, independently of whether they can be used to compress a dataset or not (see Fn. 6). We will use “string” for this purpose. In this paper we reserve “pattern” for the Dennettian notion of an aspect of a dataset that makes it compressible. Finally, we will follow L&R in using “real pattern” to refer to non-redundant (Dennettian) patterns, where non-redundancy is understood as above.¹

In Sect. 2 we introduce Dennett’s original insight, and the computer-theoretic notions on which it builds. In Sect. 3 we explain why current approaches to model selection in algorithmic information theory don’t tell the whole story about patterns. In Sect. 4 we summarize Ladyman and Ross’s RP account. Then, in Sect. 5, we present an important shortcoming of RP: we describe a model that shows that it does not abide by their own indispensability principle—sketched above, and presented more fully in Sect. 4. Finally, in Sect. 6 we propose a better definition of a real pattern, based on the notion of conditional Kolmogorov complexity, which successfully incorporates an indispensability principle. Section 7 offers some concluding remarks.

2 Algorithmic complexity and patterns

Philosophical inquiry into the role and nature of patterns in science kicks off with Daniel Dennett’s seminal “Real Patterns” (1991). Dennett’s main insight is that patternhood is linked to *algorithmic compression*. This is just the compression that contemporary everyday life has since made us familiar with: FLAC sound files or TIFF image files are *compressed*, in the sense that these files are shorter than the uncompressed WAV or bitmap originals. Speech transmitted wirelessly through cell phone networks is also heavily compressed (Rappaport and Theodore 1996). What all instances of algorithmic compression have in common, regardless of their medium or the object of compression, is the fact that a target object is faithfully reproduced by something shorter than a full bit-by-bit description.²

For a more explicit example, consider the following two objects: the first object is a list of results from a series of one million tosses of a fair coin (1 encodes heads, and 0 tails): 0100011011...; the second object is the following list of one million binary digits: 010101...010101. Since we are dealing with a fair coin, the first object is a random string of 0s and 1s. The second string, however, involves an obvious pattern: it’s a repetition of “01” half a million times. In order to describe the first string there are no substantial shortcuts to writing, digit by digit, the results of each coin toss. On the other hand, the second string can be fully reproduced via an algorithmic description that is much shorter than the full string: (roughly) “print ‘01’ half a million times”.

¹ The above informal characterization of non-redundancy will be sharpened in Sect. 4.

² Speech compression is an instance of *lossy* compression, where faithfulness of compression is judged by a certain distortion measure, or loss function (Cover and Thomas 2006, ch. 10; Shannon 1959). The main notion of compression we rely on in what follows, on the other hand, is *lossless* compression, in which the original file and the decoded version thereof are identical. TIFF, FLAC and others such as DEFLATE, typically used in zip files, are widely popular lossless algorithms. We note that the very existence of lossless compression algorithms appears to be in some tension with McAllister’s (2003a) claim that empirical datasets are incompressible—insofar, e.g., as empirical datasets can contain photographs or audio recordings. We won’t pursue this topic here.

The first, *random* string is incompressible. The second, *patterned* one incorporates redundancies that can be exploited by a compression algorithm.

Chaitin (1966), Kolmogorov (1965) and Solomonoff (1964a, b), the founders of algorithmic information theory, suggested independently that a sequence should be considered random if and only if it is incompressible. Dennett conjoined this thesis with the idea that a string has patterns to the extent that it is not purely random: “a [string] is not random—has a pattern—if and only if there is some more efficient way of describing it” (1991, p. 32). In a nutshell, *a pattern in a dataset is any aspect of the dataset that allows it to be compressed.*

As we explained above, compressibility in a string is to be thought of as affording the existence of comparatively short programs that can output the string in question.³ For some strings there will be multiple programs of different lengths that can output them, but we will be interested only in the shortest such program. The *maximum achievable algorithmic compression*, or *Kolmogorov Complexity*, $K(S)$, of a given string, S , is the *length of the shortest program that outputs S* . The Kolmogorov complexity of strings will be frequently appealed to in what follows.

A related notion is the *complexity of a string, S , conditional on another string, T , or conditional Kolmogorov complexity* $K(S|T)$. This is the length of the shortest program that outputs S , when it is allowed to use T as an input. Suppose, for example, that we have to compress a string, S , which encodes a recording of a spoken conversation. The length of the shortest program that outputs S (i.e. the program that represents the best possible compression of S) gives us $K(S)$. But suppose further that we have a computer library, T , that encodes certain speech statistics typical of recordings such as S . This will often mean that we can write a shorter program that prints S , if we use T as an additional input to the program. The shortest such program gives us $K(S|T)$.

Finally, one can quantify the amount of information an object x carries about another object y —the *mutual information* between x and y —as a measure of the reduction in the amount of descriptive effort one has to make to describe y after coming to know x . More formally, if S and T are strings, the algorithmic mutual information between them, $I(S:T)$, is:

$$I(S:T) = K(S) - K(S|T). \text{ (Li and Vitányi 2008, definition 3.9.2)}$$

3 Model selection is not pattern individuation

We claimed above that Dennett’s foundational insight was that a pattern in some dataset is any aspect of it that allows it to be compressed. “Any aspect” is, of course, rather unspecific. A full theory of patterns, if it is to help us to recognize and individuate patterns, needs to spell out in more detail what these compression-enabling aspects of

³ In formal presentations of Kolmogorov complexity (e.g. Li and Vitányi 2008, p. 107), the programs we have been alluding to are inputs to a reference universal Turing machine (UTM). For the purposes of this paper, we can just think of the reference UTM as implementing one of the very many popular Turing-complete programming languages—say, Python, or Javascript. Petersen (2018, p. 2) discusses whether the choice of UTM introduces a bias in the resulting account of patterns (for example, by making any arbitrary dataset, however big and random, compressible and hence patterned) and concludes, with Li and Vitányi (2008, p. 112), that a small enough UTM will make any such potential bias negligible.




Fig. 1 Denoising based on structure functions (from Vereshchagin and Vitányi 2010, p. 3446)

datasets amount to. In the following section we discuss Ladyman and Ross’s attempt to provide this detail (2007; Ross 2000). Beforehand, in this section, let us briefly discuss an approach in algorithmic information theory to a related question, and explain why this approach does not deliver an account of patterns, at least not as metaphysicians of science employ this notion.

Consider a noisy image; e.g. the right cross in Fig. 1 (reproduced from Vereshchagin and Vitányi 2010, p. 3446). We can intuitively analyze this object into two components: first “the information accounting for the useful regularity present in the object” (Vitányi 2006, p. 4617), which captures the noiseless cross to its left; and a meaningless one, “the information accounting for the remaining accidental information” (*ibid.*), which captures the noise.

The way algorithmic information theory approaches this analysis (following suggestions made by Kolmogorov in the early 70s and in his 1965) is by devising ways to encode the object in which the relevant code has two parts. The first part captures the meaningful information of the object (i.e. provides a *model* of the object); the second part captures the noisy remainder. In the foundational “structure functions” version of this idea (Vereshchagin and Vitányi 2006, 2010), this two-part code is implemented as follows. To encode a string x , we first identify a set S such that x is a *typical* member of S . Being a typical member of a set simply means that, in order to pick out x from the other members of S , there’s no shorter procedure than giving x ’s position in an arbitrary enumeration of S ’s members. A two-part code, then, can be constructed that first reconstructs S (this takes $K(S)$ bits, by the definition of Kolmogorov complexity), and then gives x ’s position in S (this takes $\log|S|$ bits, where $|S|$ is S ’s size). S corresponds to the best model of the data; x ’s position in S corresponds to the noisy remainder.

It can be shown (Vereshchagin and Vitányi 2006, p. 3269) that for each string x there are *optimal sets* for which the sum $K(S_{\text{optimal}}) + \log|S_{\text{optimal}}|$ is equal to $K(x)$. That is, somewhat surprisingly, the $K(S_{\text{optimal}})$ term (which captures the meaningful information in x by identifying the simplest set in which x is a typical member) together with the $\log|S_{\text{optimal}}|$ term (which captures the noisy remainder by picking out x from

a brute enumeration) add up to the Kolmogorov complexity of the original string (up to an additive constant). The best meaning-plus-noise code for x is as good as the best possible code for x .⁴

While algorithmic model selection is designed to help us distinguish meaningful from meaningless in a dataset, there are at least two respects in which it does not provide a solution to the problem real-pattern theorists are interested in. First, the two-part code idea aims at reconstructing the original string x *in its entirety*. In our case, x would correspond to an empirical dataset, and applying the procedure just sketched would leave us with a specification of *all meaningful regularities* in the dataset together with the remaining noise. But almost all patterns identified in actual scientific practice correspond to *partial* regularities in the target dataset. For one example among very many, take *CpG islands*: areas of DNA with a high concentration of the CpG dinucleotide. The abundance of CpG in a certain stretch of DNA is a clear pattern, widely studied in epigenetics (Bird 1986; Larsen, Gundersen, Lopez, and Prydz 1992). Yet, of course, full knowledge of where CpG islands are, on its own, does not allow us to reconstruct a full genome. In general, patterns in a dataset illuminate important aspects of it, without fully describing it. Model selection in the algorithmic information-theoretic tradition, as described above, offers no guidance as to how to uncover or describe patterns in this sense.

The first problem we have mentioned with the model selection approach is that it offers a way to identify *all* meaningful information in a dataset, in one go, but not a way to identify partial, incomplete portions of this information, which is what patterns are.⁵ The second respect in which model selection is not the right tool for pattern discovery stems from the fact that many of the patterns identified by scientists are partially constituted by what, arguably, is noise by the lights of model selection. Again, for one example among many, the DNA of the *C. elegans* worm contains non-coding regions of 10-base-pair periodic adenine (A)/thymine (T)-clusters (or PATCs). These are regions of non-coding DNA in which stretches of a few consecutive As and stretches of a few consecutive Ts happen more or less every ten bases, and they make up approximately 10% of the *C. elegans* genome (Frøkjær-Jensen et al. 2016). PATCs have been described as important patterns in the *C. elegans* genome, as they appear to play a role in allowing germline expression of transgenes, in regions the expression of which is otherwise silenced (*ibid.*). The fact that the clusters in PATCs happen every 10 bases (and not every 20 or 40), and that the bases involved are A and T (and not C or G) is likely to be random happenstance to a certain degree—that is, noise, in model-selection parlance. But it is PATCs themselves, their noisy ingredients included, and not just the “meaningful” core identified by structure functions, that are relevant to *C. elegans* genetics.

⁴ The foregoing few paragraphs only scratch the surface of the algorithmic approach to model selection. This is the aim of so-called *algorithmic statistics*. We point the interested reader to Gács et al. (2001), Vereshchagin and Shen (2017) and references therein for in-depth discussion and alternatives to the structure–function two-part code.

⁵ To be clear: this is a problem insofar as we want to use model selection as a method for identifying patterns. Model selection is a perfectly clear goal in algorithmic statistics, and the structure–function approach has much to recommend it, when used for its intended purpose.

An account that accommodates partial patterns (such as CpG islands), and noise-including patterns (such as PATCs) is thus in order. Before presenting our own, we turn now to describing the most developed, if ultimately unsuccessful, account of this sort.

4 Ladyman and Ross’s real patterns theory⁶

L&R’s main idea is that the “aspects” of datasets that enable compression can be captured by identifying strings that partially encode the original dataset—these strings will be the patterns in the dataset. We will soon be more precise than this; but we can already note that there can be many different strings that partially encode a dataset—indeed, there can be sets of *mutually redundant* strings, in the sense that each of them informs us of the very same aspects of the target dataset.

L&R contend that the right theory of *real* patterns should provide guidance in the process of choosing which members to recognize as real patterns, from those possible sets of mutually redundant strings. That is, a theory of patterns must help us to distinguish between potentially useful but ultimately dispensable patterns, which one can ignore without any ontological loss, and patterns such that ignoring them results in an incomplete description of the target dataset (cf. Ladyman and Ross 2007, p. 231):

Non-redundancy principle: Include in your ontology all and only those patterns that are *required* for a full (lossless) reconstruction of the target dataset.

L&R define real patterns as follows:

A pattern⁷ P is real iff

- (i) it is projectible; and
- (ii) it has a model that carries information about at least one pattern D in an encoding that has logical depth less than the bit-map encoding of D , and where D is not projectible by a physically possible device computing information about another real pattern of lower logical depth than P (Ladyman and Ross 2007, p. 233)⁸

This definition uses a number of comparatively uncommon technical terms. Our first aim in this section will be to present a version of L&R’s theory of real patterns that captures the main gist of their original definition, but is both simpler and more

⁶ We are greatly indebted to an anonymous referee for some very generous and detailed input that has much improved this section.

⁷ Keep in mind that patterns *simpliciter* for L&R are just what we have called “strings”. This is what Fn. 51 in Ladyman and Ross (2007) amounts to saying: “A mere pattern is a locatable address associated with no projectible or non-redundant object” (*ibid.*, p. 231). See also *ibid.*, p. 229: “From the ontological point of view, a non-projectible pattern exactly resembles the traditional philosophical individual.” For the meaning of “locator” and cognates in L&R’s system, see *ibid.*, p. 121ff. For the related notion of “perspective” see *ibid.*, p. 224.

⁸ We have changed the variable names to align them with the ones we use in this paper.

continuous with the rest of the literature on patterns. We will now discuss each of the technical terms in turn.

4.1 Projectibility

Projectibility is used twice in the definition. L&R say that an *entity* x *projects an entity* y (notated $x \rightarrow y$) iff it is possible to calculate y from x (*ibid.*, p. 224). Take, for example, a system of physical bodies moving about in space. We might want to predict some future positions and velocities of one particular object, given data about the current positions and velocities of a range of objects. Let x be a specification of the position and velocity of those objects at time t ; and let y be a specification of the position and velocity of our target object at a later point in time. If it is possible to effect a computation $x \rightarrow y$ which outputs y from input x (by, say, solving differential equations that correspond to some deterministic theory of gravity), we say that x projects y . If an object x projects an object y , then x doesn't merely carry some amount of information about y : it carries *all the information necessary to specify y without residue*. In algorithmic complexity terms: for any strings x and y , x projects y iff $K(y|x) = 0$, up to an additive constant independent of x and y .

According to L&R, a pattern (e.g. a differential equation, as in the example above) that makes possible a projection, $x \rightarrow y$, in the sense explained, is *projectible*—as in clause (i) above—if it affords projecting y s for multiple unobserved input x s. As L&R point out, this condition aims to “avoid trivialization of projectibility by reference to [a computer] that simply implements the one-step inference rule ‘Given input [x], output [y]’” (*ibid.*). Yet it could be argued that the appeal to multiple unobserved inputs doesn't avoid trivialization: projecting multiple y s for multiple x s is computationally equivalent to projecting a single, bigger y (say, an array of the original y s) from a single, bigger x (an array of the original x s). Again here, there is a Turing machine that calculates the bigger y from the bigger x by (paraphrasing L&R) “simply implementing the one-step inference rule ‘Given input [bigger x], output [bigger y]’”.

In our paper we are, in effect, taking “ y is projectible from x ” to mean “a trivial (very short, etc.) universal Turing machine running x as its program outputs y ”. This seems to capture what L&R want projectibility to do, while avoiding triviality. In a Kolmogorov complexity setting these constraints on universal Turing machines are standard (see Fn. 2).

4.2 Models

L&R's definition also makes reference to *models* of patterns. This has two functions in L&R's construction: first, it helps make explicit that, in order to apply algorithmic complexity theory to real-world phenomena, we need to translate them into strings. Second, it makes the definition applicable to cases in which we “may have access only to a model of the pattern in question” (e.g. when the pattern is the interior of the Sun, *ibid.* p. 233). The counterexample to RP that we will be considering in the sequel is formulated directly in terms of strings. This will allow us to sidestep these complications.

4.3 Encodings

As we saw above, there are patterns in datasets iff the latter can be compressed. The uncompressed, raw version of a dataset (for example, the raw lists of numbers and text that result from research) is what L&R call its *bit-map encoding* (see *ibid.*, p. 232). This is the string we are calling “D” in this paper. Other, more sophisticated encodings of the dataset might provide more economical (shorter) representations of it, if it is compressible.

4.4 Logical depth

Finally, the logical depth of a string is the number of steps necessary to output it from its minimal program (Bennett 1995). John Collier, co-author with Ladyman and Ross of the main chapter on real patterns in Ladyman and Ross (2007), has given reasons to opt for logical depth when discussing physical complexity in (Collier 2001; Collier and Hooker 1999). In our context, this is a relatively idiosyncratic choice: it is universally understood that the main measure of algorithmic complexity is Kolmogorov complexity (Grünwald and Vitányi 2008) and this is also the notion used in most other papers in the philosophical literature on real patterns (Dennett 1991; McAllister 2003b; Petersen 2013, 2018). We will opt for a Kolmogorov-complexity version of L&R’s definition of real patterns. In any event, the examples we will use to make our points in this paper are, in Bennett’s sense, *logically shallow*. That is, for them, Kolmogorov complexity equals logical depth (and, indeed, mostly equals length). Still, if one wishes to stick to logical depth, this is perfectly doable. Where we say “Kolmogorov complexity” one should read “logical depth”, and where we say “conditional Kolmogorov complexity” one should read “relative logical depth” (Bennett 1995, Definition 1.1).

With these clarifications in mind, we can now present a succinct version of L&R’s definition:


RP: A pattern P is a real iff there is a dataset D such that

- (i) $I(D:P) > 0$; and
- (ii) there is no Q such that $K(D|Q) = 0$ and $K(Q) < K(P)$.

The definition, first, says that for a string P to be a real pattern it has to carry information about some dataset D—which in turn means that a program that has P as input and outputs D can be shorter than a program that outputs D and has no inputs. This captures the Dennettian compressibility condition on patternhood. Second, the definition says that if P is a real pattern, then no pattern shorter than P has all the information needed to specify D without residue. This attempts to enforce the non-redundancy principle above: if there were a pattern that losslessly compressed D as a whole, while being shorter than P, then P would indeed be superfluous for the purpose of describing D.

There is an asymmetry in how this definition handles the Dennettian and the non-redundancy ingredients: on the one hand it is deemed enough that P *carry some, not all, information* about D; yet, on the other hand, redundancy is averted by appealing to *projectibility*—which, recall, requires a pattern to be *fully* informative about D, if it is to make P redundant. This combination is an understandable theoretical choice.

Fig. 2 Constructing D



First, non-redundancy cannot be implemented in terms of information-carrying. If the second part of the definition merely read

(ii)* there is no Q such that $K(D|Q) < K(D)$ and $K(Q) < K(P)$

then the unreasonable consequence would be that patterns less complex than P , that carried *less* information about D than P , would make P redundant. Second, and conversely, for the sort of reason seen in the previous section, the Dennettian insight cannot be implemented in terms of projection: it is unreasonable to rule that the only real patterns in a dataset are those patterns that *fully* reconstruct the dataset (recall CpG islands).

This combination is an understandable theoretical choice, but it is the wrong one. As we will show in the following section, the interplay between projectibility and information-carrying in clause (ii) of RP doesn't quite behave as intended, and this can be used to smuggle in very noisy patterns, and have them counted as real.

5 The problem

In RP, putatively real patterns are tested by how well they explain a dataset D . We will now present a toy model in which D is a string constructed as follows. We first take three 200 MB⁹ random binary strings, A_1 , A_2 , and A_3 . The concatenation of these three strings, we call A . Second, we construct D as the concatenation of two exact copies of A . The size of D , thus, is $(200 \times 3) \times 2 = 1200$ MB (see Fig. 2). Alongside

⁹ A megabyte (MB) is one million bytes.

D, we construct another pattern, Q, as the concatenation of A_1 and an extra 200 MB of random binary digits—i.e. A_1 plus a lot of superfluous noise.¹⁰

Now, A is a real pattern in D: it is repeated twice, verbatim, in the dataset. This is borne out by RP: first, A carries a great deal of information about D; and second, D is not projectible by any program computing information from a pattern that is shorter than A.

To see that the first condition is met (that A carries information about D) we simply show that $K(D|A) \ll K(D)$, which, as per the definition of mutual information, entails that $I(D:A) \gg 0$. D consists of a repeated 600 MB random (hence, incompressible) string, so $K(D)$ must be larger (but not much larger) than 600 MB. Now, on the other hand, if we are to write a program that outputs D, and we can use A as an extra input (as per the definition of conditional Kolmogorov complexity, presented above) the program can be as simple as concatenating two copies of A. Listing 1 is a function that does exactly that.¹¹

```
reconstruct_D(A) :
    return concatenate(A, A)
```

Listing 1

This program has a size of approximately 54 bytes.¹² This, then, is an upper bound for $K(D|A)$, and it is much smaller than $K(D) \sim 600$ MB. A, as expected, is highly informative about D: $I(D:A) = K(D) - K(D|A) \sim 600$ MB.

The second condition for A to be a real pattern is that D is not projectible by any program shorter than A—i.e. shorter than 600 MB—and, as we have already seen, it is not: A is a random, and therefore incompressible, component of D. It is impossible for any program shorter than A to produce a copy of A and, *a fortiori*, impossible for any program shorter than A to produce a copy of D.

Now, it is also intuitively compelling that long substrings in A are real patterns in the D dataset. Suppose, for example, that D-researchers haven't yet found out about the full pattern A, but have discovered that A_1 is repeated twice in D. Uncovering this fact about D would be very valuable for making sense of its structure—plausibly, the situation in genomics is relevantly analogous to this toy example. Again here, RP agrees with this intuitive assessment: A_1 is a real pattern in the RP sense. The program that reconstructs D given A_1 as an extra input has to concatenate it with A_2 and

¹⁰ For ease of discussion, we have designed our example so that redundancies in the dataset are readily apparent. Of course, in more realistic examples, sophisticated coding might be needed in order to squeeze the redundant material out of our target string. The kind of argument we develop here applies to more realistic cases as well.

¹¹ In what follows we will give our example algorithms in pseudocode—i.e. a dialect that does not correspond to any particular programming language, but can be readily translated to many of them and is tailored to maximize readability for humans. It may have occurred to some readers that the choice of coding scheme or programming language used to describe an object can condition the minimum achievable length for describing it and thus, seemingly, that different languages will introduce different Ks for the same object. This is correct, but the apparent relativity thus introduced does not affect the objectivity of K as a measure of compressibility: K is equal for every programming language up to an additive constant that is independent of the string to be compressed itself (Grünwald 2007, p. 10).

¹² At least in our laptop. The exact number will vary slightly from platform to platform.

A_3 , concatenate the resulting string, A , with itself, and then output it. A pseudocode function that does this is given by Listing 2.¹³

```
reconstruct_D_from_A1(A1):
    A2 ← "10011001..." # The whole 200MB A2 string would be here
    A3 ← "11110010..." # And again for A3
    A ← concatenate(A1, A2, A3) # concatenates all three
    substrings...
    return concatenate(A, A) # ... and everything follows as
before
```

Listing 2

This program will be just over 400 MB long, with the definitions of substrings A_2 and A_3 , each 200 MB long, taking by far the biggest chunk of this space. 400 MB is therefore a good estimate for $K(D|A_1)$: about two-thirds of $K(D)$. Again here, 400 MB is not enough to reconstruct D , which needs at least the 600 MB of A , so A_1 counts as a real pattern in D according to RP.

The problem comes with the Q pattern. This, recall, is just A_1 with an extra 200 MB of noise, so it's very clearly redundant relative to A_1 for the purpose of describing D . Yet, it too passes muster as a real pattern according to RP. First, it carries as much information about D as A_1 does. The program that reconstructs D from Q will simply keep the first 200 MB of Q , that is, A_1 , and then proceed as before:

```
reconstruct_D_from_Q(Q):
    A1 ← first_half(Q) # A1 is the first half of Q
    A2 ← "10011001..." # The whole 200MB A2 string would be here
    A3 ← "11110010..." # And again for A3
    A ← concatenate(A1, A2, A3) # ... and everything follows as
                                # before
    return concatenate(A, A)
```

Listing 3

This program has a very similar length to the one in Listing 2. That is, $K(D|A_1) \sim K(D|Q)$. And again here, while Q is double the size of A_1 , it is, at 400 MB, still much smaller than $K(D)$ —and, as we have explained, no string shorter than $K(D)$ will be able to reproduce (i.e. project) D . So Q meets the RP definition of a real pattern.

In general, any string that both carries information about D and is smaller than $K(D)$ will count as a real pattern according to RP. This, as Q shows, is an unwelcome result given the goal expressed by the non-redundancy principle: *it is trivial to multiply the redundant patterns that pass the RP test by taking a real pattern and adding to it any amount of noise such that the total length is less than the Kolmogorov complexity of the original dataset.*

¹³ In fact, if, as we said, scientists have not yet learned that A_2 and A_3 are repeated in D , these two strings should occur *twice* in listing 2, once for each repetition. This detail does not interfere with our point, and we have omitted it so as not to complicate the structure of the example.

6 A solution

How should we fix this? In the model developed in the previous section, the redundant pattern Q is longer than the non-redundant pattern A_1 . A quick—but obviously wrong—fix would be to rule that between any two strings, say L and S , that carry information about a target dataset, D , the longer one, L , should be discarded in favor of the shorter one, S . This would result in Q being discarded in favor of A_1 , as desired. The problem, of course, is that this “quick fix” pays no heed to which *particular chunks* of information the two contending strings carry about D . For example, it could be that L is identical to A_1 and S to the first half of A_2 . In that case, it would make no sense to deem L superfluous simply because it is longer than S , because L and S make complementary contributions to our understanding of D .

What we need, if we are to decide whether one putative real pattern is redundant relative to another, is a way of specifying when two patterns carry, not just the same *amount* of information, but also the same *specific pieces* of information about an object. The general question of how to express this desideratum is an open problem in information theory, explored in the *partial information decomposition* framework (Griffith et al. 2014; Williams and Beer 2010), although, as far as we are aware, it has been discussed less in algorithmic information theory. Fortunately, conditional Kolmogorov complexity allows us to arrive at the relevant notion of pattern redundancy without having to solve the general problem of partial information decomposition. We first define D -dispensability:

D-dispensability: A pattern Q that carries information about a dataset D is D -dispensable iff there is another pattern P such that:

- (i) $K(D|P, Q) = K(D|P)$
- (ii) $K(P) < K(Q)$

Clause (i) captures the idea that the information Q carries about D is *part* of the information P carries about D . This is not a claim about information quantities, but about the actual pieces of information carried by these strings: if having Q as an extra input does not allow us to shave even a few bits off the length of the program that calculates D using P as its input, then this means that all Q has to say about D is also said by P .

Notice that if we only had clause (i), it would be possible for Q to be D -dispensable in favor of P , and for P to be D -dispensable in favor of Q at the same time. In keeping with the non-redundancy principle discussed above, clause (ii) deems dispensable whichever of the contending patterns is more complex.

We will say that a pattern is *strictly D-indispensable* iff it carries information about D and is not D -dispensable. Strict D -indispensability is, we suggest, the notion that Ladyman and Ross were after in their definition of a real pattern. Thus a better definition is as simple as this:

Real Pattern: A pattern P is real iff there is at least one dataset D such that P is strictly D -indispensable.

We offer this definition as a drop-in replacement for RP, only simpler, and without the defects presented in the previous section.

7 Concluding remarks

We have shown that the theoretical desideratum captured by the non-redundancy principle—*include in your ontology all and only those patterns that are required for a full reconstruction of the target dataset*—is not successfully enforced by RP. The interplay between information-carrying and projection in this definition entails the reality of any string that is informative about D while being less complex than D.

We have proposed to define real patterns directly in terms of their contribution to reducing the conditional Kolmogorov complexity of a dataset. This captures the main insights of RP, while being simpler, and offering a better non-redundancy criterion. A further question¹⁴ is whether it is sensible to tie reality to non-redundancy, as L&R and we do. This decision is beyond the scope of this paper: we have provided clear criteria for deeming a string to be a *pattern* in a dataset (this is tied to compressibility), and also for deeming it to be a *real pattern* (this is tied to compressibility and non-redundancy). Ascertaining which, if either, of these two notions should take priority in the metaphysics of science is a matter for further research.

Our definition, like L&R's, relies on the existence of an object (a dataset in our definition, a "pattern" in theirs) which putative real patterns are more or less able to explain. It is probably possible to contrive artificial datasets that would make any desired pattern real, by any definition in the Dennettian tradition. Take any random string P, create a "dataset" D consisting of two concatenated copies of P, and—*hey presto!*—D makes P real. The lesson here is that definitions of real patterns along Dennettian lines, such as L&R's or ours, are only interesting if the datasets appealed to are *bona fide*, in the sense of coming from the world—from actual empirical research. A more explicit characterization of what should count as a *bona fide* dataset is a central question for any broadly Dennettian account of patterns, but one that must be discussed elsewhere.

Acknowledgements We would like to thank James Ladyman for his very generous discussion of the topics of this paper. We would also like to thank the participants of the reading group on Real Patterns held at the University of Barcelona in 2018. Two very detailed reviews from two anonymous referees helped us to significantly improve the paper. Abel Suñé also wishes to thank J.P. Grodniewicz for valuable discussion and Pepa Toribio for her support. Manolo Martínez would like to acknowledge research funding awarded by the Spanish Ministry of Economy, Industry and Competitiveness, in the form of grants PGC2018-101425-B-I00 and RYC-2016-20642.

References

- Andersen, H. K. (2017). Patterns, information, and causation. *The Journal of Philosophy*, 114(11), 592–622.
- Barrett, T., Wilhite, S. E., Ledoux, P., Evangelista, C., Kim, I. F., Tomashevsky, M., et al. (2013). NCBI GEO: Archive for functional genomics data sets—Update. *Nucleic Acids Research*, 41(D1), D991–D995. <https://doi.org/10.1093/nar/gks1193>.

¹⁴ Put to us by an anonymous referee.

- Bennett, C. H. (1988). Logical depth and physical complexity. In R. Herken (Ed.), *The universal Turing machine, a half century survey*, (pp. 227–257). Oxford University Press.
- Bird, A. P. (1986). CpG-rich islands and the function of DNA methylation. *Nature*, 321(6067), 209. <https://doi.org/10.1038/321209a0>.
- Bogen, J., & Woodward, J. (1988). Saving the phenomena. *The Philosophical Review*, 97(3), 303–352.
- Chaitin, G. J. (1966). On the length of programs for computing finite binary sequences. *Journal of the ACM (JACM)*, 13(4), 547–569.
- Collier, J. (2001). Dealing with the unexpected. In *Partial proceedings of CASYS 2000: Fourth international conference on computing anticipatory systems, international journal of computing anticipatory systems* (vol. 10, pp. 21–30).
- Collier, J., & Hooker, C. A. (1999). Complexly organised dynamical systems. *Open Systems and Information Dynamics*, 6(3), 241–302.
- Cover, T., & Thomas, J. (2006). *Elements of information theory (Wiley Series in Telecommunications and Signal Processing)*. New York, NY: Wiley-Interscience.
- Dennett, D. C. (1991). Real patterns. *The Journal of Philosophy*, 88(1), 27–51.
- Feigelson, E. D., & Babu, G. J. (2012). Big data in astronomy. *Significance*, 9, 22–25. <https://doi.org/10.1111/j.1740-9713.2012.00587.x>.
- Frøkjær-Jensen, C., Jain, N., Hansen, L., Davis, M. W., Li, Y., Zhao, D., et al. (2016). An abundant class of non-coding DNA can prevent stochastic gene silencing in the *C. elegans* germline. *Cell*, 166(2), 343–357. <https://doi.org/10.1016/j.cell.2016.05.072>.
- Gács, P., Tromp, J. T., & Vítányi, P. (2001). Algorithmic statistics. *IEEE Transactions on Information Theory*, 47(6), 2443–2463.
- Griffith, V., Chong, E. K., James, R. G., Ellison, C. J., & Crutchfield, J. P. (2014). Intersection information based on common randomness. *Entropy*, 16(4), 1985–2000.
- Grünwald, P. (2007). *The minimum description length principle*. Cambridge: The MIT Press.
- Grünwald, P. D., & Vítányi, P. M. B. (2008). Algorithmic Information Theory. In Adriaans, P. & van Benthem, J. (Eds.), *Philosophy of Information* (Vol. 8, pp. 281–320). Handbook of the Philosophy of Science. Amsterdam: North-Holland.
- Kolmogorov, A. N. (1965). Three approaches to the quantitative definition of information. *Problems of Information Transmission*, 1(1), 1–7.
- Ladyman, J., & Ross, D. (2007). *Every thing must go: Metaphysics naturalized*. Oxford: Oxford University Press.
- Larsen, F., Gundersen, G., Lopez, R., & Prydz, H. (1992). CpG islands as gene markers in the human genome. *Genomics*, 13(4), 1095–1107. [https://doi.org/10.1016/0888-7543\(92\)90024-M](https://doi.org/10.1016/0888-7543(92)90024-M).
- Li, M., & Vítányi, P. (2008). *An introduction to Kolmogorov complexity and its applications. Texts in computer science* (Vol. 9). New York: Springer.
- Martínez, M. (2015). Informationally-connected property clusters, and polymorphism. *Biology and Philosophy*, 30(1), 99–117.
- Martínez, M. (2017). Synergic kinds. *Synthese*. <https://doi.org/10.1007/s11229-017-1480-2>.
- McAllister, J. W. (2003a). Algorithmic randomness in empirical data. *Studies in History and Philosophy of Science Part A*, 34(3), 633–646.
- McAllister, J. W. (2003b). Effective complexity as a measure of information content. *Philosophy of Science*, 70(2), 302–307.
- McAllister, J. W. (2011). What do patterns in empirical data tell us about the structure of the world? *Synthese*, 182(1), 73–87.
- Petersen, S. (2013). Toward an algorithmic metaphysics. In D. Dowe (Ed.), *Algorithmic probability and friends: Bayesian prediction and artificial intelligence* (pp. 306–317). Berlin: Springer.
- Petersen, S. (2018). Composition as pattern. *Philosophical Studies*, 176, 1119. <https://doi.org/10.1007/s11098-018-1050-6>.
- Rappaport, T. S. (1996). *Wireless communications: Principles and practice*. Upper Saddle River, NJ: Prentice Hall PTR.
- Rissanen, J. (1998). *Stochastic complexity in statistical inquiry* (Vol. 15). Singapore: World Scientific.
- Ross, D. (2000). Rainforest realism: A Dennettian theory of existence. In D. Ross, D. Thomson, & A. Brook (Eds.), *Dennett's philosophy: A comprehensive assessment* (pp. 147–168). Cambridge, MA: The MIT Press.
- Shannon, C. E. (1959). Coding theorems for a discrete source with a fidelity criterion. Institute of Radio Engineers, International Convention Record, Part 4, 142–163.

- Solomonoff, R. J. (1964a). A formal theory of inductive inference. Part I. *Information and Control*, 7(1), 1–22.
- Solomonoff, R. J. (1964b). A formal theory of inductive inference. Part II. *Information and Control*, 7(2), 224–254.
- Sporns, O., Tononi, G., & Kötter, R. (2005). The human connectome: A structural description of the human brain. *PLoS Computational Biology*, 1(4), e42. <https://doi.org/10.1371/journal.pcbi.0010042>.
- Vereshchagin, N., & Shen, A. (2017). Algorithmic statistics: Forty years later. In A. Day, M. Fellows, N. Greenberg, B. Khossainov, A. Melnikov, & F. Rosamond (Eds.), *Computability and complexity: Essays dedicated to Rodney G. Downey on the occasion of his 60th birthday. Lecture Notes in Computer Science* (pp. 669–737). Cham: Springer. https://doi.org/10.1007/978-3-319-50062-1_41.
- Vereshchagin, N., & Vitányi, P. (2010). Rate distortion and denoising of individual data using Kolmogorov complexity. *IEEE Transactions on Information Theory*, 56(7), 3438–3454. <https://doi.org/10.1109/TIT.2010.2048491>.
- Vereshchagin, N., & Vitányi, P. (2006). On algorithmic rate-distortion function. In *2006 IEEE international symposium on information theory* (pp. 798–802). IEEE.
- Vitányi, P. M. (2006). Meaningful Information. *IEEE Transactions on Information Theory*, 52(10), 4617–4626.
- Williams, P. L., & Beer, R. D. (2010). *Nonnegative decomposition of multivariate information*. ArXiv Preprint [arXiv:1004.2515](https://arxiv.org/abs/1004.2515).

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.